

## Durham E-Theses

---

# *FPGA Based Diagnostics for the Mega-Amp Spherical Tokamak Upgrade*

CHARLES HOWARD VINCENT

### How to cite:

---

VINCENT, CHARLES HOWARD (2021) FPGA Based Diagnostics for the Mega-Amp Spherical Tokamak Upgrade. Doctoral thesis, Durham University.

### Use policy

---

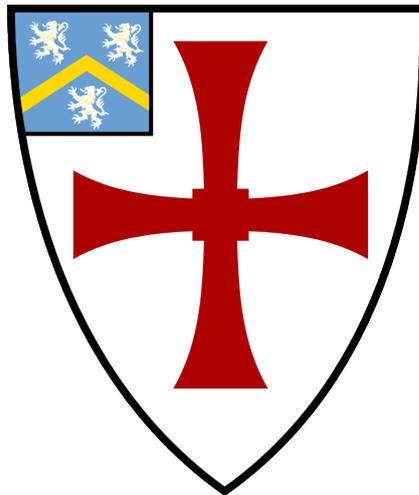


This work is licensed under a [Creative Commons Public Domain Dedication 1.0 \(CC0\)](https://creativecommons.org/licenses/by/4.0/)

# FPGA Based Diagnostics for the Mega-Amp Spherical Tokamak Upgrade

Charles Vincent

A thesis presented for the degree of  
Doctor of Philosophy



Centre for Advanced Instrumentation  
The University of Durham  
United Kingdom  
January 2021

# FPGA Based Diagnostics for the Mega-Amp Spherical Tokamak Upgrade

Charles Vincent

## Abstract

Terrestrial fusion power is a low carbon alternative to conventional power sources with reduced waste and proliferation concerns relative to fission power. The complexity of fusion research devices means that many high performance diagnostics are necessary to investigate the underlying physics of the environment. Field Programmable Gate Array technology provides a powerful and flexible option when designing bespoke instrumentation.

A low-cost FPGA solution has been implemented for the Fission Chambers for Neutron Flux Measurement data acquisition and real-time processing system. This diagnostic improves the time resolution and utility compared to the old system, with the use of FPGA technology facilitating maintainability and future customisation. Results from tests at the National Physical Laboratory show that the DC current and Campbell mode respond linearly to different neutron fluxes.

Second, a collaboration with the Plasma Science and Fusion Centre is shown in developing a proof-of-concept real-time biasing Langmuir Probe control system. This uses the same FPGA hardware and workflow as the Fission Chamber acquisition system to utilise previous experience for a short development time. The system is demonstrated to be capable of up to a 1 MHz temporal resolution of electron temperature, an order of magnitude faster than current Langmuir probe control systems currently employed on MAST-Upgrade. Open source development software and low-cost components will help facilitate application to many different plasma devices.

Finally, work to enable real-time data acquisition for the second generation Synthetic Aperture Microwave Imaging system is presented. A high bandwidth 2-D Doppler Backscattering diagnostic, the second generation SAMI instrument will utilise 30 receiving polarisation separation antennas compared to the original SAMI instrument's 8 mono-polarisation antennas. Consequently, the number of channels to be digitised has increased from 16 to 250, with accompanying increases in the data bandwidth requirements. FPGA technology has been utilised to enable data acquisition rates at up to 80 gigabits/s with potential for up to 200 gigabits/s in the future. A method for multiple FPGA board synchronisation has also been presented using the Precision Time Protocol.

Supervisors: Ray Sharples and Christopher Saunter

---

# Acknowledgements

Many thanks must go to my supervisors Ray Sharples, Roddy Vann, Graham Naylor, and Stephanie Hall who have provided invaluable help and support over the past 4 years. The work presented in chapter 3 on the Digital Mirror Langmuir Probe would also not have been possible without the close collaboration of Ted Golfinopoulos, William McCarthy and Brian LaBombard of the Plasma Science and Fusion Center at M.I.T. I would also like to thank the staff at NPL for their assistance in characterising the fission chamber acquisition system which was of significant help. Of course all the staff at Culham Centre for Fusion Energy and the York Plasma Institute deserve my gratitude for generating environments that made it a joy to come into work everyday. Special thanks must go to Scott Allan for the assistance he has provided in dealing with fission chamber organisation and planning as the fission chamber RO, and for always being happy to have a chat about next steps. I think it is safe to say that many of the results presented in chapter 2 would not have been possible without him.

Just as important have been my peers and friends in the CDT. All of them have been of great moral support over the course of the PhD and have made it a delightful experience. Special note must be made of Joe Allen as I could not have asked for a better partner over the course of the SAMI project. I would also like to name my various housemates over the years to thank them for their unending willingness to go for a drink or have a late night chat; Sam Gibson, Simon Orchard, Bhavin Patel, Will Trickey, Phil Bradford, and Chris Underwood. Caroline Lumsdon has provided the wonderful service of re-introducing me to the joy of reading, while Andrew Malcolm-Neale and Sam Ward have given me a much needed education in social issues and music respectively. Special thanks must also go to Tom Nicholas and Omkar Mytra for their companionship during the unique year that was 2020. I look forward to the many years of Christmas Carnage with you all.

Last, but definitely not least, I have endless gratitude to my parents, David and Leonora Vincent, who have been exceptional role models. My brother, Peter, also deserves mention for his apparently insatiable drive that pushes me to keep up. I would not have achieved as much without you.

The work presented in this thesis has been supported by UKRI and EPSRC. The views and opinions expressed herein do not necessarily reflect those of either organisation.

---

# Contents

<b>Declaration</b>	<b>viii</b>
<b>List of Figures</b>	<b>x</b>
<b>List of Tables</b>	<b>xix</b>
<b>Nomenclature</b>	<b>xx</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Introduction to Fusion . . . . .	1
1.1.1 Lawson Criterion . . . . .	4
1.1.2 Magnetic Confinement Fusion . . . . .	6
1.1.2.1 Tokamaks . . . . .	8
1.1.2.2 The Mega-Amp Spherical Tokamak Upgrade . . . . .	8
1.2 Instrumentation for Tokamak Diagnostics . . . . .	11
1.2.1 Passive Instrumentation . . . . .	11
1.2.2 Active Instrumentation . . . . .	12
1.2.3 Real Time Instrumentation . . . . .	12
1.2.4 Field Programmable Gate Array Technology . . . . .	13
1.2.4.1 Differences in Computing Devices . . . . .	15
1.2.5 FPGA Development Core Concepts . . . . .	17
1.2.5.1 Finite-state Machines . . . . .	19
1.2.6 The Xilinx Zynq Architecture . . . . .	20
1.2.6.1 Advanced eXtensible Interface . . . . .	20
1.3 Thesis Overview . . . . .	21
<b>2 Fission Chambers for Neutron Flux Measurements</b>	<b>22</b>
2.1 Background . . . . .	22

2.1.1	Neutrons in Fusion devices . . . . .	22
2.1.1.1	Physics Case for Neutron Flux Measurements . . . . .	23
2.1.1.2	Safety Case . . . . .	24
2.1.1.3	Fission Chamber Operation . . . . .	25
2.1.2	Analogue Electronics Front End . . . . .	26
2.1.3	Campbell Mode . . . . .	29
2.2	FPGA Control and Acquisition . . . . .	32
2.2.1	Koheron Software Development Kit . . . . .	32
2.2.2	Red Pitaya Development Board . . . . .	33
2.2.3	Peak Discrimination and Detection . . . . .	34
2.2.4	Mean and Variance Calculation . . . . .	36
2.2.5	Simulation for In-Situ Calibration . . . . .	38
2.2.6	Triggering and Clocking . . . . .	39
2.2.7	Data delivery and instrument control with Flask . . . . .	39
2.3	Fission Chamber Calibration . . . . .	40
2.3.1	National Physical Laboratory Calibration Procedure . . . . .	41
2.3.1.1	Thermal Neutron Tests . . . . .	41
2.3.1.2	Fusion Neutron Tests . . . . .	43
2.4	Implementation on MAST-Upgrade . . . . .	45
2.5	Summary . . . . .	48
2.5.1	Further Work . . . . .	49

### **3 Digital Implementation of a Mirror Langmuir Probe Control System** **52**

3.1	Background . . . . .	52
3.1.1	Langmuir Probe Theory . . . . .	53
3.1.2	Physics Case for Fast Langmuir Probes . . . . .	55
3.1.3	The Analogue Mirror Langmuir Probe . . . . .	55
3.2	Mirror Langmuir Probe Implementation Using FPGA Technology . . . . .	57
3.2.1	The Mirror Langmuir Probe Equations . . . . .	58
3.2.2	The Mirror Langmuir Probe Control System . . . . .	60
3.2.2.1	Calibration and Smoothing . . . . .	65
3.2.3	The Plasma Current Response System . . . . .	65
3.2.4	Digitisation and Hardware . . . . .	67
3.3	Mirror Langmuir Probe Emulation . . . . .	67
3.3.1	Step Change Tests . . . . .	67
3.3.2	Waveform Tests . . . . .	69
3.3.3	Emulated Real Data Tests . . . . .	70
3.4	Summary . . . . .	72

3.4.1	Further Work . . . . .	72
<b>4</b>	<b>The Synthetic Aperture Microwave Imaging Diagnostic</b>	<b>74</b>
4.1	Background . . . . .	74
4.1.1	The First Generation Synthetic Aperture Microwave Imaging Diagnostic . . . . .	78
4.1.2	Limitations to the Original SAMI Diagnostic and Digitisation and Acquisition Brief . . . . .	80
4.2	Microwaves in Tokamak Plasmas . . . . .	81
4.2.1	Image Formation with a Synthetic Aperture . . . . .	82
4.2.1.1	Beamforming . . . . .	82
4.2.2	2-D Doppler back-scattering . . . . .	83
4.2.2.1	Deriving Edge Current Density . . . . .	86
4.3	SAMI-2 Analogue Front End . . . . .	87
4.3.1	The Sinuous Antenna Design . . . . .	88
4.4	Software and Firmware Components . . . . .	89
4.4.1	10 Gigabit Ethernet . . . . .	89
4.4.1.1	Precision Time Protocol . . . . .	91
4.4.2	Intel NIC Configuration and Control . . . . .	94
4.4.3	Linux for Zynq Based Systems . . . . .	94
4.5	FPGA Based Acquisition and Control . . . . .	96
4.5.1	Xilinx 10 Gigabit Ethernet Subsystem . . . . .	96
4.5.1.1	Ethernet Frame Packaging . . . . .	98
4.5.1.2	ADC Data Conversion . . . . .	99
4.5.1.3	PTP Timestamping and Synchronisation . . . . .	101
4.5.2	Data Acquisition PC . . . . .	103
4.5.3	Network Interface Card (NIC) Driver . . . . .	105
4.5.4	Component testing . . . . .	107
4.6	Summary . . . . .	107
4.6.1	Future Work . . . . .	108
<b>5</b>	<b>Concluding Remarks</b>	<b>109</b>
5.1	Thesis Overview . . . . .	109
5.1.1	Fission Chambers for Neutron Flux Measurements . . . . .	111
5.1.2	Digital Implementation of the Mirror Langmuir Probe Control System . . . . .	111
5.1.3	The Synthetic Aperture Microwave Imaging Diagnostic . . . . .	112
5.2	Outlook . . . . .	112
5.2.1	ITER Tokamak Diagnostics . . . . .	113



---

# Declaration

The work in this thesis is based on research carried out predominantly at Culham Centre for Fusion Energy in partnership with the Centre for Advanced Instrumentation, Department of Physics, University of Durham, England. No part of this thesis has been submitted elsewhere for any other degree or qualification, and it is the sole work of the author unless referenced to the contrary in the text.

Some of the work presented in this thesis has been published in journals and conference proceedings - the relevant publications are listed below.

## Publications

C. Vincent et al, *The Digital Mirror Langmuir Probe: Field Programmable Gate Array Implementation of Real-time Langmuir Probe Biasing.*, Review of Scientific Instruments, 90(8):083504 2019, doi: 10.1063/1.510983

J. Allen et al, *Design of the Synthetic Aperture Microwave Imager Upgrade for measurement of the edge current density on MAST-U*, EPJ Web Conference, 203:03004 2019, doi: 10.1051/epjconf/20192030300

J. Allen et al, *Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic*, Proceedings of the 14th International Reflectometry Workshop for Fusion Plasma Diagnostics (IRW14) 2019

T. E. G. Nicholas et al, *Re-examining the Role of Nuclear Fusion in a Renewables-Based Energy Mix*, Energy Policy, 149:11204 2020, doi: <https://doi.org/10.1016/j.enpol.2020.11204>

W. McCarthy et al, *Real Plasma Tests of the Digital Mirror Langmuir Probe*, Prepared for publication in Review of Scientific Instruments (to be submitted January 2021)

**Copyright © 2020 by Charles Vincent.**

*“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”.*

---

# List of Figures

1.1	Figure showing the binding energy of elements relative to the number of nucleons. Used under the Creative Commons License, figure taken from “Wikimedia Commons” <sup>[1]</sup> . . . . .	2
1.2	Figure showing the potential well around an atomic nuclei. Quantum effects are utilised to overcome the Coulomb barrier such that at short range the strong atomic force dominates. Taken from “Fusion: A true challenge for an enormous reward” <sup>[2]</sup> . . . . .	3
1.3	Figure showing the fusion cross-section for different elements. Taken from “Fusion Physics” <sup>[3]</sup> . . . . .	4
1.4	Figure demonstrating the magnetic mirror plasma confinement concept. High energy particles, or particles not approaching the pinch at the correct angle, are not reflected by the tighter magnetic fields at either end and escape from the magnetic confinement. This design is therefore inherently leaky and it has been shown that the plasma losses were too high for this concept to ever produce a power plant device. Taken from “Electricity, Magnetism and Light” <sup>[4]</sup> . . . . .	6
1.5	Figure of the magnetic configuration and resulting magnetic field of a tokamak fusion experiment. Taken from the EUROfusion website <sup>[5]</sup> . . . . .	7
1.6	Diagram of the Super X divertor, a novel divertor design implemented on MAST-Upgrade. The divertor design has a much narrower baffle and longer outboard leg than classical divertor designs. The narrower baffle will prevent exhaust from re-entering the main tokamak chamber while the longer leg will increase the radiative cooling time of the exhaust before it strikes the divertor tiles. Taken from “Super-X advanced divertor design for MAST upgrade” <sup>[6]</sup> . . . . .	10
1.7	Flow chart of a basic feedback control system. The plant here refers to whatever system your feedback controller may be driving. . . . .	13

1.8	The configurable logic block (CLB) is the core component of FPGA programmable logic. A look-up table (LUT) allows the FPGA synthesiser software to implement what may be many operations into a single unit. The LUT output is connected to storage elements and the wider programmable logic (PL) system through multiplexers to facilitate greater flexibility. The outputs Q1 and Q2 are latched from O5 with signals X and I acting as clock enable or set and reset signals to give even more flexibility to the CLB. Figure taken from “User Guide UG574 UltraScale Architecture Configurable Logic Block” <sup>[7]</sup> . . . . .	14
1.9	Graphical demonstration of pipeling an addition operation. The addition of two 4 bit numbers results in a carry chain of length 4 (top). By separating this into the addition of two 2 bit numbers we have introduced an extra clock cycle of latency but the largest carry chain is now only 2 bits (bottom). If a signal can only travel the length of three carry chains before the result is latched by the clock signal the result of the (top) operation would be uncertain. . . . .	18
1.10	Representation of a Finite State Machine for a simple turnstile. These theoretical models can be used as a starting point for the implementation of control systems, communication protocols, or other automation applications. Used under the Creative Commons License, figure taken from “Wikimedia Commons” <sup>[8]</sup> . . . . .	19
2.1	Schematic of the fission chamber data acquisition implementation on MAST-Upgrade with the schematic of the front end analogue amplifier shown in figure 2.2. Shown here is only one fission chamber. While the original MAST machine ran with a single fission chamber the intention is to operate two simultaneously with the same acquisition system for MAST-Upgrade . . . . .	26
2.2	Electronics schematic of the front end electronics. The amplification factor can be switched between a gain of 10× and an inverted 100× with a two pin header on the board. The signal input can also be switched digitally from the front end inputs to a DAC output from the Red Pitaya. This feature allows testing of the system with a simulated source generated by the Red Pitaya FPGA board. Figure provided by the MAST-Upgrade Electronics Department . . . . .	27

2.3	A fast Fourier transform of the DC current from the pre-amplifier when using an internally generated signal from the Red Pitaya DAC. No significant peaks are seen from 10 Hz to 90 Hz showing that the pre-amplifier is operating as it was designed. The inset shows the full frequency range available up to 500 kHz. A peak at $\sim 366$ kHz with a modulation of $\sim 99$ kHz can be seen as the highest power peak. . . . .	28
2.4	FFT analysis of a square wave input at 61kHz. The test square wave covers half the dynamic range of the voltage input and so should serve as a reasonable benchmark for the noise element generated by the 366 kHz signal. The square wave is not entirely matched causing many harmonics within the FFT. The 366 kHz peak has a power of less than -40 dB relative to the main 61kHz peak. It has therefore been determined that no matter the source of this noise element it will have little effect on the measurements made by the final system. . . . .	29
2.5	The pulse shape of a fission event within the fission chambers. An event can be identified by the significantly larger signal when compared to the background noise, also shown in the figure. The majority of the noise floor is generated by the background alpha events. Numbers and a dashed line are used for illustrative purposes with the peak detection program in figure 2.6 . . . . .	30
2.6	A flow chart of the low latency peak detection algorithm used for the Fission Chamber acquisition system. While this is a simple method of detecting peaks the low latency means that the $1\mu\text{s}$ time resolution is preserved. Numbers are shown to pair decisions with the peak shown in figure 2.5. . . . .	34
2.7	A Gaussian pulse (FWHM 106ns) is sent by the pulse simulator at a frequency of 30.5 kHz and detected by the peak detection signal of the fission chamber acquisition system. . . . .	35
2.8	A flow chart showing the calculation of the DC and Campbell mode deliverables from the fission chamber acquisition system. As this method of calculating the current statistics from the fission chambers results in a low latency and light weight logic implementation it makes it very suitable for the Red Pitaya FPGA development board. . . . .	37
2.9	Comparison of the simulated pulse shape and a real pulse shape. The decay is slightly quicker for the simulated pulse, and the background noise has not been simulated however, this pulse shape has been sufficient for testing of the acquisition system. . . . .	38

2.10	Schematic of the fission chamber shielding assembly. Shielding layers increase the cross-section of fusion neutrons with the electrode uranium coating and attenuate gamma rays to increase the signal to noise ratio. A final thin Cadmium layer absorbs neutrons with energies that would reduce the temporal resolution of the acquisition system. . . . .	40
2.11	Diagram of the experimental setup used to test the pulse counting capabilities of the fission chamber acquisition system. Neutrons produced via a deuteron-beryllium interaction are thermalised by a graphite monitor and are then incident on the fission chamber. Neutron flux is attenuated by tilting deflection plates to reduce the number of deuterons incident on the beryllium targets. . . . .	42
2.12	Mode comparison for the fission chamber acquisition system using a thermal neutron source. This can be used to calculate the calibration parameters between pulse counting mode, DC current mode and Campbell mode. . . . .	43
2.13	A threshold scan at different neutron flux rates shows an increase in high voltage crossing points with increasing neutron flux. As these results will not actually be used for fission chamber calibration linear demand levels are shown as a relative value to demonstrate how the pulse counting mode changes at different neutron fluxes. The red cross marks 25mV above the pulse count peak to more clearly demonstrate how the maximum peaks detected changes with neutron flux demand level. . . . .	44
2.14	Number of crossing points registered for a given neutron flux at different peak detection thresholds. When measuring crossing points without pulse pile-up small threshold values would be expected to register the same number of crossing points with divergence as the threshold increases. This figure shows that there is an increase in crossing points at all threshold levels, indicating that all measurements were made with pulse pile-up occurring. . . . .	45
2.15	Diagram of the fusion neutron test setup at NPL. The fission chamber is placed in its shielding assembly and mounted on a movable arm. The height is fixed such that the proton beam and electrode centres are aligned. To measure the neutron flux the long counter <sup>[9]</sup> was placed in the same position as the fission chamber at each measurement location. A shadow cone was placed between the long counter and the neutron source to block the direct component and measure the scattered neutron flux. . . . .	46

2.16	Threshold scan using the fusion neutron source at different distances. Operation is unlikely to be in pile-up mode as the threshold count peak and FWHM is similar for each measurement distance while the crossing point count in the tail decreases with distance from the neutron source. The pulse count axis shown here is logarithmic so that the difference in threshold crossings detected is more distinct at high thresholds. While there is little difference between the number of pulses registered at 260 cm with the neutron source on or off, more pulses are clearly registered at a distance of 230 cm with more again at 79 cm. . . . .	47
2.17	Close-up comparison of a peak threshold scan comparing the fission chamber background with fission chamber response to 2.45 MeV neutrons. Acquisitions were taken over 4s at a distance of 0.79m . . . . .	48
2.18	CAD drawing and schematic of the fission chambers installed outside the MAST-Upgrade vessel. The chamber is mounted on a free standing strut to elevate it to the plasma mid-plane. Drawing provided by the MAST Drawing Office. . . . .	49
2.19	Raw DC current signal from the fission chamber installed on MAST-Upgrade. It can be inferred from the applied voltage trace that interference occurs through some coupling with the magnetic coil power supplies. The lower left inset shows applied voltage trace to a poloidal magnet. The magnet system start the voltage ramp at -0.5s which aligns well with the interference occurring in the DC current measurement. Insets show a magnified view of 3 individual periodic artefacts with a frequency of approximately 6 kHz. . . . .	50
2.20	Fast Fourier transform of the raw DC current shown in figure 2.19. Different time periods have been shown to differentiate between the coil power supply being off and on. While the DC current shows a 6 kHz artefact as being most prominent the FFT indicates that this is a modulation with a more powerful 200 kHz signal. . . . .	51
3.1	Figure showing characteristic I/V curve of a Langmuir probe. Regions of ion and electron saturation as well as the floating point voltage must be measured to accurately characterise the curve. . . . .	54
3.2	Simple circuit diagram of the transistor arrangement used in the original MLP diagnostic. This configuration of a PNP and NPN transistor is used to emulate the plasma response to an applied bias $V_P$ . Figure reproduced from “Mirror Langmuir probe: A technique for real-time measurement of magnetized plasma conditions using a single Langmuir electrode” <sup>[10]</sup> . . . . .	56

3.3	These figures show the convergence (blue regions) and non-convergence (red regions) for initial values of $T_e$ , $I_{sat}$ and $V_F$ . All figures have the same “plasma” $T_e$ and $I_{sat}$ values with three different values for “plasma” $V_F$ . Initial values are represented as scaled to the “plasma” values. Convergence of these equations is decided by running the bias state cycle until the temperature values is within 1% of the actual temperature or the result is determined to be diverging or non-physical. . . . .	60
3.4	A diagram of the Mirror Langmuir Probe test system. Shown are the connections between the Mirror Langmuir Probe (MLP) and the Plasma Current Response (PCR). Due to small sampling errors of the Red Pitaya ADC the PCR measured voltage is looped back to the MLP and this value, in conjunction with the PCR calculated value, is used to calculate the plasma temperature. . . . .	61
3.5	Diagram of the bias state sequence of the MLP control system. The calculated value and the bias value relative to temperature are annotated on each state. The first state is used to calculate a starting temperature which is then propagated through the other biases. An ideal system would have a positive bias of $0.675T_e$ and negative bias or $-3.325T_e$ but to simplify the digital logic requirements these values have been rounded to the nearest integer. . . . .	61
3.6	A flow chart of the calculation module and how it interacts with its respective look-up table and divider IP. For the ion saturation calculation the module takes an input of temperature, floating voltage and probe bias. The VHDL code that governs the parameter calculation can be seen in appendix C.1. . . . .	62
3.7	The voltage traces for the PCR (blue) and MLP (green) systems when the bias switching speed is on the same order as the PCR latency. It can clearly be seen from the neutral state $\sim 0$ that when sampling the MLP the PCR is at the wrong bias state of the cycle. . . . .	66
3.8	Figure showing the MLP response to a step change in temperature from 60 eV to 100 eV. The MLP calculated temperature is an overestimate of the applied temperature but the response is instantaneous. . . . .	68
3.9	FPGA smoothed voltage trace (blue line) of the bias states (orange markers) for a step change in temperature from 60 eV to 100 eV. . . . .	69
3.10	A demonstration of the MLP following a continuous change in temperature. The test also shows the control system break down below a 15 eV threshold. The correct temperature is again calculated once the waveform enters the operational region. . . . .	70

3.11	Triangle waveform in floating voltage at a constant temperature. The threshold for look-up table induced break down is not reached and the calculated floating voltage stays stable. While a correctly calculated floating voltage indicates good system stability, it has no direct impact on the bias state output levels and is only used in the temperature calculation cycle. . . . .	71
3.12	A demonstration of MLP behavior when C-Mod turbulence data is used to set the PCR plasma parameters. The MLP follows the turbulence well enough to facilitate accurate post-processing fit of the plasma parameters. The inset more clearly shows that while the calculated temperature by the digital MLP is not as good as the original analogue system, the general trends are presented clearly with a systematic over estimation.	72
4.1	Example graph of the pedestal formed in the pressure profile of a H-mode plasma compared to a L-mode plasma. The high edge pressure gradient suppresses turbulence at the plasma edge and leads to higher core pressures. Figure taken from “ELM control strategies and tools: Status and potential for ITER” <sup>[11]</sup> . . . . .	76
4.2	Diagram of the stability curve created by the coupling of the edge current density and pressure gradient. Three classifications of ELM are demonstrated, characterised by the destabilising mechanisms that cause them. Figure reproduced from “Magneto-hydrodynamic stability of the H-mode transport barrier as a model for edge localized modes: an overview” <sup>[12]</sup> . Currently only type I ELM theory is generally agreed, with the paper referenced here proposing a theory for the occurrence of type II and III ELMs. . . . .	77
4.3	Analysed data from shot #27969 generated by SAMI when implemented on the MAST tokamak. The red and blue zones show the regions of greatest Doppler shift as the probing microwaves are reflected off the spinning plasma surface. As we know that the magnetic field is perpendicular to the plasma rotation, the magnetic pitch angle can be inferred from the plane of the regions of maximum Doppler shift. The dotted line connects the two Doppler poles and is orthogonal to the horizontal solid lines which show the magnetic field as reconstructed by EFIT <sup>[13]</sup> . Figure taken from “2D Doppler backscattering using synthetic aperture microwave imaging of MAST edge plasmas” <sup>[14]</sup> . . . . .	79

4.4	Diagram of incident beam on a reflection grating and the back-scattered light (with diffraction orders shown as an example) resulting from corrugations along a reflecting surface. In the case of the SAMI diagnostic this surface is the iso-density cut-off surface corresponding to the frequency of the emitted beam. Corrugations along the surface are the result of turbulence within the plasma edge. . . . .	84
4.5	Back scattered power of a Doppler backscattering diagnostic. The central frequency peak is from the probing beam, after the gigahertz carrier wave has been removed (see section 4.3). While the system resolution is not sufficient to resolve the Doppler shifted peak there is still a clear asymmetry visible depending on which direction the plasma is spinning in. Colours align with those in figure 4.3, such that blue is moving away and red is moving towards the antenna. Figure taken from “2-D Doppler backscattering using synthetic aperture microwave imaging of MAST edge plasmas” <sup>[14]</sup> . . . . .	85
4.6	Simple block diagram of the SAMI-2 diagnostic showing the front end components that have been implemented on a custom PCB as described in “Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic” <sup>[15]</sup> . . . . .	87
4.7	The sinuous antenna design has two distinct arms for polarisation separation <sup>[16]</sup> . The design of the antenna allows for a relatively small cross-section for the bandwidth it is designed to respond to. It is also a planar antenna allowing for it to be printed directly on to a PCB while minimising antenna cross-talk when incorporated as part of an array. The orthogonal arms of the antenna allow to differentiate between orthogonal polarised light. Figure taken from “Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic” <sup>[15]</sup> . . . . .	89
4.8	The frame composition of an ethernet data transfer is shown here as defined by IEEE protocol 802.3. Further protocols such as TCP/IP or UDP can then be used within the packet data. In standard ethernet usage the maximum frame size is limited to 1522 bytes, including header information and error checking, but larger frame formats up to 9000 bytes can be supported if the ethernet interface is configured in “Jumbo” mode. . . . .	90
4.9	Timing diagram of a PTP synchronisation from the primary interface to a single secondary interface. Figure taken from EndRun Technologies white paper “Precision Time Protocol” <sup>[17]</sup> . . . . .	92

4.10	Graphical representation of how the Linux system components mesh to form a full operating system. . . . .	95
4.11	Flow chart of the programmable logic implemented on a Xilinx ZCU102 Evaluation Board for the SAMI-2 data delivery system. Yellow shows the IP cores developed as part of the work presented here. Green shows IP cores developed by Roddy Vann and white shows standard IP cores available as part of the Xilinx standard and proprietary library. . . . .	97
4.12	Flow chart of the ethernet frame packaging module. When the data packaging module reaches the end of a packet the last frame flag is asserted and the ethernet FIFO makes the data available to the ethernet subsystem. . . . .	99
4.13	Graphical representation of how the ADC 12 bit data sample is packed into the 16 bit output of the ADC interface module. . . . .	100
4.14	Diagram of the frame buffering loop in the ADC Data conversion module. Blue lines show points in the loop where a complete frame has been generated and output to the ethernet data packaging module while red lines show the buffering loop path. . . . .	101
4.15	Flow chart for the PTP timestamping module used to extract a timestamp generated by the acquisition PC. The timestamp is then inserted into the ethernet packet to facilitate data synchronisation across multiple FPGA acquisition boards. . . . .	102
4.16	Block diagram of the PCI-e lanes of the dual socket ASUS WS C621E SAGE motherboard. While PCI-e lanes are still multiplexed between interfaces it is possible to avoid sharing lines by carefully selecting which interfaces to populate. Figure taken from the ASUS WS C621E SAGE manual <sup>[18]</sup> . . . . .	104
B.1	Schematic of the MAST-Upgrade Operations network and how FPGA devices fit into the network. The unit program interfaces directly with the FPGA device and passes data to the scheduler code for processing and storage. . . . .	119

---

# List of Tables

1.1	Table comparing and describing the advantages and disadvantages of different computing technologies and the applications for which they are appropriate. . . . .	16
3.1	Table describing the inputs, outputs and functions of the modules shown in figure ?? . . . . .	62
4.1	Table of the PTP header information. The PTP header comes directly after the ether type (as shown in figure 4.8) as part of the packet body. Table is reproduced from “Implementing IEEE 1588v2 for use in the mobile backhaul” <sup>[19]</sup> . . . . .	92

---

# Nomenclature

**ADC** Analogue-to-Digital Converter

**ASIC** Application Specific Integrated Circuit

**AXI** Advanced eXtensible Interface<sup>[20]</sup>

**CCFE** Culham Centre for Fusion Energy

**CLB** Configurable Logic Block

**COTS** Commercial Off The Shelf

**CPLD** Complex Programmable Logic Device

**CPU** Central Processing Unit

**DAC** Digital-to-Analogue Converter

**DBS** Doppler Back-scattering

**DMA** Direct Memory Access

**DTB** Device Tree Blob

**ELM** Edge Localised Mode

**FFT** Fast Fourier Transform

**FPGA** Field Programmable Gate Array

**FSM** Finite-State Machine

**FSBL** First Stage Boot Loader

**GPU** Graphics Processing Unit

**GPIO** General Purpose Input/Output

**GUI** Graphical User Interface

**HDL** Hardware Development Language

**IC** Integrated Circuit

**ICF** Inertial Confinement Fusion

**ITER** Power plant scale tokamak, undergoing construction

**I/O** Input/Output Interface

**LP** Langmuir Probe

**LSB** Least Significant Bit

**LUT** Look-up Table

**MAC** Media Access Control

**MAST-Upgrade** Mega-Amp Spherical Tokamak Upgrade

**MCF** Magnetic Confinement Fusion

**MIT** Massachusetts Institute of Technology

**MLP** Mirror Langmuir Probe

**MSB** Most Significant Bit

**MSE** Motional Stark Effect

**NIC** Network Interface Card

**NSTX-U** National Spherical Torus Experiment Upgrade

**NBI** Neutral Beam Injector

**PSFC** Plasma Science and Fusion Center

**PFC** Plasma Facing Component

**PCR** Plasma Current Response

**PTP** Precision Time Protocol

**PCS** Plasma Control System

**PL** Programmable Logic

**POSIX** Portable Operating System Interface

**RP** Red Pitaya Development Board<sup>[21]</sup>

**SDK** Software Development Kit

**SoC** System on Chip

**SAMI** Synthetic Aperture Microwave Imaging

**SOL** Scrape-Off Layer

$U^{235}$  Uranium-235

---

# Introduction

This chapter serves as a general introduction to fusion, tokamaks and Field Programmable Gate Array (FPGA) technology. It will cover the basic plasma physics that will be referenced throughout the thesis as well as the advantages and disadvantages of using in-house designed FPGA technology over application specific custom hardware or more standard Commercial Off The Shelf (COTS) devices.

## 1.1 Introduction to Fusion

With increasing demand for low carbon energy sources it is important to continue pushing the technology envelope forward to provide as many solutions as possible<sup>[22]</sup>. Fusion power plants will provide high energy density sources with low to no carbon emissions and without much of the high level, long lasting nuclear waste or meltdown risk associated with traditional nuclear fission power<sup>[23]</sup>. However, the technological breakthroughs required to achieve a fusion power plant and the advanced technologies required to characterise and control these complex machines cannot be understated. It should also be noted that technologies developed through fusion research may contribute to other fields, as previously seen with other large scientific endeavours<sup>[24] [25] [26]</sup>. A discussion on the role fusion can play in a future energy mix can be found in appendix A.1<sup>[27]</sup>.

Nuclear fusion is the process of combining two nuclei into a single heavier nucleus. Depending on the binding energy of the source nuclei this can either be an exothermic or endothermic process. This is clearly shown in figure 1.1 which shows the average binding energy (the amount of energy needed to separate nucleons) for each element. For low nucleon elements fusion causes an increase in the binding energy per nucleon and therefore a release in energy. The opposite is true for the

high nucleon elements where fission (the act of separating an element) increases the overall binding energy.

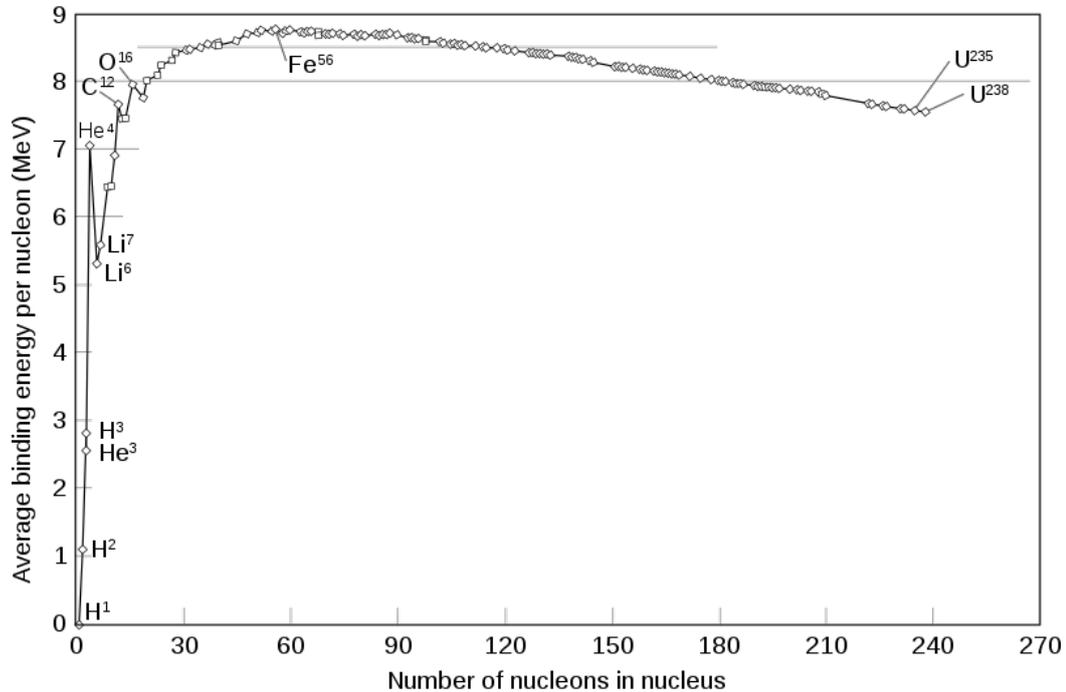


Figure 1.1: Figure showing the binding energy of elements relative to the number of nucleons. Used under the Creative Commons License, figure taken from “Wikimedia Commons”<sup>[1]</sup>.

In the core of stars the fusion process is predominantly proton-proton fusion<sup>[28]</sup>. This is not viable on earth as the proton-proton cross section is extremely low. Stars can overcome the Coulomb barrier (in part a cause of the low proton-proton cross section) between nucleons due to the large core pressures generated by their mass, and the long confinement time of the fusion fuel. On Earth, the Coulomb barrier is overcome with extremely high temperatures such that fusion reactants collide at extremely high energies, and by using fusion reactants with a high cross section, terrestrial fusion can be achieved. By themselves, both the high core pressures in stars or the high incident particle energy in terrestrial fusion, would not be sufficient to overcome the Coulomb barrier. Instead, quantum tunnelling occurs in fusion reactions to bypass the energetic summit of the Coulomb barrier. Figure 1.2 shows how quantum tunnelling can reduce the required energy necessary to overcome the electrostatic potential between two like charged particles.

Figure 1.3 shows that the deuterium-tritium reaction has the most favourable cross section with a peak collision energy at 100 keV; this is the current front runner for fuel in future fusion power plants. However, tritium is a controlled material due

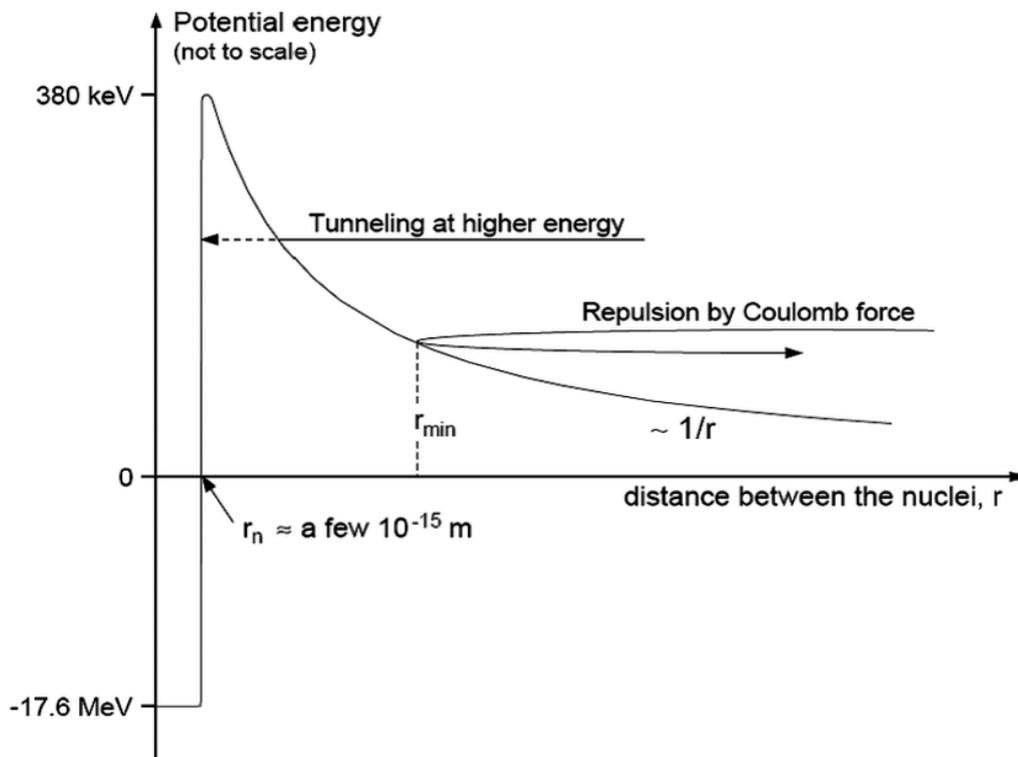


Figure 1.2: Figure showing the potential well around an atomic nuclei. Quantum effects are utilised to overcome the Coulomb barrier such that at short range the strong atomic force dominates. Taken from “Fusion: A true challenge for an enormous reward”<sup>[2]</sup>.

to its radioactivity and so is inconvenient to use for regular fusion experiments. Instead a pure deuterium is usually used and this serves as a suitable substitute to study the plasma physics and materials environment of a fusion device. From figure 1.3 we can see that very high temperatures are necessary to achieve fusion. Complicating matters further, the deuterium fusion reaction produces a high energy neutron which adds to the very harsh materials environment. The reaction equations for both a tritium-deuterium fuel mix and a pure deuterium fuel and shown in equations 1.1 and 1.2.



Finally, at the temperatures required to achieve fusion reactions the deuterium gas forms a plasma. While this means that any material coming into contact with the plasma would be destroyed, due to the electromagnetic properties of the plasma it

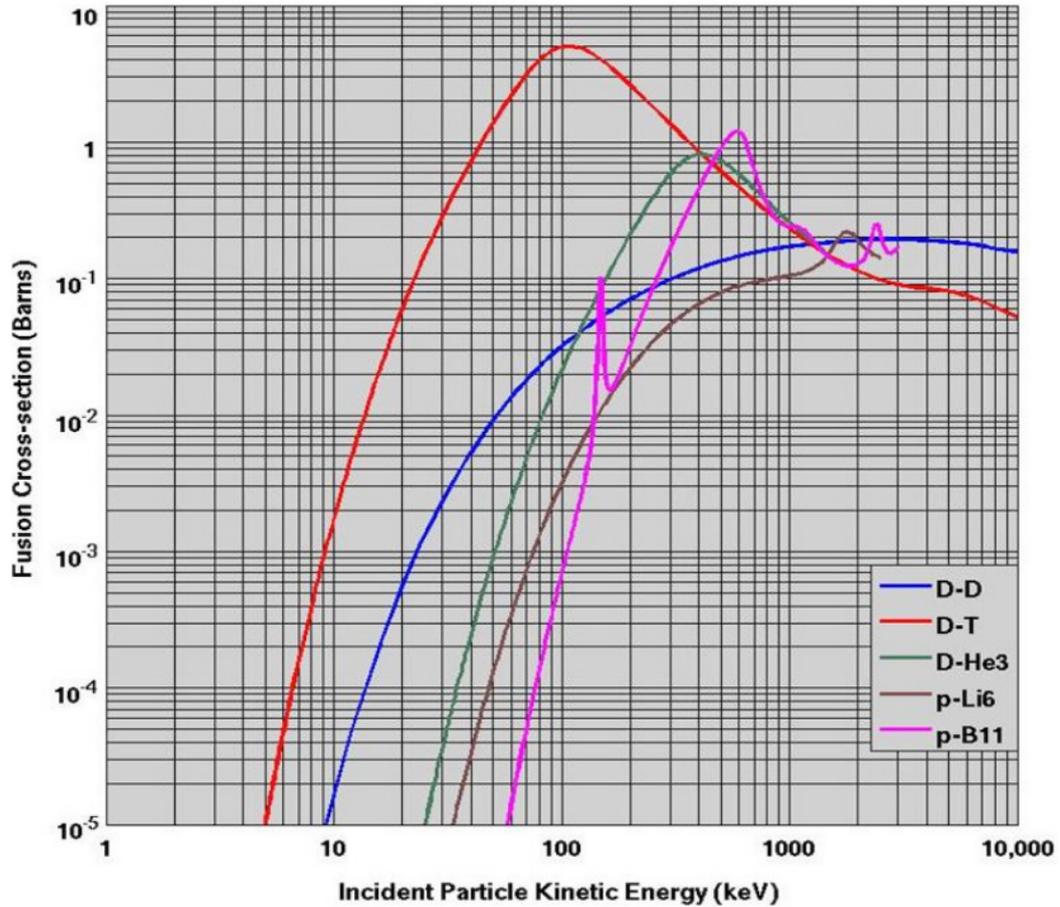


Figure 1.3: Figure showing the fusion cross-section for different elements. Taken from “Fusion Physics”<sup>[3]</sup>

can be contained by magnetic fields. All in all this means that achieving terrestrial fusion is a complicated endeavour.

### 1.1.1 Lawson Criterion

As previously mentioned, to achieve terrestrial fusion we must heat our fuel to extremely high temperatures, causing the fuel to form a plasma. To achieve a net energy output the heating necessary to maintain our fusion plasma should come from the energy emitted by the fusion reactions. It is clear that the self heating power will be proportional to the fusion reactant rate, as fusion is the dominant heat producing process, and so for maximum heating we should try to maximise this rate. For a deuterium tritium fuel mixture the power is at a maximum when equation 1.3 is true<sup>[29]</sup>, which is the result of equal number densities ( $n$ ) of tritium

and deuterium.

$$P_F = \frac{n^2}{4} \langle \sigma \nu \rangle E \quad (1.3)$$

Where the reactant rate  $\langle \sigma \nu \rangle$  has been maximised for the specific fusion reactants,  $E$  is the energy released from a reaction and  $P_F$  is the fusion power.

For the system to operate in steady state, with no external heating, the fusion power available for heating should be at least equal to the system losses. As seen in equations 1.1 and 1.2 there are two products from the fusion reaction; a Helium nucleon and a neutron. Neutrons are electromagnetically neutral and have a very small cross section with other materials; therefore the neutron generally escapes the confining medium and is not available to deposit energy within the bulk plasma. The energy available for heating is hence the fraction of energy left in the charged helium nuclei.

The plasma losses can be defined as the fraction of total plasma energy ( $W$ ) that is lost over a given timescale ( $\tau_E$ ).

$$P_L = \frac{W}{\tau_E} \quad (1.4)$$

Where  $\tau_E$  can be determined through experiment for a given machine. By balancing the power lost and heating power, a threshold can be found for the temperature relative to the product of the number density and confinement time  $\tau_E$ .

$$n\tau_E > \frac{12}{\langle \sigma \nu \rangle} \frac{T}{\epsilon_\alpha} \quad (1.5)$$

Where  $\epsilon_\alpha$  is the fraction of energy in the fusion alpha particle. Equation 1.5 can be used to calculate the minimum necessary product of number density and confinement time for a given fuel temperature. The most favourable D-T fusion fuel mixture minimises this value at a temperature of 30 keV:

$$n\tau_E > 1.5 \times 10^{20} m^{-3} s \quad (1.6)$$

This relationship is called the Lawson Criterion, the minimum criterion necessary to achieve ignition of the fusion fuel.

At the extremes of this lie two scenarios. Either very high densities over very small timescales, or low densities over much longer timescales. These two extremes have led to the development of very different technologies, Inertial Confinement Fusion (ICF) and Magnetic Confinement Fusion (MCF).

ICF technologies use high power lasers to ablate fusion fuel pellets. The resulting pressure wave from the shell ablation creates high densities in the pellet core with the objective of creating a burn front<sup>[30]</sup>. The material in this thesis concentrates on

MCF instrumentation. MCF based fusion machines use magnetic fields to confine the fusion plasma for long periods of time with external heating elements that intend to bootstrap the plasma to become “burning”<sup>[29]</sup>. That is, to a point at which the fusion power generated internally by the plasma is sufficient to achieve ignition conditions.

### 1.1.2 Magnetic Confinement Fusion

At the high temperatures required to sustain a plasma for any significant period of time, it must be contained within a vacuum. This is both because the plasma would destroy any physical container (through melting or ablation) and contact with cooler gases can very quickly quench the plasma state. As a plasma is composed of dissociated atoms and is therefore a cloud of charged nucleons and electrons, containment can be achieved with magnetic fields. An early plasma confinement design was the magnetic mirror which contained plasma using double pinched parallel magnetic field lines, as shown in figure 1.4. However, due to the limited criteria necessary for a particle to be reflected this meant the designs were extremely “leaky” and therefore incurred large plasma losses making it incompatible with ever being a viable power plant design. To fix this issue designs advanced to close the “leaks” by forming the parallel field lines into a torus.

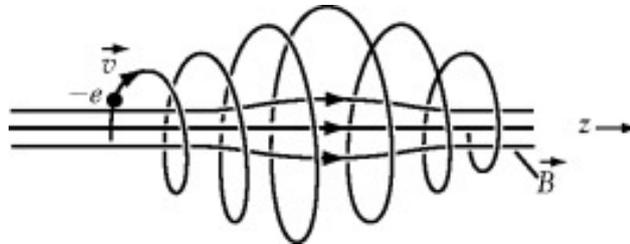


Figure 1.4: Figure demonstrating the magnetic mirror plasma confinement concept. High energy particles, or particles not approaching the pinch at the correct angle, are not reflected by the tighter magnetic fields at either end and escape from the magnetic confinement. This design is therefore inherently leaky and it has been shown that the plasma losses were too high for this concept to ever produce a power plant device. Taken from “Electricity, Magnetism and Light”<sup>[4]</sup>.

While plasma leaks have now been resolved this is still not sufficient to contain a “stable” plasma. By forming the magnetic fields into a torus the parallel field lines become non-uniform across the radius of the torus. From the Lorentz force it is not difficult to see that this would form a charge separation due to drifts induced by the magnetic field gradient<sup>[29]</sup>. A similar force induced by the field curvature also works to form a charge separation through drift for an overall drift velocity,

shown in equation 1.7.

$$v_d \propto B \times \nabla B \quad (1.7)$$

Where  $v_d$  is the drift velocity and  $B$  represents the vector magnetic field. This is called the grad B drift and introduces a third drift due to the resultant electric field, which causes the particles to leave the confinement region.

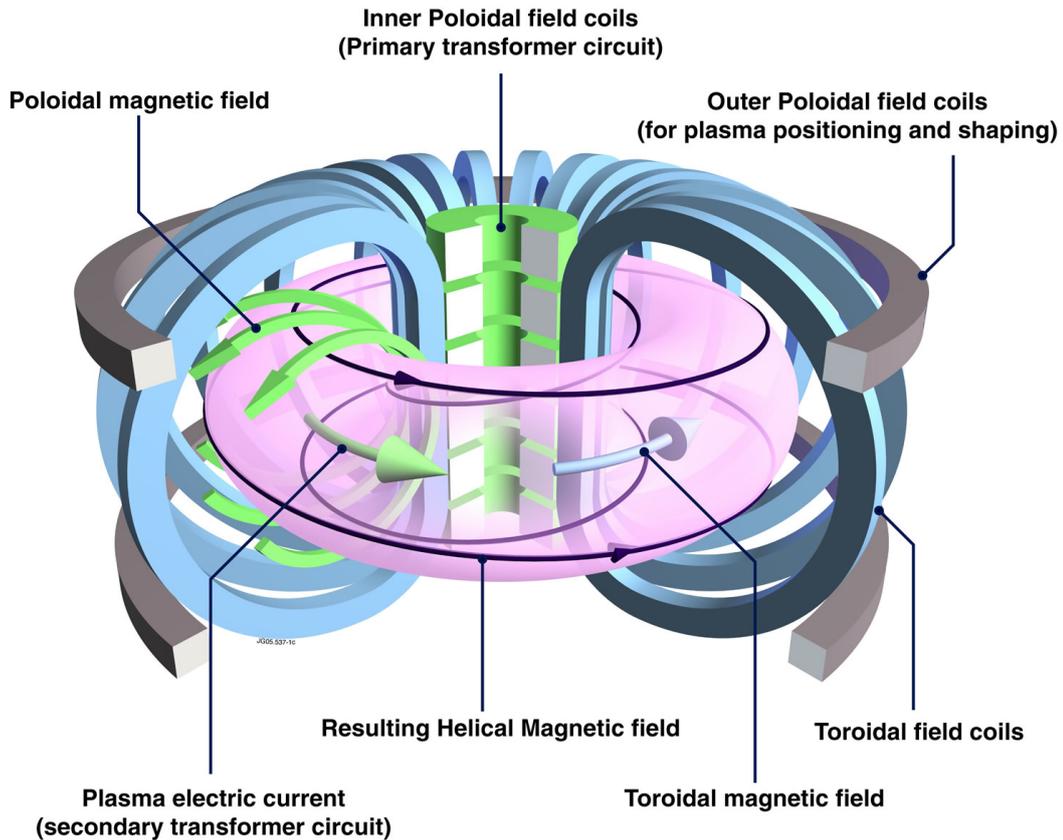


Figure 1.5: Figure of the magnetic configuration and resulting magnetic field of a tokamak fusion experiment. Taken from the EUROfusion website<sup>[5]</sup>.

To keep the particles confined the electric field must be stopped from forming. This is achieved by producing a helical field (as seen in figure 1.5), such that the average field and curvature seen by all particles is the same. The requirement to produce a helical field has resulted in two main devices for magnetic plasma confinement, the Tokamak<sup>[29]</sup> and the Stellarator<sup>[31]</sup>. Stellarator designs use a complex arrangement of magnetic fields to maintain a confined ring of fusion plasma and impose a helical field arrangement in the confining region. The focus of this thesis however is plasma diagnostics for tokamak devices. Tokamak devices use external fields and a self generated field from the spinning plasma to form the necessary constraining field as shown in figure 1.5.

### 1.1.2.1 Tokamaks

A tokamak plasma confinement system consists of two generated fields, a toroidal field and a poloidal field, to produce a helical confining field. The toroidal field is generated by external ring magnets while the poloidal field is generated by the current of the rotating internal plasma. Figure 1.5 also shows additional poloidal field coils to assist with positioning and shaping of the plasma and a central solenoid used to induce current in the plasma.

The solenoid acts like a primary core in a transformer to transfer current to the plasma which then generates the poloidal field. The solenoid itself does not deposit enough energy to achieve fusion though, and much of the fusion in tokamaks only occurs when external heating is applied<sup>[29]</sup>. The tokamak is currently the most advanced form of fusion reactor design with the (at the time of writing) JET tokamak achieving the highest recorded fusion power of 16.1 MW from 25.7 MW of heating power<sup>[32]</sup>.

Because the tokamak design generates its confining field in part from the plasma current, the external magnet design is significantly simplified compared to that of a stellarator. However, it can be easily seen from the Maxwell-Faraday equation that a tokamak can only be driven in a pulsed mode should the plasma current be solely driven by the central solenoid.

$$\Delta \times E = -\frac{\delta B}{\delta t} \quad (1.8)$$

Equation 1.8 tells us that to induce an electric field ( $E$ ) within the plasma, and therefore a current, a time varying magnetic field ( $B$ ) must be applied. This time varying magnetic field is generated by the coils of the solenoid which are driven with a voltage ramp. The voltage ramp cannot be infinite as you would quickly overload the solenoid power supplies, so the solenoid, and therefore the tokamak, can only be operated in a pulsed manner.

This is an issue for future reactor designs as a varying thermal load on plasma facing components induces stress that would not be present in steady state. From the perspective of energy generation it is also unfavourable as pulsed operation inherently includes reactor downtime, reducing the return on investment. Other forms of current drive exist such that the solenoid may be only used for startup, or not at all, but these also come with associated complications<sup>[33] [34] [35]</sup>.

### 1.1.2.2 The Mega-Amp Spherical Tokamak Upgrade

The Mega-Amp Spherical Tokamak Upgrade (MAST-Upgrade) tokamak is an upgrade to the UK operated MAST fusion reactor. The original MAST experiment

was built to test the novel design concept of a tight aspect ratio (ratio of major to minor plasma radii) fusion reactor<sup>[36]</sup>. Unlike conventional fusion reactors the MAST design has a ratio of major and minor radii that is close to one, hence the term spherical tokamak. The increased proximity to the magnets means that higher containment can be achieved for equivalent magnetic field strength. As a lot of the cost of a fusion reactor is due to the magnets, being able to reduce the necessary field makes spherical tokamaks economically attractive. As the solenoid is also closer to the plasma reduced power is necessary for equivalent current drive.

The advantages of spherical tokamaks are expressed by the beta parameter, which quantifies the efficiency of plasma confinement<sup>[29]</sup>.

$$\beta = \frac{p}{B_0^2/2\mu_0} \quad (1.9)$$

The beta, as seen in equation 1.9 is the ratio of plasma pressure ( $p$ ) to magnetic energy density ( $B_0^2/2\mu_0$ ). The beta for spherical tokamaks is significantly improved such that START, the first spherical tokamak to be operated at Culham Centre for Fusion Energy (CCFE) broke the beta factor world record by a factor of 2 when it was first built<sup>[37]</sup>. As higher pressures generally correspond to a higher fusion rate, and beta can be seen as quantifying the pressure efficiency of a tokamak, a high beta is a good indicator of performance for tokamak designs.

The effects of a spherical tokamak are not all positive however. While the small size does mean reduced costs it also means that sensitive components, such as magnets and the solenoid, are at greater proximity to the plasma. This exposes them to a much larger neutron flux without the required space for sufficient shielding to avoid significant damage. In a continuous operation scenario this would mean a comparatively shorter lifetime for some of the most expensive parts of a plant. While methods of heating and inducing current radiatively have been studied, such that the central solenoid is no longer required, much of these technologies are still in their infancy<sup>[38]</sup> <sup>[34]</sup>.

To exploit the enhanced damage potential of a spherical tokamak the original remit of the MAST machine was to explore the possibility of a fusion component test facility<sup>[39]</sup>. Such a facility would be able to test novel component designs quickly without building an entirely new machine and therefore enhance the pace of tokamak development. While not nominally a component test facility itself, MAST-Upgrade has continued to build on this remit by implementing a new divertor design called the Super-X<sup>[6]</sup>.

The Super X divertor concept, shown in figure 1.6 is a novel divertor design that improves on classical designs in two ways. The divertor baffle, separating the main

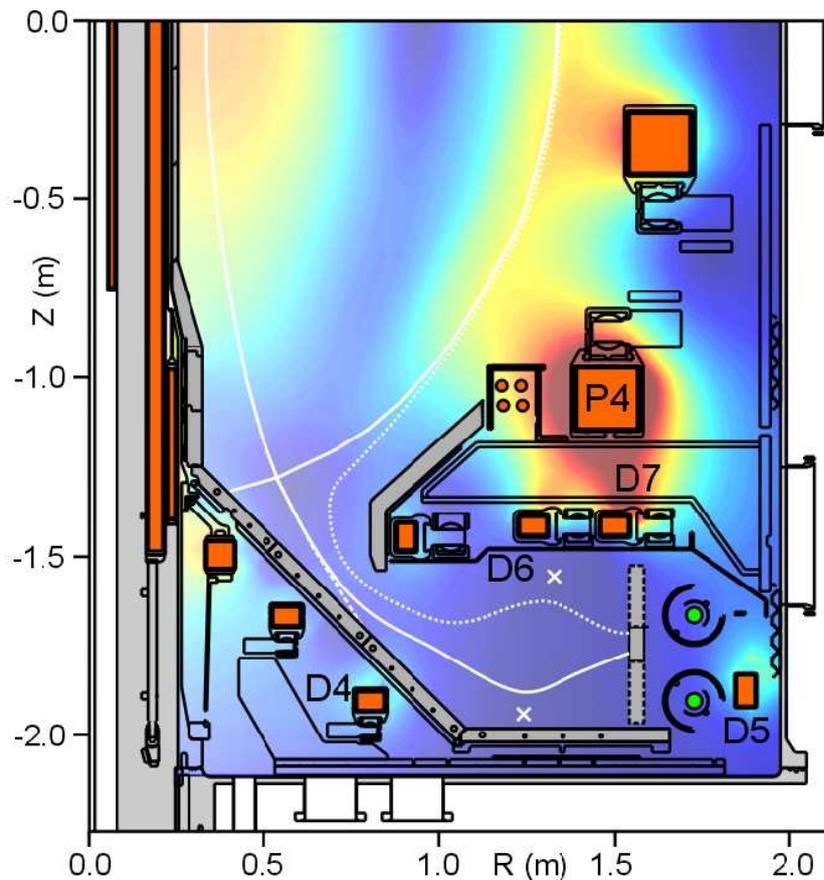


Figure 1.6: Diagram of the Super X divertor, a novel divertor design implemented on MAST-Upgrade. The divertor design has a much narrower baffle and longer outboard leg than classical divertor designs. The narrower baffle will prevent exhaust from re-entering the main tokamak chamber while the longer leg will increase the radiative cooling time of the exhaust before it strikes the divertor tiles. Taken from “Super-X advanced divertor design for MAST upgrade”<sup>[6]</sup>

tokamak chamber and the exhaust region is narrower than the original MAST design and will prevent neutralised exhaust gases from re-entering the main chamber. It is planned that the reduced neutral gas in the edge region will reduce cooling and therefore improve the fusion plasma efficiency. Secondly, the novel divertor design lengthens the exhaust leg considerably to increase the radiative cooling time before the plasma strikes the divertor. It is important to explore whether or not this will significantly reduce divertor heat loads as for power plant tokamak designs the divertor strike points would wear very quickly and require a lot of plant downtime to replace the divertor material.

## 1.2 Instrumentation for Tokamak Diagnostics

This thesis concerns itself with the development of several diagnostics for tokamak devices. While the helical field geometry of the tokamak design solves the problem of particle drifts there are still plenty of instabilities present that make the issue of achieving a burning plasma very difficult. A tokamak plasma can be separated into two regions of interest, the core and the edge. The plasma edge can be cool enough to measure with physical probes as it consists of a mixture of neutral particles at low temperatures and ionised particles outside the plasma core. The plasma core has an extremely high temperature and so cannot be probed physically. Core diagnostics therefore use probing beams and measure the changes in spectrum and diffraction as they pass through the plasma.

Diagnostics on experimental devices like MAST-Upgrade and JET not only help with the operation of the tokamak but also inform future designs for power plant scale machines. As the majority of wall space in a fusion power plant will be allocated to heat generating neutron blankets (converting neutron kinetic energy into heat through collisions), leaving little room for diagnostic ports, it is important to determine and formulate schemes for which diagnostics are essential for tokamak operation. It is also important to note that diagnostics that can operate on smaller machines such as MAST-Upgrade may not be applicable to larger machines such as ITER<sup>[40]</sup> or SPARC<sup>[41]</sup>. The increased neutron flux from these machines means that some components (especially optical components such as lenses and mirrors) would degrade too quickly for them to be viable.

Plasma diagnostics vary widely in their electronic complexities, from completely analogue systems (e.g. activation foils), part analogue with data digitisation (e.g. fission chambers) to almost completely digitised end-to-end (e.g. camera systems). With the ubiquitous use of computers for scientific research in modern times almost all systems need some form of digital data acquisition. This can be done through specially built hardware (Application Specific Integrated Circuit (ASIC)) or programmable COTS (FPGA, Complex Programmable Logic Device (CPLD)) devices. The following material will look at different classifications of diagnostic equipment and cover the basics of FPGA based instrumentation, the core focus of this thesis.

### 1.2.1 Passive Instrumentation

Passive instrumentation defines a class of diagnostic that does not interfere with the material that is being observed. This means that passive diagnostics consist solely

of detectors for some form of radiation or field emitted from the experiment. In the case of tokamaks, this generally consists of neutrons, external magnetic fields, and different wavelengths of emitted electromagnetic radiation.

The advantage of passive diagnostics over active diagnostics is that you can be sure that you are not interfering with your experiment in some way and therefore affecting your measurements. However, it can be difficult to measure quantities of interest meaning that plasma parameters must sometimes be inferred through the use of complicated models. An example of this is the reconstruction software EFIT<sup>[13]</sup> that uses input from many diagnostics to calculate internal plasma quantities that are difficult to measure directly due to the harsh environment conditions.

### 1.2.2 Active Instrumentation

Active diagnostics are defined by their probing capabilities. This means interfering with the experiment medium, whether through a physical probe or through the emission of some form of radiation into the plasma. Active diagnostics will generally consist of two major components; the probing instrumentation and a return signal detector.

While active diagnostics can usually measure experimental parameters more directly, they do have an effect on their environment. This can mean that the effect of the “probe” must be calibrated for through modelling, or that they can only be used in certain scenarios. For example, the field of Langmuir probe modelling is very large with lots of ongoing research even today, although the original Langmuir probe concept has been around since 1926<sup>[42]</sup>. An example of a diagnostic that uses a probing medium that is already present for tokamak operation would be the Motional Stark Effect (MSE)<sup>[43]</sup> diagnostic which utilises the Neutral Beam Injectors (NBI), used for injecting the majority of the heat in current tokamak designs, as its probing beam. While this means that the MSE diagnostic does not have secondary effects on the plasma that must be accounted for, it can only be used during periods of NBI operation.

### 1.2.3 Real Time Instrumentation

Real-time diagnostics describes a subset of instrumentation that analyse data for use in feedback control systems. This is useful for rapidly changing unstable systems where human reactions would be too slow to account for experimental changes<sup>[44]</sup>. The two core components of real-time instrumentation are the detector and the actuator. The detector will measure the quantity of interest, process the

data and then present this to the actuator. The actuator will then take action to reduce the difference between the reference and the actual value of the measured parameter. A simple flow diagram represents this concept in figure 1.7. Obviously this becomes a more challenging problem on shorter timescales and real-time diagnostics must usually make a balance between their latency and their bandwidth.

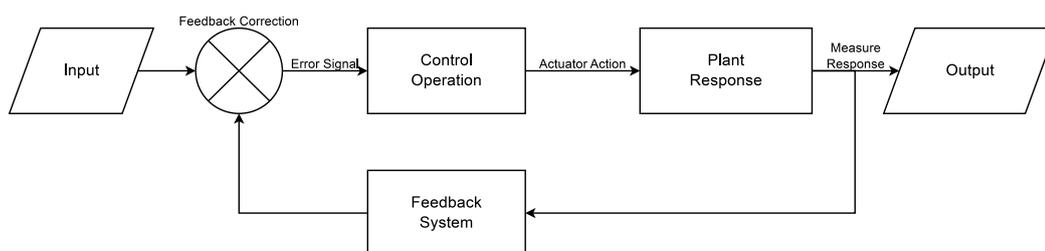


Figure 1.7: Flow chart of a basic feedback control system. The plant here refers to whatever system your feedback controller may be driving.

In real-time instrumentation systems there are two quantities of significance. Latency is the amount of time between a quantity being measured and the effect it is possible to produce from the associated actuator. The latency of a system must be low enough such that the actuator may reliably affect the experimental quantity it is trying to control. This must be carefully balanced with bandwidth, the amount of data that can be collected with a given amount of time, sometimes also referred to as bit rate or data rate (this is not to be confused with bandwidth when used in reference to electromagnetic detectors, in which case bandwidth would refer to the region of frequencies that the detector is sensitive to). The bandwidth of a detector must be sufficient to be able to measure the output of a system accurately enough to provide a reliable feedback signal. For many detectors a larger bandwidth or data rate requires a longer measuring time therefore increasing latency. These two quantities therefore can directly impact each other and experimentation and modelling are usually required to determine the correct balance.

#### 1.2.4 Field Programmable Gate Array Technology

Due to its flexibility and speed, FPGA technology is attractive for developing fast data acquisition instrumentation for use in real-time control and high data bandwidth applications. An FPGA is an Integrated Circuit (IC) that can be programmed and reconfigured continually after manufacture. This makes them very useful for applications that have unique but changing requirements. In research environments FPGAs are crucial for the continual development of diagnostics whose

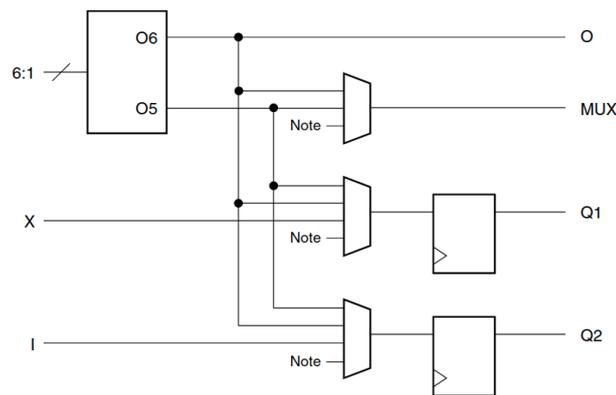


Figure 1.8: The configurable logic block (CLB) is the core component of FPGA programmable logic. A look-up table (LUT) allows the FPGA synthesiser software to implement what may be many operations into a single unit. The LUT output is connected to storage elements and the wider programmable logic (PL) system through multiplexers to facilitate greater flexibility. The outputs Q1 and Q2 are latched from O5 with signals X and I acting as clock enable or set and reset signals to give even more flexibility to the CLB. Figure taken from “User Guide UG574 UltraScale Architecture Configurable Logic Block” [7].

requirements may change over time, or for which a bespoke piece of equipment would be too expensive to commission. By developing instrumentation in-house, research organisations are able to reduce maintenance costs and have fine tuned control over diagnostic development.

Figure 1.8 shows a configurable logic block (CLB) which constitutes the majority of the programmable logic in FPGAs. The main component of a CLB is the 6 input look-up table (LUT). Programming an FPGA constitutes three main stages of code development, synthesis, and implementation. The synthesis of a Hardware Development Language (HDL) code requires a synthesis program to translate what has been written into a collection of interacting CLBs. The 6 input LUT allows the synthesiser to generate one LUT for what may be quite a few operations within the code. The CLB interacts with the greater network of programmable logic using multiplexed outputs and storage elements to facilitate complex operations with a single CLB. Also included in some CLBs is dedicated carry logic to improve arithmetic performance. After synthesis of the HDL code implementation software will use algorithmic methods to implement the logic network onto the required hardware. The implementation software then produces a file called a bitstream that can be loaded on to the FPGA.

While FPGAs can be extremely useful in certain situations, the steep learning curve and complexity of design compared to more mainstream software design means that the initial time investment can be quite high. This can be costly in time sensitive

industries and so the advantages of the technology and the requirements of the application must be weighed appropriately. Table 1.1 and the following section (1.2.4.1) compare the four main technologies currently available and give examples of where each technology can most successfully be implemented.

#### **1.2.4.1 Differences in Computing Devices**

Alternatives to FPGA technology exist and may be more appropriate depending on the given application. ASICs are pre-configured ICs that have been specifically designed for a particular application. While this makes them very efficient at their job, ASICs are difficult to develop due to their inflexibility and the complexity of designing ICs. ASICs become useful in high volume applications as the cost reductions gained from the individual unit after development will outweigh the high initial development costs.

Complex Programmable Logic Devices (CPLD) are similar to FPGAs in that they are re-configurable after manufacture. Much of the difference between modern FPGAs and CPLDs is a matter of scale. CPLDs generally consist of much fewer configurable blocks and have a much lower interconnect to logic ratio. While this somewhat constrains their flexibility it considerably reduces the complexity of their programming. CPLDs are also configured with some form of permanent storage solution such that they are programmed as soon as they are turned on. This is in contrast to FPGAs that store their configuration in volatile memory, and must be re-programmed after every power cycle. A CPLD may be preferable for applications that have limited access and reduced complexity.

As FPGA usage moves into the machine learning field, classical architectures like Graphics Processing Unit (GPU)s become direct competition. Due to their specialised nature and ease of programming, much of the work in GPU development for machine learning applications currently surpasses FPGA implementations. However, due to their flexibility the power-to-performance metrics for FPGAs versus classical architectures can be very competitive in certain applications, such as those with non-traditional bit-widths<sup>[45]</sup>.

Table 1.1: Table comparing and describing the advantages and disadvantages of different computing technologies and the applications for which they are appropriate.

Technology	Pros	Cons	Application
Central Processing Unit	Good for general compute applications	Is slower than other options for specific applications and does not handle custom IO	Personal computers, web servers, massively parallel High Performance Computing
Graphics Processing Unit	Efficient at processing standard width floating point numbers for specific applications	Highly specialised hardware only suitable for a few particular tasks	Graphics processing such as audio/visual editing and video games. Increasingly used for Machine Learning applications
Field Programmable Gate Array and Complex Programmable Logic Devices	Highly customisable to any specific process with very general IO for integration with varied peripherals	High setup cost and low efficiency compared to custom built circuits. Low accessibility compared to CPUs and GPUs	Custom instrumentation with real-time and low latency requirements. Specialised Machine Learning applications with non-standard data formats.
Application Specific Integrated Circuit	Highly efficient at the task it is designed for.	Low development time and high initial investment.	High volume digital signal processing (DSP) applications such as cryptocurrency mining or digital audio processing.

## 1.2.5 FPGA Development Core Concepts

While FPGA technology is extremely useful, the fundamentals of programmable logic hardware mean that some core concepts in traditional software development are no longer valid<sup>[46]</sup>. Some core functionality missing in FPGA development when contrasted to classical software is the lack of a division operator. As all mathematical operations must consist of bit shifts, additions, or subtractions, operators like division, logarithm and exponent become non-trivial to implement.

Depending on the application this can be circumvented through the use of look-up tables. Look-up tables are attractive because they can provide a low latency implementation of many practically difficult calculations. However, if the application requires high accuracy or range then they can become quite resource intensive. As seen later on in the thesis, the ideal solution is usually some mix of a logic implementation and look-up table to fit within the resource and timing constraints of the problem.

Another core concept in the programming of FPGAs, not usually seen in traditional software development, is that of pipe-lining. Pipe-lining is the process of separating a large operation, that has the requirement of completing instantly, into many smaller operations that happen sequentially. This is necessary due to the path length parameter of many operations within FPGA fabric. The path length of an operation is the amount of time necessary for an operation to get from its start point to its end point. For the programmed application to work correctly, each operation must get from start to finish in less time than one clock cycle. This is because the operation start value is updated every clock cycle. If the path length of the previous operation is longer than this, the start value will be undefined, propagating unknown values through the system. Figure 1.9 demonstrates graphically how the process of pipelining can reduce the longest carry chain and hence why it is a necessary concept during FPGA development.

An example of an operation requiring pipelining would be a large multiplication. Multiplications implemented on FPGA fabric consist of large bit carry chains. This means that the path-length of a multiply operation can quickly exceed the timescale of the given clock frequency. By separating the multiplication into smaller bit widths, the operations can now occur within the specified window and the values propagated through the system will be fully defined. The side effect of pipe-lining is of course increased latency due to the multiple smaller operations that must happen sequentially. Latency may be reduced through clever programming, such as a carry cascade, but it is up to the programmer to decide which solution will fit best to the application.

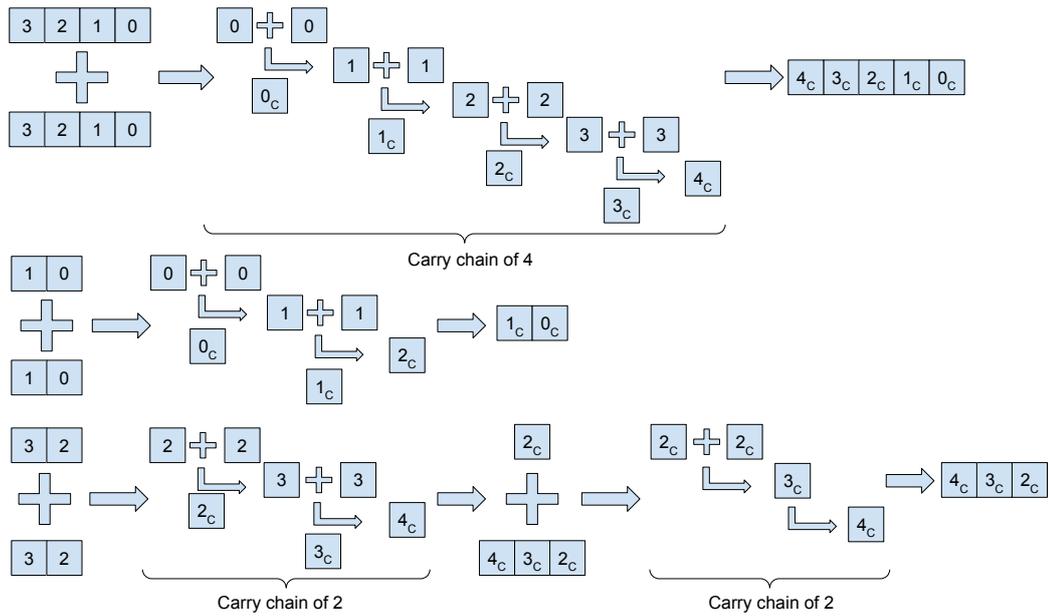


Figure 1.9: Graphical demonstration of pipeling an addition operation. The addition of two 4 bit numbers results in a carry chain of length 4 (top). By separating this into the addition of two 2 bit numbers we have introduced an extra clock cycle of latency but the largest carry chain is now only 2 bits (bottom). If a signal can only travel the length of three carry chains before the result is latched by the clock signal the result of the (top) operation would be uncertain.

This also shows one of the main advantages of FPGAs over the more conventional commercially available technologies, and that is the manipulation of non-standard bit widths. While the more commonly used CPU's and GPU's must use standard length memory arrays of exponent two length, the logic of an FPGA can be designed to use bit lengths of any size. This means that the logic resources can be used much more efficiently and has actually led to FPGA based technologies exceeding the watt-to-performance metrics of GPU's in machine learning applications<sup>[47]</sup> <sup>[45]</sup>.

Finally, when developing FPGA technologies we must introduce the concept of IP cores. IP refers to the intellectual property intrinsic to an FPGA module designed to accomplish a certain task. An FPGA design will consist of many interfacing IP cores that range in complexity. While normal software can be written in a similar manner (and some may argue this would be good practice) this is not the general case. FPGA IP modules range quite broadly in complexity, from full Finite-state Machines (see section 1.2.5.1) and complex I/O interfaces (see section 4.5.1) to very simple IP cores designed to separate or concatenate data words or multiplex data streams<sup>[7]</sup>.

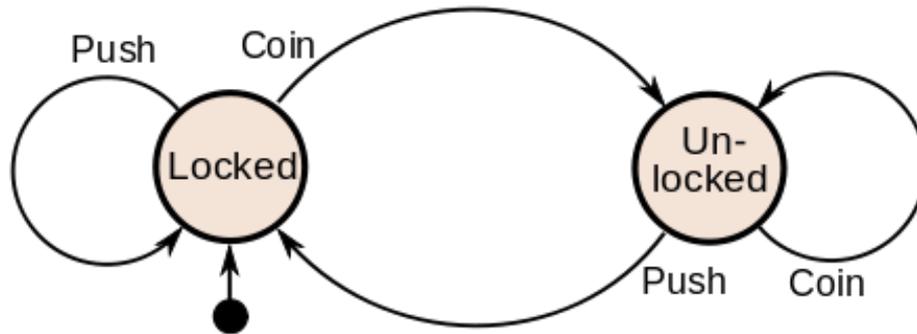


Figure 1.10: Representation of a Finite State Machine for a simple turnstile. These theoretical models can be used as a starting point for the implementation of control systems, communication protocols, or other automation applications. Used under the Creative Commons License, figure taken from “Wikimedia Commons”<sup>[8]</sup>.

### 1.2.5.1 Finite-state Machines

Finite-state Machines (FSM) are a mathematical model of an abstract computer used to help design solutions to problems in automation, communication protocols and other areas of digital electronics. Examples of FSMs can be seen throughout this body of work either explicitly, such as figure 4.12, or implicitly such as figure ???. FSMs are a core concept of FPGA development<sup>[46][48]</sup> due to their usefulness in designing iterative control systems and the ease of implementing an FSM using the “CASE” keyword provided as part of both the VHDL<sup>[49]</sup> and Verilog<sup>[50]</sup> development languages. A graphical representation of a simple FSM is shown in figure 1.10.

The behaviour of an FSM is defined by a list of states, of which only one can be held by the state machine at any time. While in each state, the FSM will sample various properties of the system it controls to decide the next state. Producing a theoretical state machine before delving into FPGA development helps to define critical decisions and system states of the problem being solved. The use of FSMs with the “CASE” keyword also helps to improve legibility of HDL code, meaning that the end system is easier to debug, maintain and adapt should modifications be required.

## 1.2.6 The Xilinx Zynq Architecture

The Xilinx Zynq architecture<sup>[51]</sup> is a core component of the work produced here as it has allowed diagnostic systems to be built that would have been significantly more complex just ten years ago. Before the introduction of Zynq technology FPGA based solutions that required Linux had to either operate a soft processor on chip<sup>[52]</sup> or communicate with a processor using the available FPGA IO. Zynq technology was therefore introduced to address embedded applications where a full processing unit and programmable logic are required, with a high-speed integrated data channel connecting the two. The System-on-Chip design of the Zynq architecture offers several advantages over how similar solutions would have been implemented before its introduction.

First, the processor component no longer requires space on the programmable logic and as the transistors per pound (£) are much higher for ARM processors than FPGAs the cost of the system is significantly lower for the same or better performance. Second, a logic chip specifically designed to run as a CPU is more efficient and capable of processing data faster than a processor implemented in PL. Lastly, memory interfaces between the processor can now be integrated on chip using high bandwidth specially built interfaces as opposed to the generic FPGA IO. This is especially significant with regards to the processor system memory as, within the Zynq architecture, it now has a direct interface to the FPGA allowing implementation of Direct Memory Access (DMA) protocols for fast writing and reading of data natively on the memory bus (see chapter 4.5), or slower protocols (as used in the Koheron SDK, see section 2.2.1) that are generally less complex to implement.

### 1.2.6.1 Advanced eXtensible Interface

The Advanced eXtensible Interface (AXI) is a high speed data transfer interface defined as part of the ARM Advanced Microcontroller Bus Architecture<sup>[53]</sup> [20]. As many of the Xilinx IP cores are built with native AXI4 interfaces it is important to understand how to use this protocol. By standardising the interface protocols for standard modules, designing with Xilinx based FPGAs is significantly simplified.

AXI4 interfaces come in three flavours:

- **AXI4** Designed for use with memory mapped interfaces, the full AXI4 interface provides functionality for reading and writing to different memory addresses. This interface type is generally used for IP core configuration as

the memory map functionality allows reading and writing to various configuration registers within the module.

- **AXI4-Lite** Similar to the AXI4 interface but for simpler and lower performance applications. The main difference between the lite and standard interface is the lack of support for bursting, a high speed method of quickly writing to many contiguous registers.
- **AXI4-Stream** The AXI4-Stream interface is a high speed data streaming interface. In contrast to the standard and lite version of AXI4 the stream interface operates in a continuous burst mode for high speed delivery of data.

The core of the standardised AXI protocol is the interaction of valid and ready flags. When a module has data prepared a valid flag is set high, notifying the receiving interface that data is ready to be read. Similarly, modules that receive data have a standard ready flag to be set high when the module is ready to be written to. It is only when both flags are high for a given interface that the data is marked as used and the modules can move on to the next data word.

For both the AXI4 and AXI4-Lite interfaces data is accompanied by a memory address that the data is written to or read from. The full AXI4 and AXI4-Stream interfaces support operating in burst mode, where both the valid and ready flags are held high as the address and data variables are incremented through their values.

### 1.3 Thesis Overview

The thesis presents work on the development of three separate FPGA based instruments for tokamak research. All three systems utilise the Xilinx Zynq architecture to implement programmable logic with a Linux system for instrument configuration and control. The fission chamber acquisition and processing system and real-time biasing Langmuir probe control system highlight the use of the open source Kohern SDK<sup>[54]</sup> for a low cost easily maintainable fusion diagnostic. While showing that this is certainly a viable route for future fusion instrumentation, both instruments highlight how the Zynq-7000<sup>[51]</sup> SoC requires some compromises in terms of system performance compared to more expensive hardware platforms.

The Synthetic Aperture Microwave Imaging diagnostic shows how the Zynq architecture can be applied to high performance instrumentation and the cutting edge of networking technology. However, the resulting instrument is significantly more complex than systems based on the lower end Zynq-7000 architecture, as well as added complexity in developing linux for a custom application.

---

# Fission Chambers for Neutron Flux Measurements

Fission chambers are used to measure the neutron flux of neutron emitting devices. This chapter will describe the design and implementation of a new fission chamber acquisition system for MAST-Upgrade. Test results from the National Physics Laboratory will be shown where measurements were made to verify correct operation. Preliminary results from implementation on the MAST-Upgrade tokamak are also presented which shows that some external issues due to electrical interference must still be solved.

## 2.1 Background

### 2.1.1 Neutrons in Fusion devices

Both a deuterium-deuterium and a tritium-deuterium fusion reaction produce a helium isotope and a neutron. As the charged helium nucleus is trapped by the magnetic fields, the neutron flux from fusion experiments is therefore a key direct measurement of the fusion rate. However, due to the small interaction cross-section with materials and neutral charge of the neutron it is a difficult particle to detect.

Neutron diagnostics transfer the kinetic energy of a neutron through collisions with some detection medium to create a detectable electrical signal. There are multiple neutron diagnostics that will operate on MAST-Upgrade, each with their advantages and disadvantages. Fission chambers allow the measurement of time resolved neutron flux and can be placed sufficiently far from a neutron source to assume that the neutron flux is isotropic. This is in contrast to neutron cameras that can deliver localised but spectrally resolved neutron data<sup>[55]</sup>, and activation

foils which are used for total flux measurements but are not time resolved<sup>[56]</sup>. Together, all three types of neutron diagnostics will be able to fully characterise the tokamak plasma as a neutron source.

Two fission chambers will be used on MAST-Upgrade, at different radial locations around the vacuum vessel, as demonstrated in figure 2.18. Offset radial positions of the chambers will verify that the isotropic approximation is reliable as the neutron beam heating sources create highly localised regions of fusion reactions. Should the central column be between a fission chamber and the fusion region the fission chamber may measure a significantly lower fusion rate than is being produced. However, it is worth noting that the original MAST tokamak operated with only one fission chamber and this was considered sufficient for fulfilling the dose monitoring safety requirement. Both fission chambers will feed their signal current to the same acquisition unit, where both will be acquired and processed in real-time, simultaneously. Data will then be sent to the MAST-Upgrade DATAC system (an overview of which is given in appendix B.2) over ethernet using a Flask (see 2.2.1) implemented webserver.

#### **2.1.1.1 Physics Case for Neutron Flux Measurements**

When operating a tokamak as a fusion experimental device the fusion product is obviously an important quantity that must be characterised. As each fusion reaction produces a neutron, the neutron flux of a tokamak plasma gives a measurement of the total number of fusion reactions occurring. The neutron flux can therefore be used as a measurement of fusion power and will be an important factor for real-time control in future fusion power plants.

As previously mentioned (see section 1.1) the fusion reaction produces a charged helium nucleus and a neutron. As the charged nucleus is trapped by the confining fields, and is necessary for heat deposition in a burning plasma, the neutron is the majority factor in extracting energy. In a fusion power plant a blanket of lithium, and usually some neutron multiplier<sup>[57][58]</sup>, will cover the exterior of the vessel such that when a neutron passes through it, it will deposit its energy and create tritium fuel. A power plant will need to keep precise control of the blanket temperature and a real-time measurement of neutron flux will help to achieve this.

For current research machines the neutron flux can help with diagnosis of other problems. Neutral beams are a way of depositing energy into fusion plasmas. In MAST-Upgrade the majority of the fusion reactions will happen during Neutral Beam Injector (NBI) operation. The neutron flux while the NBI is operating will

allow measurements of the neutral beam efficiency as well as validation of NBI modelling codes<sup>[59]</sup>.

Disruptions of the fusion plasma can also be accompanied by runaway electrons<sup>[60]</sup>. Runaway electron beams can sustain a current that is a significant fraction of the total plasma current with individual electron energies in the  $MeV$  range<sup>[61]</sup>. The beams can therefore cause significant damage to plasma facing components (PFC) and are a failure risk for power plant scale machines like ITER. When a runaway electron beam interacts with a PFC the resulting gamma ray and hard X-ray radiation can result in photo-neutrons. Spikes in measured neutron flux (a side effect of incident gamma rays on the detector) after a plasma disruption can therefore be used to characterise and investigate runaway electron mechanisms.

Also important to note is that in all these cases the neutron spectrum is important in diagnosing the plasma mechanisms involved in generating neutrons. Other neutron diagnostics, such as neutron cameras must therefore be used in conjunction with neutron flux monitors such as fission chambers. While the implementation for MAST-Upgrade utilises two reasonably sized fission chambers, the system planned for ITER will deploy many micro-fission chambers around the vessel for a more granular resolution of neutron flux<sup>[40]</sup>.

### **2.1.1.2 Safety Case**

As well as a physics case for neutron flux measurement a safety case exists in the form of structural integrity and a radiation bio-hazard. When a neutron collides with an atom inside a solid object it can knock the atom out of place. This causes defects within the crystalline structure of metals called dislocations. Dislocations can be either voids or extra atoms and cause the structure to become more brittle, decreasing the tolerable stress<sup>[62]</sup>.

Instead of colliding with an atom, a neutron may be captured by a nucleus. The new, heavier isotope may then undergo beta decay to transmute into a new element. The materials composing the fusion reactor will continue to be irradiated over the lifetime of the device and these impurities will build up. As before, the presence of a new element can change the properties of the material, reducing the integrity of the tokamak structure.

By recording the neutron flux over the lifetime of machines like MAST-Upgrade and JET, the material structure transformation can be extrapolated to power plant scale devices like ITER. This then advises material choices for different components and shielding requirements for particularly affected regions.

Also necessary for a nuclear site is radiation dosage monitoring. As a nuclear research lab with licensing for the use of tritium on site, the Culham Centre for Fusion Energy has a radiation dosage limit that cannot be exceeded. Neutron radiation is damaging to organic tissue and for this reason a bio shield consisting of a 1 metre<sup>[63]</sup> thick concrete wall surrounds the MAST-Upgrade reactor to protect people within the vicinity. For both these reasons it is essential to track the neutron flux so that the yearly dosage allocation is not surpassed and to ensure that the current bio-shielding arrangements are sufficient.

### 2.1.1.3 Fission Chamber Operation

The fission chambers used in this application are FC765 chambers designed and built by Centronic<sup>[64]</sup>. Each chamber consists of an aluminium casing containing two internal electrodes electrically insulated from the casing. One electrode provides the high voltage bias and the other is used for the current signal. The fission chamber can be used for pulse counting, current monitoring and Campbell mode<sup>[65]</sup> operation.

Both electrodes are coated in a fissile coating, in this case Uranium-235 ( $U^{235}$ ). Neutron flux incident on the chamber causes fission of the coating with the resulting high energy products inducing  $\beta$  activity in the filling gas (a mixture of Argon and Nitrogen). As these charged beta products are quickly drawn to, and neutralised, by the signal electrode a fission event produces a current pulse that can be detected. When individual current pulses can be distinguished it is therefore easy to infer that the number of pulses is proportional to the fission rate, and therefore the incident neutron flux.

While it is expected that fission events produce the majority of the signal electrode current, other factors also have measurable effects. Random alpha decays of the fissile coating in the chamber cause a low level background that the acquisition electronics must take care to discriminate against. Gamma radiation can cause the production of high energy  $\beta$  particles through direct ionisation of the containing gas<sup>[66] [67] [68]</sup>. This also causes detectable pulses above the background alpha current.

As the induced current pulses have a finite width, increasing neutron flux can result in an elevated DC current (due to overlapping pulses) from the signal electrode. Therefore when the neutron flux is sufficiently high the fission chamber acquisition electronics provides both a DC current and a Mean Square Voltage measurement (described in section 2.1.3).

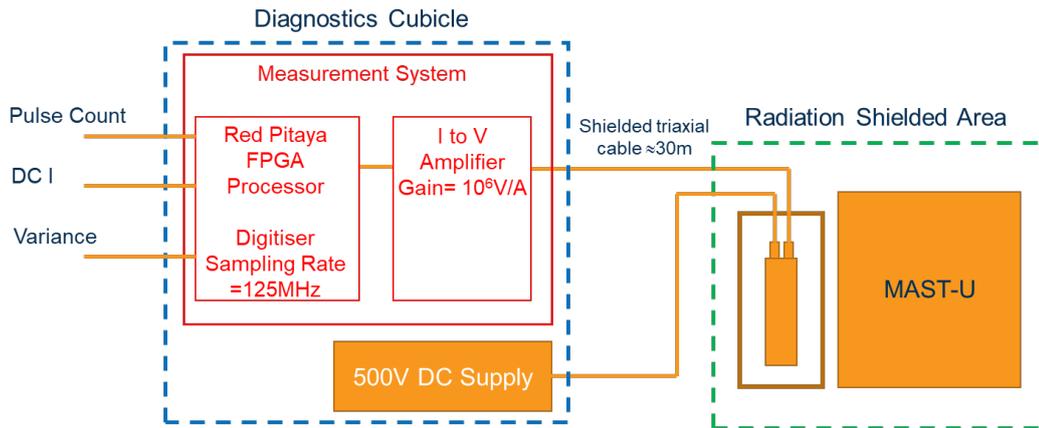


Figure 2.1: Schematic of the fission chamber data acquisition implementation on MAST-Upgrade with the schematic of the front end analogue amplifier shown in figure 2.2. Shown here is only one fission chamber. While the original MAST machine ran with a single fission chamber the intention is to operate two simultaneously with the same acquisition system for MAST-Upgrade

Figure 2.1 shows how a fission chamber and data acquisition and processing system will be implemented on MAST-U. The fission chamber current signal is fed into a pre-amplifier before being sampled by an ADC integrated into a Red Pitaya development board<sup>[21]</sup>. The integrated ADC samples the voltage out of the pre-amplifier at 125 MSPS with a 14 bit precision over a 2 volt range. The FPGA then processes this signal and uses standard AXI interfaces (see section 1.2.6.1) provided by the Koheron SDK<sup>[54]</sup> to transfer the data to an ARM based CPU. The Flask micro-framework is then used to run a light-weight web server which integrates into the MAST-Upgrade DATAC system (see appendix B.2) for instrument control and data retrieval.

## 2.1.2 Analogue Electronics Front End

Before the fission chamber signal current reaches the Red Pitaya for acquisition and processing it is amplified by the front end electronics. The amplification can be switched from  $10\times$  to an inverted  $100\times$  using a 2 pin header on the amplification board. Also designed into the front end amplifier is a loopback path for a DAC output from the Red Pitaya board. This allows for a source to be simulated, fed through the amplifier, and then sampled as a normal signal to test the electronics and processing system without a physical source. A schematic of the front end electronics is shown in figure 2.2.

The front end electronics is also designed to filter out frequency harmonics generated by the 50 Hz AC power input. Figure 2.3 shows a fast Fourier transform of

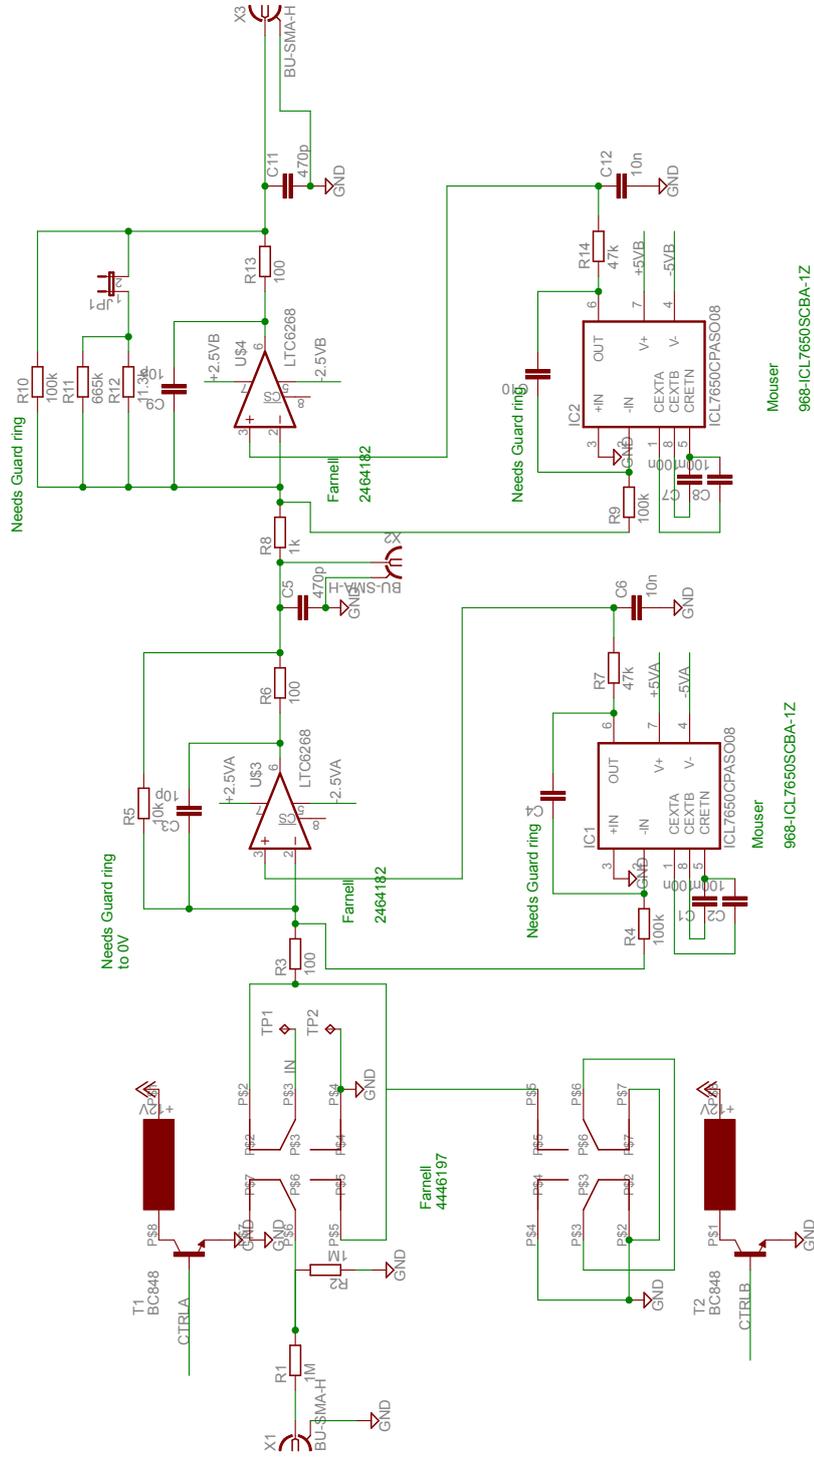


Figure 2.2: Electronics schematic of the front end electronics. The amplification factor can be switched between a gain of  $10\times$  and an inverted  $100\times$  with a two pin header on the board. The signal input can also be switched digitally from the front end inputs to a DAC output from the Red Pitaya. This feature allows testing of the system with a simulated source generated by the Red Pitaya FPGA board. Figure provided by the MAST-Upgrade Electronics Department

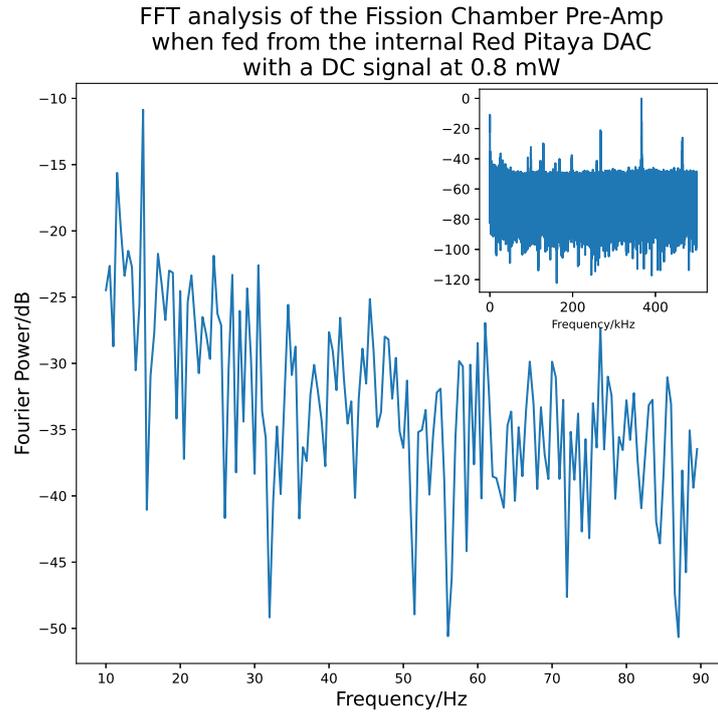


Figure 2.3: A fast Fourier transform of the DC current from the pre-amplifier when using an internally generated signal from the Red Pitaya DAC. No significant peaks are seen from 10 Hz to 90 Hz showing that the pre-amplifier is operating as it was designed. The inset shows the full frequency range available up to 500 kHz. A peak at  $\sim 366$  kHz with a modulation of  $\sim 99$  kHz can be seen as the highest power peak.

the fission chamber current and variance signal in a neutral environment. No significant peaks can be seen in the 10 Hz to 90 Hz range showing that the pre-amplifier is working as designed.

The FFT power has been normalised relative to the peak at  $\sim 366$  kHz shown in the inset. It is unclear where the main noise peak at 366 kHz is being generated as this is not a commonly used frequency. Peaks at 267 kHz and 465 kHz may be side-bands of this peak resulting from a modulation with the peak at 99 kHz. A possible source may be the AC (switch-mode) power supply within the acquisition unit.

To quantify whether or not these noise elements are likely to cause issues when deploying on MAST-Upgrade the pre-amplifier was fed a square wave at 61 kHz (the slowest possible frequency with the available input RAM to the DAC) and an amplitude of 0.5 V. Figure 2.4 shows the FFT analysis of the square wave with a signal power of 0.4 mW. A typical neutron induced current pulse, such as that

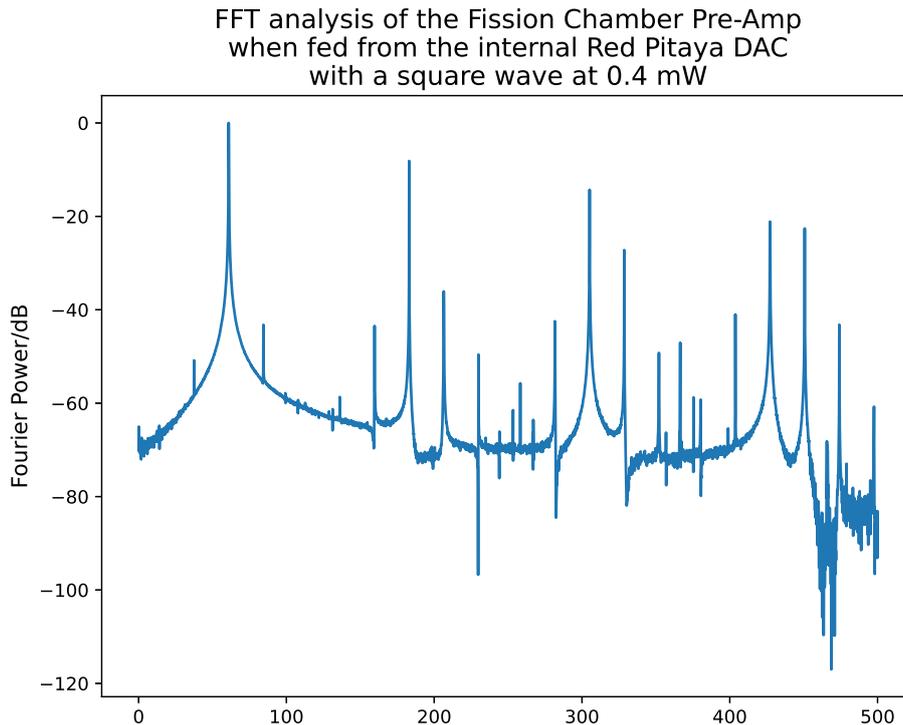


Figure 2.4: FFT analysis of a square wave input at 61kHz. The test square wave covers half the dynamic range of the voltage input and so should serve as a reasonable benchmark for the noise element generated by the 366 kHz signal. The square wave is not entirely matched causing many harmonics within the FFT. The 366 kHz peak has a power of less than -40 dB relative to the main 61kHz peak. It has therefore been determined that no matter the source of this noise element it will have little effect on the measurements made by the final system.

shown in figure 2.5, has a peak power of a few Watts. At a frequency of 366 kHz and a duty cycle (approximating the neutron pulse as a square wave) of 0.2, the average power of neutron induced pulses would be on the order of 0.5 W. The power of the 366 kHz signal seen in figure 2.3 is over 40 dB down on the 0.4 mW test square wave. Assuming a duty cycle of 0.5 the amplitude of the signal would be approximately 0.5 mV, below the 10 mV DC level of the alpha noise. The 366 kHz signal is therefore not considered as an issue for operation of the fission chamber acquisition system on MAST-Upgrade.

### 2.1.3 Campbell Mode

When a neutron strikes the uranium coated electrodes within the chamber, a  $U^{235}$  fission event occurs and the resulting products ionise the internal gas. A pulse event above the background noise of the chamber therefore signifies one neutron count. However, the generated pulse has some finite width (determined by the

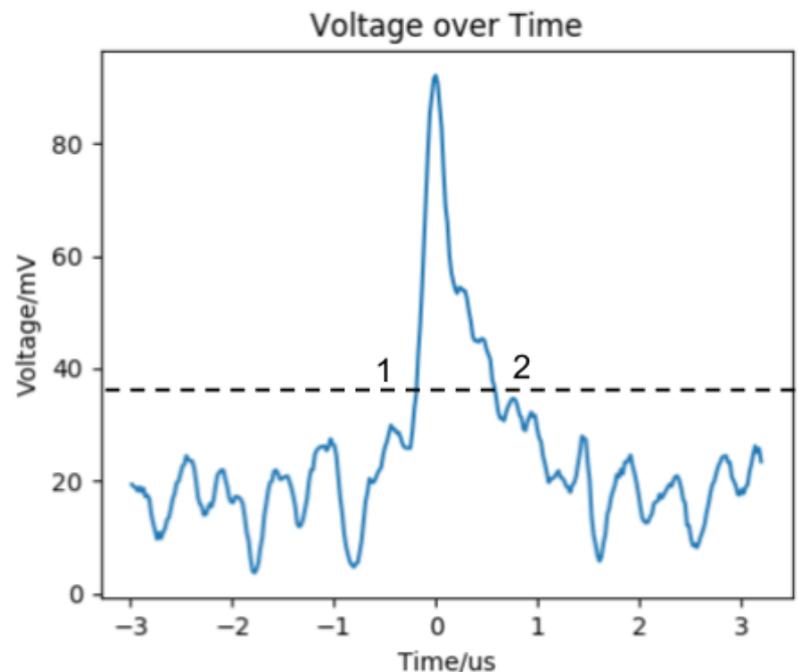


Figure 2.5: The pulse shape of a fission event within the fission chambers. An event can be identified by the significantly larger signal when compared to the background noise, also shown in the figure. The majority of the noise floor is generated by the background alpha events. Numbers and a dashed line are used for illustrative purposes with the peak detection program in figure 2.6

current decay characteristics of the fission chamber and the front end amplifier electronics), as shown in figure 2.5 which results in a phenomenon called pile-up. Pile-up is when many events occur faster than the pulse from an event can decay or the pulse can be distinguished by the digitisation system. This results in the events stacking, and a signal current consisting of a DC signal that is proportional to the number of pulses per unit time.

However, it is important to note that fission chamber events may not always be due to a neutron interaction. Gamma rays can also cause fission of the  $U^{235}$  coating and therefore erroneously increase the perceived neutron flux from the measured DC output current.

To mitigate the impact of gamma events, fission chambers can be operated in Campbelling mode. Originally developed for quantifying electronic noise of valves and circuits<sup>[65]</sup> the method has been successfully applied to fission chamber systems for neutron flux measurement<sup>[69][70]</sup>. Based upon the statistics of random events, the Campbell mode theory is used to calculate the number of neutron events once pile-up has caused the fission chamber signal current to appear as a DC signal.

The core postulate of the Campbell theorem is that electronic noise consists of randomly generated events within the circuit. When a fission chamber signal current is no longer operating in a pulse detection regime, that is pile-up makes individual pulses undetectable, the neutron and gamma ray fission events can be considered to constitute the majority of the “noise” of the signal. For the theorem to be valid for fission chambers then the noise event must dissipate over time and be agnostic to previous events.

$$S = \sum_i n_i s(iL) \quad (2.1)$$

In this case the signal can be written as in equation 2.1 where  $S$  is the cumulative signal at a time  $t$  of the  $n_i$  events that have previously occurred in interval  $L$ ,  $i$  intervals since time  $t_0$ .  $s(iL)$  is the discrete time signal of events that occurred in period  $iL$ . For events that decay this is obviously a nonphysical quantity as measuring  $s(iL)$  at a specific time  $t$  would result in sampling on a single time in the period  $iL$ . Therefore for this derivation can imagine  $s(iL)$  as storing the cumulative signal in the period  $iL$ .

We are interested in how these noise events affect the mean-square of our signal current. The mean current of our signal is relatively simple to derive and is shown in equation 2.2 where we have taken the limit of our signal  $S$  as the time interval  $L$  goes to 0 and  $i$  goes to infinity such that,

$$\begin{aligned} \bar{S} &= \lambda \int_0^\infty s(t) dt \\ t = iL, L &\longrightarrow dt, i \longrightarrow \infty \end{aligned} \quad (2.2)$$

and we have also used the fact that the average number of events ( $\bar{n}$ ) in time  $L$  is equal to the rate of events, notated by  $\lambda$ .

The other term of interest for the mean-square signal is the mean of the square of our signal  $S$ . We can construct the square of our signal using two time independent elements as in equation 2.3.

$$S^2 = \sum_i \sum_j n_i n_j s(iL) s(jL) \quad (2.3)$$

Similar to our derivation of equation 2.2 from equation 2.1 we must first determine the mean  $\overline{n_i n_j}$ , which has two distinct components.

$$\overline{n_i n_j} = \bar{n}_i \bar{n}_j + \overline{(n_i - \bar{n}_i)(n_j - \bar{n}_j)} \quad (2.4)$$

As before we know that the mean of  $n$  for a randomly occurring event is the probability of that event times the length of period  $L$ , such that the first term in equation 2.4 is just equal to,

$$\bar{n}_i \bar{n}_j = (\lambda L)^2 \quad (2.5)$$

The second term can be partitioned into the cases where  $i = j$  and  $i \neq j$ . For uncorrelated events random in time then probability theory tells us that<sup>[65]</sup>,

$$\begin{aligned} \overline{(n_{i=j} - \bar{n}_{i=j})^2} &= \lambda L \\ \overline{(n_{i \neq j} - \bar{n}_{i \neq j})(n_{j \neq i} - \bar{n}_{j \neq i})} &= 0 \end{aligned} \quad (2.6)$$

We can now assemble our resulting equation for the mean of the square of our signal,  $\bar{S}^2$ , as shown in equation 2.7.

$$\begin{aligned} (\bar{S})^2 &= (\lambda L)^2 \left( \sum_i s(iL) \right)^2 \\ \bar{S}^2 &= (\bar{S})^2 + \lambda L \left( \sum_i s(iL) \right)^2 \end{aligned} \quad (2.7)$$

Using the limits from equation 2.2 and the fact that the variance of a randomly distributed variable  $x$  is  $Var(x) = \bar{x}^2 - (\bar{x})^2$  we find that the signal variance due to a series of random events is,

$$Var(S) = \lambda \int_0^\infty s(t)^2 dt \quad (2.8)$$

From equation 2.8 we can see that the variance of signal  $S$  is proportional to both the frequency of events and the square of the magnitude of individual events  $s$ . As the current generated by a gamma event is less than that of a neutron event the variance signal inherently discriminates against gamma events. The measurement of this variance quantity is called Campbell mode and is now a core component of any fission chamber data processing system.

## 2.2 FPGA Control and Acquisition

The system presented here has been developed for implementation on the MAST-Upgrade fusion experiment and shares no components with the previous system. It has been built in-house and uses open source software development tools for easier cross-experiment implementation and maintenance. Improvements over the old system are increased time resolution, simultaneous operation of pulse counting, DC, and Campbell modes and reduced cost. The relative ease and reduced maintenance costs of a system developed in-house are also a major advantage over the original bespoke, proprietary system. This is due to the lack of proprietary software and the use of Commercial-Off-The-Shelf hardware components that be easily replaced.

### 2.2.1 Koheron Software Development Kit

The Koheron Software Development Kit is an open source Software Development Kit (SDK) for use with Zynq based System on Chip (SoC) devices<sup>[54]</sup>. Zynq based

SoC's are produced by Xilinx as all-in-one FPGA and CPU integrated systems<sup>[51]</sup>. Designing a chip like this allows Xilinx to provide standardised IP modules for FPGA to CPU integration (see section 1.2.6). However, to utilise a chip to its full potential it is still necessary to build a software system around it. As will be seen later (see section 4.4.3), building an operating system and the associated software to utilise a system is a non-trivial task.

The Koheron SDK<sup>[54]</sup> is an open source project to take Zynq development standardisation one step further. Designed for specific commercially available Zynq boards the project provides useful IP (such as ADC and DAC control, configuration registers and digital signal analysers) and a standard Ubuntu based operating system. To interface with the FPGA the Koheron SDK provides a Python package that can read and write registers over a network interface. This significantly simplifies the usual process of accessing registers via memory addressing using a low level language directly on the system. Also included is a web interface to create a network accessible GUI using commonly available browsers.

While there are many advantages to the Koheron SDK, the standardised workflow does mean that designs using it must fit within its constraints. Control and status registers must all be allocated in 32 bit blocks which can be difficult to change in further design revisions. Mostly due to the use case of the Zynq chip, Koheron SDK compatible boards are generally low power with relatively small amounts of programmable logic, so high performance designs must be light weight and well optimised.

In contrast to the Xilinx recommended workflow for development of Zynq-based systems the Koheron SDK significantly simplifies and speeds up the process of getting a system running. The Koheron SDK has been used for the development of both the fission chamber acquisition and Langmuir probe control and acquisition systems presented in this thesis (see section 3.2.2).

## 2.2.2 Red Pitaya Development Board

The Red Pitaya Development Board is a Xilinx Zynq-based platform<sup>[51]</sup> consisting of a dual core processor and programmable logic chip. The processor can support 32-bit Linux and has access to 512 MB of on-board RAM. Two ADC's and two DAC's with a sampling rate of 125 MHz are attached to the board for signal processing applications, and two 26 pin extension connectors provide slower ADC's, GPIO and board power. General IO accessible to the processor includes a 1 Gbit Ethernet and two USB-A ports.

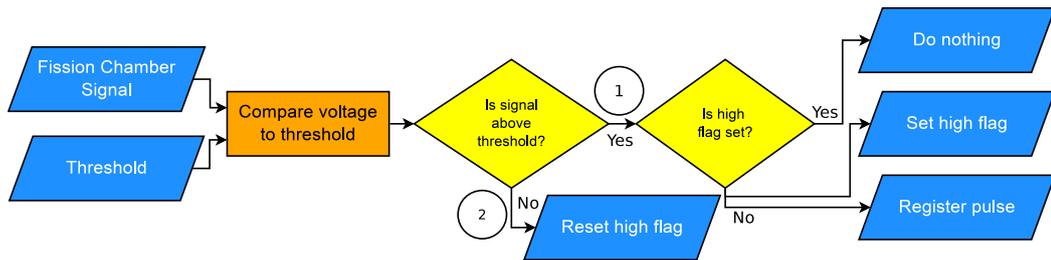


Figure 2.6: A flow chart of the low latency peak detection algorithm used for the Fission Chamber acquisition system. While this is a simple method of detecting peaks the low latency means that the  $1\mu\text{s}$  time resolution is preserved. Numbers are shown to pair decisions with the peak shown in figure 2.5.

The Red Pitaya company also provides software to operate the board using a browser interface. A part of this interface is a Python scripting environment, but more complex operations such as programming the FPGA are not possible through this interface. The Koheron SDK supports the Red Pitaya board and allows full access to adding user generated IP and control through the integrated registers.

### 2.2.3 Peak Discrimination and Detection

A low latency peak detection algorithm has been designed for operating the acquisition system in pulse counting mode. A threshold is set relative to a measured baseline and when the ADC voltage crosses this threshold twice (exceeding and then falling below) a peak is registered. A flow chart demonstrating the algorithm operation is shown in figure 2.6

As mentioned previously, issues with peak detection can occur due to pile-up (see 2.1.3). A digital filter could be used to more clearly define the peaks and suppress background noise, but this would introduce latency to the system. Several factors have determined that a digital filter will not be used as part of this system. From a practical point of view the Zynq 7010 SoC used in the Red Pitaya is a relatively small resource ADC and so only a short filter, specifically up to twenty coefficients, could be implemented. This would not be sufficient to implement a measurable difference in the peak distinctiveness. Secondly, a twenty point filter would introduce latency that would be almost a fifth of the total sampling time. While for the moment this is not necessarily an issue as the fission chamber operates significantly faster than most diagnostics, it is not good practice for the measurement output to be 20 percent of the sample time behind the sample timestamp. Last but not least, by looking at the measurements of the neutron flux on the original MAST machine and when the original fission chamber operated usefully in peak detection mode (that is the DC and Campbell modes could not accurately determine neutron flux)

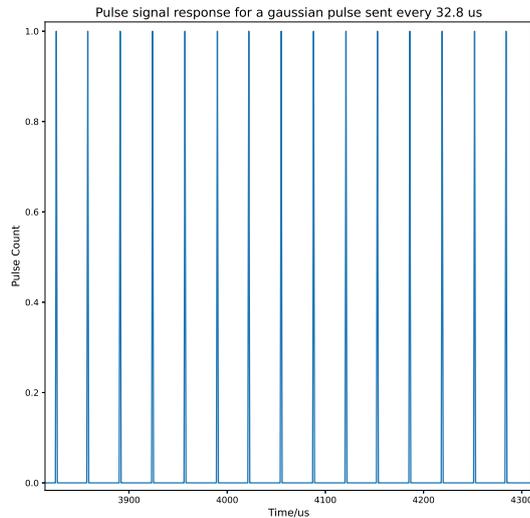


Figure 2.7: A Gaussian pulse (FWHM 106ns) is sent by the pulse simulator at a frequency of 30.5 kHz and detected by the peak detection signal of the fission chamber acquisition system.

it can be determined that greater cross-over between the three modes would not materially affect the precision of the system in measuring neutron flux.

While this peak detection algorithm is extremely simple, it does have the advantage of being extremely low cost with very little latency. Synthetic tests with a pulse generator, as shown in figure 2.7, and as shown later in section 2.3.1 in tests with low flux neutron sources, a sensibly set threshold is able to accurately discriminate and count voltage peaks above the average background.

A side effect of this method of peak detection is the ability to locate and measure the background noise of the acquisition electronics. As the peak detection algorithm detects crossing points between a threshold and a configurable baseline, when the threshold is zero a “pulse” is registered when the voltage signal crosses the baseline level. By sweeping the baseline across the range of the ADC the pre-amplifier noise profile can be measured. The noise from the detectors in a neutral environment can then be evaluated using the same method and comparing to the unconnected pre-amplifier profile. This gives a measurement of the alpha noise, the current generated by random fission events of the electrode uranium coating.

Examples of how threshold sweeping has been used to characterise the fission chambers can be seen in section 2.3.1 in both the thermal neutron and fusion neutron tests. Figure 2.13 shows how a threshold sweep changes when the chambers are operating in a pile-up regime. The width of the curve is significantly broadened

and the peak location is also shifted. Figure 2.16 shows how the threshold scans change when the fission chamber operates in a regime without pile-up. This helps inform what the pulse threshold should be set to in order to ensure that pulses from background events are not counted within the fission chamber pulse parameter.

## 2.2.4 Mean and Variance Calculation

As discussed in section 1.2.5 the implementation of mathematical operations such as divisions, exponents, and logarithms, in FPGA fabric can be logic intensive and introduce large amounts of latency. Both the DC value and the Campbell mode measurement of a 1 MHz signal from a 125 MSPS ADC require, at first glance, division and square root operations. The variance of a discrete random variable is shown in equation 2.9 where  $x_i$  is the current sample and  $\mu$  is the mean of the samples from time  $t_0$  to  $t_n$  (in this case  $n = 125$ ). In this form there is not only a division by  $n$  but a storage requirement of each sample  $x_i$  to find the difference with  $\mu$  that can only be calculated at time  $t_n$  (once  $\mu$  has been calculated).

$$Var(x) = \sigma^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2 \quad (2.9)$$

The variance can be expanded as shown in equation 2.10 to remove the storage requirement of each individual sample in equation 2.9.

$$n\sigma^2 = \sum_{i=1}^n x_i^2 - 2\mu \sum_{i=1}^n x_i + n\mu^2 \quad (2.10)$$

The sample mean can then be expressed as a sum, shown in equation 2.11, where  $n\mu$  now becomes the DC current signal deliverable.

$$\mu = \frac{1}{n} \sum_{i=1}^n x_i \quad (2.11)$$

Terms with respect to the mean in equation 2.10 can be expressed as shown in equations 2.12 and 2.13.

$$2\mu \sum_{i=1}^n x_i = \frac{2}{n} \left( \sum_{i=1}^n x_i \right)^2 \quad (2.12)$$

$$n\mu^2 = \frac{1}{n} \left( \sum_{i=1}^n x_i \right)^2 \quad (2.13)$$

And substituting these equations into the original variance expansion and multiplying by  $n^2$  results in equation 2.14.

$$n^2\sigma^2 = n \sum_{i=1}^n x_i^2 - \left( \sum_{i=1}^n x_i \right)^2 \quad (2.14)$$

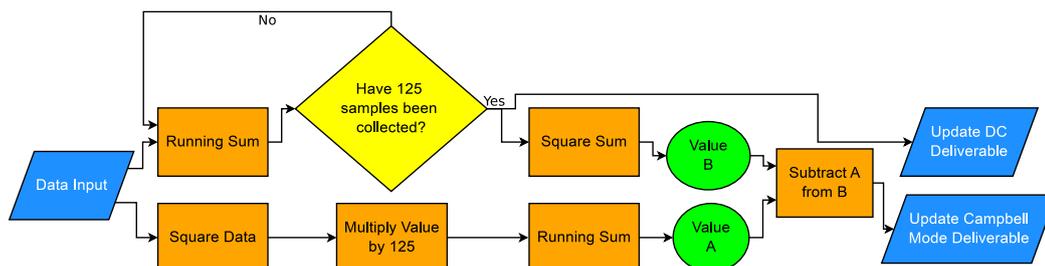


Figure 2.8: A flow chart showing the calculation of the DC and Campbell mode deliverables from the fission chamber acquisition system. As this method of calculating the current statistics from the fission chambers results in a low latency and light weight logic implementation it makes it very suitable for the Red Pitaya FPGA development board.

Figure 2.8 shows a flow chart following the equations described in this section. This method of calculating the DC current and Campbell mode of the fission chambers results in a low latency and light-weight solution suitable for implementation on low power boards like the Red Pitaya (described in section 2.2.2), and situations where high temporal resolution is required.

Implementation of these equations was not a trivial task. When designing FPGA instrumentation it is critical that the programmable logic on the chip “passes timing”. Once the design has been finalised, the implementation software that translates the Hardware Development Language (HDL) code into a bitstream that can be loaded onto the FPGA also calculates the path lengths of each operation in the PL design. If the path length is too long, the signal cannot get from one end to the other before it is latched by the end point clock, and the path has failed timing. This means that it is not guaranteed that the digital signal will propagate through the logic before it is sampled at the output. As the logic itself is unable to determine whether or not the value it is measuring is correct, the end point signal of the path becomes undefined. This could then potentially propagate through the logic resulting in undefined behaviour of the entire system

Within the Fission Chamber acquisition logic, the 14 bit ADC value cannot be processed as a single quantity, as the number of carry bits increases the path length beyond the speed limitations of the device. The ADC value must therefore be split into two segments which are processed in parallel and then recombined at a later stage. Snippets of the code used to achieve this can be seen in appendix B.1.

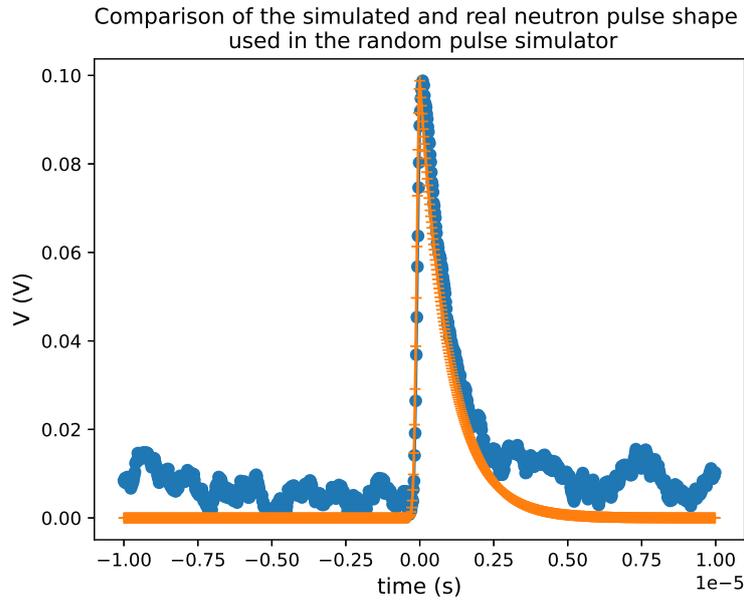


Figure 2.9: Comparison of the simulated pulse shape and a real pulse shape. The decay is slightly quicker for the simulated pulse, and the background noise has not been simulated however, this pulse shape has been sufficient for testing of the acquisition system.

### 2.2.5 Simulation for In-Situ Calibration

Also included within the programmable logic is a system for generating a fission chamber pulse simulation signal. The signal consists of fixed height pulses at random intervals, but with an overall user-defined frequency. This is achieved with a 32-bit random number generator and a threshold, such that if the generated number is larger than the threshold a pulse is generated and output over the DAC. The pulse shape can be user defined, but for the purposes of simulating the response to a fission chamber input this would always be similar to the pulse shape seen in figure 2.5. This pulse shape is approximated by the convolution of a Gaussian pulse and an exponential decay, with a comparison between the real and simulated pulse shown in figure 2.9.

This simulated pulse signal has been used to test the acquisition system when a neutron source has been unavailable. It has allowed characterisation of the analogue electronics (which can be seen in section 2.1.2) and in future, once the system has been implemented on MAST-U, will allow testing of new firmware implementations while installed on the machine.

## 2.2.6 Triggering and Clocking

The acquisition system can utilise both internal and external triggering and clocking sources. External sources are provided by the MAST-Upgrade Plasma Control System (PCS) to allow precise synchronisation with other diagnostics. The MAST-Upgrade PCS provides a 10 MHz free-running reference clock for all the installed diagnostics. An automatic switching module detects when the 10 MHz clock has been provided and automatically switches to it over the internal 125 MHz clock. A time-stamping module then uses this signal to generate a 1 MHz pulse that enables the data reader for the statistics module and the pulse counter.

The internal 125 MHz clock is still used to clock the ADC's and the fission chamber output modules, but aligning to the 10 MHz clock means that the sample period can be correlated with other diagnostics to the closest  $\mu\text{s}$ . Triggering can be provided in software using either the Python or Flask web server interface. The system accepts an "arm" signal from the MAST-Upgrade PCS to prepare the data acquisition for data retrieval, but this is merely a software arm as the system can be triggered in hardware and will collect and store data regardless. When implemented as part of the MAST-Upgrade DATAC system the ARM will act as a reset for any data stored from the previous shot that has not been collected.

## 2.2.7 Data delivery and instrument control with Flask

The MAST-Upgrade data acquisition system uses http *GET* and *POST* requests to view and set diagnostic parameters. Diagnostics implemented on MAST-Upgrade must therefore be running some form of interface for such requests. As the Red Pitaya used for the acquisition PC in the fission chamber system is low power, it was decided that a light weight web server implementation was necessary.

Flask is a web application microframework for Python<sup>[71]</sup>. This means it is a lightweight code library designed to allow developers to simply construct web applications using Python. Both its low resource usage and the fact that it allows a web server to be easily written and implemented using Python<sup>[72]</sup>, a commonly used programming language, meant that Flask was ideal for this application.

The web server acts as a black box to a diagnostic interface program running on a central data acquisition PC on the MAST-Upgrade network. This interface program can remotely view and set the control and status parameters of the fission chamber system as well as remotely trigger and reset the system.

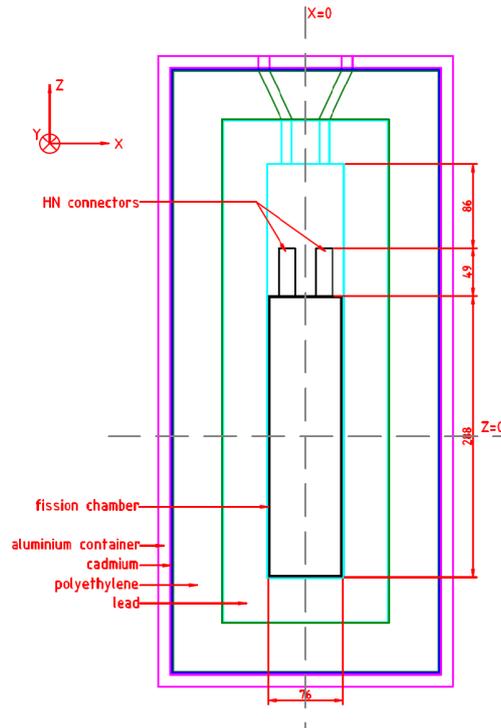


Figure 2.10: Schematic of the fission chamber shielding assembly. Shielding layers increase the cross-section of fusion neutrons with the electrode uranium coating and attenuate gamma rays to increase the signal to noise ratio. A final thin Cadmium layer absorbs neutrons with energies that would reduce the temporal resolution of the acquisition system.

## 2.3 Fission Chamber Calibration

The two fission chambers will be placed at different radial positions on the outside of the MAST-Upgrade vessel (see figure 2.18). To increase the cross-section of the fusion neutrons and to reduce the impact of gamma rays, each fission chamber has several shielding layers. A schematic of the shielding layers is shown in figure 2.10. The inner shielding layer constitutes 45mm of lead to block gamma rays. A sandwiched layer of 50mm thick high density polyethylene (HDPE) is used to thermalise incoming neutrons (attenuation to less than 1eV) to increase their cross-section with the uranium coating. The outer layer of 0.5mm of Cadmium blocks neutrons whose energies have been attenuated too much to maintain the temporal resolution of the system (this is explored in appendix B).

The shielding layers mean that not all fusion neutrons that are incident on the chamber housing will register as a current pulse, and the effect of pile-up also means that simple pulse counting cannot be used to measure the neutron flux. Instead, the chambers and associated shielding assemblies must be calibrated to known neutron

fluxes. In this instance calibration parameters have been achieved with sources of known flux, comparisons with activation coils during NBI commissioning and neutron sources with a known half-life.

### 2.3.1 National Physical Laboratory Calibration Procedure

The National Physical Laboratory (NPL) is a radiation and measurement lab, and as such has the facilities for calibrating fission chambers for use in nuclear applications. Measurements at NPL were made using thermal (less than 1eV) and Deuterium fusion neutrons (2.45 MeV). Both fission chambers were tested using the thermal neutron source to confirm an identical response in pulse counting mode. The original intention was to use the data gathered at NPL to calibrate the fission chamber response to a neutron flux. However, changes made to the system between when these tests were completed and when the system is implemented on MAST-Upgrade mean that these tests merely act as proof of system functionality.

The following tests are standard procedures formulated by NPL for the calibration of fission chambers. They are described in sufficient detail so as to understand that the fission chamber output is a result of incident neutrons, not as a description of novel experimental procedure presented in this thesis.

#### 2.3.1.1 Thermal Neutron Tests

The thermal neutron source emitted a flux of  $10^7$  neutrons  $\text{cm}^{-2}\text{s}^{-1}$ , as measured at the NPL facility. To test the pulse counting capabilities of the acquisition system at different flux levels a fission chamber was placed at different distances from the source. A diagram of the experimental setup is shown in figure 2.11.

Figure 2.11 shows a Van de Graaff generator that produces charged deuteron particles. These are then accelerated towards beryllium targets set in a graphite moderator. The graphite moderator thermalises the neutrons for detection by the fission chamber, which has been lowered into a cavity.

Deflection plates were used to control the number of deuterons that are incident on the beryllium targets, and therefore adjust the neutron flux to the fission chamber. To calculate the maximum flux available gold activation foils were placed at different heights within the cavity. Subsequent flux levels were calculated as a fraction of this maximum, corresponding to the reduced number of deuterons striking the beryllium targets.

Figure 2.12 shows the three calibration lines of the fission chamber output. As expected each parameter increases linearly with increasing neutron flux. The sig-

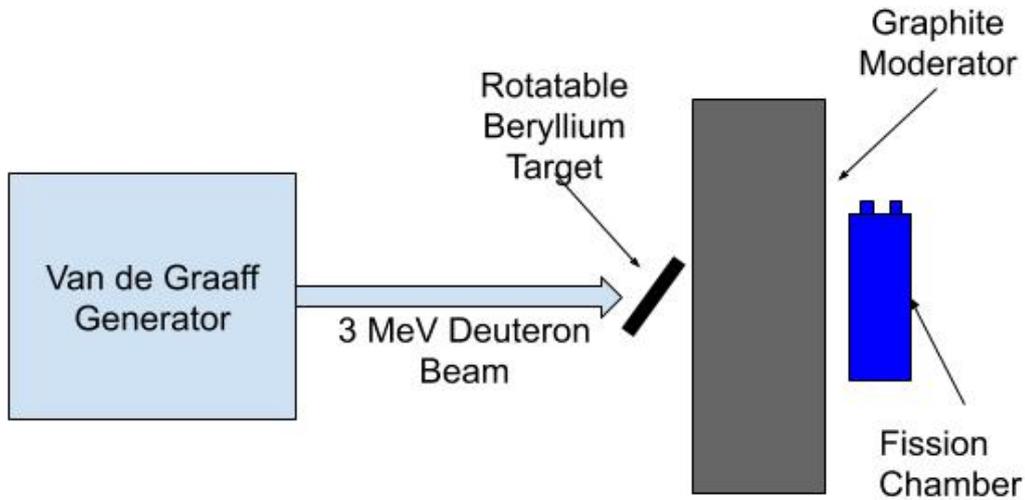


Figure 2.11: Diagram of the experimental setup used to test the pulse counting capabilities of the fission chamber acquisition system. Neutrons produced via a deuteron-beryllium interaction are thermalised by a graphite moderator and are then incident on the fission chamber. Neutron flux is attenuated by tilting deflection plates to reduce the number of deuterons incident on the beryllium targets.

nificant increase in DC current means that the pulse counting mode is operating in a pile-up regime. However, as the peak detection algorithm is counting crossing points of the voltage threshold, and from the Campbell theorem we know that signal variance increases with neutron flux, it is expected that the number of registered “pulses” should increase when pile-up is occurring.

Figure 2.13 shows how the number of registered peaks changes with different neutron flux. Due to the neutron pile-up the full-width half maximum of the curve increases with flux, which is in contrast to threshold scans of the 2.45 MeV source, as seen in figures 2.16 and 2.17, where the reduced neutron count does not result in pile-up.

Figure 2.14 demonstrates the change in pulse count at different thresholds with increasing neutron flux. Behaviour of the pulse counting system is as expected with a reduction in registered peaks as either threshold or neutron flux increases. The demand level is representative of the relative neutron flux. A future trip to NPL will use accurate measurements of the neutron flux to calibrate the fission chamber response absolutely.

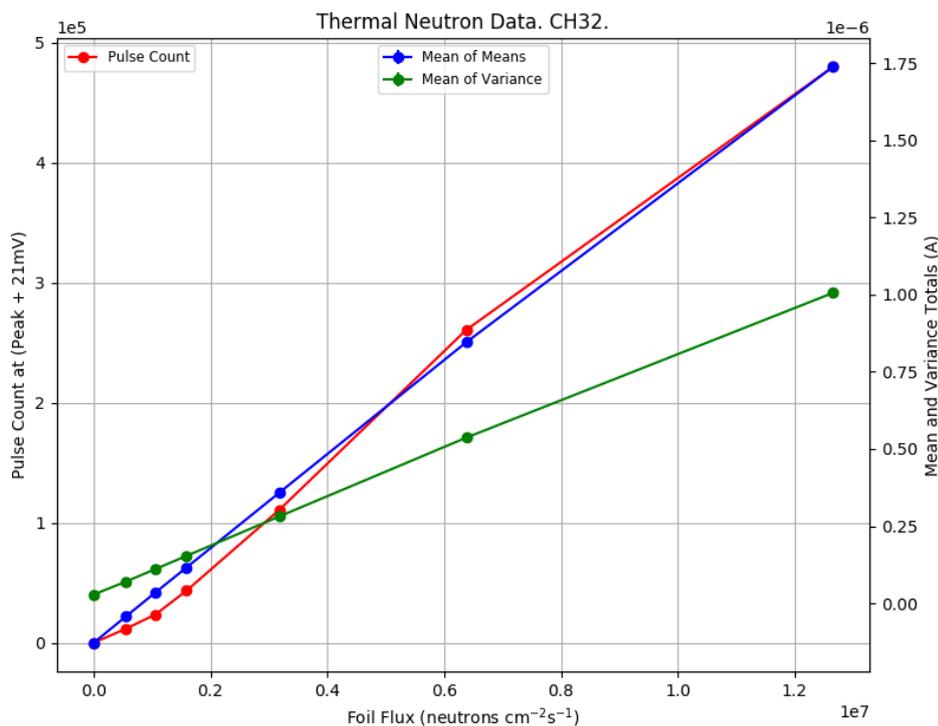


Figure 2.12: Mode comparison for the fission chamber acquisition system using a thermal neutron source. This can be used to calculate the calibration parameters between pulse counting mode, DC current mode and Campbell mode.

### 2.3.1.2 Fusion Neutron Tests

Fusion neutrons were produced using a Tritium-Proton interaction. Similar to the thermal neutron test, a Van de Graff generator produces 3.27 MeV protons which are fired into a tritium doped target. Neutrons are produced at 2.45 MeV equivalent to the energy of neutrons produced by the deuterium-deuterium fusion reaction. Figure 2.15 shows a cartoon of the experimental setup to test the response of the fission chamber and shielding assembly to fusion neutrons.

The fission chamber and shielding assembly were placed on a remotely movable arm so the electrode angle of incidence and the distant to the source could be varied without entering the experimental hall. The chamber height was fixed such that the electrode centres and neutron beam are aligned. To reduce the effect of scattered neutrons the hall was clear at the neutron beam height and the minimum distance to a wall from the fission chamber was 6m. A long counter<sup>[9]</sup> (or flat response detector, an energy independent neutron detector) was used to precisely measure the neutron flux at each radial distance that measurements were made

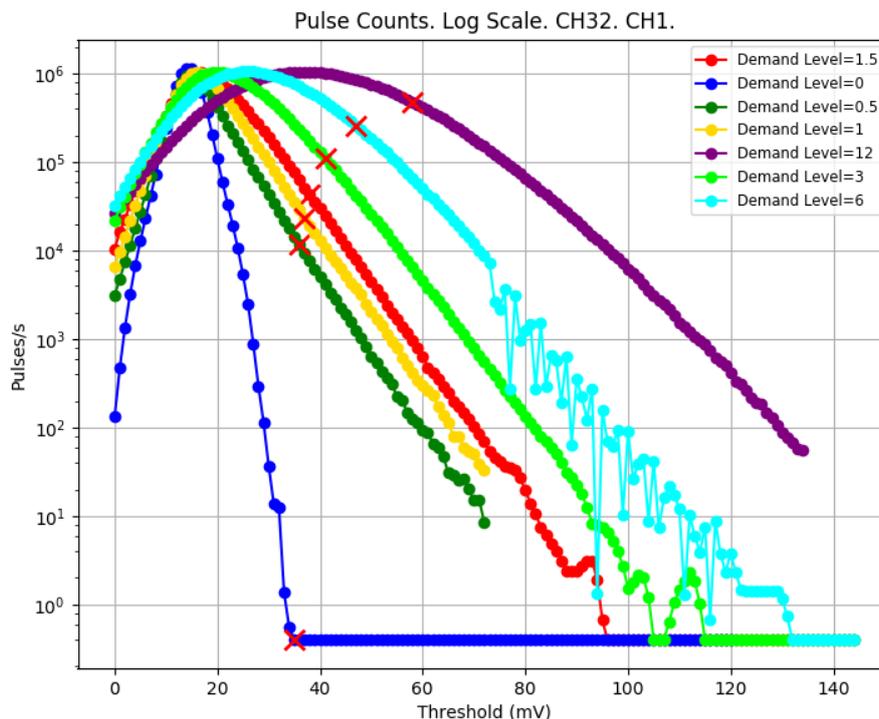


Figure 2.13: A threshold scan at different neutron flux rates shows an increase in high voltage crossing points with increasing neutron flux. As these results will not actually be used for fission chamber calibration linear demand levels are shown as a relative value to demonstrate how the pulse counting mode changes at different neutron fluxes. The red cross marks 25mV above the pulse count peak to more clearly demonstrate how the maximum peaks detected changes with neutron flux demand level.

with the fission chamber assembly. Using a shadow cone<sup>[73]</sup> to block the direct neutron beam component to the long counter the scattered neutron flux could be measured separately to the direct neutron flux. Only the full neutron flux was measured at a distance of 79 cm as the separation between source and detector was not large enough to be able to fit the shadow cone. Three radial distances were used to acquire calibration measurements, 79 cm, 230 cm and 260 cm, as well as a measurement at 260 m with the source off as a control.

Figure 2.16 shows the threshold scans for the 2.45 MeV source with the fission chamber at different distances. The reduced count rate, due to the shielding and lower fluxes available with the source meant that pile-up did not occur, as seen by the stationary position of the count peak. Figure 2.17 shows a fine scan of the difference in pulse counts between the 2.45 MeV source and fission chamber back-

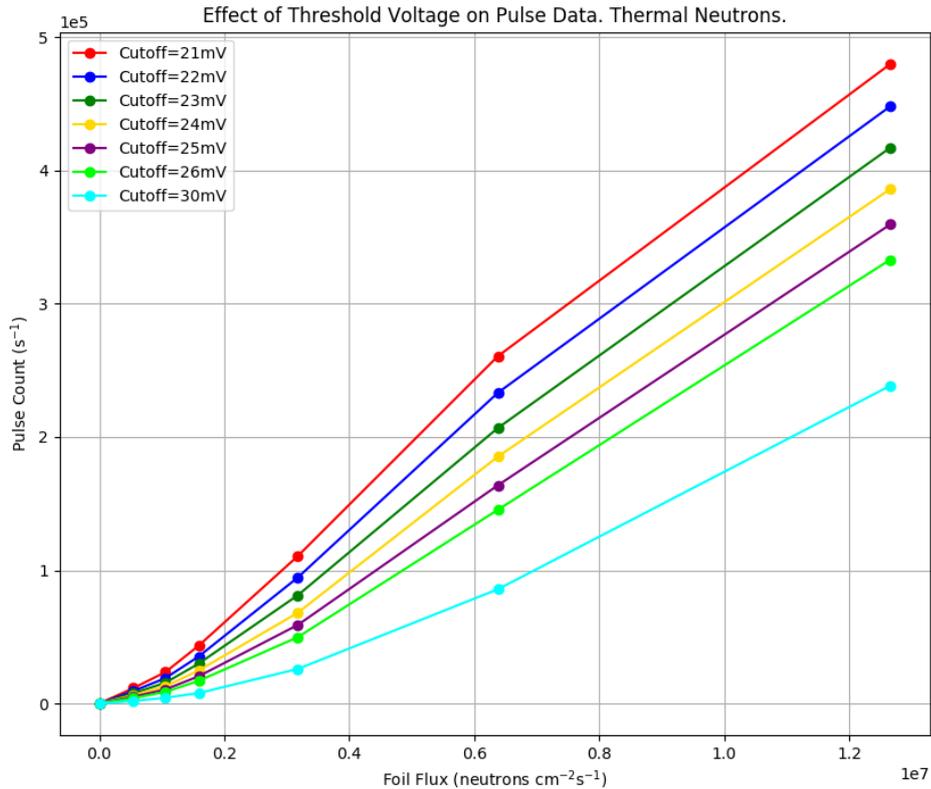


Figure 2.14: Number of crossing points registered for a given neutron flux at different peak detection thresholds. When measuring crossing points without pulse pile-up small threshold values would be expected to register the same number of crossing points with divergence as the threshold increases. This figure shows that there is an increase in crossing points at all threshold levels, indicating that all measurements were made with pulse pile-up occurring.

ground at an assembly distance of 79 cm to the source. The number of registered peaks per second is significantly lower than the thermal neutron source and meant that comparisons could not be made to the DC and Campbell modes. To correctly characterise the cross-over point of the three modes a source producing  $10^6$  counts per second at threshold of 35 mV would be ideal. This would correspond to approximately 1 registered pulse per  $\mu\text{s}$  sample.

## 2.4 Implementation on MAST-Upgrade

Before the start of the experimental campaign at MAST-Upgrade, a fission chamber and accompanying acquisition system was installed during the commissioning of the Neutral Beam Injectors in November 2020. This will act as a first test of the acquisition system and in conjunction with activation foils placed inside the vessel

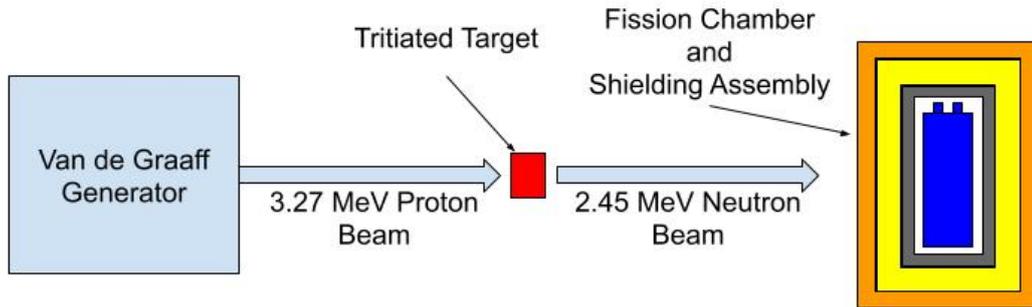


Figure 2.15: Diagram of the fusion neutron test setup at NPL. The fission chamber is placed in its shielding assembly and mounted on a movable arm. The height is fixed such that the proton beam and electrode centres are aligned. To measure the neutron flux the long counter<sup>[9]</sup> was placed in the same position as the fission chamber at each measurement location. A shadow cone was placed between the long counter and the neutron source to block the direct component and measure the scattered neutron flux.

and around the experimental hall gives an idea of the fission chamber response to future fusion plasmas. While a full plasma will not be present during the neutral beam commissioning, a low density plasma is necessary for ensuring that the NBI system does not damage the vessel and PFCs. The majority of fusion neutrons are expected to be produced by deuterium-deuterium reactions between the beam particles and adsorbed deuterium in the beam dump.

Figure 2.18 shows how the fission chambers will be mounted outside the vessel wall of the MAST-Upgrade reactor. Figure 2.18 also shows a CAD model of the fission chamber mounting system. The fission chamber is elevated to the plasma mid-plane by a free standing strut that supports the 150 kg weight of the chamber and shielding.

The proximity to the magnets means that currents induced in the signal and power cables are detected by the acquisition system as electrical interference. Figure 2.19 shows the raw DC current values for shot 42663 where significant interference is visible from 0.8 seconds before time zero until the end of the shot. Insets on the figure show the artefacts magnified with a frequency of approximately 6 kHz. The lower left insert shows the applied voltage to a poloidal magnet which aligns well with the start of the interference. There is little to no interference visible on the unconnected second channel so it is safe to assume that this signal is due to electrical pickup of the signal cable.

Figure 2.20 shows a Fast Fourier Transform (FFT) of two different time periods

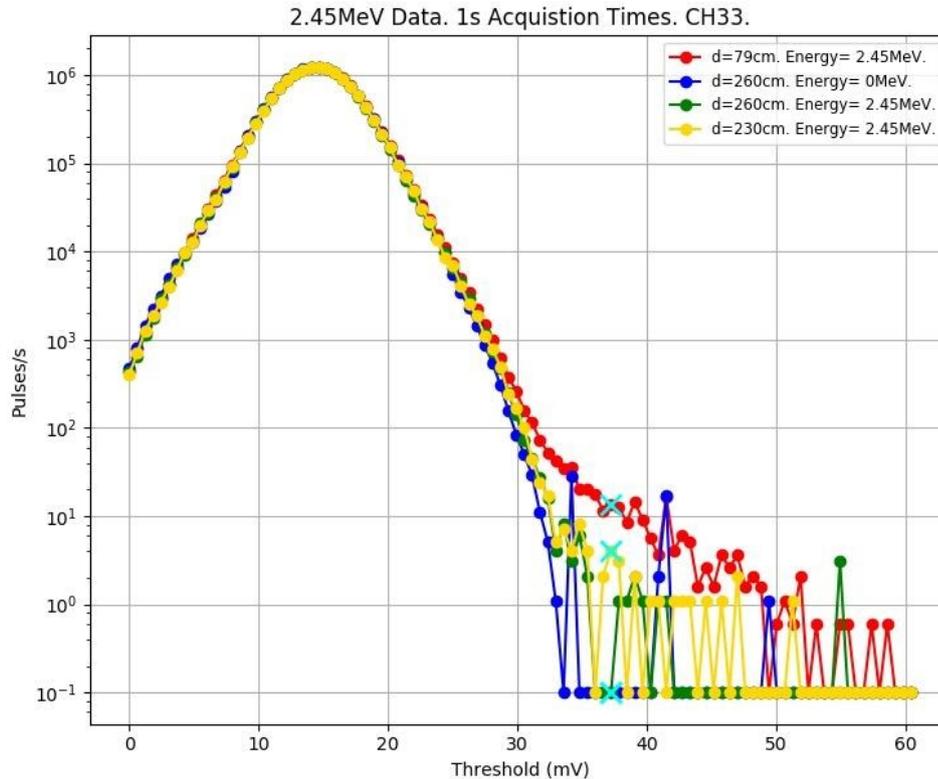


Figure 2.16: Threshold scan using the fusion neutron source at different distances. Operation is unlikely to be in pile-up mode as the threshold count peak and FWHM is similar for each measurement distance while the crossing point count in the tail decreases with distance from the neutron source. The pulse count axis shown here is logarithmic so that the difference in threshold crossings detected is more distinct at high thresholds. While there is little difference between the number of pulses registered at 260 cm with the neutron source on or off, more pulses are clearly registered at a distance of 230 cm with more again at 79 cm.

during the fission chamber acquisition. The FFT shows that the main noise element is a 200 kHz signal which has been modulated with the 6 kHz signal that is shown in the inset in figure 2.19.

It is currently unclear how the magnet coil voltage couples to the fission chamber signal line. One theory is that the triax cable shielding has been grounded in such a way as to be coupled to the magnetic coil power supply. This is obviously a significant issue as the periodic artefacts occupy a significant portion of the signal bandwidth. The oscillations also saturate the variance parameter making it unusable for neutron flux measurement purposes. The noise elements also only occur when specific power supplies are on, and as all the power supplies are located outside the main torus hall this also would indicate some issue with mixed earths. However, while this is an interesting problem, an investigation of the noise source

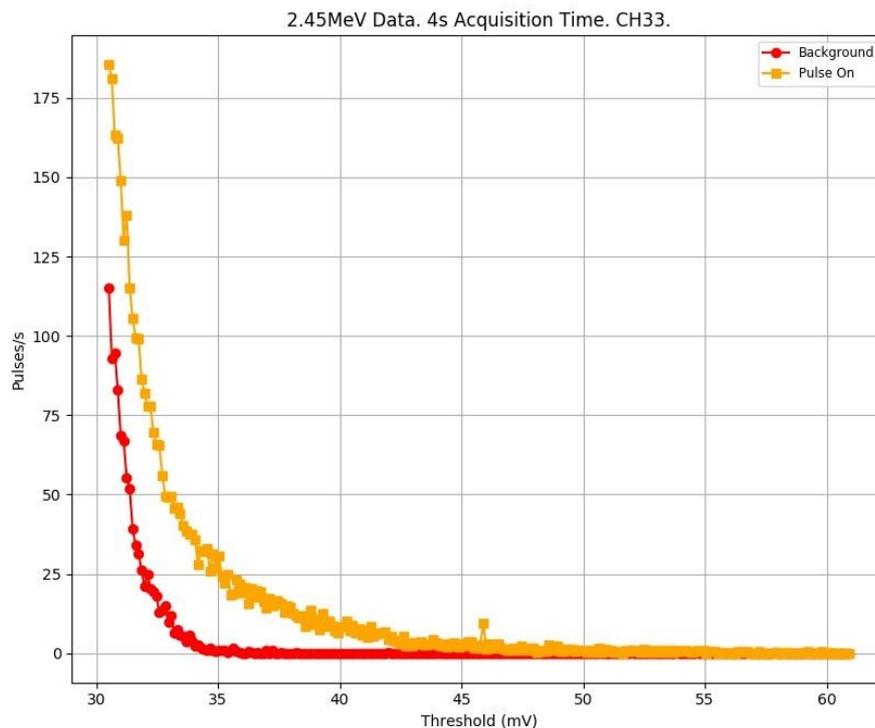


Figure 2.17: Close-up comparison of a peak threshold scan comparing the fission chamber background with fission chamber response to 2.45 MeV neutrons. Acquisitions were taken over 4s at a distance of 0.79m

and solutions is outside the scope of this work.

## 2.5 Summary

A real-time processing and data acquisition system has been designed for two fission chambers. The fission chambers will be mounted for operation during the upcoming MAST-Upgrade campaign for the purpose of neutron flux measurements. This is a novel system and improves on the original implementation by an order of magnitude in time resolution and the capacity to run in pulse counting, DC current and Campbell mode simultaneously. An original logic implementation of the variance calculation is used to provide a low latency, real-time Campbell output at 1 MHz.

Tests at the National Physical Laboratory have shown that the data acquisition system responds linearly to applied neutron flux as expected. While the neutron flux threshold for pile-up is very low, an appropriately selected current pulse threshold is able to measure differences in small applied neutron fluxes. Installation

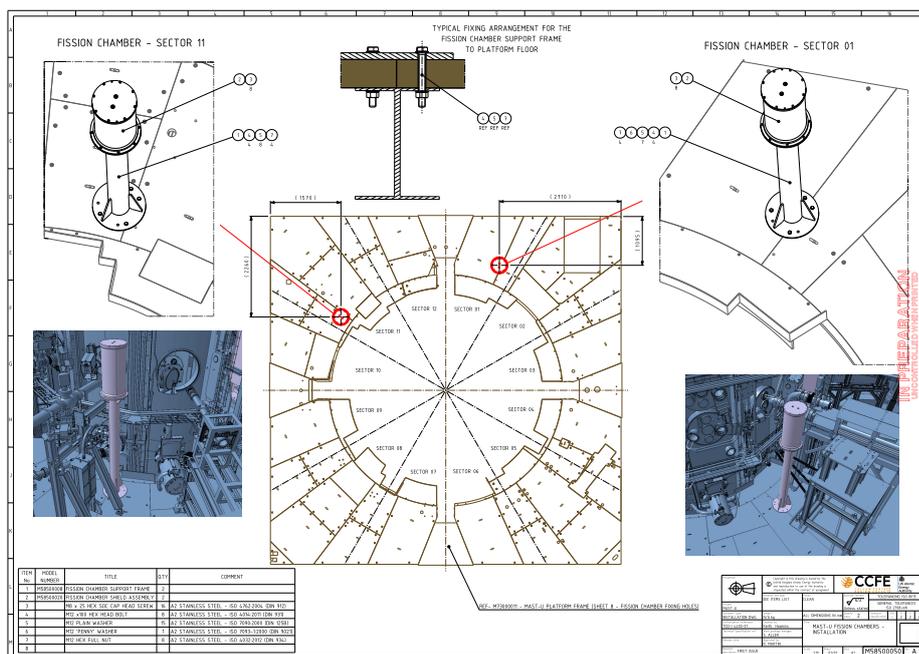


Figure 2.18: CAD drawing and schematic of the fission chambers installed outside the MAST-Upgrade vessel. The chamber is mounted on a free standing strut to elevate it to the plasma mid-plane. Drawing provided by the MAST Drawing Office.

on MAST-Upgrade has shown extreme interference through some interaction with the magnetic coil power supplies.

### 2.5.1 Further Work

Further work on the fission chamber acquisition system will concentrate on providing a real-time output of the neutron flux. This will be useful for real-time monitoring and could potentially be incorporated into the MAST-Upgrade Plasma Control System to assist in disruption prediction<sup>[60]</sup>. It would also be helpful to further develop the web interface to provide a network accessible GUI to adjust the simulated pulse counting. The ability to view the pulse counting mode response in real-time would also help with diagnostic control. The relatively low cost and agnostic nature of the acquisition and processing system means that implementation on other machines, with different fission chamber hardware, should also be a fairly simple procedure.

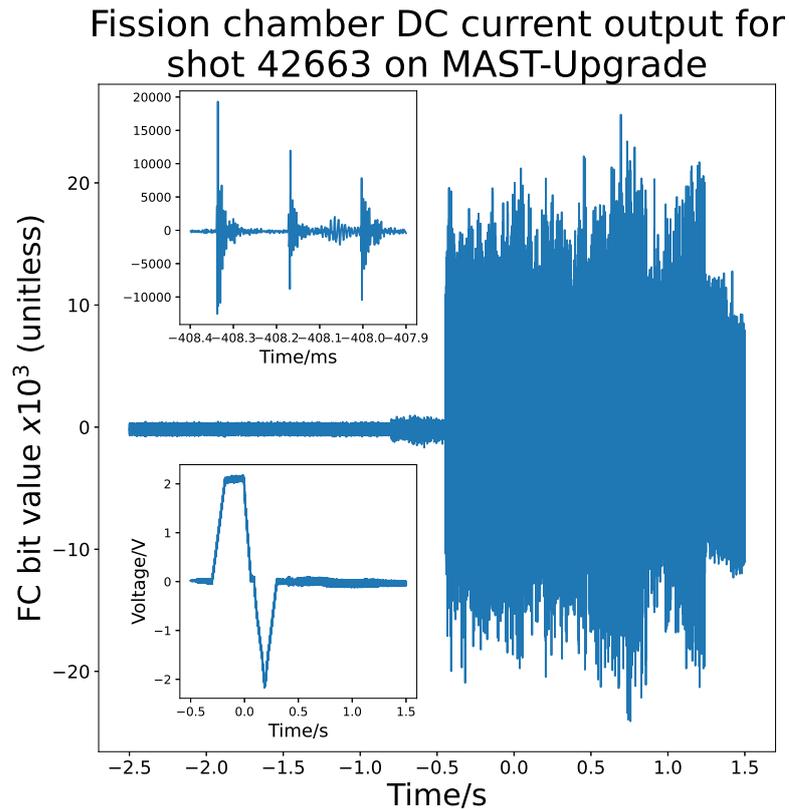


Figure 2.19: Raw DC current signal from the fission chamber installed on MAST-Upgrade. It can be inferred from the applied voltage trace that interference occurs through some coupling with the magnetic coil power supplies. The lower left inset shows applied voltage trace to a poloidal magnet. The magnet system start the voltage ramp at -0.5s which aligns well with the interference occurring in the DC current measurement. Insets show a magnified view of 3 individual periodic artefacts with a frequency of approximately 6 kHz.

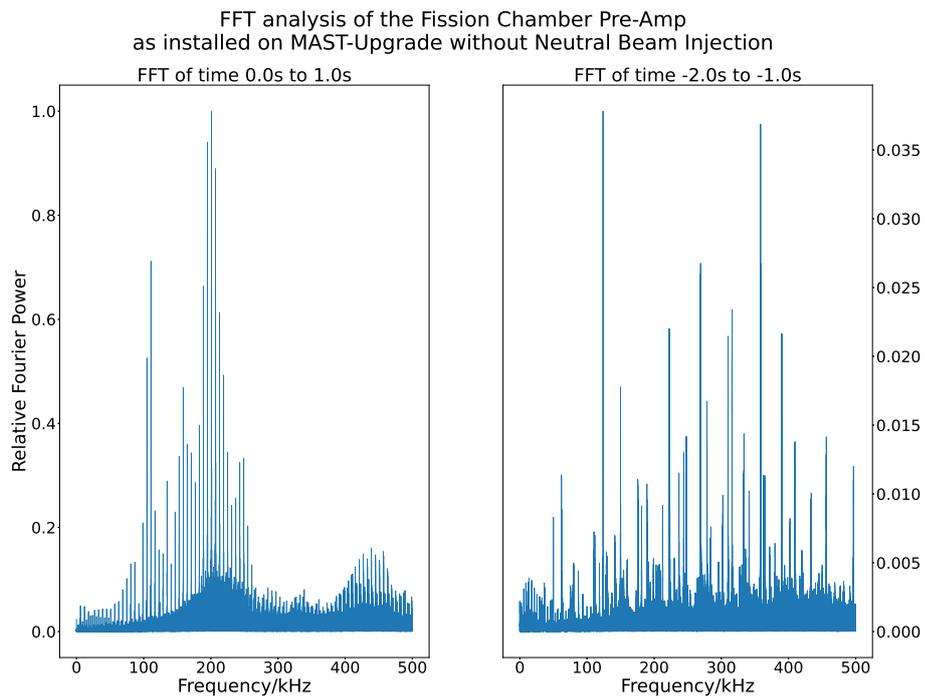


Figure 2.20: Fast Fourier transform of the raw DC current shown in figure 2.19. Different time periods have been shown to differentiate between the coil power supply being off and on. While the DC current shows a 6 kHz artefact as being most prominent the FFT indicates that this is a modulation with a more powerful 200 kHz signal.

---

# Digital Implementation of a Mirror Langmuir Probe Control System

This chapter covers the development of a prototype control system for real-time biasing of Langmuir probes. The prototype control scheme was developed in partnership with the Plasma Science and Fusion Center (PSFC), a department of Massachusetts Institute of Technology (MIT). While this is theoretically a continuation of the original MLP project<sup>[10]</sup> the control scheme is very different, and as such the "Mirror" aspect is no longer an accurate descriptor of the system. A more accurate description of the FPGA implementation of the MLP would be as a time-multiplexed triple probe.

Presented here is the initial proof-of-concept development consisting of FPGA based control and test systems. This control system will allow the measurement of electron temperature and density in plasma experiments with significantly higher time resolution than most current systems. Its implementation using FPGA technology will also facilitate porting the control system to other plasma devices. The results presented in this chapter show that the core concept of the control system, the capability to alter the Langmuir probe bias in real time to fit fluctuating plasma parameters, is fully functional.

Part of this work has been published in Review of Scientific Instruments<sup>[74]</sup>.

## 3.1 Background

Much of the research on fusion experiments is concentrated in understanding the physics of plasma transport both in the core and the edge. To study transport at the edge it is necessary to take measurements of plasma temperature, density and potential at different edge depths and at the time scales of the turbulence

driven transport<sup>[75][76][77]</sup>. As a thin conductor, Langmuir probes can measure the conductive properties of the plasma and how they change with varying plasma parameters, making them an essential diagnostic for fusion experiments<sup>[78][79]</sup>. A scanning Langmuir probe can be inserted into the plasma edge and its position adjusted to take measurements at different depths. However, current techniques for Langmuir probe operation are unable to resolve fast transport phenomena due to either their low temporal resolution (single probe operation), or low spatial resolution (multi-probe techniques)<sup>[80]</sup>.

The MLP biasing system, first deployed at Alcator C-Mod, was the first Langmuir probe control system to make measurements at both high temporal and high spatial resolutions<sup>[78]</sup>. By implementing real-time voltage biasing of a Langmuir probe the analogue MLP was able to measure plasma parameters at a rate of 1.1 MSPS. The Alcator C-mod team used the MLP to great effect in characterising the quasi-coherent mode (QCM) of H mode discharges due to higher levels of detail available<sup>[78]</sup>. This has since contributed significantly to studies of the physics of the plasma edge.

The system implemented on Alcator C-Mod consisted of a complicated arrangement of analogue electronics that made implementation and maintenance of the diagnostic very complex. By introducing a digital electronics version the Langmuir probe control system described here shows that similar results can be achieved with significantly simplified setup and maintenance costs in both time and money.

### 3.1.1 Langmuir Probe Theory

A Langmuir probe is a conductive piece of material that can interact with the electric fields generated by a plasma. When working with fusion research plasma both surface mounted and scanning probes (mounted on a reciprocating arm) are useful. Due to the high temperatures, scanning probes can only operate in the plasma edge region without worrying about probe ablation or melting. The response of a Langmuir probe exposed to a plasma is shown in figure 3.1. By varying the voltage bias applied to the probe a characteristic I-V curve can be generated which is then used to calculate plasma parameters<sup>[42]</sup>.

$$I_{LP} = I_{sat} \left( e^{\frac{V_P - V_F}{T_e}} - 1 \right) \quad (3.1)$$

The equation to describe the ion saturation region of the Langmuir Probe (LP) characteristic is shown in equation 3.1<sup>[81]</sup>, where  $I_{LP}$  and  $I_{sat}$  are the probe current and the ion saturation current respectively,  $V_P$  and  $V_F$  are the probe voltage and floating voltage respectively, and  $T_e$  is the electron temperature of the plasma. The full LP characteristic curve can be seen in figure 3.1.

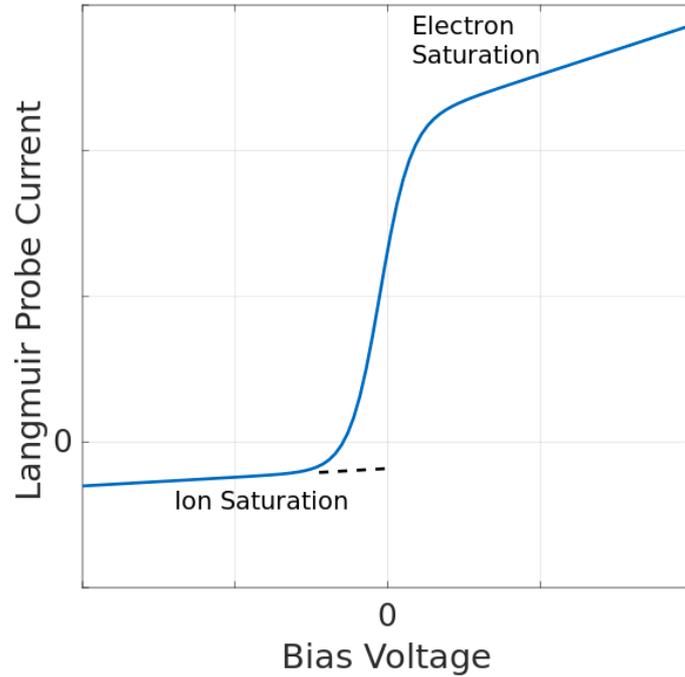


Figure 3.1: Figure showing characteristic I/V curve of a Langmuir probe. Regions of ion and electron saturation as well as the floating point voltage must be measured to accurately characterise the curve.

To obtain the plasma parameters  $T_e$ ,  $V_F$ , and  $I_{sat}$  equation 3.1 is fitted to a curve obtained by scanning the bias applied to the LP. It can be seen from equation 3.1 that the electron temperature ( $T_e$ ) and the floating voltage ( $V_F$ ) set the bounds of the characteristic, and as these quantities are unknown then a traditional probe control system must sweep the probe bias over a wide range of voltages to ensure the curve is fully characterised. This sets an upper limit to the speed of conventional single probe control systems as many more data samples than is strictly necessary must be collected and analysed to extract the required plasma parameters. A note must be made that while this does mean a low temporal resolution compared to what is theoretically possible, the many samples means that error analysis can be performed, which is not possible with the three bias of the MLP system.

To compensate for the sweep speed of a single probe some LP systems utilise a double or triple probe configuration<sup>[80]</sup>. However, this results in a low spatial resolution and complexity in the analysis as electromagnetic effects from adjacent probes must be taken into consideration.

### 3.1.2 Physics Case for Fast Langmuir Probes

Cross-field transport in the Scrape-Off Layer (SOL) of magnetic confinement fusion devices is an important area of study, as it is believed to be a major contributing factor in setting the SOL width and in turn contributing to divertor power loads, fuelling and plasma wall interactions<sup>[82][83]</sup>. A major field of study within cross-field transport is the nature of filaments.

It is known that filamentary structures contribute significantly to both material and energy transport out of the core<sup>[84][85]</sup>, but their formation and exact dynamics have still been elusive. A lot of progress has been made with the use of fast cameras<sup>[86][87]</sup> but, due to the similar radiation broadening effects, the separate temperature and density dynamics are entangled and therefore still unknown. As thin filaments are ubiquitous in both turbulent and H-mode (where turbulence is generally suppressed) plasmas, their characterisation would be useful in understanding cross-field transport<sup>[84]</sup>

A filament of diameter 1 cm with a velocity of 1 km/s<sup>[76]</sup> would cross the surface of a thin probe within the order of  $10\mu s$ . To fully characterise the change in temperature and density that would result from this type of filamentary transport a temporal resolution on the order of 1 MHz would be necessary. To examine these filamentary structures it is therefore necessary to have a fast, single probe system. These types of measurements could then be used with fast camera data to fully characterise these thin filaments and therefore provide insights into the resultant transport.

### 3.1.3 The Analogue Mirror Langmuir Probe

The primary feature of the original mirror Langmuir probe was the generation of real-time analog signals corresponding to the electron temperature, ion saturation current and floating current of a single electrode in contact with a tokamak plasma. A three stage voltage waveform drives the plasma electrode and a “Mirror Langmuir Probe” electronic device. By actively altering the output current using a feedback loop from the single electrode the MLP can mimic the electrical response on the real Langmuir probe in contact with the plasma. The feedback loop incorporates scaled quantities equivalent to the plasma parameters and these are iterated to generate the correct output voltage, that being a voltage equal to that measured from the plasma. This effectively means that the MLP is calculating a “best fit” to the plasma I/V characteristic in real-time.

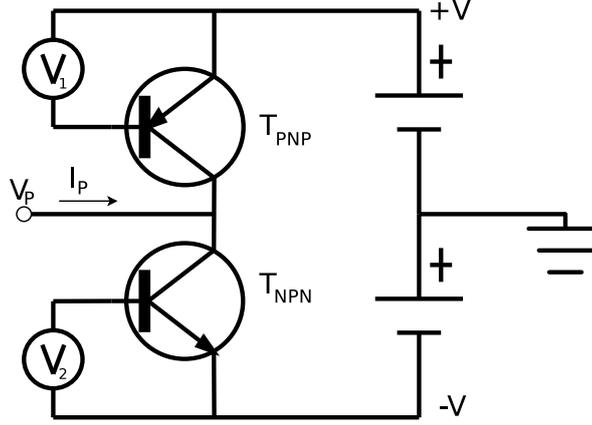


Figure 3.2: Simple circuit diagram of the transistor arrangement used in the original MLP diagnostic. This configuration of a PNP and NPN transistor is used to emulate the plasma response to an applied bias  $V_P$ . Figure reproduced from “Mirror Langmuir probe: A technique for real-time measurement of magnetized plasma conditions using a single Langmuir electrode”<sup>[10]</sup>

The core concept of the MLP circuit is that by appropriately scaling the inputs to a bipolar transistor, its I/V response is equivalent to that of a single electrode Langmuir probe. The emitter current ( $I_C$ ) of a bipolar transistor with base emitter bias ( $V_{BE}$ ) is given by the Ebers-Moll equation<sup>[88]</sup>:

$$I_C = I_S(e^{\alpha V_{BE}} - 1) \quad (3.2)$$

Where the saturation current ( $I_S$ ) is a characteristic of the transistor and the factor  $\alpha$  is dependent on the electron temperature in the transistor. At room temperature  $\alpha$  is approximately equal to 40 Volts<sup>-1</sup>. An arrangement of pnp and npn transistors in series, biased with the appropriate voltages, can reproduce equation 3.1.

The transistor configuration shown in figure 3.2 is designed to emulate the plasma response to an applied probe bias  $V_P$ . Summing the transistor currents using the Ebers-Moll equation (equation 3.2) results in:

$$I_P = I_S^{npn} \exp[\alpha_{npn} V_2] - I_S^{pnp} \exp[\alpha_{pnp} V_1] \quad (3.3)$$

Where we have used the fact that the base emitter component,  $e^{\alpha V_{BE}}$  is generally very large and so the subtraction of 1 can be ignored. As  $V_P$  is the probe voltage then  $V_1$  and  $V_2$  must be dependent on the three plasma parameters that describe the Langmuir Curve,  $I_{sat}$ ,  $V_F$  and  $T_E$ . To start, an initial voltage  $V_i$  can be set across transistor  $T_{pnp}$ . This would result in the base emitter potential of  $T_{npn}$  summing to include our other parameters.

$$\begin{aligned} V_1 &= V_i \\ V_2 &= V_i + V_B + \frac{V_P - V_F}{40T_e} \end{aligned} \quad (3.4)$$

Equations 3.4 show the substituted voltages, where the factor of 40 is used to balance the transistor  $\alpha$  value at room temperature. The voltage  $V_B$  can be adjusted to select a specific saturation current  $I_{sat}$ , with  $V_F$  and  $T_e$ , being held constant depending on the plasma parameters being modelled.

By expressing  $V_i$  and  $V_B$  in terms of the saturation currents of our transistors we can form an equation for  $I_P$  similar in form to the Langmuir probe equation. Using

$$V_i = \ln\left\{\frac{I_{sat}}{I_s^{pnp}}\right\}/\alpha_{pnp} \quad (3.5)$$

and

$$V_B = \ln\left\{\frac{I_{sat0}}{I_s^{nnp}}\right\}/\alpha_{nnp} + \ln\left\{\frac{I_s^{pnp}}{I_{sat0}}\right\}/\alpha_{pnp} \quad (3.6)$$

the transistor pair equation in 3.3 simplifies to a version of equation 3.7

$$I_P = I_{sat0}\left\{\left[\left(\frac{I_{sat}}{I_{sat0}}\right)^{\frac{\alpha_{nnp}}{\alpha_{pnp}}} e^{\alpha_{nnp} \frac{V_P - V_F}{T_e}} - 1\right]\right\} \quad (3.7)$$

We can see that setting the  $\alpha$  parameter for both transistors to 40 reduces equation 3.7 to exactly the Langmuir probe equation.

While the concept of the transistor pair may seem simple, the values for  $V_F$ ,  $V_B$ , and  $V_i$  must all be continually adjusted using a feedback loop comparing the real plasma potential to  $V_P$ . This arrangement of analogue electronics is quite complex and an in-depth description of the original Mirror Langmuir Probe control system, and helpful figures for understanding its operation, is available in, “Mirror Langmuir probe: A technique for real-time measurement of magnetized plasma conditions using a single Langmuir electrode”, *B. LaBombard and L. Lyons*<sup>[10]</sup>. While an extremely useful system<sup>[78]</sup>, the complexity has kept it from being implemented on other machines. It would therefore be preferable to show that similar levels of performance are possible using modern technology for simplified application to other plasma devices.

## 3.2 Mirror Langmuir Probe Implementation Using FPGA Technology

By implementing a real-time biasing control scheme on FPGA technology similar results to those given by the original Mirror Probe system can be achieved with reduced implementation costs and complexity. Due to the embedded system-on-chip nature of many available FPGA systems the hardware requirements are considerably simplified and integration with other diagnostic systems as well as customisation of the device becomes more manageable.

### 3.2.1 The Mirror Langmuir Probe Equations

Section 3.1.1 described the I/V characteristic of a plasma and how characterisation of this relationship measures the temperature and density. To obtain these parameters quickly the sweep speed of the curve must be maximised by sampling as little as possible. To characterise a curve the minimum number of points that must be sampled is equal to the number of degrees of freedom. In the case of equation 3.1 this would be three. To fully capture the physics of the curve these points must measure the ion saturation current ( $I_{sat}$ ), the floating voltage ( $V_F$ ) and the curvature above the floating voltage (set by  $T_e$ ). However, as these values change in time the optimal bias voltages that measure them must also change in time.

The real-time adjustment of these control values can be simplified somewhat by capacitively coupling the probe driver to the probe electrode. This ensures that one state can always be a zero state, with the voltage difference equal to the floating voltage. To ensure that the driver is capacitively coupled to the electrode then the positive bias and negative bias states must draw equal but opposite currents.

We can now express what the probe bias should be relative to the temperature expressed in electron volts. By setting our negative and positive bias so that they draw the same current we can form the following relationship from the LP characteristic (equation 3.1).

$$\exp\left(\frac{V_+ - V_F}{T_e}\right) + \exp\left(\frac{V_- - V_F}{T_e}\right) = 2 \quad (3.8)$$

Where  $V_+$  and  $V_-$  are the positive and negative probe biases respectively. To make sure that the full range of the I/V curve is captured we can see that the difference between  $V_+$  and  $V_-$  is large relative to  $T_e$ .

$$V_+ - V_- = 4T_e \quad (3.9)$$

By substituting equation 3.9 into equation 3.8 we find a value for  $V_+$  and  $V_-$  relative to  $T_e$  and  $V_F$ .

$$V_+ = 0.675T_e + V_F \quad (3.10)$$

$$V_- = -3.325T_e + V_F \quad (3.11)$$

Where  $V_F$  is approximated as 0 for the driving voltage as we assume ideal capacitive coupling with the probe electrode. The probe bias can now be expressed as:

$$V_P = V_{\pm,0} + V_C \quad (3.12)$$

Where  $V_{\pm,0}$  is the driving voltage and  $V_C$  is the voltage of the coupling capacitor (equal to  $V_F$  under ideal conditions).

Now that we have  $V_{\pm}$  in terms of  $T_e$  we can see that if we can calculate  $T_e$  in real-time we can set the probe bias in real time. Fitting routines to the I/V curve would be too slow (as discussed in 3.1.1) but by re-arranging the Langmuir probe equation 3 times for each bias state the electron temperature can be found in an iterative way.

- *Positive Bias  $T_e$* : The positive state is used to calculate the electron temperature necessary for adjusting the bias. Either default values for  $I_{sat}$  and  $V_F$  are used or values calculated from the previous bias cycle.

$$T_e = (V_P - V_F) / \ln\left(\frac{I_{LP}}{I_{sat}} + 1\right) \quad (3.13)$$

- *Negative Bias  $I_{sat}$* : The negative state is assumed to be sufficiently biased that it is collecting the ion saturation current. Re-arranging the Langmuir probe equation gives a correction to account for the slight gradient of the ion saturation current region.

$$I_{sat} = I_{LP} / \left(\exp\left(\frac{V_P - V_F}{T_e}\right) - 1\right) \quad (3.14)$$

- *Floating Bias  $V_F$* : The floating state is held at a potential voltage of 0V and if the capacitance assumption is true the probe should therefore be drawing minimal current. The floating voltage of the plasma can therefore be calculated by setting  $V_P$  to 0 and substituting in the previously calculated  $T_e$  and  $I_{sat}$  values.

$$V_F = 0_V - T_e \ln\left(\frac{I_{LP}}{I_{sat}} + 1\right) \quad (3.15)$$

So long as the equations can be iterated at a rate faster than a possible change in  $T_e$ , and the equations converge, then the temperature can be calculated. The order of the bias states is shown in figure 3.5. It is not necessarily obvious that iterating through these equations would converge on a temperature value.

Figure 3.3 shows the regions where iterating over the re-arranged Langmuir probe equations converges to the correct plasma temperature. While the interaction of the bias equations create interesting patterns (and have therefore been shown in several configurations) these figures have only been used as a preliminary check to confirm the equations can converge under specific conditions, and therefore inform the choice of default bias values should the instrument react in an unstable manner. The conclusion drawn from these figures is that the equations are generally likely to converge when  $T_e$  is significantly larger than the actual floating voltage of the plasma  $V_F$ .

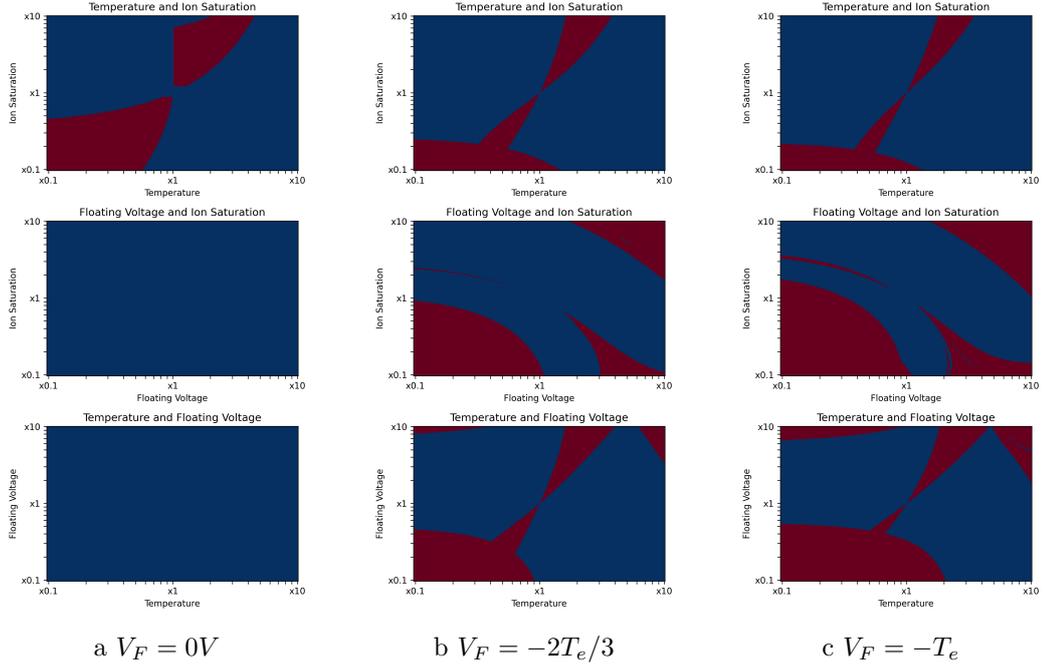


Figure 3.3: These figures show the convergence (blue regions) and non-convergence (red regions) for initial values of  $T_e$ ,  $I_{sat}$  and  $V_F$ . All figures have the same “plasma”  $T_e$  and  $I_{sat}$  values with three different values for “plasma”  $V_F$ . Initial values are represented as scaled to the “plasma” values. Convergence of these equations is decided by running the bias state cycle until the temperature values is within 1% of the actual temperature or the result is determined to be diverging or non-physical.

### 3.2.2 The Mirror Langmuir Probe Control System

The MLP control system is implemented on a Red Pitaya development board. This board has been previously described in 2.2.2. The dual 125 MSPS ADCs and DACs made it an ideal development environment for the project. The current sampling and calculation were achieved on FPGA while the data storage and instrument interface were handle using the Koheron SDK described in 2.2.1.

The MLP has separate calculation modules for each bias parameter detailed in equations 3.15, 3.14, and 3.13. Each value is calculated in sequence with the resulting output bias state voltage demonstrated in figure 3.5

Each calculation module consists of a multiplication step, a Xilinx IP division block<sup>[89]</sup> and a look-up table. An example of the calculation module for the  $I_{sat}$  parameter is shown in figure 3.6. The look-up table allows fast implementation of the logarithms and exponents found in the bias state equations. While complex functions are possible on FPGA<sup>[90][91]</sup> the required logic is expensive and introduces unwanted latency to the application. The Xilinx divider IP<sup>[89]</sup> was evaluated for

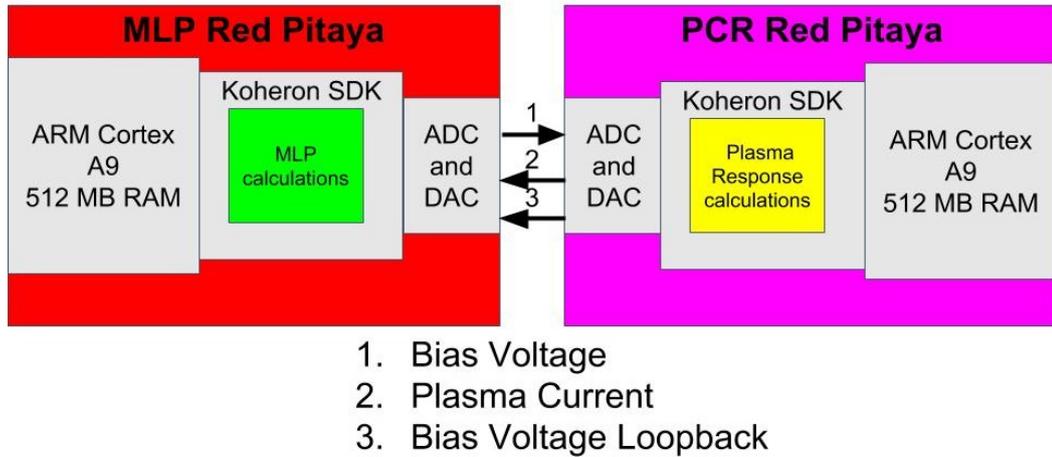


Figure 3.4: A diagram of the Mirror Langmuir Probe test system. Shown are the connections between the MLP and the PCR. Due to small sampling errors of the Red Pitaya ADC the PCR measured voltage is looped back to the MLP and this value, in conjunction with the PCR calculated value, is used to calculate the plasma temperature.

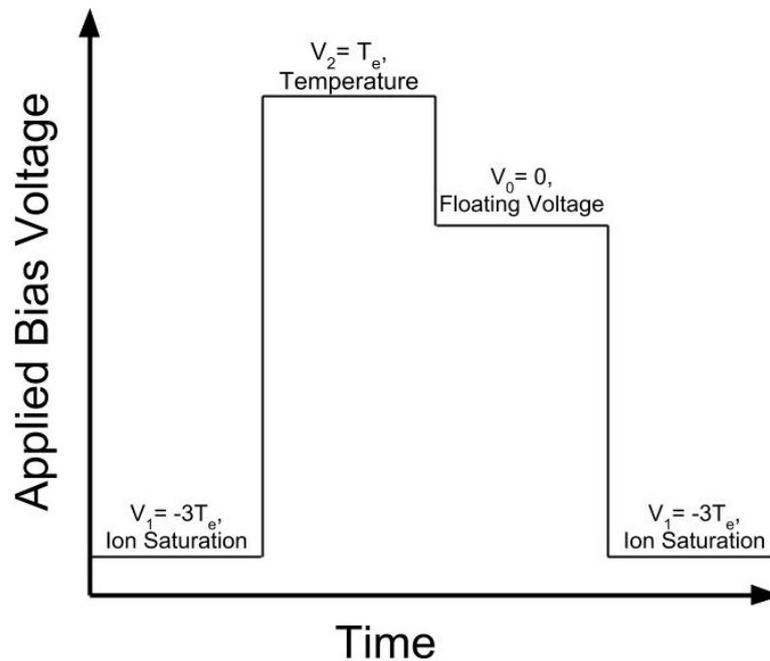


Figure 3.5: Diagram of the bias state sequence of the MLP control system. The calculated value and the bias value relative to temperature are annotated on each state. The first state is used to calculate a starting temperature which is then propagated through the other biases. An ideal system would have a positive bias of  $0.675T_e$  and negative bias or  $-3.325T_e$  but to simplify the digital logic requirements these values have been rounded to the nearest integer.

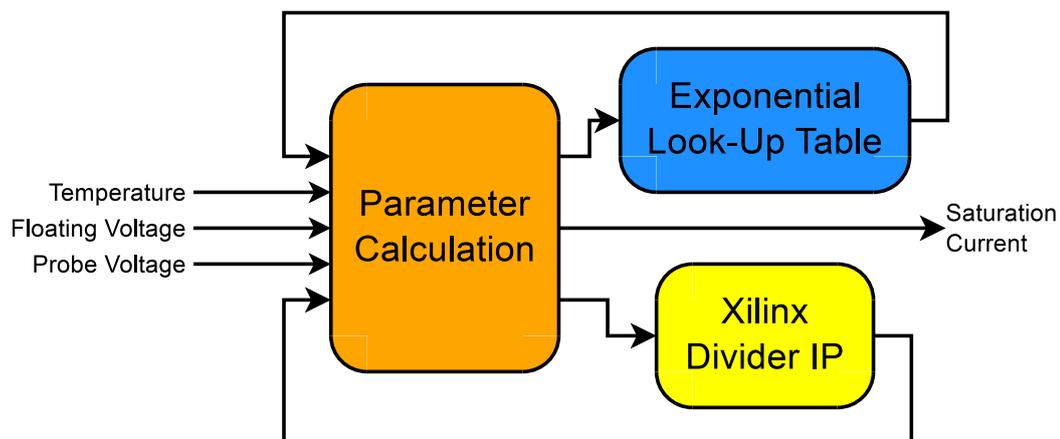


Figure 3.6: A flow chart of the calculation module and how it interacts with its respective look-up table and divider IP. For the ion saturation calculation the module takes an input of temperature, floating voltage and probe bias. The VHDL code that governs the parameter calculation can be seen in appendix C.1.

its speed and resource usage and found to be suitable for this implementation. The plug-and-play nature of Xilinx provided IP meant that the complexity of the design was significantly reduced and the limited block RAM available on the Red Pitaya could be used for the look-up tables, increasing the range and precision of the logarithm and exponent functions.

Table 3.1: Table describing the inputs, outputs and functions of the modules shown in figure ??.

Module	Inputs	Outputs	Function
Trigger	N/A	All other modules	Generates a core enable for all modules and a 1 MHz timestamp so that data can be accurately aligned
Calibration	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Probe Current</li> <li>• Monitor Current</li> </ul>	<ul style="list-style-type: none"> <li>• Smoothing</li> </ul>	Uses measured values to correct the input values from the ADCs

Continued on next page

**Table 3.1 – continued from previous page**

Module	Inputs	Outputs	Function
Smoothing	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Calibration</li> </ul>	<ul style="list-style-type: none"> <li>• Ion Saturation</li> <li>• Temperature</li> <li>• Floating Voltage</li> <li>• Data Collector</li> </ul>	The smoothing module is used to provide a stable input by mitigating the background noise from the ADCs
Ion Saturation	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Smoothing</li> <li>• Temperature</li> <li>• Floating Voltage</li> <li>• Bias Voltage</li> </ul>	<ul style="list-style-type: none"> <li>• Temperature</li> <li>• Floating Voltage</li> <li>• Data Collector</li> </ul>	Ion Saturation calculation module as described in figure 3.6
Temperature	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Smoothing</li> <li>• Floating Voltage</li> <li>• Ion Saturation</li> <li>• Bias Voltage</li> </ul>	<ul style="list-style-type: none"> <li>• Floating Voltage</li> <li>• Ion Saturation</li> <li>• Data Collector</li> </ul>	Temperature calculation module as described in figure 3.6
Floating Voltage	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Smoothing</li> <li>• Ion Saturation</li> <li>• Temperature</li> <li>• Bias Voltage</li> </ul>	<ul style="list-style-type: none"> <li>• Ion Saturation</li> <li>• Temperature</li> <li>• Data Collector</li> </ul>	Floating Voltage calculation module as described in figure 3.6

Continued on next page

**Table 3.1 – continued from previous page**

Module	Inputs	Outputs	Function
Bias Voltage	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Temperature</li> </ul>	<ul style="list-style-type: none"> <li>• Ion Saturation</li> <li>• Temperature</li> <li>• Floating Voltage</li> <li>• Data Collector</li> </ul>	The Bias Voltage module set the output bias voltage from the the input temperature value. It also times how long a bias state has been active and triggers the calculation of each parameter at the appropriate time.
Data Collector	<ul style="list-style-type: none"> <li>• Trigger</li> <li>• Smoothing</li> <li>• Ion Saturation</li> <li>• Temperature</li> <li>• Floating Voltage</li> <li>• Bias Voltage</li> </ul>	<ul style="list-style-type: none"> <li>• Zynq Processing System</li> </ul>	Collects the data for each of the required values and amalgamates them into a 64 bit number. This value is then passed to the Zynq Processing System using the infrastructure provided by the Koheron SDK.
Zynq Processing System	<ul style="list-style-type: none"> <li>• Data Collector</li> </ul>	N/A	An ARM based SoC that can directly interface to the FPGA logic using AXI interfaces <sup>[20]</sup> . The data can then be written to file using python or C++ or as in this case handled by a Flask implemented webserver for delivery to MAST-Upgrade DATAC (described in Appendix B.2).

### 3.2.2.1 Calibration and Smoothing

A simple module within the control system can be used to scale the output and input voltages to the Red Pitaya by calibration constants. These calibration constants can be obtained by scanning the DAC output over the full range of -1 V to 1 V. This may be necessary in more complex applications as the calibration errors introduced by the DAC would alter the calculated settling time of the probe bias due to cable capacitance. As the settling time from cable capacitance can introduce dead-time on the order of the required cycle time, errors in the output voltage due to imprecise DACs may mean that the cycle time becomes too slow to sample accurate probe voltages<sup>[10]</sup>.

Also important in practical applications is smoothing of the temperature parameter used to set the probe bias. Sudden fluctuations in temperature of arcs from the probe to the plasma have the potential to require the MLP control system to set probe bias outside of the range of the control electronics. By limiting the rate of change of the temperature parameter this behaviour can be mitigated against. The application of such a smoothing module can be seen in figure 3.9. The smoothing algorithm, shown in equation 3.16, used in this case is a basic real-time applicable algorithm that has a low logic requirement and acts as a simple low-pass filter.

$$X_n = X_{n-1} + \frac{x_i - X_{n-1}}{\alpha} \quad (3.16)$$

When setting the scaling factor ( $\alpha$ ) to 1 in equation 3.16 the smoothing output ( $X_n$ ) follows the input ( $x_i$ ) exactly. As  $\alpha$  is increased the high frequency response of  $X_n$  is reduced. As previously mentioned (see section 1.2.5), division operations are resource intensive and introduce latency, therefore the scaling factor is implemented as a right bit shift which can be implemented essentially instantaneously, though it does limit the scaling factor to exponents of two.

### 3.2.3 The Plasma Current Response System

In order to test and validate the MLP system a separate Red Pitaya development board was used to implement a Plasma Current Response (PCR) system. By utilising the Langmuir probe equation (as shown in equation 3.1) the PCR can feed back a current to the MLP for a given “applied” voltage. The PCR can then be loaded with preset temperature wave-forms which will alter the response current.

The configuration of the instruments is shown in figure 3.4 in a what we chose to term “duelling” mode. The MLP should be agnostic to the plasma source it is connected to and so the PCR is designed to be completely independent of the

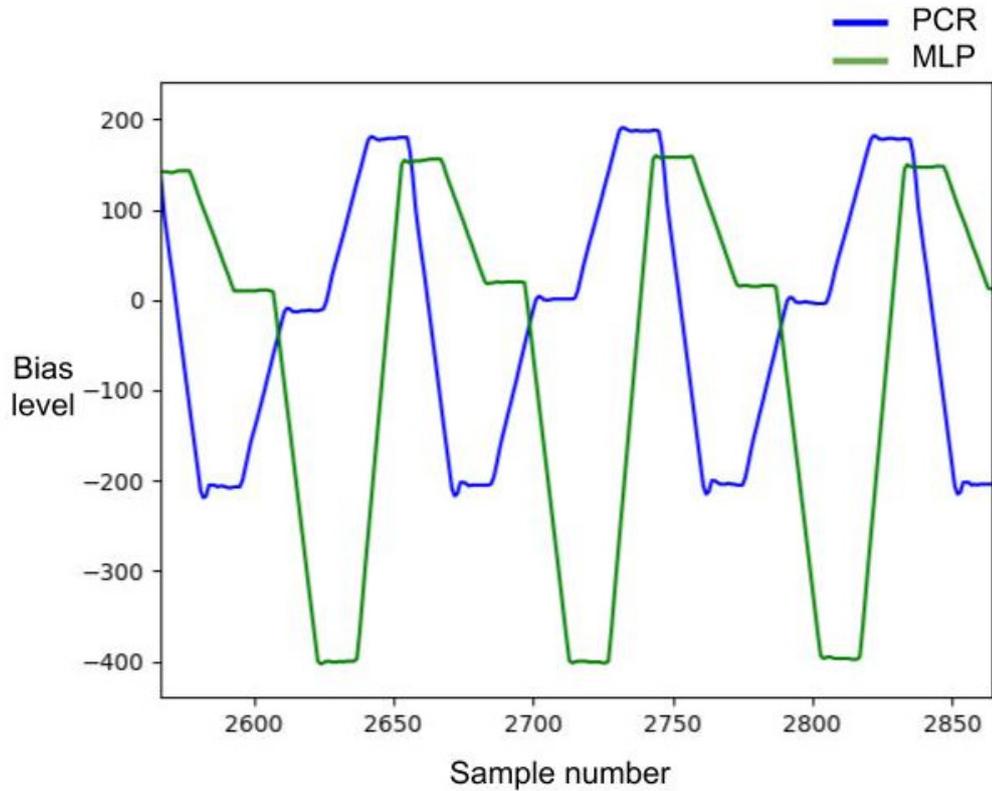


Figure 3.7: The voltage traces for the PCR (blue) and MLP (green) systems when the bias switching speed is on the same order as the PCR latency. It can clearly be seen from the neutral state  $\sim 0$  that when sampling the MLP the PCR is at the wrong bias state of the cycle.

MLP system, other than the single “voltage bias” input. The Koheron SDK allows the use of arbitrary waveforms to be used as inputs into the PCR so tests can be performed using unrealistic temperature changes or past real machine data.

Unlike the MLP the PCR runs a calculation stream, such that each input has a corresponding latency-separated output. This latency in the data stream can affect the MLP calculation at high speeds as the sampled PCR output would no longer be well aligned with the transmitted voltage bias as demonstrated in figure 3.7.

A feature seen in the original MLP system but not incorporated here is the capacitive coupling between MLP and the plasma source. Capacitance in cables has the affect of introducing time delays into the applied biases as well as causing current reflections. As such the MLP bias is the full probe bias applied to the PCR as opposed to a pre-gain value as would be used in full scale systems. The  $V_c$  term in equation 3.12 is therefore neglected.

### 3.2.4 Digitisation and Hardware

The test setup is presented as a cartoon in 3.4. The Red Pitaya boards are connected to each other using SMA terminated cables and the on-board 125 MSPS ADCs and DACs. Each input is terminated with a  $50\Omega$  load to reduce current reflections, critical to operating at high voltage switching speeds. Evidence of these reflections can be seen in figure 3.7 at the start of each MLP bias state.

To demonstrate the “all-in-one” advantage of the RP board, digitisation of the signals is done on the MLP and transferred to storage over ethernet using the Koheron SDK. To mitigate against sampling error of the ADC the PCR has an output of the measured input voltage to the MLP along with the calculated plasma response voltage. As mentioned in 3.2.3 there is a calculation latency between the PCR input and output so the looped back current is also delayed by the same period. At each bias state the MLP takes samples at five equally spaced intervals of the output voltage and the return voltage so that current stability can be verified in post processing. However only the final measurement at the end of the bias state (at which point it is assumed that the voltage has settled to the correct value) is used to calculate the next value in the sequence. This is deemed to be preferable to taking an average of the bias state as the voltage variance at high switching speed could cause it to be an unstable value.

## 3.3 Mirror Langmuir Probe Emulation

The objective of the proof of concept digital Mirror Langmuir Probe was to emulate the real-time biasing functionality of the original Mirror Langmuir Probe. To demonstrate this the PCR was used to with different temperature waveforms, both synthetic and from previous shots on Alcator C-Mod.

### 3.3.1 Step Change Tests

The preliminary test of a step change in temperature was designed to demonstrate the core functionality of the Langmuir probe control system. That is, the capability to change the Langmuir probe bias states in real-time. While this is an unrealistic test as the temperature would not change at this rate within a real plasma it does give an indication of the reaction time of the MLP to changes in temperature.

Figure 3.8 shows an applied step change in temperature, the temperature calculated by the MLP, and the calculated temperature from a standard Langmuir probe fitting routine applied to both the MLP input current and the PCR output current.

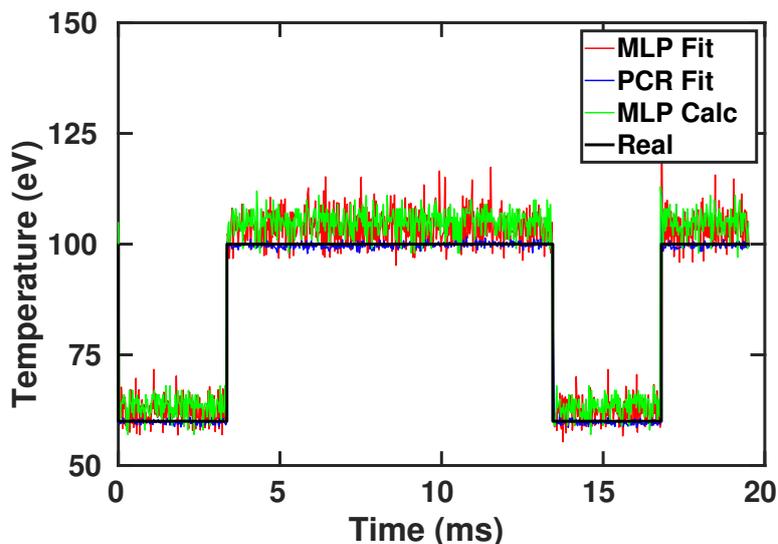


Figure 3.8: Figure showing the MLP response to a step change in temperature from 60 eV to 100 eV. The MLP calculated temperature is an overestimate of the applied temperature but the response is instantaneous.

The figure shows that while there is a systematic overestimation of the temperature by the MLP calculation it reacts instantaneously (within one calculation cycle) to the large change in both an increasing and decreasing direction. It is important to note that the systematic offset of the MLP calculated temperature does not affect the calculated post-processed parameters. The overestimated temperature means that the voltage levels that the DAC must switch between are larger than necessary, but it is preferable to an underestimation as we can be sure that we are still collecting within the electron and ion saturation regions. It is unclear as to why this systematic offset occurs as there are many places during the system it could be introduced. For example, it could be caused by an incorrect calibration of the Digital-to-Analogue Converter (DAC), the inherently inaccurate nature of a look-up table based calculation or some aspect of how the PCR and MLP interact.

The action of the bias states can be seen in figure 3.9 where the blue line denotes the voltage trace between biases and the orange marks show the bias steady state value. Also seen is the effect of the smoothing algorithm shown in section 3.2.2.1. It should be noted that figure 3.8 and 3.9 are from the same data set, and that while the applied bias are smoothed the calculated temperature is still at the correct value.

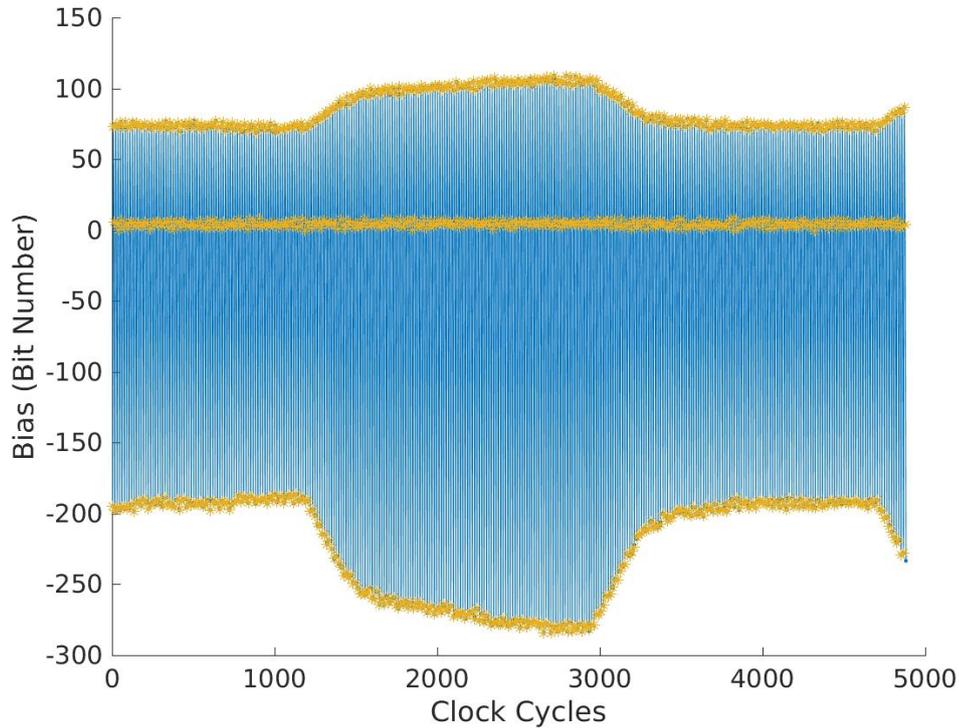


Figure 3.9: FPGA smoothed voltage trace (blue line) of the bias states (orange markers) for a step change in temperature from 60 eV to 100 eV.

### 3.3.2 Waveform Tests

While again un-physical, triangular waveforms in temperature and floating voltage show that the MLP control system can follow continuous changes. Figure 3.10 shows the MLP response to a continuous change in temperature over a range of 0 eV to 100 eV. A break down in the control system is apparent below a 15 eV threshold. It was found that due to the look-up table implementation in the PCR the value for the returned current for small  $T_e$  becomes much larger than if calculated analytically. The look-up table size is limited by the logic resources of the FPGA, and due to the exponential nature of equation 3.1, the fact that it must become more dense for small values of  $T_e$ . Tests with a real plasma would show the better defined limits of the control system, which would likely be due to the ADC precision.

The control system break down does show however that the MLP can re-lock to the correct temperature once it re-enters the operation region. To facilitate re-locking when the MLP is deemed to be calculating incorrect temperatures (negative values, or values that are too big) the calculated temperature is set to a default value until

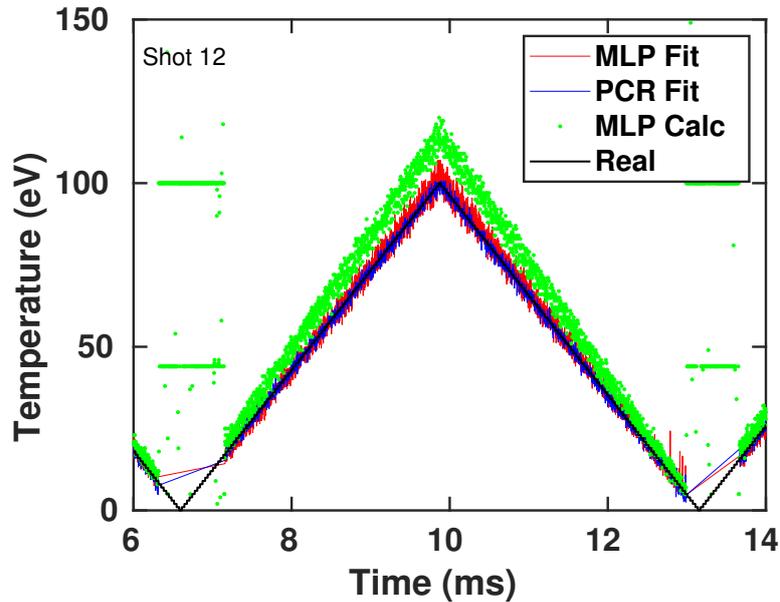


Figure 3.10: A demonstration of the MLP following a continuous change in temperature. The test also shows the control system break down below a 15 eV threshold. The correct temperature is again calculated once the waveform enters the operational region.

it becomes stable again.

Figure 3.11 shows a triangular waveform of the floating voltage at a constant temperature. As the probe current relationship to floating voltage is inverse to that of temperature the calculation is stable throughout the the range. While it is useful to know that the plasma floating voltage can be calculated correctly it is important to note that the floating voltage and ion saturation values have no direct effect on the bias state levels.

Figure ?? shows the bias states for the temperature waveform in figure 3.10. The high switching speed of the bias states results in the voltage trace of the MLP output appearing as a blue block. Yellow markers bound and bisect the blue block representing the bias steady state probe voltage. When the calculated temperature becomes erratic the bias levels are set to default states equivalent to a calculated temperature of 100 eV, and the voltage trace is not recorded. The probe trace is then recorded again once the calculated temperature becomes stable.

### 3.3.3 Emulated Real Data Tests

The closest to full MLP emulation achievable with this test system is the use of real C-mod shot data to set the plasma parameters of the PCR. Due to the

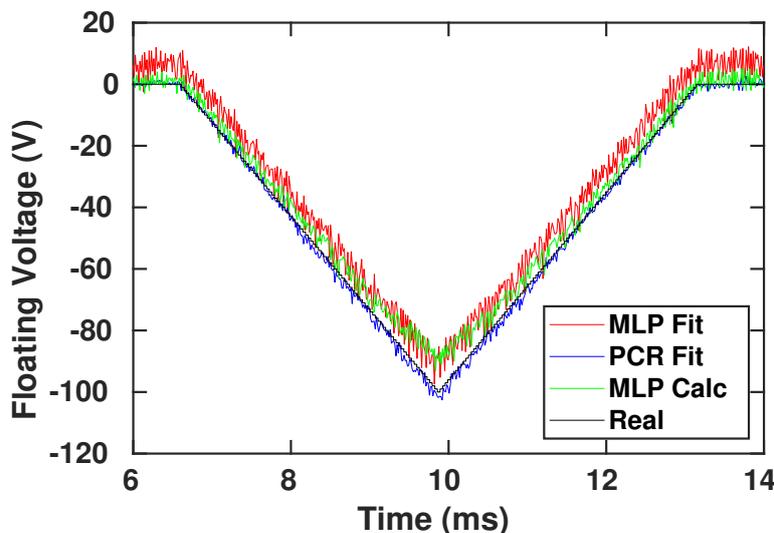


Figure 3.11: Triangle waveform in floating voltage at a constant temperature. The threshold for look-up table induced break down is not reached and the calculated floating voltage stays stable. While a correctly calculated floating voltage indicates good system stability, it has no direct impact on the bias state output levels and is only used in the temperature calculation cycle.

latency issues mentioned in section 3.2.3 the C-mod data was slowed down by a factor of 2 and the bias cycle rate of the MLP was set to 500 kHz, so that relative to each other the time resolution was similar to the original MLP. Without this adjustment the PCR latency would mean that the input and output signals were significantly de-synchronised relative to each other so that the MLP would not be sampling the equivalent “Langmuir probe” response for its applied bias. In real fusion applications, so long as the cycle time is slow enough that the equations can propagate through the PL, the real-time biasing algorithm would be agnostic to the speed at which it is operated. This does not mean that it can be always operated up to the calculation limit however, as the switching speed of the electronics delivering the probe bias must be taken into account.

Time series data of temperature, floating voltage, and ion saturation current were used from run day 160628 of C-Mod operation. Figure 3.12 shows that the general trends of the temperature are followed well with a systematic over-calculation of the temperature in almost all cases. As mentioned previously MLP over-calculation of temperature is not an issue if the post processed parameters can produce a good fit. In this case the post-processed parameters follow the data precisely with a less than 10% error, less than the error usually associated with Langmuir probe measurements. Figure 3.12 also shows an inset of a 10 ms window to better view the behaviour of the MLP calculated temperature. While the calculated temperature

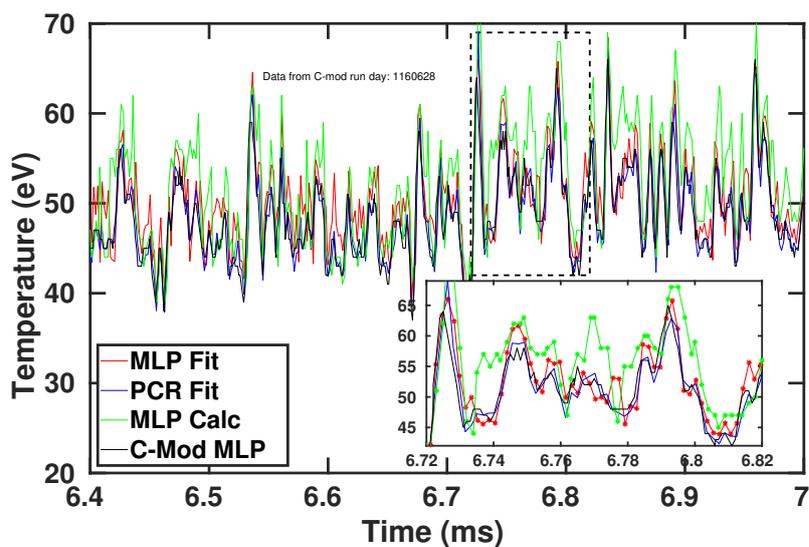


Figure 3.12: A demonstration of MLP behavior when C-Mod turbulence data is used to set the PCR plasma parameters. The MLP follows the turbulence well enough to facilitate accurate post-processing fit of the plasma parameters. The inset more clearly shows that while the calculated temperature by the digital MLP is not as good as the original analogue system, the general trends are presented clearly with a systematic over estimation.

of the original analogue system manages to follow the post-processed parameters more closely, the digital MLP still captures the large temperature spike features.

### 3.4 Summary

This work showed the first proof-of-concept for a digital Langmuir probe control system that can alter probe biases in real-time. Through the implementation of this control system on digital electronics its application to real plasma experiments is significantly simplified. While FPGA's have been used before to implement Langmuir probe control systems<sup>[79]</sup> new advancements in the size and flexibility of implementation of FPGA technology means that advanced instruments can be developed and deployed quickly. The real-time biasing scheme is also not limited to this application and could be applied to other FPGA based Langmuir probe systems to reduce the amount of scanning time spent in irrelevant regions.

#### 3.4.1 Further Work

Work has been undertaken by W. McCarthy et al<sup>[92]</sup> at the Plasma Science and Fusion Centre to demonstrate the application of the real-time digital Mirror Lang-

muir Probe control scheme on the linear plasma device DIONISIS<sup>[93]</sup>. These tests showed that the real-time control system was functional with a temporal resolution up to 140 kHz and was able to follow fluctuations in plasma parameters such that the post-processed values accurately calculated electron temperature, ion saturation current and plasma potential. All components used were Commercial-Off-The-Shelf showing promise for quick and easy adaptation to other devices. While the 140 kHz temporal resolution is less than what has been demonstrated as part of the proof-of-concept system it is still significantly faster than most Langmuir probe control systems in use<sup>[79] [94] [95]</sup>.

Work will also be continued on MAST-Upgrade through a PhD studentship. This project will work to implement a real-time Langmuir probe biasing system for implementation as part of a reciprocating probe. While the inspiration for the project is based on the work presented here, the programmable logic will be redesigned from the ground up.

---

# The Synthetic Aperture Microwave Imaging Diagnostic

This chapter presents work on the data acquisition system for the second generation Synthetic Aperture Microwave Imaging (SAMI) diagnostic. Background physics, motivation for the project, and the technological challenges to overcome that are necessary to produce the physics deliverables are covered. An overview of the FPGA implementation and server side drivers used to surmount these challenges is described along with how these will be used to aid physics studies once the diagnostic is implemented on MAST-Upgrade.

## 4.1 Background

Tokamak plasmas have many operating regimes characterised by their capacity to retain heat, edge plasma pressure gradient and plasma stability. High confinement mode (H mode) is characterised by a plasma edge transport barrier that creates a temperature and pressure pedestal. The pedestal increases the core temperature and pressure such that significantly more fusion occurs in the core region than when the plasma is in low confinement (L mode) conditions. Core pressure increase induced by the transport barrier is shown in figure 4.1.

A lot of the physical mechanisms behind the creation of this transport barrier have been studied in depth<sup>[96] [97] [98]</sup> (these citations represent a small sample of previous research into the transport barrier and steady state H mode), and entering H mode is now a relatively standard procedure, achieved through the injection of heat into a low impurity ohmic plasma<sup>[99]</sup>. The fusion rate is significantly increased by the presence of this transport barrier, and is accompanied by the suppression of

edge turbulence. For this reason H mode operation has been decided as the ITER baseline operation scenario<sup>[100]</sup>.

While H mode is favourable for the fusion rate of a tokamak, the ITER baseline operation scenario is currently expected to include the presence of Edge Localised Modes (ELM)<sup>[101]</sup>. ELMs are quasi-periodic expulsions of energy and matter from the plasma and can comprise of a significant portion of the total energy stored in the plasma. While ELMs in smaller devices can be useful in removing plasma impurities from the core, as tokamaks increase in size they can become problematic to tokamak integrity<sup>[102]</sup>. The periodic expulsions of energy and matter increase the heat flux to plasma-facing components (PFCs) and for power plant scale experiments could cause failure of the machine<sup>[103][104]</sup>. In addition to causing damage, requiring the replacement of PFCs, ELMs also limit the pressure and temperature pedestal height and therefore reduce the plasma efficiency. For these reasons it is generally recognised that a good quantitative understanding of ELM processes is necessary for their prediction and control.

A proposed theory for ELM mechanisms is peeling-ballooning mode theory<sup>[12]</sup>. Peeling-ballooning mode theory states that coupling of the edge pressure and current works to limit the pedestal height and destabilise the edge pressure gradient, leading to the onset of ELMs. Figure 4.2 shows the cyclic nature of the coupling pressure and current modes. Type II and type III ELMs occur when the edge current density and edge pressure gradient respectively become large enough to be unstable against ballooning or peeling modes<sup>[12]</sup>. Both of these types of ELM result in relatively small expulsions of energy. If both variables rise too quickly then type I ELMs can occur. Type I ELMs can cause significant damage to plasma facing components, due to the amounts of energy and matter they expel, and it is therefore critical to understand the conditions that lead to their occurrence.

To correctly classify the stable and unstable regions of figure 4.2 both the edge pressure gradient and edge current density must be accurately characterised. While the pressure gradient can be accurately measured by current diagnostics, such as the Thomson scattering diagnostic (electron density and temperature), reciprocal Langmuir probes (see section 3.1.1) and charge-exchange spectroscopy (ion temperature), high temporal and spatial resolution measurements of the edge current density are lacking. For this reason the current density is usually inferred, using reconstruction programs such as EFIT<sup>[13]</sup>, from either the magnetic pitch angle or estimated using the force balance equation in relation to the pressure gradient.

Diagnostics capable of measuring the magnetic pitch angle are the Motional Stark Effect (MSE)<sup>[43][105]</sup> diagnostic or Zeeman polarimetry of lithium during lithium pellet or lithium beam injection<sup>[106]</sup>. Both these methods have drawbacks however.

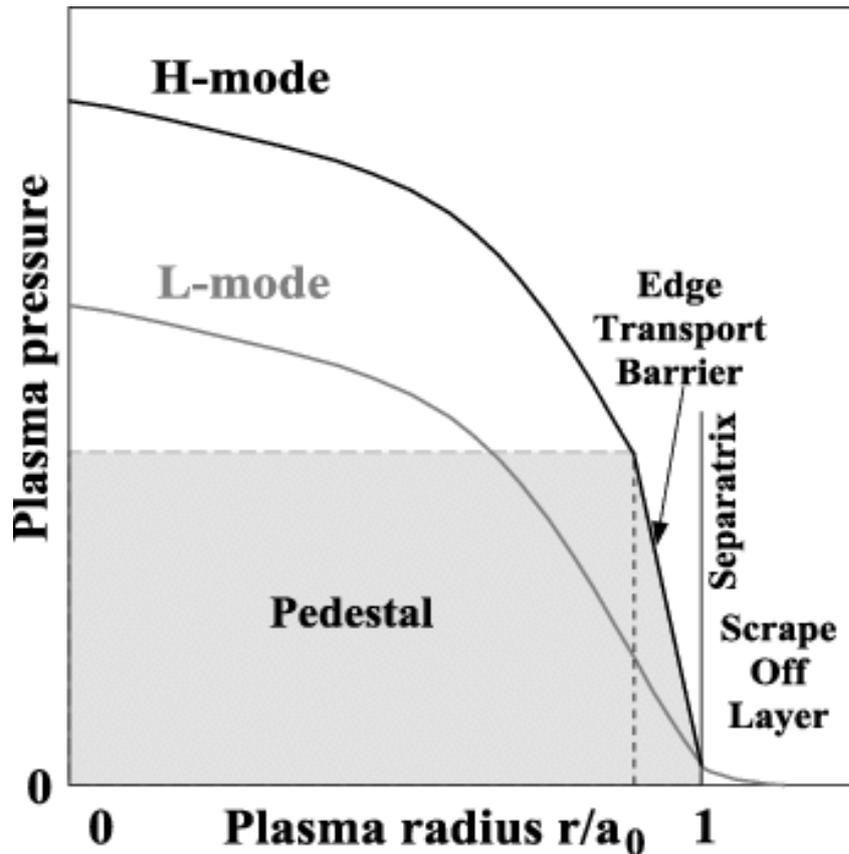


Figure 4.1: Example graph of the pedestal formed in the pressure profile of a H-mode plasma compared to a L-mode plasma. The high edge pressure gradient suppresses turbulence at the plasma edge and leads to higher core pressures. Figure taken from “ELM control strategies and tools: Status and potential for ITER”<sup>[11]</sup>

As lithium must be handled with care and is therefore not utilised on most machines, the MSE diagnostic has therefore become a standard part of the diagnostic suite in order to measure the magnetic pitch angle.

MSE is an optical diagnostic technique that measures the line splitting of injected neutral deuterium atoms due to their high velocities in a magnetic field. While MSE gives good spatial resolution the drawback in this case is that a neutral beam is required for the MSE to function and as such it cannot be used in operating scenarios where NBI (see section 2.1.1.1) is not required. Also, while not necessarily a concern for ELM studies, as previously discussed the application of the NBI can cause the plasma to transition to H-mode, and as such the edge current density in L-mode is not well understood. The limited suite of diagnostics for edge current density, peeling mode limits, and how they lead to the occurrence of ELMs means that a diagnostic capable of studying these quantities is desirable.

The following material will introduce the original SAMI diagnostic and its capab-

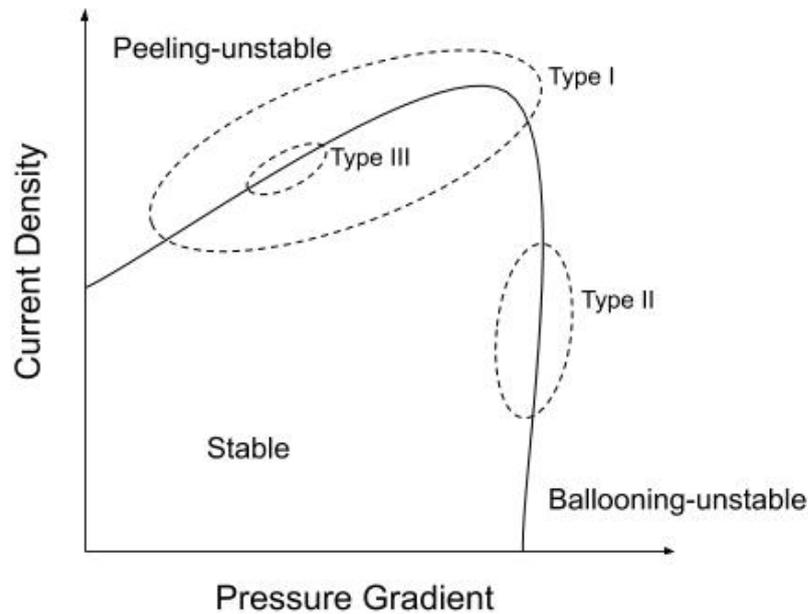


Figure 4.2: Diagram of the stability curve created by the coupling of the edge current density and pressure gradient. Three classifications of ELM are demonstrated, characterised by the destabilising mechanisms that cause them. Figure reproduced from “Magneto-hydrodynamic stability of the H-mode transport barrier as a model for edge localized modes: an overview”<sup>[12]</sup>. Currently only type I ELM theory is generally agreed, with the paper referenced here proposing a theory for the occurrence of type II and III ELMs.

ilities, as well as the results from its implementation on the spherical tokamaks MAST and NSTX. Presented next will be the background physics of microwaves in plasma so that the requirements of the diagnostic are fully defined. This includes showing how a synthetic aperture can be used to image the 2-D plasma surface and how Doppler back-scattering (DBS) enables us to infer the magnetic pitch angle. A brief description of the front-end electronics necessary to receive the microwaves emitted by the plasma will follow.

An overview of current 10 Gigabit ethernet technologies will be given to appreciate why it is the preferable communication protocol, as well as the required Linux systems needed to control and operate the FPGA boards and acquisition PC’s. The second part of the chapter will cover the work done as part of the data delivery system to enable these technologies within the second generation SAMI instrument (SAMI-2) and therefore allow the diagnostic to operate over a full shot length with an increased number of polarisation separating antennas compared to the original SAMI diagnostic. The capability for SAMI-2 to acquire this data in real-time over an arbitrary shot length is a significant improvement over the original SAMI diagnostic. The use of a standard network communication protocol (10 gigabit

ethernet) means that SAMI-2 is agnostic to the device it is installed on and could potentially be implemented on other fusion experiments in the future with relative ease.

#### 4.1.1 The First Generation Synthetic Aperture Microwave Imaging Diagnostic

The first generation Synthetic Aperture Microwave Imaging (SAMI) diagnostic was a novel imaging diagnostic studying the tokamak plasma edge. Using a synthetic aperture technique (similar to that used in radio astronomy and ultrasound) with antennas tuned to the microwave range, SAMI can resolve the electromagnetic radiation characteristics of a plasma surface corresponding to a given microwave frequency<sup>[107]</sup>. Probing the edge in this manner is preferable to optical diagnostics as the antennas can be made of metal which is much more resilient to neutron bombardment than glass. While not an issue for smaller machines like MAST-Upgrade, the neutron flux on ITER-like devices will mean that optical instruments will degrade very quickly.

The original SAMI was implemented on MAST and NSTX and used to measure the magnetic pitch angle in the plasma edge. It was able to achieve this by illuminating the plasma with a microwave beam and measuring the bandwidth of the reflected signal. The synthetic aperture approach to resolving the return signal meant that SAMI was able to build up an image of the returned microwave frequency across the plasma surface, as shown in figure 4.3, and locate the regions of greatest Doppler shift due to the rotating plasma surface.

As the regions of Doppler shift seen in figure 4.3 are aligned with respect to the plasma rotation we know the velocity vector of the charge carrier. From section 1.1.2 we also know that the magnetic field is perpendicular to the velocity of the charge carriers it is coupled with, and we can therefore infer the magnetic pitch angle. This method of measuring the pitch angle is called 2-D Doppler back-scattering (DBS). A probing beam is required to illuminate the plasma surface and produce the back-scattered signal, but as this can be generated by the diagnostic itself this method can therefore operate in all plasma operating scenarios.

DBS itself is a well understood technique but had previously only been used in one dimension to measure fluctuations in density<sup>[108]</sup>. The application of this method in 2-D to measure magnetic pitch angle is unique and novel to the SAMI diagnostic. Unfortunately, deriving the edge current density would require an edge profile of the magnetic pitch angle and this was not within the capacity of the original SAMI diagnostic.

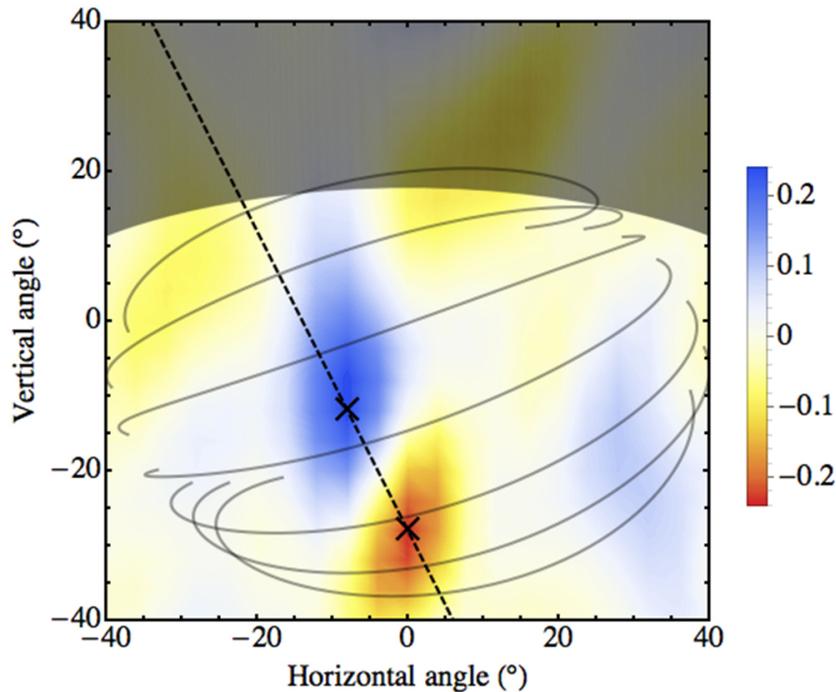


Figure 4.3: Analysed data from shot #27969 generated by SAMI when implemented on the MAST tokamak. The red and blue zones show the regions of greatest Doppler shift as the probing microwaves are reflected off the spinning plasma surface. As we know that the magnetic field is perpendicular to the plasma rotation, the magnetic pitch angle can be inferred from the plane of the regions of maximum Doppler shift. The dotted line connects the two Doppler poles and is orthogonal to the horizontal solid lines which show the magnetic field as reconstructed by EFIT<sup>[13]</sup>. Figure taken from “2D Doppler backscattering using synthetic aperture microwave imaging of MAST edge plasmas”<sup>[14]</sup>

Due to the novel nature of the SAMI instrument, FPGA technology was used as the primary data acquisition and processing hardware. The ADC’s used to digitize the antenna signals were operated by a pair of Xilinx ML605’s in a primary-secondary configuration. The primary board implements a MicroBlaze soft processor to control the ethernet interface and also generates a global clock for both of the FPGA boards. Synchronisation of both boards is critical to ensure that the antenna data is aligned correctly, as misalignment would mean that the synthetic aperture approach would not be able to produce a correct image. Synchronisation is accomplished by cloning the programmable logic architecture for both boards but only using the global clock from the primary. This, in theory, means that the electrical paths of the programmable logic are identical with the same clock source. Pre-shot calibration is also used to make sure that data is correctly aligned. The trigger is also only supplied to the primary board and then distributed to the acquisition logic on both boards using external I/O.

The acquired data is streamed directly into RAM over a small-outline-dual-inline memory interface on the FPGA board. The board could support a maximum of 2 GB of data over this interface, and the transfer rate was limited to 4 GB/s. The data was then be streamed post shot into an acquisition PC using a User Datagram Protocol (UDP), specifically designed for SAMI, over ethernet on a Local Area Network (LAN). The combined limits to data capacity and write speed mean that the maximum acquisition time of the original SAMI system was 500 ms. For the implementation on MAST and NSTX this was fine as the maximum shot length was also 500 ms.

The Linux system used to operate the diagnostic was compiled to occupy the lowest 120 MB of the memory-space available and allowed full control of the FPGA peripherals, such as the digitisation core and ethernet interface. Since the high data rate meant that the MicroBlaze core would not have been able to handle the data throughput, multiplexers were used to shift the memory control I/O pins between the soft processor and the digitisation logic as required.

#### **4.1.2 Limitations to the Original SAMI Diagnostic and Digitisation and Acquisition Brief**

From this brief description of the original SAMI diagnostic some limitations are apparent, due to technology limitations at the time and the novelty of the instrument, when looking to implement on the new MAST-Upgrade machine. The two major limitations of the original instrument are the limited acquisition time and the inability to separate the incoming microwave O and X mode polarisations (see appendix D.1). To solve these issues the new SAMI-2 diagnostic will use current FPGA technology to implement high-speed data streaming directly to an acquisition PC built with COTS components. Separating the received microwave signal into orthogonal polarisation components is achieved using a new antenna design with custom built analogue electronics.

SAMI-2 also benefits from an increase in the number of receiving antennas from eight to thirty. An additional two probing antennas makes for an array of thirty-two antennas as described in J. Allen et al, *Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic*<sup>[15]</sup>. For each antenna the received signal is split into vertical and horizontal polarisation components, the two frequencies from the emitted probing beam are then separated, and finally each frequency is divided into its in-phase and quadrature components. This results in a total of 240 channels that require digitisation. To achieve this SAMI-2 will deploy ten Xilinx ZCU102 devel-

opment boards, each digitising and processing 24 channels of data, equivalent to the response from 3 antennas.

## 4.2 Microwaves in Tokamak Plasmas

To understand the motivation behind development of the SAMI diagnostic we will first cover the physics of microwaves in plasma and how they are relevant to the edge current density problem. SAMI-2 is intended primarily as a 2-D DBS diagnostic and so the theory presented here will focus on motivating the hardware requirements for this purpose. Other applications of the SAMI diagnostic could be the passive measurement of Electron-Bernstein waves, which would require a more in-depth treatment of plasma wave interactions, but this is outside the scope of this thesis. The following therefore assumes that the plasma regimes SAMI-2 will be observing follow the cold plasma approximation.

In a plasma the electrons are sufficiently energetic that they are no longer bound to nuclei. However, due to the large mass difference between the electron and a nucleus the cold plasma approximation assumes that the velocity of the electrons is relatively larger than the ions such that the ions are essentially stationary relative to the electrons. Within the reference frame of an ion we can imagine some background cloud of free electrons. When this cloud is separated some small distance from the ion, electrostatic attraction will draw the electrons back and the electrons will oscillate around the ion. This mechanism results in a plasma parameter called the plasma frequency which is shown in equation 4.1<sup>[109]</sup>, where  $n_0$  is the electron number density.

$$\omega_p = \left(\frac{n_0 e^2}{\epsilon_0 m_e}\right)^{1/2} \quad (4.1)$$

A fundamental property of the plasma, the plasma frequency is only dependent on the plasma density. As plasma density in a tokamak is related to radial position, plasma edge probing depth can be selected by the corresponding plasma frequency at that specific density.

When an electromagnetic wave is incident on the plasma, the oscillating electric and magnetic fields perturb the unconstrained electrons<sup>[29]</sup>. The interaction between the incoming wave and the plasma leads to a dispersion relationship dependent on the plasma frequency.

$$\omega^2 = k^2 c^2 + \omega_p^2 \quad (4.2)$$

Equation 4.2 shows that as the waves penetrate deeper into the plasma the density rises to a point where  $\omega = \omega_p$  and the wave vector,  $k$ , becomes purely imaginary.

This therefore defines a cut-off density where the wave can no longer propagate through the plasma medium. A wave that reaches its cut-off density is most likely to be reflected or scattered, though other effects can occur such as mode conversion or tunnelling<sup>[110]</sup>.

Reflected and scattered waves can be imaged to measure properties of the plasma surface from which they occur. Receiving hardware for microwaves are generally metallic antennas where the electromagnetic wave is converted to an electrical signal. For long time-frames this signal must then be digitised for storage and analysis. The rest of this chapter covers how this is achieved for the SAMI-2 instrument.

## 4.2.1 Image Formation with a Synthetic Aperture

For a spatially incoherent source, constructive interference only occurs from a single point on the object being imaged, and the source intensity distribution can be imaged with a collection of receiving antennas<sup>[111]</sup>. This is called synthetic aperture imaging and is exploited by the SAMI-2 diagnostic to image the plasma surface at microwave wavelengths.

The original SAMI diagnostic, as implemented on MAST, used a far-field approximation which assumes that the source is sufficiently distant that waves incident on the detector are plane-parallel. Described below is the beam-forming technique used to facilitate the 2-D DBS method for measuring the magnetic pitch angle. While more complex methods exist, incorporating adjustments such as near-field effects, an overview of the beamforming technique will give sufficient context for synthetic aperture image formation as used in this application.

### 4.2.1.1 Beamforming

Originally developed for ultrasound imaging, the beamforming technique is now widely used in many areas, such as radar, telecommunications and astronomy<sup>[112][113]</sup>. As mentioned above, assuming that a source is spatially incoherent, constructive interference will only occur along a single line of sight between the source and the receiving antenna. For a multiple antenna array, a line of sight can be selected by applying a phase shift between antenna pairs such that constructive interference can only occur from that direction.

$$s_1(\underline{\alpha}, f) = \sum_{j=1} W_j s_j(\underline{\alpha}, f) s^{\frac{i2\pi f}{c} \underline{\sigma}_{j,1} \underline{\alpha}} \quad (4.3)$$

Equation 4.3, when applied to the received signal, suppresses all signals along direction  $\alpha$ , with frequency  $f$ , to the source from reference antenna 1. The constant  $W_j$  acts as a scaling factor calculated from the voltage response function of antenna  $j$ , with  $\sigma_{j,1}$  denoting the vector between antenna 1 and antenna  $j$ . Equation 4.3 uses the far field approximation to assume that vectors  $\sigma$  and  $\alpha$  are orthogonal. The intensity of signal  $s_{\alpha}$  can be calculated by applying the inverse Fourier transform as shown in equation 4.4.

$$I_1(\underline{\alpha}) = \int_{-f_{Nyq}}^{f_{Nyq}} |FT^{-1}[s_1(\underline{\alpha}, f)]|^2 df \quad (4.4)$$

Where we have accounted for the digitisation discretisation by integrating between the Nyquist limits of our digitisation electronics. Equations 4.3 and 4.4 can be evaluated for any direction  $\alpha$  and as such a 2-D image may be formed by “scanning” across the detector’s field-of-view (FOV).

Equation 4.3 shows that by increasing the number of antennas, the number of sight lines ( $\alpha$ ) is increased. The resolution of the system is actually proportional to the antennas squared<sup>[114]</sup>, so while adding more antennas does burden the digitisation system, they pay-off is generally worth it. As FPGA’s allow parallel signal lines, meaning that it is relatively easy to duplicate the digital signal processing for one signal to many others, then the extra digitisation requirements are relatively minor so long as there are sufficient logic resources on the chosen hardware platform.

## 4.2.2 2-D Doppler back-scattering

We have now described how electromagnetic waves are reflected or scattered at specific densities relative to their frequencies and that a 2-D image of the returned microwave intensity distribution can be formed from an array of antennas. Using this information we can formulate a diagnostic technique for forming a 2-D image of a plasma density surface. By illuminating the plasma with a single frequency microwave source and constructing receiving electronics to receive at that frequency a detector can measure the intensity of back-scattered light. This technique can be used not just for DBS, as in this application, but also reflectometry to measure the distance to, and therefore radius of, the plasma<sup>[115]</sup>. Radius measurements are useful for characterising turbulence present in the plasma edge region as well as being a useful real-time parameter for feedback control of the constraining magnetic field.

A tokamak plasma iso-density surface is neither flat nor stationary, with density corrugations along the propagation path occurring due to turbulence in the edge region. These density corrugations act to modulate the refractive index and therefore

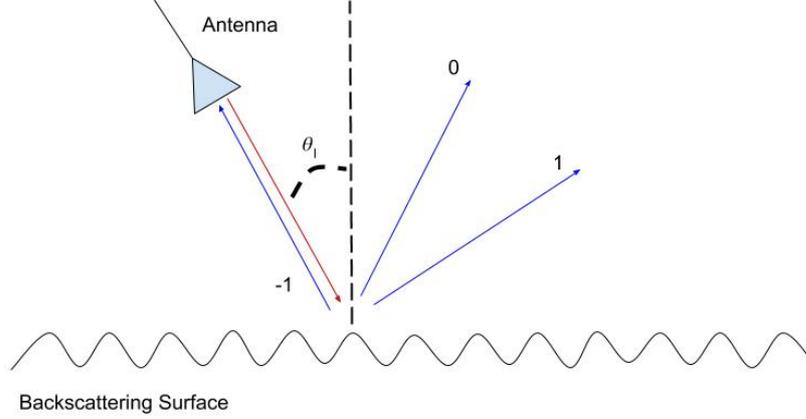


Figure 4.4: Diagram of incident beam on a reflection grating and the back-scattered light (with diffraction orders shown as an example) resulting from corrugations along a reflecting surface. In the case of the SAMI diagnostic this surface is the iso-density cut-off surface corresponding to the frequency of the emitted beam. Corrugations along the surface are the result of turbulence within the plasma edge.

reflected light from the cut-off is scattered back in the direction of the source<sup>[116]</sup>. By imagining a microwave beam incident on a grating in a vacuum, we can form a model for how the back-scattered light relates to the scale of the turbulence causing the density corrugations.

Figure 4.4 presents a diagram of an incident beam on a reflection grating that is moving with some velocity  $v$ . In contrast to conventional Doppler reflectometry the beam is incident on the surface at an oblique angle. This model represents the turbulent plasma surface so that we can apply the Bragg condition, shown in equation 4.5.

$$K_T = 2k_I \sin \theta_I \quad (4.5)$$

The Bragg condition relates the grating spacing to the angle of incidence. By scanning the beam, with wave number  $k_I$ , across the corrugation the wave number spectrum ( $K_T$ ) of the grating is determined, in this case generated by the different scales of plasma turbulence. The scattering surface is moving with velocity  $v_T$ , the phase velocity of the turbulent structure relative to the plasma surface, resulting in a Doppler shift shown in equation 4.6

$$\delta\omega = vk_I \quad (4.6)$$

An example of the returned power spectrum can be seen in figure 4.5 where, while the resolution is not sufficient to resolve the backscatter Doppler peak, there is

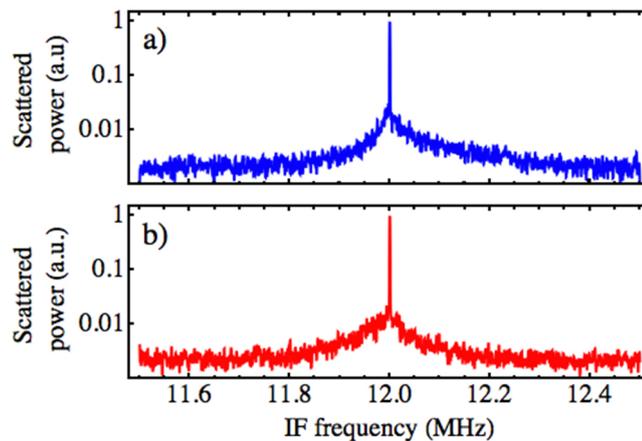


Figure 4.5: Back scattered power of a Doppler backscattering diagnostic. The central frequency peak is from the probing beam, after the gigahertz carrier wave has been removed (see section 4.3). While the system resolution is not sufficient to resolve the Doppler shifted peak there is still a clear asymmetry visible depending on which direction the plasma is spinning in. Colours align with those in figure 4.3, such that blue is moving away and red is moving towards the antenna. Figure taken from “2-D Doppler backscattering using synthetic aperture microwave imaging of MAST edge plasmas”<sup>[14]</sup>

a clear asymmetry in the backscattered power depending on the direction that the plasma is spinning in. The standard DBS technique is used to investigate the turbulent structure velocity distribution which can give insight into edge transport, the L-H mode transition and H-mode pedestal structure<sup>[117][108]</sup>.

Advantages of the DBS technique over other edge turbulence velocity diagnostics such as Beam Emission Spectroscopy<sup>[118]</sup> and Gas Puff Imaging<sup>[119]</sup> are listed below:

- DBS antennas can be made out of neutron resistant materials, the diagnostic system does not include mirrors or focusing optics<sup>[120]</sup> which are more susceptible to neutron damage.
- While it is an active diagnostic, the probing beam of DBS has little to no affect on the plasma conditions, unlike the beam penetration of BES<sup>[120]</sup> or the injected pellet of GPI<sup>[121]</sup>.
- Future tokamaks will not be able to operate with port access as this plasma facing area will be used by breeding blankets. Implementations of the DBS technique on such devices could utilise wave-guides to de-localise receiving electronics behind neutron shielding.

While the intricacies of DBS, and how it compares to other instruments of its type, are out of the scope of this thesis the advantages stated here also apply to the SAMI diagnostic.

#### 4.2.2.1 Deriving Edge Current Density

2-D Doppler back-scattering is a technique for viewing the Doppler shift distribution of an emitted frequency of light across an iso-density plasma surface. Receiving electronics are therefore designed to filter out the central frequency of a probing beam and extract the side-bands where the Doppler shifted signal is greatest. Signal side bands, as demonstrated in figure 4.5, occur due to a shift in frequency of the probing beam when reflected off the moving plasma surface. As the field of view is scanned across the plasma by the beamforming algorithm the beam power in the side-bands will change. An image of the plasma surface velocity can therefore be produced. This in turn can be used to determine the magnetic pitch angle at that iso-density surface.

The Doppler shift of the back-scattered beam can be attributed to two conditions: (i) the turbulent phase velocity relative to the iso-density background surface can be random in direction and so results in a simple broadening of the back-scattered signal; (ii)  $\underline{E} \times \underline{B}$  or pressure-force effects can drive an underlying velocity that results in an overall directional Doppler shift relative to the lab frame. The received power spectrum is therefore asymmetrical, depending on the direction of the motion relative to the antenna. A complete analysis of the DBS signal should also include a discussion of the efficiency of the back-scattering process in relation to turbulence structures<sup>[116] [122] [123]</sup> and while a full discussion of this concept is outside the scope of this thesis, it would conclude that back-scattering efficiency peaks perpendicular to the magnetic field.

It can now be seen that combining the underlying surface velocity and the back-scattering efficiency we expect a 2-D DBS image of the plasma surface at a given frequency to form Doppler shifted poles. These poles would be aligned perpendicular to the magnetic field lines and thus allow measurement of the magnetic pitch angle. The intention of the SAMI-2 diagnostic is to simultaneously probe the plasma edge with two separate frequencies, and thus measure the pitch angle at two different locations. Measurement of the pitch angle can then be used to derive the edge current density using Amperes Law, shown in equations 4.7 and 4.8.

$$\mu_0 \mathbf{j} = \nabla \times \mathbf{B} \tag{4.7}$$

$$\mu_0 j_\phi = \frac{\delta B_Z}{\delta R} - \frac{\delta B_R}{\delta Z} \tag{4.8}$$

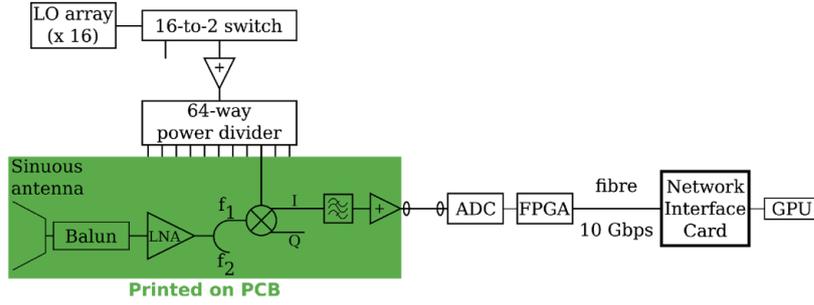


Figure 4.6: Simple block diagram of the SAMI-2 diagnostic showing the front end components that have been implemented on a custom PCB as described in “Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic”<sup>[15]</sup>

The pitch angle,  $\gamma$ , is related to the ratio of  $B_\theta$  (poloidal field component) and  $B_\phi$  (toroidal field component) and at the tokamak mid-plane where  $\theta = 0$  then it follows that:

$$\gamma = \arctan \frac{B_Z}{B_\phi} \quad (4.9)$$

Equation 4.9 can be used to calculate  $B_Z$  from the measured pitch angle and a value for  $B_\phi$  obtained from a basic equilibrium. The term  $\frac{\delta B_R}{\delta Z}$  in equation 4.8 can also be calculated without in-depth equilibrium reconstruction<sup>[124]</sup>. This results in a value for  $j_\phi$  (toroidal current density), which when calculated at multiple locations in the plasma edge will determine the edge current density gradient.

The requirement to receive two simultaneous frequencies does not materially affect the complexity of the digitisation system, other than doubling the number of signal lines that must be processed. This is not an issue, as previously discussed in section 4.2.1.1, as the signal processing logic can essentially just be copied for each additional signal line. It does however impact the resolution of the data that can be delivered, as the two ten gigabit ethernet outputs must now service double the number of channels.

### 4.3 SAMI-2 Analogue Front End

As mentioned in section 4.2.2.1 the SAMI-2 diagnostic must be able to measure the DBS signal from two probing frequencies simultaneously. A custom PCB has been designed, described in J. Allen et al, *Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic*<sup>[15]</sup>, to achieve this. A simple block diagram of the front end components attached to the data acquisition back-end is shown in figure 4.6

The antennas used for SAMI-2 (described in section 4.3.1) generate a signal from two orthogonal polarisations. These signals carry the information from the two probing beams and so these frequencies must be separated. To achieve this the signal is first fed through two mixers with Local Oscillators (LO) corresponding to the probing beam frequencies. Once the probing beam carrier wave has been subtracted the In-phase and Quadrature (I/Q) components are separated to allow the main beam side-bands to be distinguished.

Finally each signal is digitised by a 12-bit, 125 MHz ADC that is connected to the FPGA through a custom interface board. Polarisation, dual frequencies, and I/Q components, mean a total of eight signal channels must be digitised per antenna. The ADC-to-FPGA interface board is designed for 24 digital channels meaning that each FPGA is responsible for processing the data from three antennas.

#### 4.3.1 The Sinuous Antenna Design

As previously mentioned (see section 4.1.2) a key improvement of SAMI-2 over the original instrument is the capacity to separate polarisations and hence be able to distinguish between the O-mode and X-mode cut-off surfaces (for a brief description of O and X-mode waves see appendix D.1). The antenna must also be able to operate at up to 40 GHz so as to penetrate into the MAST-Upgrade pedestal as far as possible. As the resolution of the system scales with the square of the number of antennas, the more antennas that can be fitted within the available port space is also important, therefore the antenna cross-section would preferably be small. The original SAMI diagnostic used antennas that protruded toward the plasma surface causing interference through internal array reflections. Therefore a significant improvement would be the use of planar antennas flush to the array surface, eliminating the possibility of such inter-antenna reflections.

To fulfil these requirements a sinuous antenna printed on PCB has been chosen. The sinuous antenna design<sup>[125]</sup> is a broad spectrum planar antenna with dual arms for polarisation separation<sup>[16]</sup>. The designs of the antenna arms give it a small cross-section relative to the minimum frequency it is designed to receive while also minimising cross-talk between the polarisations. Figure 4.7 shows the sinuous antenna design with the different polarisation arms highlighted with different colours. This antenna design has been shown to be effective for applications where a large broadband and good polarisation separation is required<sup>[125]</sup>.



Figure 4.7: The sinuous antenna design has two distinct arms for polarisation separation<sup>[16]</sup>. The design of the antenna allows for a relatively small cross-section for the bandwidth it is designed to respond to. It is also a planar antenna allowing for it to be printed directly on to a PCB while minimising antenna cross-talk when incorporated as part of an array. The orthogonal arms of the antenna allow to differentiate between orthogonal polarised light. Figure taken from “Dual-polarisation broadband sinuous antenna and microstrip power divider design for the Synthetic Aperture Microwave Imager-2 diagnostic”<sup>[15]</sup>

## 4.4 Software and Firmware Components

### 4.4.1 10 Gigabit Ethernet

Defined by the Institute of Electrical and Electronics Engineers (IEEE) as protocol 802.3, ethernet describes the physical and first digital layer of the hardware and software required to transfer information between nodes on a network. Nodes of a network may be a computing device such as a personal computer or server, or parts of the network infrastructure such as network switches or repeaters. The ethernet protocol can be split into two distinct standards; (i) the physical layer (the system defining the physical medium that the data is transmitted over); (ii) the data link-layer (the system defining how data is encoded digitally and transferred across the

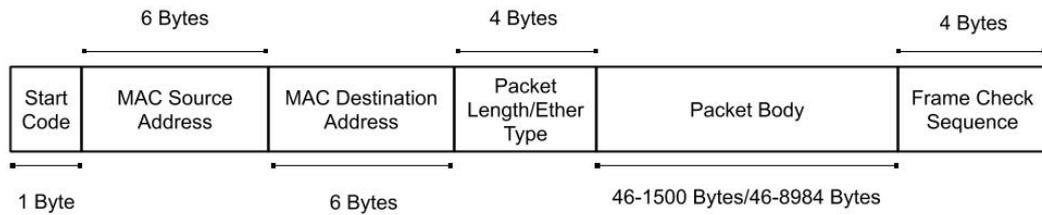


Figure 4.8: The frame composition of an ethernet data transfer is shown here as defined by IEEE protocol 802.3. Further protocols such as TCP/IP or UDP can then be used within the packet data. In standard ethernet usage the maximum frame size is limited to 1522 bytes, including header information and error checking, but larger frame formats up to 9000 bytes can be supported if the ethernet interface is configured in “Jumbo” mode.

physical connecting medium)<sup>[126]</sup>.

The physical layer is fairly broad and defines the electrical or optical properties of the network hardware. Many connection standards have been defined and used since the introduction of ethernet in the 1980’s with the most common connector types being RJ45 for copper cabling and SFP for optical fibres. Since the introduction of fast ethernet, when the data throughput achievable reached 100 Mb/s ethernet networks have employed a star topology, where devices in a network are connected through a central repeater or switch. While this approach may require more cabling than a direct connect implementation it introduces redundancy to the network, as if one cable fails a maximum of only one device is taken offline, as well as advanced routing techniques for greater network flexibility.

The 802.3 protocol also defines the specification for the data link layer, the first software layer for the transfer of packets between systems. The data link layer relates to how information is passed between nodes in a network and so sets out methods for determining data packet start and end as well as the possibility of including methods for error correction. The data link layer only cares about point to point communication, that is adjacent nodes in a network, and so it is indifferent to the final destination of the packet information.

The format of an ethernet packet as defined in IEEE protocol 802.3 is shown in figure 4.8. The first 14 bytes consist of the ethernet header, which contains the packet destination, source and ethernet type or length. For a standard ethernet frame the size is limited to 1522 bytes including the 14 byte header frame header and 4 byte error checking sequence at the end. Larger frames up to 9000 bytes can be supported if an interface is configured to send and receive them, with the packet length bytes in the header determining the packet frame size. A jumbo frame is

limited to 9000 bytes as larger packet labels would conflict with tags for different ether types. It is useful in this application to utilise the jumbo frame capabilities of the ethernet protocol as this reduces the number of interrupts that must be sent to the acquisition PC CPU telling it to collect data. Reducing the number of interrupts in turn increases the number of bytes than can be processed in any given amount of time.

Ether types are tags for denoting special forms of ethernet frames that a network interface must treat differently to a standard network transfer. While in theory it would be possible to define an ethernet tag for SAMI frames specifically this would require editing of the core ethernet driver on top of the development of the acquisition driver already being written. This would introduce unnecessary complexity as the jumbo frame size limit should be enough for SAMI-2 purposes.

#### **4.4.1.1 Precision Time Protocol**

One of these ether types is the Precision Time Protocol (PTP). This protocol is a standardised method for synchronising clocks across a network to enable high precision, time sensitive operations<sup>[127]</sup>. A PTP-enabled network requires hardware capable of generating high accuracy timestamps. While a specification exists for software time-stamping, the accuracy is low and this method would not be recommended for applications where sub nano-second precision is required.

Figure 4.9 shows a diagram of the message exchange between a primary and secondary interface. Before the clocks in a PTP enabled network can be synchronised the primary clock must be established. This is achieved using a standardised best Master Clock Algorithm which will determine the highest accuracy clock on the network. An announce message is used to state the clock statistics of each device and then, knowing its own local statistics, the devices can determine whether they are the primary or a secondary device. Once the primary interface has been decided it can now send the various messages that constitute PTP operations. A PTP message starts with a 34 byte header, demonstrated in table 4.1, which determines the type and contents of the message. The main types of message are synchronisation, follow-up and delay messages. For a version 2 type PTP network to initialise a PTP exchange the primary interface will send a synchronise message to its connected secondary devices. The synchronise message contains a zero timestamp, and is followed by a follow-up message which contains the timestamp of when the synchronise message was sent.

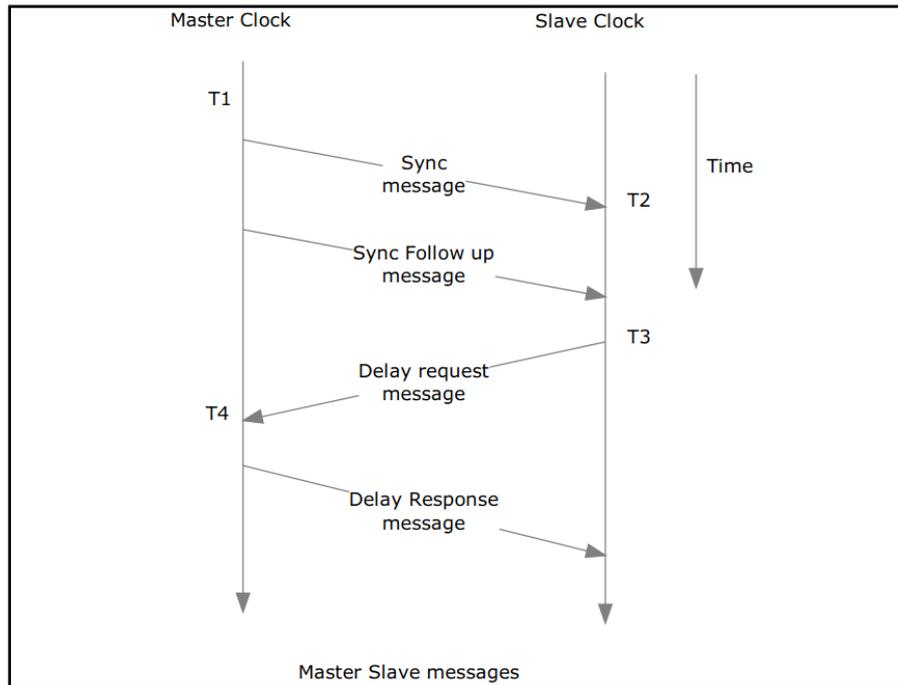


Figure 4.9: Timing diagram of a PTP synchronisation from the primary interface to a single secondary interface. Figure taken from EndRun Technologies white paper “Precision Time Protocol”<sup>[17]</sup>

Table 4.1: Table of the PTP header information. The PTP header comes directly after the ether type (as shown in figure 4.8) as part of the packet body. Table is reproduced from “Implementing IEEE 1588v2 for use in the mobile backhaul”<sup>[19]</sup>

Content	Length (Bytes)	Bytes from start	Description
Message Type	1	0	Defines the message type, examples are sync, follow-up, delay request etc.
PTP Version	1	1	Declares the PTP version, defining whether the sync message timestamp is delivered with the sync message or the follow-up message.
Message Length	2	2	Declares the length of the PTP message

Continued on next page

Table 4.1 – continued from previous page

Content	Length (Bytes)	Bytes from start	Description
Domain Number	1	4	Defines the clock domain that the message belongs to. An interface may be on the same network as multiple master clocks but will be grouped with systems all linked to the same master clock in a domain.
Reserved	1	5	Header information that is not user accessible.
Flags	2	6	Status flags to be read by the receiving interface.
Correction Field	8	8	Value in nanoseconds of the correction that must be applied to the secondary interface clock.
Reserved	4	16	Header information that is not user accessible.
Source Port Identity	10	20	PTP source port identity address.
Sequence ID	2	30	Contains the message sequence number.
Control Field	1	32	Similar to the message type the Control Field values depends on the message being sent.
Log Message Interval	1	33	Header field whose contents depend on the type of message being sent

### 4.4.2 Intel NIC Configuration and Control

The SAMI-2 acquisition PC will use Intel X710 Network Interface Cards (NIC) to acquire data from the FPGAs. To acquire data at the necessary rate the NICs selected operate at speeds of up to 10 gigabits per second. As a PCI-e 3.0 x8 slot operates at a maximum of 40 gigabits per second, to obtain the maximum number of data lines to the acquisition PC the 4 port version of the PCI-e card was selected.

The Intel X710 NICs are standard COTS devices and a Linux driver is provided as part of the main Linux kernel. This driver tells the system how to initialise the NICs and where to store and retrieve the ethernet frames received or transmitted by the interface. Direct Memory Access (DMA) is a protocol used to allow peripheral devices, such as PCI-e expansion cards, to directly access system memory for reading and writing. This is in contrast to serving an interrupt to the CPU and waiting until a thread is available to process the memory transaction.

Once DMA has been used to insert frames into memory the standard ethernet library functions can be used to read the collected packets and transfer the data to user available memory locations. The receive buffers, that are filled by DMA functions, can be up to 4000 frames in size which means that the CPU has plenty of time to be served an interrupt and process the memory transaction. As the read and write speed of system memory is an order of magnitude faster than that of the 10 gigabit ethernet connection this should give plenty of time for the CPU to complete the transfer operation.

### 4.4.3 Linux for Zynq Based Systems

Setting up Linux for an embedded system requires many components. A first stage boot loader (FSBL), stored in non-volatile memory, is used to bootstrap the system in a programmable state. A second stage boot loader is then required to set up system memory and initialise the operating system. It may take some user inputs and prepares system memory and basic networking. Once a second stage boot loader has set up the system memory a kernel image can be loaded with a root file system, and the device can now be booted into the available operating system. A graphical representation of how the different system components relate to each other is shown in figure 4.10.

The first stage boot loader is the first component in running Linux on an embedded system. Loaded onto non-volatile memory, attached directly to the CPU, the first stage boot loader will define the primary partition table of the device and tell the CPU what is and isn't available for writing to. It will also tell the CPU how

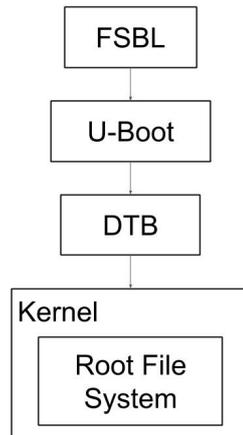


Figure 4.10: Graphical representation of how the Linux system components mesh to form a full operating system.

to initialise and address basic input and output hardware of the device such as memory and storage interfaces and network interfaces. The final operation of the first stage boot loader is loading the second stage boot loader in preparation for acceptance of a kernel image.

The second stage does a similar job to the first stage boot loader, but can contain more complex information, such as configuring and executing an operating system, as it is allowed to be much larger. Second stage boot loaders may give users access to basic hardware controls such as network interfaces and basic memory interfaces, but they must operate in a driver-less state and therefore do not have much functionality beyond that required to load the desired operating system. Some boot loaders, such as “Das U-Boot” operate as a first and second stage boot loader and are designed generically so as to be implementable on many embedded systems. Size of the boot loader is limited due not only to memory restrictions on embedded systems but also to speed up the transition of the system from down to operational.

As embedded systems can come attached with any number of varying peripherals an operating system must be told, what and where these peripherals are. To do so, embedded system developers must generate a device tree blob (DTB). A DTB is a file containing descriptions of the peripherals attached and available to the CPU. The description consists of clock information, the memory address used to configure and operate the peripheral, and other parameters specific to the particular piece of hardware.

The final stage of any fully functioning system is the root file system. The root file system contains the directories for storing the programs necessary within a modern

operating system. This includes the configuration files and various driver programs for operation of the attached hardware, as well as the directory maps linking these files.

Once the various software components have been generated they can be written to memory for use by the system-on-chip of the embedded hardware platform. For a lot of consumer technology this memory is generally in the form of a solid state device soldered to the platform. In the case of FPGAs it is more usual to load the firmware and software components to an SD card which can then be inserted into the FPGA hardware platform.

## 4.5 FPGA Based Acquisition and Control

The data acquisition and delivery system on the SAMI-2 diagnostic consists of two parts. An FPGA board interfaces with the receiving analogue electronics to acquire data from the Analogue-to-Digital converters. The FPGA then uses 10 Gigabit ethernet technology to transfer the data to an acquisition PC that uses a bespoke driver to store the raw ethernet frames in memory before eventual transfer to a further storage PC.

The FPGA acquisition system itself breaks down into two further segments. An ADC interface collects the frame and data signal from the ADC boards and generates a 128 bit data stream for delivery to on-board RAM using the Direct Memory Access (DMA) protocol. The processing system of the FPGA is limited to 8 gigabytes of memory, and with digitisation of twenty-four 12 bit data streams at 125 MHz this means only 1.8 seconds of raw data can be stored. This is further exacerbated by the fact that the AXI interface requires data to be packed into words that are a multiple of 8 bits, so the raw data must be packed to 16 bits. As shots on MAST-Upgrade are intended to run to 5 seconds this is not sufficient. Instead, an ethernet interface to the ADC acquisition modules has been designed to stream a reduced data rate at 10 Gbits/s for the entire length of a MAST-Upgrade shot. The work in this chapter concerns itself with the ethernet streaming interface and the associated receiving driver.

### 4.5.1 Xilinx 10 Gigabit Ethernet Subsystem

The SAMI-2 FPGA implementation uses the Xilinx 10G/25G Ethernet Subsystem IP<sup>[128]</sup> to provide two high speed 10 Gigabit ethernet interfaces for data transmission. The ethernet interfaces operate over the Small Form-factor Pluggable+

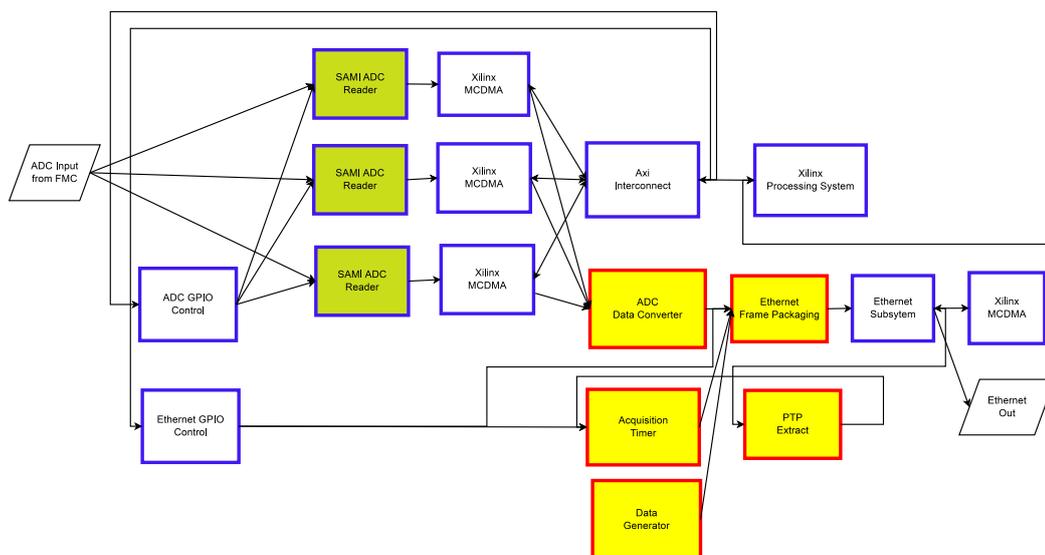


Figure 4.11: Flow chart of the programmable logic implemented on a Xilinx ZCU102 Evaluation Board for the SAMI-2 data delivery system. Yellow shows the IP cores developed as part of the work presented here. Green shows IP cores developed by Roddy Vann and white shows standard IP cores available as part of the Xilinx standard and proprietary library.

(SFP+) cages for 10 gigabit over fibre optic cable operation. These interfaces use gigabit transceivers, special hardware logic blocks on the FPGA, to operate at the super-gigabit speeds not achievable by the standard hardware available as part of the FPGA programmable logic.

The ethernet subsystem can be operated by AXI interface blocks to the ARM processor. Separate interfaces are required for the subsystem configuration and data delivery. As the subsystem only requires configuration on power-up this can be solely dealt with by the ARM connected AXI interface. This is in contrast to the previous generation 10 Gigabit Ethernet<sup>[129]</sup> subsystem that required specific configuration for operation with jumbo packets and so would have required intervention from the programmable logic.

For data delivery directly from the ADC interface module there must obviously be a direct connection to the ethernet subsystem. To allow the interface to be usable from the ARM processor, data must therefore be multiplexed between the two sources. The data interface to the ethernet subsystem follows the AXI4-Stream protocol<sup>[20]</sup> and is connected to the ARM processor through a Multi Cast DMA (MCDMA) Xilinx IP core. This allows the PL to write directly to the single 8 Gb RAM module attached to the FPGA board. The MCDMA core communicates with the ethernet subsystem data interface through FIFOs that perform the clock-crossing operation from the MCDMA core (on the system clock) to the data

interface (on the ethernet data clock). It is the FIFO interfaces that are multiplexed between the MCDMA core and the user generated core that packages data from the ADC interface.

Data to be sent over the ethernet subsystem must be pre-packaged into an ethernet frame, as demonstrated in figure 4.8. This is accomplished with a custom written module, described in section 4.5.1.1. The various clocks for the subsystem are generated internally by the IP, derived from a crystal oscillator designed for use with gigabit transceivers and a 150 MHz system clock. The generated receive and transmit clock are then looped back into the system to clock the various data interfaces. Only one ethernet subsystem has to be instantiated which supports two independent ethernet configuration and data interfaces. A PTP extraction module, described in section 4.5.1.3 is multiplexed with the receive FIFO to facilitate synchronisation of multiple FPGA boards.

#### 4.5.1.1 Ethernet Frame Packaging

The ethernet frame packaging module receives data from the ADC conversion module (described in section 4.5.1.2) and prepares it for transmission using the ethernet subsystem. Data entering the ethernet subsystem must be presented as part of a complete ethernet packet. This means that an ethernet source and destination address must be attached at the beginning of the frame and then followed by the frame length. The ethernet subsystem itself will add the ethernet packet start and stop code, but without the standard ethernet header information (as shown in figure 4.8) the ethernet packet will not be received by the destination interface. A state machine (as described in 1.2.5.1) is used to control the ethernet frame section that is to be generated. A block diagram of the FSM used in this module is shown in figure 4.12.

Components of the ethernet frame are generated in 64 bit (8 byte) segments, as this is the largest segment block accepted by the ethernet subsystem. As a MAC address is 48 bits long and the length component (also ether type, see section 4.4.1) is 16 bits long this means that two 64 bit frames of of the ethernet packet must be generated before data can be inserted. It is therefore necessary to buffer the incoming ADC data so that no data is lost while waiting for the address and length components to be output from the module.

Once the ethernet header has been generated there are still 16 bits (2 bytes) of data left within the 64 bit word that is used for the data interface. These 16 bits are used to hold the timestamp on which the acquisition trigger is received, generated by the PTP timestamping module (described in section 4.5.1.3). Once that the

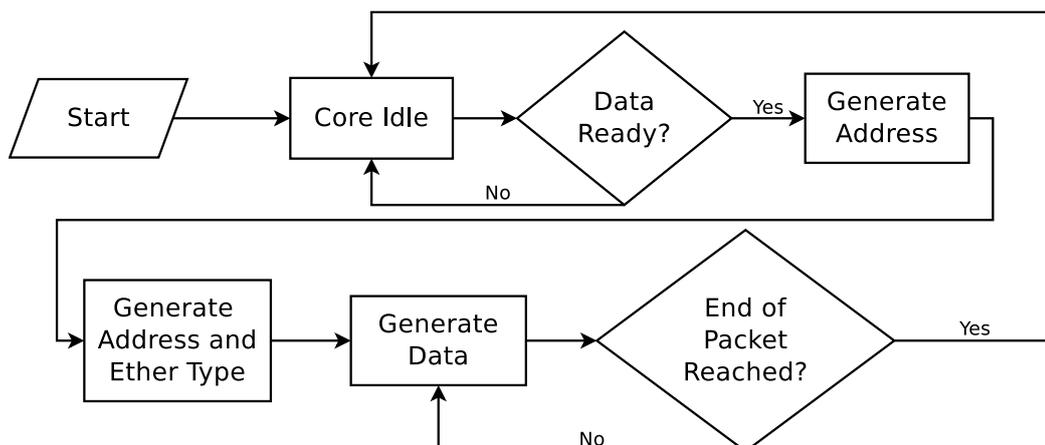


Figure 4.12: Flow chart of the ethernet frame packaging module. When the data packaging module reaches the end of a packet the last frame flag is asserted and the ethernet FIFO makes the data available to the ethernet subsystem.

ethernet header and timestamp have been delivered to the ethernet input FIFO, data can be written into the packet by the ethernet subsystem. The number of bytes delivered must be counted to ensure that the packet length is correct, else the receiving interface will register an error, or not be able to receive the packet at all.

To ensure correct operation of the ethernet subsystem the input FIFO is operated in packet mode. This means that while the data packaging module continuously fills the FIFO on the 150 MHz system clock, the FIFO waits until the last byte flag of the ethernet packet is made high before making the data available to the ethernet subsystem. As the FIFO supports asynchronous clocks the entire ethernet packet can be delivered at once on the subsystem clock, eliminating the requirement for clock crossing logic from the data clock to the subsystem clock.

#### 4.5.1.2 ADC Data Conversion

Data is generated by the ADC at a rate of 125 MHz with a 12 bit resolution. For each channel of data the ADC generates two signal lines of 6 bits in serial. This is then parallelised by an ADC interface. The ADC interface (developed by Roddy Vann) outputs the data as a 16 bit word to a DMA module for delivery to the ARM processor over a high-speed AXI4 interface (see section 1.2.6.1). The 16 bit word is produced by packing the 6 bit ADC segments with two zero bits and then concatenating the two 8 bit words. Each 16 bit data segment is therefore organised as shown in figure 4.13 and output to the DMA module on the ADC clock of 125 MHz.

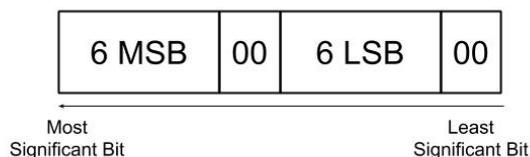


Figure 4.13: Graphical representation of how the ADC 12 bit data sample is packed into the 16 bit output of the ADC interface module.

As previously discussed (see section 4.1.2), each FPGA of the SAMI-2 data acquisition system will process 24 channels of data. At 12 bits per channel at a rate of 125 MHz this equates to 36 Gigabits of data, which is obviously too much for the two 10 gigabit ethernet data links. To allow headroom within the 10 gigabit ethernet protocol it has been decided that at first the data will be downsampled to 4 bits at 125 MHz for a data bandwidth of 6 gigabits per ethernet link.

However, packaging the data into 64 bit words for use with the ethernet subsystem is now more complicated. This is because while it is easy enough to see that the 24 data channels can be split into 12 channels for each ethernet stream, simply packaging the 4 bits from each of the 12 channels and transmitting that would result in only 48 bits of the 64 bit word containing data. It would also result in sending a 64 bit word at a rate of 125 MHz for a data bandwidth of 8 gigabits per second, defeating the purpose of reducing the data resolution to 4 bits to allow headroom in the first place. Instead a buffering system has been created to ensure that the full 64 bit word is used and therefore an ethernet line rate of 6 gigabits per second is achieved. Figure 4.14 shows a graphical representation of how the buffering loop is achieved.

As each data segment is buffered multiple times within this module it also facilitates crossing from the 125 MHz clock to the 150 MHz system clock that the ethernet packaging module operates on (described above in section 4.5.1.1). However, the uniqueness of the buffering loop makes it very particular to this data configuration and major modifications would be necessary for operation with 5 bit and 6 bit downsampling. Neither 5 or 6 are a factor of 64 meaning that while the 4 bit downsampling can be achieved by splitting the channel data into 4 bit words and concatenating, an implementation for 5 or 6 bit words would also include splitting of individual elements. Not only does this complicate the implementation of a buffering loop, it would also complicate the collation of data on the acquisition PC as ethernet frames are stored as 32 bit characters (again, neither 5 or 6 being factors of 32).

Synthetic aperture imaging has been shown to work with as little as 1-bit resolu-

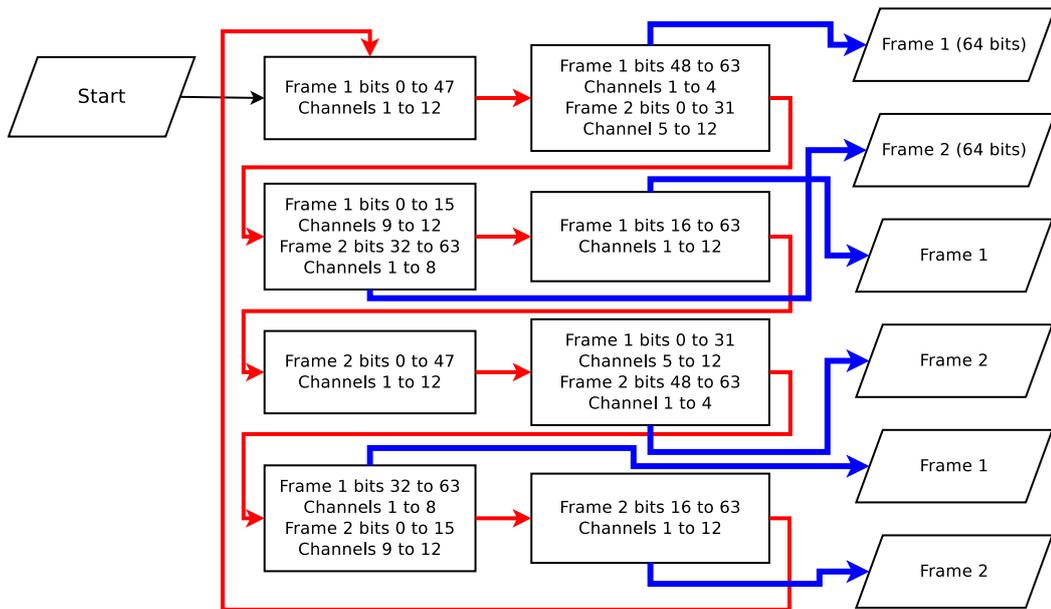


Figure 4.14: Diagram of the frame buffering loop in the ADC Data conversion module. Blue lines show points in the loop where a complete frame has been generated and output to the ethernet data packaging module while red lines show the buffering loop path.

ution<sup>[130]</sup>. While it is not clear what effect downsampling will have on the final resolution the method is still expected to work with a 4-bit bandwidth. To investigate the effect of downsampling, the maximum amount of raw data will also be stored and compared against the streamed downsampled data.

### 4.5.1.3 PTP Timestamping and Synchronisation

A PTP timestamping module has been implemented to extract data from incoming PTP packets (described in section 4.4.1.1) and generate a synchronised timestamp across the separate FPGA acquisition boards. The acquisition trigger for all the FPGA boards is delivered from the MAST-Upgrade control system along a single fibre line and then distributed by an internal break-out board. While in theory this should mean that all the FPGAs receive the trigger simultaneously, in practice it is highly likely that the distributed trigger for each board will have different delays due to differences in electrical paths.

The beam forming technique described in section 4.2.1.1 relies on cross-correlation of antenna data. This means that antenna signals must be synchronised precisely to be able to form accurate images of the plasma surface. Knowing exactly, relative to each other, when the acquisition boards receive the trigger is therefore essential for the beam forming technique.

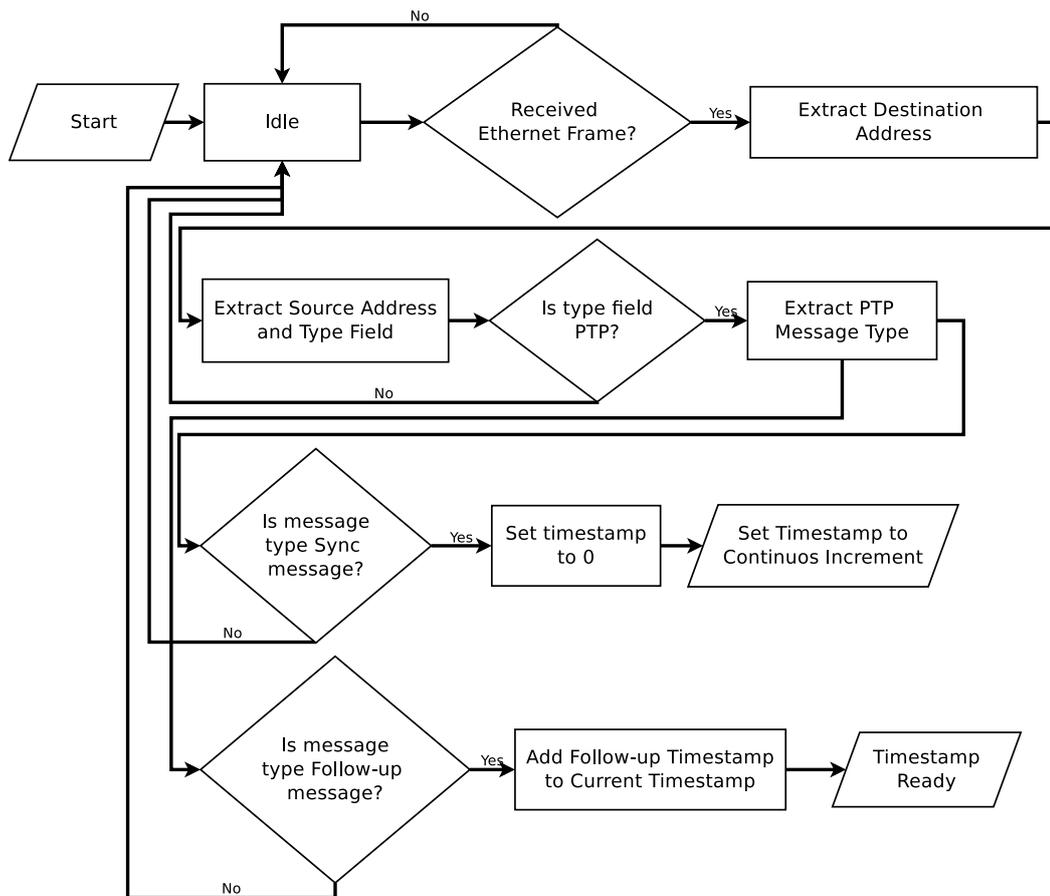


Figure 4.15: Flow chart for the PTP timestamping module used to extract a timestamp generated by the acquisition PC. The timestamp is then inserted into the ethernet packet to facilitate data synchronisation across multiple FPGA acquisition boards.

As mentioned in section 4.4.1.1 the PTP handshaking protocol consists primarily of synchronisation and follow-up messages. The PTP module FSM, as shown in figure 4.15 is therefore designed to reset the timestamp to zero on a synchronisation message and increment the time counter. When the follow-up message is received the new timestamp is added to the free running timestamp of the PTP module, and it is only once the follow-up message has been delivered to all the boards that they can be assumed to be synchronised. The PTP timestamp at the trigger time is then inserted into the ethernet packet to be transmitted so that when the data arrives at the acquisition PC it can be aligned relative to the other packets.

The PTP module uses a 300 MHz clock to increment the timestamp counter, which enables the sub-8 ns accuracy of the PTP protocol. To allow ethernet functionality within the Linux system the ethernet subsystem data outputs are simultaneously connected to the DMA module that interfaces with the ARM processor and the

PTP module. While the ethernet subsystem does contain logic to enable native PTP timestamping, this would only be beneficial for attaching a timestamp to the outgoing packets. This is because the native PTP logic does not determine if a packet is a PTP packet and only presents the potential timestamp as an output. Therefore even using the native PTP logic, the current PTP packet filtering would still have to be used to determine when synchronisation and follow-up packets have been received.

## 4.5.2 Data Acquisition PC

The technical specification for the data acquisition PC is constrained, mostly by the high data throughput that requires multiple ethernet-over-fibre add-in cards. As mentioned multiple times in the preceding material, the full SAMI-2 data acquisition system will require twenty 10 gigabit fiber pairs for data delivery from the FPGA boards to the acquisition PC. The Intel X710 NIC (described in section 4.4.2) has a 4 port variant and so five of these PCI-e add-in cards will be used for data acquisition. Each NIC requires eight PCI-e lanes to operate at the full 40 gigabit bandwidth. This means that the CPU and motherboard combination used for the SAMI-2 acquisition PC must support a total of forty PCI-e lanes and expansion slots for at least 5 full width PCI-e expansion cards.

While almost all Intel Xeon CPUs<sup>[131]</sup> support up to forty four PCI-e lanes as standard, motherboards with more than three full length PCI-e lanes are rare. Further compounding the constraints is that to fulfil all the PCI-e requirements that a COTS system delivers, the PCI-e lanes of single socket motherboards are distributed around many peripherals such that the lanes connecting the PCI-e interface to the CPU are multiplexed with other high-speed interfaces. Multiplexing PCI-e lanes from different cards could become a potential bottleneck and so it would be preferable to avoid this. Therefore the SAMI-2 acquisition PC uses a dual socket motherboard, specifically the ASUS WS C621E SAGE<sup>[18]</sup>, that supports forty non-multiplexed PCI-e lanes divided between the two CPUs. The block diagram for the motherboard is shown in figure 4.16. It can be seen that by carefully selecting PCI-e slots 1, 2, 3, 5, and 6, PCI-e multiplexing can be avoided.

Choice of CPU was fairly straight forward as the Intel Xeon Silver 4208 was the lowest cost CPU compatible with the motherboard with hyper-threading available. Hyper-threading means that each core (the CPU component that executes instructions) can operate two simultaneous threads for greater multi-tasking capacity. While the threads share the low-level, high-speed, L3 cache and so are not completely distinct, they can generally be viewed as being able to execute distinct



Finally a standard 750 watt ATX PC power supply<sup>[132]</sup> provides headroom for the expected 300 watt load and leaves possibility for adding a GPU in the future. An In-Win R400-01N 4U<sup>[133]</sup> standard rack mount server case allows installation of the acquisition PC within the server racks available on MAST-Upgrade with additional room for adding storage in the future.

### 4.5.3 Network Interface Card (NIC) Driver

To acquire data from the FPGA a program has been written to acquire data concurrently from all the 10 gigabit ethernet interfaces. This has been achieved using standard C++ libraries and it is hoped that the program layout is flexible for ease of customisation in the future. The program consists of three files, the main program file, a thread function file and a socket class file. These are compiled using the “GNU make” framework<sup>[134]</sup> such that small changes to the individual files do not mean a complete re-compilation of the entire program.

To interface directly with a socket the “sys/socket.h” and “sys/types.h” standard libraries must be used to instantiate and configure each necessary ethernet interface<sup>[135]</sup>. To reduce any processing overhead, and therefore be capable of reaching the maximum bandwidth, the ethernet interfaces are configured using the RAW protocol. A list of the protocols and their descriptions can be found as part of the socket library documentation<sup>[135]</sup>. Using ethernet interfaces with the RAW protocol means that no processing is done on the packet when it is transferred from the receive buffer to main memory. The recall from the receive buffer to main memory is accomplished with the “recvfrom”<sup>[136]</sup> function, also part of the “sys/socket.h” and “sys/types.h” libraries.

A socket configuration class has been written to standardise socket operations such as socket configuration and sending and receiving packets. The socket class means that instantiation of socket objects can be automated and hence easily adjusted for the number of ethernet interfaces that need to be operated. This is useful functionality as the preliminary application of the SAMI-2 diagnostic will operate with a reduced number of antennas, compared to the full complement of thirty that the system will eventually run with. Socket class has methods for:

- Configuration  
Sockets are configured to operate using the “SOCK\_RAW” protocol. This is in contrast to configuration with the (more generally used) “SOCK\_DGRAM” protocol that would pre-process incoming packets to remove the source and destination address to leave just the ethernet packet body. While it may seem

preferable to automatically strip the ethernet header information to prevent having to do it manually, it is not trivial to pass the correct source and address information the “recvfrom” function for correct packet identification.

Once a socket has been instantiated it is configured with the required interface MAC address and the promiscuous mode flag is set. Promiscuous mode requires the interface to receive all packets that arrive instead of just those that are addressed to the socket. Alternatively, once the system is in an operational stage the point-to-point flag can be set with the specific destination MAC address.

- Send Packets

The send packet method is used for testing purposes so that the bandwidth limits of the interfaces can be determined. This method is not expected to be used during diagnostic operation as the FPGA Linux system should be able to interface with the ethernet packets directly and therefore use data transfer methods that are normally used between Linux systems.

- Receive Packets

The receive packet method utilises the “recvfrom” function to transfer individual packets from the receive buffer (described in section 4.4.2) into system memory. The real-time acquisition packets are generated by the ethernet packaging module (described in section 4.5.1.1) and so a different MAC address can be used to the Linux MAC address from the FPGA to distinguish the packets. Linux systems also send loopback packets through an interface and so each incoming packet is checked for the correct source address from the corresponding FPGA interface. Once the ethernet header information has been removed from the packet buffer the data is transferred to the thread function for concatenation with the full data array.

- Send PTP Packet

The send PTP packet method constructs a PTP synchronisation and follow-up packet using the formatting described in section 4.4.1.1. This is used by the FPGA, as described in section 4.5.1.3, for synchronisation of the FPGA boards so the cross-correlations between the antennas can be calculated accurately.

To achieve the maximum possible data bandwidth, each interface must be operated independently so that ethernet frames can be transferred from their respective receive buffers to system memory uninterrupted. The POSIX thread library functions<sup>[137]</sup> were chosen to instantiate a separate instruction thread for each interface.

To operate the thread function takes a thread object, a function to be executed by the thread and arguments for the thread function. A thread function has been written that takes a socket class object as an argument and then executes the receive packets class method for that socket object. This means that each interface is operated identically and that different class methods can be executed within the thread function while abstracting the method code to the socket class. Compartmentalising the program into these separate files makes the individual scripts that are used in the program more readable and helps keep individual file sizes small which is useful when debugging.

#### **4.5.4 Component testing**

While it is not possible to test the entire system end-to-end until all the hardware components are ready, it is possible to test individual parts in isolation. Each of the FPGA modules that has been described has been simulated using the Xilinx Vivado development suite to ensure that the logic is operating as intended. Modules are then connected as they would be in the final system and it has been verified that the components also produced the expected behaviour when working in concert.

Separately, the send and receive functions described as part of the NIC software driver can be used with each other to check that they operate correctly. Data transfer speeds of up to nine gigabits per second have been achieved using this method. While this is less than what should be possible for the interface, it is sufficient for the six gigabit rate that is expected of the initial SAMI-2 implementation.

## **4.6 Summary**

A real-time data acquisition system has been designed for the SAMI-2 diagnostic to overcome the on-board data limitations of the FPGA for a maximum acquisition time of five seconds. This corresponds to the maximum shot length that is expected on the MAST-Upgrade tokamak. The SAMI-2 diagnostic will help in exploring the edge current density to further understanding of the L-to-H mode transition as well as other plasma edge related phenomena.

FPGA modules have been designed to downsample data in real-time across a clock boundary, package data for the state-of-the-art 10 gigabit ethernet protocol and synchronise data across multiple boards using the Precision Time Protocol. Not only are the data rates shown here some of the highest currently used in fusion but a solution for FPGA board synchronisation has been shown using a standard network

protocol. Once development of the front end electronics has been completed the real-time data acquisition system shown here can be tested. The diagnostic will then be mounted on the MAST-Upgrade tokamak.

#### **4.6.1 Future Work**

The first improvement to the system that can be made is a reduction of the down-sampling to better utilise the full 10 gigabit bandwidth. This will mean alterations to the ADC conversion module (described in section 4.5.1.2), though the current frame buffering system can still be utilised. More sophisticated downsampling and filtering techniques can be implemented, such as frequency downsampling and the application of digital filters to retain the full bit width by reducing the temporal resolution.

Currently only two SFP+ ports are utilised of the four available on the ZCU102 Development board. Significant future work on the SAMI system may include the use of the other two SFP+ ports to a dedicated analysis box containing more FPGAs. This would facilitate real-time processing of the magnetic pitch angle as well as applications to gyrotron steering for radiative heating.

---

## Concluding Remarks

What follows is a summary of the material presented in this thesis, with particular note made of the novel aspects that have been contributed to fusion research. Next is an overview of fusion diagnostics in general and how they will continue to be utilised going forward. Finally, a brief review is made of FPGA technology as applied to the ITER tokamak, the next major step in realising a commercial fusion power plant.

### 5.1 Thesis Overview

It is clear that diagnostic instrumentation is a critical part of any scientific research. The scale and complexity of fusion experiments means that this is especially true with a wide range of diagnostic techniques being employed. To facilitate research it is also important that diagnostic systems have comprehensive data acquisition, processing, and diagnostic control methods. Presented within this body of work are three distinct diagnostic instruments that nonetheless have used similar workflows in the development of their acquisition, processing, and control systems. This is important as the reduced development time that comes with the similar workflows means reduced cost and greater time spent on innovative research. Shared architectures mean that advantages are also to be found in system maintenance and improvement.

Field Programmable Gate Array technology continues to provide a flexible platform for custom instrumentation with unique I/O and data processing requirements. This makes FPGAs ideal for research environments where bespoke one-off instrumentation is regularly required. The re-programmability also makes in-house developed FPGA based instrumentation easy to improve as the functionality can be adjusted as the diagnostic requirements change over time. In short time-scales dia-

agnostics may also be initially implemented with reduced functionality and upgraded as and when the opportunity presents itself. Finally, while in-house development of instrumentation can take more time than commissioning from an outside source, the advantages in maintenance cannot be understated. In-house developed instruments can be maintained on site due to the available knowledge base and this can reduce costs in maintenance and reduce downtime, critical to research environments where experiment time is limited.

Zynq is the architecture name for the integration of an FPGA and a Central Processing Unit as a System-On-Chip developed by Xilinx. Alternative architectures are available from Intel in the form of their Agilex, Arria, and some Stratix FPGA technology families<sup>[138]</sup>. Packaging these two technologies in such an integrated form factor reduces the complexity and increases the performance of FPGA instrument control with more traditional computer software systems (accessible via the ARM CPU). To facilitate communication between ARM and FPGA components the AXI interface standard has been defined by the IEEE. By integrating this protocol into IP cores, ARM to FPGA interactions can be standardised making collaboration and resource sharing significantly easier compared to previous applications of Linux for FPGA control.

The Fission Chambers for Neutron Flux Measurements acquisition system and the real-time biasing Langmuir probe prototype use the low-cost Red Pitaya development board and the Koheron SDK to implement adaptable state-of-the-art instrumentation. Utilising the open-source Koheron SDK allows for a quick and simple workflow for designing high performance diagnostics. While still making use of the free version of the Vivado toolkit, the Koheron SDK allows complex FPGA development with the associated support software without buying into the expensive Vivado software eco-system. The Red Pitaya development board is a low-cost Zynq based FPGA system that provides high speed Analogue-to-Digital and Digital-to-Analogue converters making it a good choice for high performance electronic signal sampling applications. Both instruments have used the Red Pitaya FPGA and Koheron SDK workflow to great effect in creating high performance instrumentation with significant advantages over previous systems.

While ARM processors are generally regarded as components for low power systems (such as phones, IoT devices and light-weight embedded systems), the FPGA incorporated as part of the Zynq SoC allows integration of state-of-the-art networking and high speed I/O usually only possible with high-end bespoke instrumentation. The original SAMI diagnostic employed a processor in PL using the MicroBlaze architecture<sup>[52]</sup>. By offloading the Linux control system to the attached ARM processing cores the entire FPGA fabric is available for data acquisition and pro-

cessing. The ARM processors can also run the Linux based control software faster and more reliably than a MicroBlaze implementation, using a standard Ubuntu image with all the advantages in access to software that that entails.

### 5.1.1 Fission Chambers for Neutron Flux Measurements

Presented here has been the development of a new digital electronic acquisition system for dual fission chambers to be implemented on the Mega Amp Spherical Tokamak Upgrade fusion experiment. This system improves the time resolution of neutron flux measurements compared to the old system and brings development in-house to improve diagnostic flexibility and reduce maintenance costs. The acquisition system also incorporates a web server implemented with Flask for delivery of data and diagnostic configuration over the MAST-U network. This is an added feature over the previous fission chamber acquisition system. Firmware and software components are backed-up and available using the on-site Gitlab version control system. The acquisition system has been tested and calibrated at the National Physical Laboratory with multiple sources and data from the commissioning of the Neutral Beam Injectors is also presented.

While the fission chamber signal processing presented here is not the first digital implementation of real-time DC current and Campbell mode calculation<sup>[139]</sup>, it does show a novel, low latency, light-weight method for calculating the signal variance<sup>[140]</sup>. The expected linear response has been verified with testing at the National Physical Laboratory and will be the first COTS built digital fission chamber acquisition system implemented on a fusion experimental device.

### 5.1.2 Digital Implementation of the Mirror Langmuir Probe Control System

Presented here is the development of a proof-of-concept real-time biasing Langmuir probe control system. Much of this work has been previously published<sup>[74]</sup>. The proof-of-concept shows real-time biasing capability of a Langmuir probe with respect to plasma temperature. This will facilitate high temporal resolution measurements of plasma temperature fluctuations to provide new insights into the physics of plasma edge turbulence and filamentary structures.

The use of FPGAs allowed this project to be implemented quickly and at low cost. The proof-of-concept firmware is open source and available as a public Github repository to facilitate collaboration and further development. While the FPGA concepts used to build the system are not new, they have been applied to create an

advanced diagnostic that should be easy to implement on plasma devices. In fact, this has already been shown by further work completed by W. McCarthy and will be presented in a future publication prepared for submission to Review of Scientific Instruments.

### **5.1.3 The Synthetic Aperture Microwave Imaging Diagnostic**

The Synthetic Aperture Microwave Imaging 2 diagnostic will measure the plasma edge magnetic pitch angle simultaneously in two locations so as to infer the edge current density. This will facilitate physics studies into many edge phenomena such as ELMs, the L-H mode transition and edge turbulence. The work that has been presented here enables the SAMI-2 diagnostic by using FPGA technology to stream in real-time the acquired data from 30 polarised antennas into an acquisition PC using the standard ethernet protocol and commercially available digital electronics.

This is a continuation of the initial SAMI instrument that used 8 non-polarisation discriminating antennas to measure the magnetic pitch angle of the plasma edge at a single depth. The technology used for the new SAMI-2 instrument is cutting edge with an initial data bandwidth far exceeding other instrumentation installed on the MAST-Upgrade diagnostic. Challenges included configuring the FPGA for 10 Gigabit networking and designing the software for data capture and storage on the data acquisition PC. Due to the extremely high data bandwidth requirements and novelty of the project this would not have been possible within this timescale without the use of FPGA technology. Also presented here are innovations in generating low-latency downsampled data for a 64-bit AXI interface as well as a novel implementation of the PTP timestamping protocol for multiple FPGA board synchronisation.

## **5.2 Outlook**

It is clear that as FPGA technology becomes cheaper and more accessible it will continue to play a critical role in the development of instrumentation for fusion experimental devices. Open source workflows will also contribute to increasing the accessibility of FPGA development, hopefully leading to greater innovation in the field. Standardisation of hardware and software components will facilitate collaboration as instruments become easier to adapt and implement for a variety of machines.

A major advance of the past few years is the integration of programmable logic and conventional computer system components into a System-on-Chip architec-

ture. This symbiosis will continue to proliferate into fusion research as increasing complexity in diagnostics and control systems drives more complex software requirements. Also to consider is the use of Local Area Networks for holistic diagnostic management. Data security and the increasing number of digital diagnostic control systems mean that LANs are a critical aspect of Tokamak control and management. Zynq-like architectures allow diagnostics to present web interfaces to the wider instrument network without requiring a separate control PC. Not only does this reduce the space and power requirements of the experiment overall, but it also means that control and data acquisition protocols can be standardised across devices, reducing the complexity of running the experiment.

However, although there are many advantages in speed and energy efficiency as gate technology reduces in size, components become more susceptible to radiation damage. This is magnified by the increase in surface area that accompanies the increased number of components on silicon when creating a Zynq-like SoC. Further work is needed on how FPGA based components within the bio-shield degrade over time with respect to performance and reliability.

### 5.2.1 ITER Tokamak Diagnostics

The ITER tokamak is the next step in tokamak evolution<sup>[141]</sup>. It has been designed to be the first fusion device to reach thermal break-even with an eventual thermal output of 500 MW for input heating energy of 50 MW. FPGA technology will be critical part of many planned diagnostics on ITER, used as part of the neutron diagnostic suite<sup>[142]</sup>, the real-time plant protection system<sup>[143]</sup>, X-ray diagnostics<sup>[144]</sup>, and many others. Such is the pervasive use of FPGA technology that the researchers working as part of the ITER project have undertaken studies into how the harsh radiation environment will affect the reliability and longevity of the various FPGA hardware platforms<sup>[145]</sup>. The prolific use of FPGA based instrumentation throughout the ITER ecosystem shows that FPGA technology will continue to be a major part of instrumentation development within the fusion community for the foreseeable future.

---

# Introduction

## A.1 Future of Fusion Energy

This appendix serves as an overview of a paper<sup>[27]</sup>, written with other members of the fusion Centre for Doctoral Training, on wider aspects of fusion energy research outside of the main focus of this thesis. The motivation behind the work presented in the paper is to examine in depth the role fusion power plants will play when the technology becomes widely available, with the aim of advising fusion research policy. Policy suggestions are made that might facilitate the penetration of future fusion power plants into a future carbon neutral energy grid.

As is often quoted when talking about fusion in public, “Fusion is only 30 years away, and always will be”. This phrase, that has plagued the credibility of the fusion community for the past 50 years, shows that we must be careful in our messaging going forward<sup>[146][147]</sup>. The current motivations for fusion power are its clean energy credentials and the low-level waste criterion. This criterion states that after 100 years from decommissioning a fusion power plant will not leave any long lived nuclear waste (> 100 years) of a radioactivity level greater than “intermediate” classification. By examining current research this paper aims to show that these constraints can and should be relaxed to enable the penetration of fusion power as an energy source.

Many proponents of fusion power state that its low carbon credentials mean that fusion based power plants can be a solution to climate change and the issue of increasing CO<sub>2</sub> and green-house gas emissions<sup>[23]</sup>. While this is theoretically true, it is only relevant if climate change is still an issue by the time commercial fusion power becomes available. The IPCC has estimated that pollution causing global warming must reach net-zero by 2050<sup>[22]</sup> to avoid risking irreversible effects. Given the development time involved for a fusion power plant it is reasonable to ask

therefore if fusion can usefully contribute to a net-zero carbon solution within that time-frame.

The EU fusion roadmap states that a power plant that can demonstrate electricity to the grid will be operational approximately 20 years after the current European fusion research machine, ITER, has demonstrated energy gain from fusion<sup>[148]</sup>. As ITER is not slated to operate at full power until 2035 this puts the first demonstration of fusion power to the grid in 2055, a full 5 years after the global transition to a net-zero carbon energy infrastructure. It is therefore highly unlikely that fusion based power plants will contribute to an initial de-carbonised grid.

Instead fusion power would enter a market most likely dominated by renewable energy<sup>[27]</sup>. To fit within this market fusion would most likely have to demonstrate some capability over the currently assumed role of a “base-load” power source. This could be accomplished with some form of load following or the application of fusion power to other industries such as heat generation for either commercial or industrial applications.

The second tenet of fusion power is the low-level waste criterion. Set by fusion programs in the 1970s to improve fusion’s image against fission, the low-level waste criterion states that neutron irradiated material from a fusion power plant must be able to be disposed of in low-level waste repositories 100 years after decommissioning. However, recent research has shown that the structural steel specifically designed to meet this specification, EUROFER97, will in fact produce intermediate-level waste, as defined by the UK government, with a half-life of over a thousand years.

As the mass of structural material used within a plant is quite high it is possible that a future fusion power plant built using EUROFER97 would produce more intermediate waste than an equivalent nuclear fission based reactor. While the regulations defining radioactive waste disposal are not universal this does pose a significant problem. One solution would be for a set of global regulations to be defined for the safe construction and decommissioning of fusion power plants, a solution in line with the extremely co-operative research and development that fusion has experienced over the last couple of decades.

Overall the paper finds that while there will still be a place for fusion in future energy markets thanks to its high energy density and low proliferation attributes, significant strides can be made to more accurate messaging and research into how fusion power can be most usefully applied. It may also be possible and necessary to relax the nuclear waste requirements while still maintaining a significantly smaller environmental impact relative to conventional nuclear power.

---

## Fission Chambers for Neutron Flux Measurements

Thermal Neutron Latency in Fission Chamber shielding ]Electron Bernstein waves can be characterised as having wavelengths of the order of four times the gyro-radius, and must be sufficiently high frequency such that the ion motion can be ignored

Thermal Neutron Latency in Fission Chamber shielding The neutrons generated by fusion reactions must travel from the plasma where they are generated to the fission chamber electrode to be detected. During this travel neutrons may be scattered and lose energy, and therefore speed. To precisely measure the neutron flux for a  $1 \mu\text{s}$  period then it is important to ensure that the neutron travel time is less than the sampling time, preferably by an order of magnitude or more. This is so that the neutron flux measured by the fission chamber system can be reliably correlated to the corresponding timestamp.

For an averaging time of  $1 \mu\text{s}$  it has already been established (see section 2.2.3) that a delay of greater than 200 ns would have an undesirable effect on the temporal resolution of the system. It is relatively trivial to set a lower limit on the neutron energy that would definitely introduce more than 200 ns of latency. The fission chambers are approximately 2m from the centre column, which to first order we can assume is the average distance to the plasma bulk. For a neutron to travel 2 m in 200 ns or less it must have an energy of approximately 500 keV or greater. While this is a significantly greater energy than the Cadmium layer is designed to adsorb<sup>[149]</sup>, modelling using the neutron code MCNP<sup>[150]</sup> has shown that the time resolution is sufficiently maintained.

## B.1 Code snippets from the Fission Chamber Acquisition system

### B.1.1 Mean and Variance Calculation Code

The following code snippet shows how the mean and variance parameters reported by the fission chamber data processing system are calculated using the VHDL programming language. An example of the pipelining technique described in section 1.2.5 can be seen in processes “square\_proc\_1”, “square\_err\_proc”, and “cross\_err\_data\_proc”, where individual components of the sum of the squares are calculated and then summed in “stat\_proc\_2”.

```
— Process to sum the initial downsampled data and store the
— sample period sum at the end of the period
stat_proc_1 : process(adc_clk)
begin
  if rising_edge(adc_clk) then
    if clk_en = '1' or acquire_pulse = '1' then
      — incorporating the error component
      mean_sum_hold <= mean_sum + errSum(12 downto 5);
      — Starting the initial sum
      mean_sum <= to_signed(to_integer(cutData), mean_sum'length);
    else
      — Adding consequent samples
      mean_sum <= mean_sum + to_signed(to_integer(cutData), mean_sum'length);
    end if;
  end if;
end process;

square_proc_1 : process(adc_clk)
begin
  if rising_edge(adc_clk) then
    — Calculating the square of the Most Significant Bits for the variance parameter
    square_hold <= to_unsigned(to_integer(cutData*cutData), square_hold'length);
  end if;
end process;

square_err_proc : process(adc_clk)
begin
  if rising_edge(adc_clk) then
    — Calculating the square of the Least Significant Bits for the variance parameter
    err_squ_hold <= to_unsigned(to_integer(errData*errData), err_squ_hold'length);
  end if;
end process;

cross_err_data_proc : process(adc_clk)
begin
  if rising_edge(adc_clk) then
    — Calculating the cross multiplication for the variance parameter
    cross_err_hold <= shift_left(to_signed(to_integer(errData)*to_integer(cutData),
    cross_err_hold'length), 1);
  end if;
end process;

— Summing the square of the samples for the variance parameter
stat_proc_2 : process(adc_clk)
begin
  if rising_edge(adc_clk) then
    — Counter times the end calculation relative to the end of the sample period
    if proc_count = 1 then
      square_sum_hold <= square_sum;
      square_sum <= to_unsigned((to_integer(square_hold) +
```

```

        to_integer(err_squ_hold(9 downto 5)) +
        to_integer(cross_err_hold(13 downto 5)), square_sum'length);
    else
        square_sum <= square_sum +
        to_unsigned((to_integer(square_hold) + to_integer(err_squ_hold(9 downto 5)) +
        to_integer(cross_err_hold(13 downto 5))), square_sum'length);
    end if;
end if;
end process;

-- Multiplying by the number of samples for the variance parameter
calc_squares_1 : process(adc_clk)
begin
    if rising_edge(adc_clk) then
        if proc_count = 2 then
            sample_part <= to_unsigned(to_integer(stat_counter_hold)*to_integer(square_sum_hold),
            sample_part'length);
        end if;
    end if;
end process;

--Calculating the square of the means for the variance parameter
calc_squares_2 : process(adc_clk)
begin
    if rising_edge(adc_clk) then
        if proc_count = 1 then
            sum_part <= to_unsigned(to_integer(mean_sum_hold*mean_sum_hold), sum_part'length);
        end if;
    end if;
end process;

-- Integrating the variance paramater components into the final value
var_calc : process(adc_clk)
begin
    if rising_edge(adc_clk) then
        if proc_count = 3 then
            var_hold <= to_unsigned(to_integer(sample_part) - to_integer(sum_part),
            var_hold'length);
        end if;
    end if;
end process;

-- Setting the module outputs and ready signal for parameter delivery
set_vars : process(adc_clk)
    variable start_count : std_logic := '0';
begin
    if rising_edge(adc_clk) then
        calcDone <= '0';
        if clk_en = '1' then
            mean_buf <= std_logic_vector(mean_sum(15 downto 0) + errSum(12 downto 5));
            start_count := '1';
        end if;
        if start_count = '1' then
            proc_count <= proc_count + 1;
            if proc_count = 4 then
                start_count := '0';
                var_buf <= std_logic_vector(var_hold(27 downto 12));
                calcDone <= '1';
                proc_count <= 0;
            end if;
        end if;
    end if;
end process;

```

## B.2 MAST-U DATA C

Every FPGA based device is connected to the MAST-Upgrade Operations network and has a corresponding Unit program that collects and then distributes data from the device. Figure B.1 shows a schematic of how the FPGA device sits on the network, along with the unit program, plasma control system, scheduler codes and central storage.

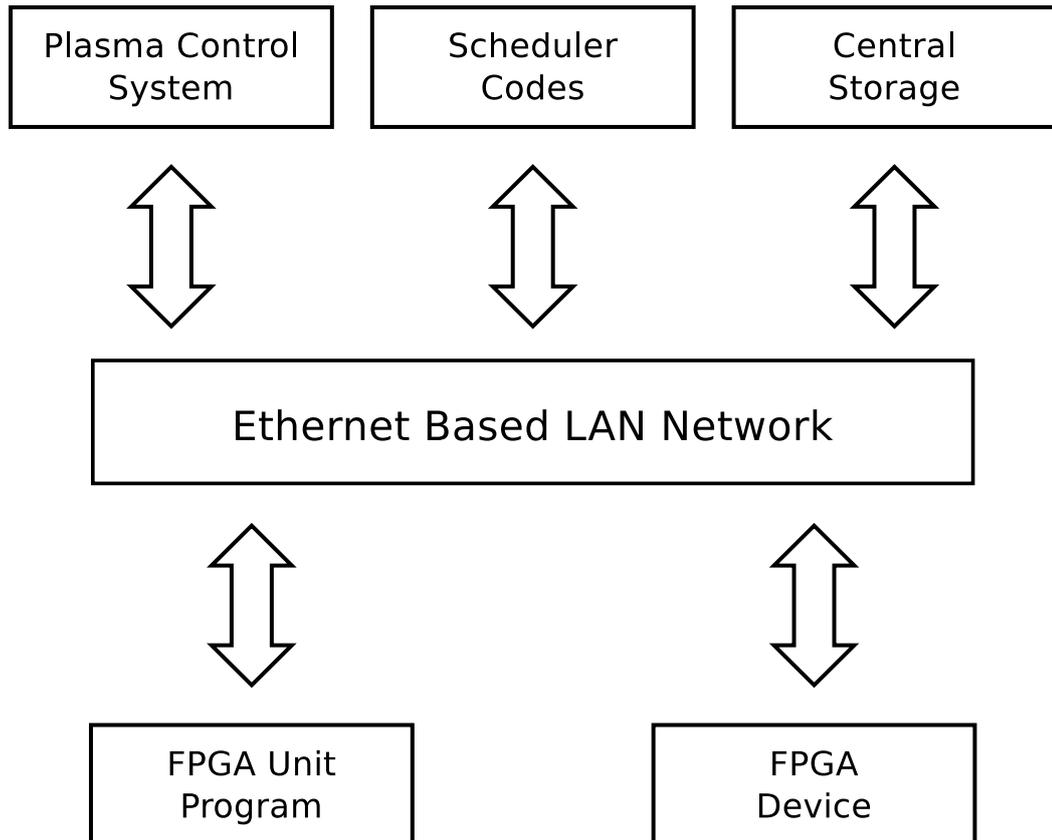


Figure B.1: Schematic of the MAST-Upgrade Operations network and how FPGA devices fit into the network. The unit program interfaces directly with the FPGA device and passes data to the scheduler code for processing and storage.

FPGA based devices are only directly interfaced by the unit program, which collects and distributes data from the diagnostic. Before a shot the unit program will ready the device for data collection. In the case of the fission chamber acquisition system this sets a software arm that clears the previous shot data in preparation for receiving the next shot. After a shot has been run the unit program will then collect data from the instrument using the HTTP “GET” protocol. The diagnostic then holds this data in memory, to be re-collected in the case of errors, until the software arm has been reset.

Once the unit program has collected data it is processed into the netcdf-4 format and sent to the scheduler code for routine post shot processing. Currently the analysed data for the fission chamber system consists of converting the raw signal into the real DC current signal and the real variance values. Once calibration parameters have been obtained a value for the neutron flux will also be returned.

---

# Digital Implementation of the Mirror Langmuir Probe Control System

## C.1 Code snippets from the Real-time Langmuir Probe Biasing System

The following is a code snippet for the temperature calculation module in the Mirror Langmuir Probe control system. The other modules follow basically the same layout with minor adjustments depending on the specific equation. Of note is the timing process to set the divider inputs and the process to retrieve data from the divider and convert it into a look-up table address.

```

-- purpose: Process to do core reset
-- type   : sequential
-- inputs : adc_clk, clk_rst, iSat_guess
-- outputs: Temp
reset_proc : process (adc_clk) is
begin -- process reset_proc
  if rising_edge(adc_clk) then -- rising clock edge
    if clk_rst = '1' then -- synchronous reset (active high)
      Temp_proxy <= to_signed(Temp_guess, 16);
    else
      if output_trigger = '1' then
        if Temp_mask > to_signed(1200, Temp_mask'length) then
          Temp_proxy <= to_signed(1200, 16);
        elsif Temp_mask < to_signed(20, Temp_mask'length) then
          if Temp_mask < to_signed(0, Temp_mask'length) then
            Temp_proxy <= to_signed(Temp_guess, 16);
          else
            Temp_proxy <= to_signed(20, 16);
          end if;
        else
          Temp_proxy <= Temp_mask(15 downto 0);
        end if;
      end if;
    end if;
  end if;
end process;

```

```

        data_valid <= '1';
    else
        data_valid <= '0';
    end if;
end if;
end if;
end process reset_proc;

-- purpose: Process to calculate Temperature value
-- type   : combinational
-- inputs : adc_clk
-- outputs: saturation current
Temp_proc : process (adc_clk) is
begin
    if rising_edge(adc_clk) then
        if exp_en = '1' then
            if calc_switch = "01" then
                Temp_mask <= shift_right(difference_hold * signed(BRAMret), 9);
            elsif calc_switch = "10" then
                Temp_mask <= shift_right(to_signed(to_integer(difference_hold) *
                    to_integer(unsigned(BRAMret)), 32), 13);
            elsif calc_switch = "00" then
                Temp_mask <= shift_right(difference_hold * signed(BRAMret), 9);
            end if;
            output_trigger <= '1';
        else
            output_trigger <= '0';
        end if;
    end if;
end process Temp_proc;

-- purpose: Process to calculate the difference between bias and floating voltage
-- type   : sequential
-- inputs : adc_clk, waitBRAM, storeSig, storeSig2
-- outputs: difference_hold
difference_proc : process (adc_clk) is
begin -- process difference_proc
    if rising_edge(adc_clk) then -- rising clock edge
        if waitBRAM = '1' then -- synchronous reset (active high)
            difference_hold <= storeSig - storeSig2;
        end if;
    end if;
end process difference_proc;

-- purpose: process to set the divisor and dividend for the divider
-- type   : combinational
-- inputs : adc_clk
-- outputs: divisor, dividend, tUser
div_proc : process (adc_clk) is
    variable divisor_mask : signed(13 downto 0) := (others => '0');
begin -- process diff_proc
    if rising_edge(adc_clk) then
        if clk_en = '1' then
            -- Setting the variables to go into the division
            divisor_mask := signed(iSat(13 downto 0));
            if divisor_mask /= to_signed(0, 14) then
                divisor_tdata <= "00" & std_logic_vector(divisor_mask);
            else
                divisor_tdata <= "00" & std_logic_vector(to_signed(iSat_guess, 14));
            end if;
            dividend_tdata <= "00" & std_logic_vector(shift_right(signed(volt_in), 0));
            dividend_tvalid <= '1';
            divisor_tvalid <= '1';
            storeSig <= signed(volt2);
            storeSig2 <= signed(vFloat);
        else

```

```

        dividend_tvalid <= '0';
        divisor_tvalid  <= '0';
    end if;
end if;
end process div_proc;

-- purpose: process to set the BRAM address for data retrieval.
-- type   : combinational
-- inputs : adc_clk
-- outputs: BRAM_addr, waitBRAM
BRAM_proc : process (adc_clk) is
    variable divider_int : integer range -8191 to 8191 := 0;
    variable divider_rem : integer range -2047 to 2047 := 0;
    variable addr_mask   : integer range 0 to 16383   := 0;
begin -- process BRAM_proc
    if rising_edge(adc_clk) then
        if index = '1' then
            -- Extracting the integer part and the fractional part returned by the
            -- divider core to use in the bram address mapping
            divider_rem := to_integer(signed(divider_tdata(11 downto 0)));
            divider_int := to_integer(signed(divider_tdata(25 downto 12)));
            int_store   <= divider_int;
            rem_store   <= divider_rem;
            case divider_int is
                when -1 =>
                    addr_mask := 0;
                    calc_switch <= "01";
                when 0 =>
                    addr_mask := 2048 + divider_rem;
                    if addr_mask < 2048 then
                        calc_switch <= "01";
                    else
                        calc_switch <= "00";
                    end if;
                when 1 =>
                    addr_mask := 4096 + divider_rem;
                    calc_switch <= "10";
                when 2 =>
                    addr_mask := 6144 + divider_rem;
                    calc_switch <= "10";
                when 3 =>
                    addr_mask := 8192 + divider_rem;
                    calc_switch <= "10";
                when 4 =>
                    addr_mask := 10240 + divider_rem;
                    calc_switch <= "10";
                when 5 =>
                    addr_mask := 12288 + divider_rem;
                    calc_switch <= "10";
                when 6 =>
                    addr_mask := 14336 + divider_rem;
                    calc_switch <= "10";
                when others =>
                    if divider_int < to_signed(-1, 14) then
                        addr_mask := 0;
                        calc_switch <= "01";
                    elsif divider_int >= to_signed(7, 14) then
                        addr_mask := 16383;
                        calc_switch <= "10";
                    end if;
            end case;
            addr_mask_store <= addr_mask;
            BRAM_addr      <= std_logic_vector(to_unsigned(addr_mask, 14));
            waitBRAM       <= '1';
        else
            waitBRAM <= '0';
        end if;
    end if;
end process BRAM_proc;

```

```
    end if;
  end if;
end process BRAM_proc;

-- purpose: process to collect bram data after address is set by division module
-- type   : combinational
-- inputs : adc_clk
-- outputs: exp_ret, exp_en
collect_proc : process (adc_clk) is
begin -- process collect_proc
  -- Setting a collection tick to get the right block ram memory back once
  -- the address has been assigned.
  if rising_edge(adc_clk) then
    if waitBRAM = '1' then
      exp_count <= exp_count + 1;
    end if;
    if exp_count = 1 then
      exp_count <= exp_count + 1;
    elsif exp_count = 2 then
      exp_en <= '1';
    end if;
    if exp_en = '1' then
      exp_count <= 0;
      exp_en <= '0';
    end if;
  end if;
end process collect_proc;
```

---

# The Synthetic Aperture Microwave Imaging Diagnostic

## D.1 Electron Bernstein Wave X-O Mode Conversion

Electron Bernstein waves are electrostatic waves that occur in a plasma. They are defined as high frequency longitudinal oscillations and can be derived from the plasma wave equations in the absence of a magnetic field<sup>[151]</sup>. Electron Bernstein waves can be characterised as having wavelengths of the order of four times the gyro-radius, and must be sufficiently high frequency such that the ion motion can be ignored

O and X-mode waves correspond to orthogonal polarised light that propagates perpendicular to the magnetic field. Due to differences in the dispersion relation, the cut-off density, and therefore the edge penetration depth, for O and X mode waves is slightly different. This difference in penetration depth could be significant enough to cause interference in resolving the backscattered image as described in section 4.2.1.1, hence the need for polarisation separation antennas so that the different backscattered locations can be distinguished.

The SAMI diagnostic was originally intended to be used in a passive mode to detect emission of O and X-mode waves that had been generated by electron Bernstein waves<sup>[152]</sup>. Electron Bernstein waves can be mode converted and escape the plasma in the form of microwaves. Depending on the exact plasma conditions at the point of mode conversion, either an O or X-mode wave will be emitted. These conditions can be determined from the dispersion relationships that result from electromagnetic wave propagation perpendicular to the magnetic field through a plasma. As the mode conversion conditions are determined by the local plasma parameters, including the magnetic pitch angle, the SAMI diagnostic could measure

the edge current density without an active probing beam.

This appendix is to further explain the need for polarisation antennas, and is not intended as an in-depth explanation of Electron-Bernstein waves and their mode conversion properties. For a full discussion of mode conversion of microwaves in the plasma edge and how they apply to the application of the SAMI diagnostic when operated in a passive mode, see the theses of Simon Freethy<sup>[153]</sup>, David Thomas<sup>[154]</sup>, and Matthew Thomas<sup>[110]</sup>. The many papers written on the application of the previous SAMI diagnostic would also be highly informative<sup>[14] [107] [152]</sup>.

---

## Bibliography

- [1] Fastfission. Binding energy curve - common isotopes.svg. URL [https://commons.wikimedia.org/wiki/File:Binding\\_energy\\_curve\\_-\\_common\\_isotopes.svg](https://commons.wikimedia.org/wiki/File:Binding_energy_curve_-_common_isotopes.svg). Accessed Jun. 15, 2020, generated with data from <https://www.einstein-online.info/wp-content/uploads/table-of-nuclear-data.txt>.
- [2] Jef Ongaena. Fusion: A true challenge for an enormous reward. *EPJ Web of Conferences*, 98:05004, 01 2015. doi: 10.1051/epjconf/20159805004.
- [3] M. Kikuchi, K. Lackner, and M. Q. Tran. *Fusion Physics*. International Atomic Energy Agency, 2012. ISBN 978-92-0-130410-0. Page 22.
- [4] Wayne. M. Saslow. *Electricity, Magnetism and Light*. Elsevier Inc., 2002. ISBN 978-0-12-619455-5. Page 44.
- [5] EUROfusion. URL [https://www.euro-fusion.org/fileadmin/user\\_upload/Archive/wp-content/uploads/2011/09/jg05-537-1c.jpg](https://www.euro-fusion.org/fileadmin/user_upload/Archive/wp-content/uploads/2011/09/jg05-537-1c.jpg). Accessed Jun. 29, 2020.
- [6] S Lisgo, Marko Istenic, J Canik, R Buttery, I Katramados, M Kotschenreuther, M Kovari, S Mahajan, M Shannon, D Taylor, and P Valanju. Super-x advanced divertor design for mast upgrade. *36th EPS Conference on Plasma Physics 2009, EPS 2009 - Europhysics Conference Abstracts*, 33, 01 2009.
- [7] Xilinx. Ultrascale architecture configurable logic block. Xilinx User Guide, UG574 (v1.5) February 28, 2017.
- [8] Chetvorno. Turnstile state machine colored.svg. URL [https://commons.wikimedia.org/wiki/File:Turnstile\\_state\\_machine\\_colored.svg](https://commons.wikimedia.org/wiki/File:Turnstile_state_machine_colored.svg). Accessed Dec. 23, 2020.

- [9] M.F L'Annunziata. *Handbook of Radioactivity Analysis*. Elsevier Inc., 2004. ISBN 978-0-12-436603-9.
- [10] B. LaBombard and L. Lyons. Mirror Langmuir probe: A technique for real-time measurement of magnetized plasma conditions using a single Langmuir electrode. *Review of Scientific Instruments*, 78(7):073501, 2007. doi: 10.1063/1.2754392. URL <https://doi.org/10.1063/1.2754392>.
- [11] P. Lang, Alberto Loarte, G. Saibene, Larry Baylor, M. Becoulet, M. Cavinato, S. Clement-Lorenzo, Edward Daly, T. Evans, M.E. Fenstermacher, Y. Gribov, L.D. Horton, C. Lowry, Y.R. Martin, Olaf Neubauer, N. Oyama, Michael Schaffer, Derek Stork, W. Suttrop, and Oliver Schmitz. Elm control strategies and tools: Status and potential for iter. *Nuclear Fusion*, 53:043004, 03 2013. doi: 10.1088/0029-5515/53/4/043004.
- [12] H R Wilson, S C Cowley, A Kirk, and P B Snyder. Magneto-hydrodynamic stability of the h-mode transport barrier as a model for edge localized modes: an overview. *Plasma Physics and Controlled Fusion*, 48(5A):A71–A84, apr 2006. doi: 10.1088/0741-3335/48/5a/s06. URL <https://doi.org/10.1088/0741-3335/48/5a/s06>.
- [13] M. Fitzgerald, L.C. Appel, and M.J. Hole. EFIT tokamak equilibria with toroidal flow and anisotropic pressure using the two-temperature guiding-centre plasma. *Nuclear Fusion*, 53(11):113040, oct 2013. doi: 10.1088/0029-5515/53/11/113040. URL <https://doi.org/10.1088/0029-5515/53/11/113040>.
- [14] D.A. Thomas, K.J. Brunner, S.J. Freethy, B.K. Huang, V.F. Shevchenko, and R.G.L. Vann. 2d doppler backscattering using synthetic aperture microwave imaging of MAST edge plasmas. *Nuclear Fusion*, 56(2):026013, jan 2016. doi: 10.1088/0029-5515/56/2/026013. URL <https://doi.org/10.1088/0029-5515/56/2/026013>.
- [15] J. Allen, C. H. Vincent, F. J. Hambly, S. Bale, and R. Vann. Dual-polarisation broadband sinuous antenna and microstrip power divider design for the synthetic aperture microwave imager-2 diagnostic. In *Proceedings of the 14th International Reflectometry Workshop for Fusion Plasma Diagnostics (IRW14)*, 2019. URL [https://www.epfl.ch/research/domains/swiss-plasma-center/wp-content/uploads/2019/05/IRW14\\_Book\\_Of\\_Abstracts.pdf](https://www.epfl.ch/research/domains/swiss-plasma-center/wp-content/uploads/2019/05/IRW14_Book_Of_Abstracts.pdf).
- [16] R.H. DuHammell. Dual polarized sinuous antennas, 1985.

- [17] EndRun Technologies. Precision time protocol. Technical report.
- [18] ASUSTeK COMPUTER INC. Ws c621e sage series. Technical report, IEEE, February 2020. URL [https://dlcdnets.asus.com/pub/ASUS/mb/Socket3647/WS\\_C621E\\_SAGE/Manual/E16279\\_WS\\_C621E\\_SAGE\\_UM\\_V7\\_WEB.pdf](https://dlcdnets.asus.com/pub/ASUS/mb/Socket3647/WS_C621E_SAGE/Manual/E16279_WS_C621E_SAGE_UM_V7_WEB.pdf).
- [19] Calnex Solutions Ltd. Implementing ieee 1588v2 for use in the mobile backhaul. Technical report.
- [20] Advanced microcontroller bus achitecture overview. <https://developer.arm.com/architectures/system-architectures/amba>. Accessed: 23rd March 2021.
- [21] Red Pitaya Organisation. <https://www.redpitaya.com/>, 2018. Accessed: 23rd March 2021.
- [22] V. Masson-Delmotte, P. Zhai, H.O. Pörtner, D. Roberts, J. Skea, P.R. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J.B.R Matthews, Y. Chen, X. Zhou, M.I. Gomis, E. Lonnoy, T. Maycock, M. Tignor, and T. Waterfield. IPCC, 2018: Global warming of 1.5°C. an IPCC special report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. *In Press.*, 2018.
- [23] C. Bustreo, U. Giuliani, D. Maggio, and G. Zollino. How fusion power can contribute to a fully decarbonized european power mix after 2050. *Fusion Engineering and Design*, 146:2189 – 2193, 2019. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2019.03.150>. URL <http://www.sciencedirect.com/science/article/pii/S0920379619304831>. SI:SOFT-30.
- [24] Nasa technology transfer program. URL <https://technology.nasa.gov/>. Accessed: 23rd March 2021.
- [25] Eso technology transfer. URL <https://www.eso.org/public/industry/techtrans/>. Accessed: 23rd March 2021.
- [26] J. Binks. Fascination, issue 11. Science and Technology Facilities Council.
- [27] T.E.G. Nicholas, T.P. Davis, F. Federici, J. Leland, B.S. Patel, C. Vincent, and S.H. Ward. Re-examining the role of nuclear fusion in a renewables-based energy mix. *Energy Policy*, 149:112043, 2021. ISSN 0301-4215. doi: <https://doi.org/10.1016/j.enpol.2021.112043>.

- [//doi.org/10.1016/j.enpol.2020.112043](https://doi.org/10.1016/j.enpol.2020.112043). URL <http://www.sciencedirect.com/science/article/pii/S0301421520307540>.
- [28] H. A. Bethe. Energy production in stars. *Phys. Rev.*, 55:434–456, Mar 1939. doi: 10.1103/PhysRev.55.434. URL <https://link.aps.org/doi/10.1103/PhysRev.55.434>.
- [29] J. Wesson. *Tokamaks*. Clarendon Press-Oxford, 2004. ISBN 0-19-8509227. Third Edition.
- [30] J. Lindl. Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Physics of Plasmas*, 2(11):3933–4024, 1995. doi: 10.1063/1.871025. URL <https://doi.org/10.1063/1.871025>.
- [31] J. L. Johnson. The stellarator approach to toroidal plasma confinement. *Nuclear Technology - Fusion*, 2(3):340–361, 1982. doi: 10.13182/FST82-A20768. URL <https://doi.org/10.13182/FST82-A20768>.
- [32] The JET Team (presented by F X Söldner). Shear optimization experiments with current profile control on JET. *Plasma Physics and Controlled Fusion*, 39(12B):B353–B369, dec 1997. doi: 10.1088/0741-3335/39/12b/027. URL <https://doi.org/10.1088/0741-3335/39/12b/027>.
- [33] D.F.H. Start. Non-inductive current drive in tokamaks (report on the IAEA technical committee meeting,culham, 1983). *Nuclear Fusion*, 23(10):1411–1417, oct 1983. doi: 10.1088/0029-5515/23/10/016. URL <https://doi.org/10.1088/0029-5515/23/10/016>.
- [34] D.J. Battaglia, M.W. Bongard, R.J. Fonck, and A.J. Redd. Tokamak startup using outboard current injection on the pegasus toroidal experiment. *Nuclear Fusion*, 51(7):073029, jun 2011. doi: 10.1088/0029-5515/51/7/073029. URL <https://doi.org/10.1088/0029-5515/51/7/073029>.
- [35] Qingquan Yu, Daqing Zhang, and Xuemin Shen. Hybrid inductive current drive—non-inductive current drive for a tokamak reactor. *Fusion Engineering and Design*, 26(1):619 – 622, 1995. ISSN 0920-3796. doi: [https://doi.org/10.1016/0920-3796\(94\)00228-Y](https://doi.org/10.1016/0920-3796(94)00228-Y). URL <http://www.sciencedirect.com/science/article/pii/092037969400228Y>. Proceedings of the Fifth International Toki Conference on Plasma Physics and Controlled Nuclear Fusion.
- [36] A. Sykes, J.-W. Ahn, R. Akers, E. Arends, P. G. Carolan, G. F. Counsell, S. J. Fielding, M. Gryaznevich, R. Martin, M. Price, C. Roach, V. Shevchenko,

- M. Tournianski, M. Valovic, M. J. Walsh, and H. R. Wilson. First physics results from the mast mega-amp spherical tokamak. *Physics of Plasmas*, 8(5):2101–2106, 2001. doi: 10.1063/1.1352595. URL <https://doi.org/10.1063/1.1352595>.
- [37] M. Gryaznevich, R. Akers, P. G. Carolan, N. J. Conway, D. Gates, A. R. Field, T. C. Hender, I. Jenkins, R. Martin, M. P. S. Nightingale, C. Ribeiro, D. C. Robinson, A. Sykes, M. Tournianski, M. Valovic, and M. J. Walsh. Achievement of record  $\beta$  in the start spherical tokamak. *Phys. Rev. Lett.*, 80:3972–3975, May 1998. doi: 10.1103/PhysRevLett.80.3972. URL <https://link.aps.org/doi/10.1103/PhysRevLett.80.3972>.
- [38] JaeChun Seol, B.H. Park, S.S. Kim, J.Y. Kim, and Yong-Su Na. Electron cyclotron heating during ECRH-assisted pre-ionization in a tokamak. *Nuclear Fusion*, 50(10):105008, aug 2010. doi: 10.1088/0029-5515/50/10/105008. URL <https://doi.org/10.1088/0029-5515/50/10/105008>.
- [39] G.M. Voss, S. Davis, A. Dnestrovskij, A. Kirk, P.J. Knight, M. Loughlin, M.H. O'Brien, D. Sychugov, A. Tabasso, and H.R. Wilson. Conceptual design of a component test facility based on the spherical tokamak. *Fusion Engineering and Design*, 83(10):1648 – 1653, 2008. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2008.05.002>. URL <http://www.sciencedirect.com/science/article/pii/S0920379608000914>. Proceedings of the Eight International Symposium of Fusion Nuclear Technology.
- [40] M. Yamauchi, T. Nishitani, K. Ochiai, Y. Morimoto, J. Hori, K. Ebisawa, S. Kasai, and C. Walker. Development of in-vessel neutron monitor using micro-fission chambers for iter. *Review of Scientific Instruments*, 74(3):1730–1734, 2003. doi: 10.1063/1.1534397. URL <https://doi.org/10.1063/1.1534397>.
- [41] R.A. Tinguely, A. Rosenthal, R. Simpson, S.B. Ballinger, A.J. Creely, S. Frank, A.Q. Kuang, B.L. Linehan, W. McCarthy, L.M. Milanese, K.J. Montes, T. Mouratidis, J.F. Picard, P. Rodriguez-Fernandez, A.J. Sandberg, F. Sciortino, E.A. Tolman, M. Zhou, B.N. Sorbom, Z.S. Hartwig, and A.E. White. Neutron diagnostics for the physics of a high-field, compact,  $q \geq 1$  tokamak. *Fusion Engineering and Design*, 143:212 – 225, 2019. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2019.03.148>. URL <http://www.sciencedirect.com/science/article/pii/S0920379619304818>.

- [42] H. M. Mott-Smith and Irving Langmuir. The Theory of Collectors in Gaseous Discharges. *Phys. Rev.*, 28:727–763, Oct 1926. doi: 10.1103/PhysRev.28.727. URL <https://link.aps.org/doi/10.1103/PhysRev.28.727>.
- [43] N. C. Hawkes, K. Blackler, B. Viacoz, C. H. Wilson, J. B. Migozzi, and B. C. Stratton. Design of the joint european torus motional stark effect diagnostic. *Review of Scientific Instruments*, 70(1):894–897, 1999. doi: 10.1063/1.1149415. URL <https://doi.org/10.1063/1.1149415>.
- [44] A Murari, E Joffrin, R Felton, D Mazon, L Zabeo, R Albanese, P Arena, G Ambrosino, M Ariola, O Barana, M Bruno, L Laborde, D Moreau, F Piccolo, F Sartori, F Crisanti, E de la Luna, J Sanchez, and EFDA-JET Contributors. Development of real-time diagnostics and feedback algorithms for JET in view of the next step. *Plasma Physics and Controlled Fusion*, 47(3):395–407, feb 2005. doi: 10.1088/0741-3335/47/3/001. URL <https://doi.org/10.1088/0741-3335/47/3/001>.
- [45] Griffin Lacey, Graham W. Taylor, and Shawki Areibi. Deep learning on fpgas: Past, present, and future, 2016.
- [46] B. Mealy and F. Tappero. *Free Range VHDL*. Creative Commons Attribution-ShareAlike UnportedLicense, 2018.
- [47] A. Boutros, S. Yazdanshenas, and V. Betz. Embracing diversity: Enhanced dsp blocks for low-precision deep learning on fpgas. In *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*, pages 35–357, 2018.
- [48] Hana Kubátová. *Finite State Machine Implementation in FPGAs*, pages 175–184. Springer US, Boston, MA, 2005. ISBN 978-0-387-28327-2. doi: 10.1007/0-387-28327-7\_15. URL [https://doi.org/10.1007/0-387-28327-7\\_15](https://doi.org/10.1007/0-387-28327-7_15).
- [49] Institute of Electrical and Electronics Engineers. IEEE 1076-2019 - IEEE Standard for VHDL Language Reference Manual. Technical report, IEEE, September 2019. URL <https://standards.ieee.org/standard/1076-2019.html>.
- [50] Ieee standard for verilog hardware description language. *IEEE Std 1364-2005 (Revision of IEEE Std 1364-2001)*, pages 1–590, 2006. doi: 10.1109/IEEESTD.2006.99495.
- [51] Inc Xilinx. Zynq-7000 soc - technical reference manual. Technical report, July 2018. URL [https://www.xilinx.com/support/documentation/user\\_guides/ug585-Zynq-7000-TRM.pdf](https://www.xilinx.com/support/documentation/user_guides/ug585-Zynq-7000-TRM.pdf). UG585.

- [52] Inc Xilinx. Microblaze processor reference guide. Technical report, May 2019. URL [https://www.xilinx.com/support/documentation/sw\\_manuals/xilinx2019\\_1/ug984-vivado-microblaze-ref.pdf](https://www.xilinx.com/support/documentation/sw_manuals/xilinx2019_1/ug984-vivado-microblaze-ref.pdf). UG984.
- [53] Inc Xilinx. Axi reference guide. Technical report, March 2011. URL [https://www.xilinx.com/support/documentation/ip\\_documentation/ug761\\_axi\\_reference\\_guide.pdf](https://www.xilinx.com/support/documentation/ip_documentation/ug761_axi_reference_guide.pdf). UG761.
- [54] Koheron SDK. <https://www.koheron.com/software-development-kit>. Accessed: 23rd March 2021.
- [55] M. Cecconello, M. Turnyanskiy, S. Conroy, G. Ericsson, E. Ronchi, S. Sangaroon, R. Akers, I. Fitzgerald, A. Cullen, and M. Weiszflog. A neutron camera system for mast. *Review of Scientific Instruments*, 81(10):10D315, 2010. doi: 10.1063/1.3479038. URL <https://doi.org/10.1063/1.3479038>.
- [56] T. Vasilopoulou, I.E. Stamatelatos, P. Batistoni, A. Colangeli, D. Flaminio, N. Fonnesu, S. Loreti, B. Obryk, M. Pillon, and R. Villari. Improved neutron activation dosimetry for fusion. *Fusion Engineering and Design*, 139:109 – 114, 2019. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2019.01.002>. URL <http://www.sciencedirect.com/science/article/pii/S092037961930002X>.
- [57] Seungyon Cho, Mu-Young Ahn, Cheol Woo Lee, Eung Seon Kim, Yi-Hyun Park, Youngmin Lee, and Dong Won Lee. Feasibility study of fusion breeding blanket concept employing graphite reflector. *Fusion Engineering and Design*, 98-99:1747 – 1750, 2015. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2015.02.033>. URL <http://www.sciencedirect.com/science/article/pii/S0920379615001167>. Proceedings of the 28th Symposium On Fusion Technology (SOFT-28).
- [58] A.R Raffray, M Akiba, V Chuyanov, L Giancarli, and S Malang. Breeding blanket concepts for fusion and materials requirements. *Journal of Nuclear Materials*, 307-311:21 – 30, 2002. ISSN 0022-3115. doi: [https://doi.org/10.1016/S0022-3115\(02\)01174-1](https://doi.org/10.1016/S0022-3115(02)01174-1). URL <http://www.sciencedirect.com/science/article/pii/S0022311502011741>.
- [59] H. Weisen, Hyun-Tae Kim, J. Strachan, S. Scott, Y. Baranov, J. Buchanan, M. Fitzgerald, D. Keeling, D.B. King, L. Giacomelli, T. Koskela, M.J. Weisen, C. Giroud, M. Maslov, W.G. Core, K.-D. Zastrow, D.B. Syme, S. Popovichev, S. Conroy, I. Lengar, L. Snoj, P. Batistoni, and M. Santala and. The ‘neutron deficit’ in the JET tokamak. *Nuclear Fusion*, 57(7):076029, jun

2017. doi: 10.1088/1741-4326/aa6dcc. URL <https://doi.org/10.1088/1741-4326/aa6dcc>.
- [60] V Riccardo and JET EFDA contributors. Disruptions and disruption mitigation. *Plasma Physics and Controlled Fusion*, 45(12A):A269–A284, nov 2003. doi: 10.1088/0741-3335/45/12a/018. URL <https://doi.org/10.1088/0741-3335/45/12a/018>.
- [61] R.D Gill, B Alper, A.W Edwards, L.C Ingesson, M.F Johnson, and D.J Ward. Direct observations of runaway electrons during disruptions in the JET tokamak. *Nuclear Fusion*, 40(2):163–174, feb 2000. doi: 10.1088/0029-5515/40/2/302. URL <https://doi.org/10.1088/0029-5515/40/2/302>.
- [62] M. Kiritani. Defect structure evolution from radiation damage with d-t fusion neutrons. *Journal of Nuclear Materials*, 133-134:85 – 91, 1985. ISSN 0022-3115. doi: [https://doi.org/10.1016/0022-3115\(85\)90116-3](https://doi.org/10.1016/0022-3115(85)90116-3). URL <http://www.sciencedirect.com/science/article/pii/0022311585901163>.
- [63] IARC Working Group on the Evaluation of Carcinogenic Risks to Humans. *Radiation*. International Agency for Research on Cancer, 2012. ISBN 978-92-832-0136-6.
- [64] Centronic LTD. Fission chambers - type fc765. Technical report, Centronic LTD, King Henry’s Drive, Croydon, Surrey, CR9 0BG, United Kingdom, April 2005. URL <http://www.centronic.co.uk/products/10/fission-chambers>. Ref 313/01.
- [65] Norman Robert Campbell and V. J. Francis. A theory of valve and circuit noise. *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, 93:45–52, 1946.
- [66] A. Sari, C. Thiam, K. Boudergui, F. Carrel, R. Coulon, J. Dumazert, C. Frangville, H. Hamrita, F. Lainé, M. Trocmé, B. Krausz, R. Pissarello, and R. Delalez. Investigation of fission chamber response in the frame of fuel debris localization measurements at fukushima daiichi. *Radiation Measurements*, 130:106223, 2020. ISSN 1350-4487. doi: <https://doi.org/10.1016/j.radmeas.2019.106223>. URL <https://www.sciencedirect.com/science/article/pii/S1350448719305098>.
- [67] Zs Elter, Christian Jammes, Imre Pázsit, L. Pál, and Philippe Filliatre. Performance investigation of the pulse and campbelling modes of a fission chamber using a poisson pulse train simulation code. *Nuclear Instruments and*

- Methods in Physics Research Section A Accelerators Spectrometers Detectors and Associated Equipment*, 774, 02 2015. doi: 10.1016/j.nima.2014.11.065.
- [68] Christian Jammes, Philippe Filliatre, Benoit Geslot, L. Oriol, F. Berhouet, J-F Villard, and Ludo Vermeeren. Research activities in fission chamber modeling in support of the nuclear energy industry. *Nuclear Science, IEEE Transactions on*, 57:3678 – 3682, 01 2011. doi: 10.1109/TNS.2010.2075939.
- [69] Zeinalov, Shakir, Sidorova, Olga, Sedyshev, Pavel, Shvetsov, Valery, Lee, Youngseok, and Nam, Uk-Won. Thermal neutron intensity measurement with fission chamber in current, pulse and campbell modes. *EPJ Web Conf.*, 231:05009, 2020. doi: 10.1051/epjconf/202023105009. URL <https://doi.org/10.1051/epjconf/202023105009>.
- [70] J. Alarcón, L. Marzano, S. Thorp, and C. Verrastro. Fpga based wide range neutron flux monitoring system using campbell mode. In *2019 X Southern Conference on Programmable Logic (SPL)*, pages 77–81, 2019.
- [71] Flask microframework. URL <https://flask.palletsprojects.com/en/1.1.x/>. Accessed Sept. 7, 2020.
- [72] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009. ISBN 1441412697.
- [73] Sang In Kim, Bong Hwan Kim, Jang Lyul Kim, and Jung Il Lee. A review of neutron scattering correction for the calibration of neutron survey meters using the shadow cone method. *Nuclear Engineering and Technology*, 47(7):939 – 944, 2015. ISSN 1738-5733. doi: <https://doi.org/10.1016/j.net.2015.07.005>. URL <http://www.sciencedirect.com/science/article/pii/S173857331500176X>.
- [74] C. Vincent, W. McCarthy, T. Golfinopoulos, B. LaBombard, R. Sharples, J. Lovell, G. Naylor, S. Hall, J. Harrison, and A. Q. Kuang. The digital mirror langmuir probe: Field programmable gate array implementation of real-time langmuir probe biasing. *Review of Scientific Instruments*, 90(8):083504, 2019. doi: 10.1063/1.5109834. URL <https://doi.org/10.1063/1.5109834>.
- [75] N. Smick and B. LaBombard. Wall scanning probe for high-field plasma measurements on Alcator C-Mod. *Review of Scientific Instruments*, 80(2): 023502, 2009. doi: 10.1063/1.3069290. URL <https://doi.org/10.1063/1.3069290>.
- [76] A. Kirk, A. J. Thornton, J. R. Harrison, F. Militello, N. R. Walkden, the MAST Team, and the EUROfusion MST1 Team. L-mode filament

- characteristics on MAST as a function of plasma current measured using visible imaging. *Plasma Physics and Controlled Fusion*, 58(8):085008, 2016. doi: 10.1088/0741-3335/58/8/085008. URL <http://stacks.iop.org/0741-3335/58/i=8/a=085008>.
- [77] P. W. Terry. Suppression of turbulence and transport by sheared flow. *Rev. Mod. Phys.*, 72:109–165, Jan 2000. doi: 10.1103/RevModPhys.72.109. URL <https://link.aps.org/doi/10.1103/RevModPhys.72.109>.
- [78] B. LaBombard, T. Golfopoulos, J. L. Terry, D. Brunner, E. Davis, M. Greenwald, and J. W. Hughes. New insights on boundary plasma turbulence and the quasi-coherent mode in Alcator C-Mod using a Mirror Langmuir Probe. *Physics of Plasmas*, 21(5):056108, 2014. doi: 10.1063/1.4873918. URL <https://doi.org/10.1063/1.4873918>.
- [79] J. Lovell, R. Stephen, S. Bray, G. Naylor, S. Elmore, H. Willett, M. Peterka, M. Dimitrova, A. Havranek, M. Hron, and R. Sharples. A compact, smart Langmuir Probe control module for MAST-Upgrade. *Journal of Instrumentation*, 12(11):C11008, 2017. URL <http://stacks.iop.org/1748-0221/12/i=11/a=C11008>.
- [80] Sin-Li Chen and T. Sekiguchi. Instantaneous Direct-Display System of Plasma Parameters by Means of Triple Probe. *Journal of Applied Physics*, 36(8):2363–2375, 1965. doi: 10.1063/1.1714492. URL <https://doi.org/10.1063/1.1714492>.
- [81] P. Stangeby. *The Plasma Boundary of Magnetic Fusion Devices*. CRC Press, 2000.
- [82] B. Lipschultz, X. Bonnin, G. Counsell, A. Kallenbach, A. Kukushkin, K. Krieger, A. Leonard, A. Loarte, R. Neu, R.A. Pitts, T. Rognien, J. Roth, C. Skinner, J.L. Terry, E. Tsitrone, D. Whyte, S. Zweben, N. Asakura, D. Coster, R. Doerner, R. Dux, G. Federici, M. Fenstermacher, W. Fundamenski, P. Ghendrih, A. Herrmann, J. Hu, S. Krasheninnikov, G. Kirnev, A. Kreter, V. Kurnaev, B. LaBombard, S. Lisgo, T. Nakano, N. Ohno, H.D. Pacher, J. Paley, Y. Pan, G. Pautasso, V. Philipps, V. Rohde, D. Rudakov, P. Stangeby, S. Takamura, T. Tanabe, Y. Yang, and S. Zhu. Plasma–surface interaction, scrape-off layer and divertor physics: implications for ITER. *Nuclear Fusion*, 47(9):1189–1205, aug 2007. doi: 10.1088/0029-5515/47/9/016. URL <https://doi.org/10.1088/0029-5515/47/9/016>.
- [83] N.R. Walkden, F. Militello, J. Harrison, T. Farley, S. Silburn, and J. Young. Identification of intermittent transport in the scrape-off layer

- of mast through high speed imaging. *Nuclear Materials and Energy*, 12: 175 – 180, 2017. ISSN 2352-1791. doi: <https://doi.org/10.1016/j.nme.2016.10.024>. URL <http://www.sciencedirect.com/science/article/pii/S2352179116300667>. Proceedings of the 22nd International Conference on Plasma Surface Interactions 2016, 22nd PSI.
- [84] J. A. Boedo, D. L. Rudakov, R. A. Moyer, G. R. McKee, R. J. Colchin, M. J. Schaffer, P. G. Stangeby, W. P. West, S. L. Allen, T. E. Evans, R. J. Fonck, E. M. Hollmann, S. Krasheninnikov, A. W. Leonard, W. Nevins, M. A. Mahdavi, G. D. Porter, G. R. Tynan, D. G. Whyte, and X. Xu. Transport by intermittency in the boundary of the diii-d tokamak. *Physics of Plasmas*, 10(5):1670–1677, 2003. doi: 10.1063/1.1563259. URL <https://doi.org/10.1063/1.1563259>.
- [85] B D Dudson, N Ben Ayed, A Kirk, H R Wilson, G Counsell, X Xu, M Umansky, P B Snyder, and B LLoyd and. Experiments and simulation of edge turbulence and filaments in MAST. *Plasma Physics and Controlled Fusion*, 50(12):124012, nov 2008. doi: 10.1088/0741-3335/50/12/124012. URL <https://doi.org/10.1088/0741-3335/50/12/124012>.
- [86] S.B. Ballinger, J.L. Terry, S.G. Baek, K. Tang, A. von Stechow, C. Killer, D. Nicolai, G. Satheeswaran, P. Drews, and O. Grulke. Fast camera imaging of plasmas in alcator c-mod and w7-x. *Nuclear Materials and Energy*, 17:269 – 273, 2018. ISSN 2352-1791. doi: <https://doi.org/10.1016/j.nme.2018.11.015>. URL <http://www.sciencedirect.com/science/article/pii/S2352179118301121>.
- [87] Jordan Cavalier, Nicolas Lemoine, Frederic Brochard, Vladimir Weinzettl, Jakub Seidl, Scott Silburn, Patrick Tamain, Renaud Dejarnac, Jiri Adamek, and Radomir Panek. Tomographic reconstruction of tokamak edge turbulence from single visible camera data and automatic turbulence structure tracking. *Nuclear Fusion*, 59(5):056025, apr 2019. doi: 10.1088/1741-4326/ab0d4c. URL <https://doi.org/10.1088/1741-4326/ab0d4c>.
- [88] P. Horowitz and W. Hill. *The Art of Electronics, Third Edition*. Cambridge university Press, 2015.
- [89] Xilinx divider ip. URL <https://www.xilinx.com/products/intellectual-property/divider.html>. Accessed: 23rd March 2021.
- [90] Salvador Tropea. Fpga implementation of base-n logarithm. pages 27 – 32, 03 2007. ISBN 1-4244-0606-4. doi: 10.1109/SPL.2007.371719.

- [91] Nikolaos Alachiotis and Alexandros Stamatakis. Fpga optimizations for a pipelined floating-point exponential unit. In Andreas Koch, Ram Krishnamurthy, John McAllister, Roger Woods, and Tarek El-Ghazawi, editors, *Reconfigurable Computing: Architectures, Tools and Applications*, pages 316–327, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg. ISBN 978-3-642-19475-7.
- [92] W. McCarthy, T. Golfinopoulos, K.B Woller, C. Vincent, A.Q. Kuang, and B. LaBombard. Real plasma tests of the digital mirror langmuir probe. 2021. Submitted for publication in Review of Scientific Instruments.
- [93] F. Bedoya, K. B. Woller, and D. G. Whyte. Study of the properties of thin Li films and their relationship with He plasmas using ion beam analysis in the DIONISOS experiment. *Review of Scientific Instruments*, 89(10):10J106, 2018. doi: 10.1063/1.5034240. URL <https://doi.org/10.1063/1.5034240>.
- [94] F Militello, P Tamain, W Fundamenski, A Kirk, V Naulin, and A H Nielsen and. Experimental and numerical characterization of the turbulence in the scrape-off layer of MAST. *Plasma Physics and Controlled Fusion*, 55(2):025005, jan 2013. doi: 10.1088/0741-3335/55/2/025005. URL <https://doi.org/10.1088/0741-3335/55/2/025005>.
- [95] S Elmore, S Y Allan, A Kirk, G Fishpool, J Harrison, P Tamain, M Kočan, R Gaffka, R Stephen, and J W Bradley and. Upstream and divertor ion temperature measurements on MAST by retarding field energy analyser. *Plasma Physics and Controlled Fusion*, 54(6):065001, may 2012. doi: 10.1088/0741-3335/54/6/065001. URL <https://doi.org/10.1088/0741-3335/54/6/065001>.
- [96] D J Campbell, S Ali Arshad, A Gondhalekar, and P R Thomas. H-modes under steady-state conditions in JET. *Plasma Physics and Controlled Fusion*, 36(7A):A255–A260, jul 1994. doi: 10.1088/0741-3335/36/7a/037. URL <https://doi.org/10.1088/0741-3335/36/7a/037>.
- [97] H R Wilson, S C Cowley, A Kirk, and P B Snyder. Magneto-hydrodynamic stability of the h-mode transport barrier as a model for edge localized modes: an overview. 48(5A):A71–A84, apr 2006. doi: 10.1088/0741-3335/48/5a/s06. URL <https://doi.org/10.1088/0741-3335/48/5a/s06>.
- [98] F. Wagner, G. Fussmann, T. Grave, M. Keilhacker, M. Kornherr, K. Lackner, K. McCormick, E. R. Müller, A. Stäbler, G. Becker, K. Bernhardt, U. Ditte, A. Eberhagen, O. Gehre, J. Gernhardt, G. v. Gierke, E. Glock, O. Gruber,

- G. Haas, M. Hesse, G. Janeschitz, F. Karger, S. Kissel, O. Klüber, G. Lisitano, H. M. Mayer, D. Meisel, V. Mertens, H. Murmann, W. Poschenrieder, H. Rapp, H. Röhr, F. Ryter, F. Schneider, G. Siller, P. Smeulders, F. Söldner, E. Speth, K. H. Steuer, Z. Szymanski, and O. Vollmer. Development of an edge transport barrier at the h-mode transition of asdex. *Phys. Rev. Lett.*, 53:1453–1456, Oct 1984. doi: 10.1103/PhysRevLett.53.1453. URL <https://link.aps.org/doi/10.1103/PhysRevLett.53.1453>.
- [99] F. Wagner, G. Fussmann, T. Grave, M. Keilhacker, M. Kornherr, K. Lackner, K. McCormick, E. R. Müller, A. Stäbler, G. Becker, K. Bernhardt, U. Ditte, A. Eberhagen, O. Gehre, J. Gernhardt, G. v. Gierke, E. Glock, O. Gruber, G. Haas, M. Hesse, G. Janeschitz, F. Karger, S. Kissel, O. Klüber, G. Lisitano, H. M. Mayer, D. Meisel, V. Mertens, H. Murmann, W. Poschenrieder, H. Rapp, H. Röhr, F. Ryter, F. Schneider, G. Siller, P. Smeulders, F. Söldner, E. Speth, K. H. Steuer, Z. Szymanski, and O. Vollmer. Development of an edge transport barrier at the h-mode transition of asdex. *Phys. Rev. Lett.*, 53:1453–1456, Oct 1984. doi: 10.1103/PhysRevLett.53.1453. URL <https://link.aps.org/doi/10.1103/PhysRevLett.53.1453>.
- [100] A.C.C. Sips, J. Schweinzer, T.C. Luce, S. Wolfe, H. Urano, J. Hobirk, S. Ide, E. Joffrin, C. Kessel, S.H. Kim, P. Lomas, I. Nunes, T. Pütterich, F. Rimini, W.M. Solomon, J. Stober, F. Turco, P.C. de Vries, , , , and and. Assessment of the baseline scenario at  $q \sim 95 \sim 3$  for ITER. *Nuclear Fusion*, 58(12):126010, sep 2018. doi: 10.1088/1741-4326/aade57. URL <https://doi.org/10.1088/1741-4326/aade57>.
- [101] A. Loarte, B. Bazylev, D. J. Campbell, S. Pestchanyi, T. Eich, T. E. Evans, O. Schmitz, W. Suttrop, M. Becoulet, R. Sartori, M. Fenstermacher, E. Luna, Y. Gribov, M. J. Schaffer, R. Pitts, H. R. Wilson, J. Canik, I. Landman, Jochen Linke, E. Nardon, V. Podkovyrov, P. Lang, G. Saibene, N. Klimov, A. Zhitlukhin, B. Riccardi, H. Thomsen, Y. Liang, A. Kirk, and R. Maingi. Iter-elm control requirements, elm control schemes and required rd. 2010.
- [102] Michiya Shimada. The present status of impurity transport investigations. *Fusion Engineering and Design*, 15(4):325 – 341, 1991. ISSN 0920-3796. doi: [https://doi.org/10.1016/0920-3796\(92\)90018-Y](https://doi.org/10.1016/0920-3796(92)90018-Y). URL <http://www.sciencedirect.com/science/article/pii/092037969290018Y>.
- [103] R.P. Wenninger, M. Bernert, T. Eich, E. Fable, G. Federici, A. Kallenbach, A. Loarte, C. Lowry, D. McDonald, R. Neu, T. Pütterich, P. Schneider, B. Sieglin, G. Strohmayer, F. Reimold, and M. Wischmeier. DEMO divertor limitations during and in between ELMs. *Nuclear Fusion*, 54(11):114003, nov

2014. doi: 10.1088/0029-5515/54/11/114003. URL <https://doi.org/10.1088/0029-5515/54/11/114003>.
- [104] Yu Igitkhanov and Boris Bazylev. Erosion of pfc in demo due to elm impact. pages 1–6, 06 2013. ISBN 978-1-4799-0169-2. doi: 10.1109/SOFE.2013.6635280.
- [105] Alex Thorman, Clive Michael, John Howard, Brian Victor, Chris Holcomb, and Steve Allen. Motional stark effect imaging first results on the diii-d tokamak. *Review of Scientific Instruments*, 89(10):10D124, 2018. doi: 10.1063/1.5038969. URL <https://doi.org/10.1063/1.5038969>.
- [106] K. Kamiya, T. Fujita, A. Kojima, and H. Kubo. Zeeman polarimetry measurement for edge current density determination using li-beam probe on jt-60u. *Review of Scientific Instruments*, 81(3):033502, 2010. doi: 10.1063/1.3309793. URL <https://doi.org/10.1063/1.3309793>.
- [107] R. G. L. Vann, K. J. Brunner, R. Ellis, G. Taylor, and D. A. Thomas. Preliminary measurements of the edge magnetic field pitch from 2-d doppler backscattering in mast and nstx-u (invited). *Review of Scientific Instruments*, 87(11):11D902, 2016. doi: 10.1063/1.4962253. URL <https://aip.scitation.org/doi/abs/10.1063/1.4962253>.
- [108] Jon Hillesheim, W.A. Peebles, Hendrik Meyer, and Neal Crocker. Doppler backscattering measurements for mast. *40th EPS Conference on Plasma Physics, EPS 2013*, 1:201–204, 01 2013.
- [109] F.F Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer, Cham, 1974. ISBN 978-3-319-22309-4.
- [110] Matthew Bryn Thomas. *3D full-wave modelling of microwave interactions with plasma density fluctuations*. PhD thesis, York University, November 2017.
- [111] Laurent Ferro-Famil and Eric Pottier. 1 - synthetic aperture radar imaging. In Nicolas Baghdadi and Mehrez Zribi, editors, *Microwave Remote Sensing of Land Surface*, pages 1 – 65. Elsevier, 2016. ISBN 978-1-78548-159-8. doi: <https://doi.org/10.1016/B978-1-78548-159-8.50001-3>. URL <http://www.sciencedirect.com/science/article/pii/B9781785481598500013>.
- [112] C. E. Burckhardt, P. . Grandchamp, and H. Hoffmann. An experimental 2 mhz synthetic aperture sonar system intended for medical use. *IEEE Transactions on Sonics and Ultrasonics*, 21(1):1–6, 1974. doi: 10.1109/T-SU.1974.29783.

- [113] M. Younis, C. Fischer, and W. Wiesbeck. Digital beamforming in sar systems. *IEEE Transactions on Geoscience and Remote Sensing*, 41(7):1735–1739, 2003. doi: 10.1109/TGRS.2003.815662.
- [114] A. Farina. Fundamentals of multisite radar systems: multistatic radars and multiradar systems [book review]. *IEEE Aerospace and Electronic Systems Magazine*, 16(4):44–44, 2001. doi: 10.1109/MAES.2001.918029.
- [115] C Laviron, A J H Donné, M E Manso, and J Sanchez. Reflectometry techniques for density profile measurements on fusion plasmas. *Plasma Physics and Controlled Fusion*, 38(7):905–936, jul 1996. doi: 10.1088/0741-3335/38/7/002. URL <https://doi.org/10.1088/0741-3335/38/7/002>.
- [116] M Hirsch, E Holzhauser, J Baldzuhn, B Kurzan, and B Scott. Doppler reflectometry for the investigation of propagating density perturbations. *Plasma Physics and Controlled Fusion*, 43(12):1641–1660, oct 2001. doi: 10.1088/0741-3335/43/12/302. URL <https://doi.org/10.1088/0741-3335/43/12/302>.
- [117] J.C. Hillesheim, N.A. Crocker, W.A. Peebles, H. Meyer, A. Meakins, A.R. Field, D. Dunai, M. Carr, and N. Hawkes and. Doppler backscattering for spherical tokamaks and measurement of high-k-density fluctuation wavenumber spectrum in MAST. *Nuclear Fusion*, 55(7):073024, jun 2015. doi: 10.1088/0029-5515/55/7/073024. URL <https://doi.org/10.1088/0029-5515/55/7/073024>.
- [118] W Mandl, R C Wolf, M G von Hellermann, and H P Summers. Beam emission spectroscopy as a comprehensive plasma diagnostic tool. *Plasma Physics and Controlled Fusion*, 35(10):1373–1394, oct 1993. doi: 10.1088/0741-3335/35/10/003. URL <https://doi.org/10.1088/0741-3335/35/10/003>.
- [119] S. J. Zweben, J. L. Terry, D. P. Stotler, and R. J. Maqueda. Invited review article: Gas puff imaging diagnostics of edge plasma turbulence in magnetic fusion devices. *Review of Scientific Instruments*, 88(4):041101, 2017. doi: 10.1063/1.4981873. URL <https://doi.org/10.1063/1.4981873>.
- [120] G. McKee, R. Ashley, R. Durst, R. Fonck, M. Jakubowski, K. Tritz, K. Burrell, C. Greenfield, and J. Robinson. The beam emission spectroscopy diagnostic on the diii-d tokamak. *Review of Scientific Instruments*, 70(1):913–916, 1999. doi: 10.1063/1.1149416. URL <https://doi.org/10.1063/1.1149416>.

- [121] I. Shesterikov, Y. Xu, M. Berte, P. Dumortier, M. Van Schoor, M. Vergote, B. Schweer, and G. Van Oost. Development of the gas-puff imaging diagnostic in the textor tokamak. *Review of Scientific Instruments*, 84(5):053501, 2013. doi: 10.1063/1.4803934. URL <https://doi.org/10.1063/1.4803934>.
- [122] G R Tynan, A Fujisawa, and G McKee. A review of experimental drift turbulence studies. *Plasma Physics and Controlled Fusion*, 51(11):113001, oct 2009. doi: 10.1088/0741-3335/51/11/113001. URL <https://doi.org/10.1088/0741-3335/51/11/113001>.
- [123] E Z Gusakov and A V Surkov. Spatial and wavenumber resolution of doppler reflectometry. *Plasma Physics and Controlled Fusion*, 46(7):1143–1162, jun 2004. doi: 10.1088/0741-3335/46/7/012. URL <https://doi.org/10.1088/0741-3335/46/7/012>.
- [124] C.C. Petty, W.R. Fox, T.C. Luce, M.A. Makowski, and T. Suzuki. Analysis of current drive using MSE polarimetry without equilibrium reconstruction. *Nuclear Fusion*, 42(9):1124–1133, aug 2002. doi: 10.1088/0029-5515/42/9/311. URL <https://doi.org/10.1088/0029-5515/42/9/311>.
- [125] Karim Mohammad Pour Aghdam, Reza Faraji-Dana, and Jalil Rashed-Mohassel. The sinuous antenna-a dual polarized feed for reflector-based searching systems. *AEU - International Journal of Electronics and Communications*, 59(7):392 – 400, 2005. ISSN 1434-8411. doi: <https://doi.org/10.1016/j.aeue.2004.11.046>. URL <http://www.sciencedirect.com/science/article/pii/S1434841104000639>.
- [126] Institute of Electrical and Electronics Engineers. IEEE 802.3-2018 - IEEE Standard for Ethernet. Technical report, IEEE, August 2018.
- [127] Institute of Electrical and Electronics Engineers. IEEE standards interpretations for IEEE Std 1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems. Technical report, IEEE, December 2011.
- [128] Inc Xilinx. 10g/25g high speed ethernet subsystem v3.2. Technical report, September 2020. URL [https://www.xilinx.com/support/documentation/ip\\_documentation/xxv\\_ethernet/v3\\_2/pg210-25g-ethernet.pdf](https://www.xilinx.com/support/documentation/ip_documentation/xxv_ethernet/v3_2/pg210-25g-ethernet.pdf). PG210.
- [129] Inc Xilinx. 10g ethernet subsystem v3.2. Technical report, July 2020. URL [https://www.xilinx.com/support/documentation/ip\\_documentation/axi\\_10g\\_ethernet/v3\\_1/pg157-axi-10g-ethernet.pdf](https://www.xilinx.com/support/documentation/ip_documentation/axi_10g_ethernet/v3_1/pg157-axi-10g-ethernet.pdf). PG157.

- [130] X. Dong and Y. Zhang. A map approach for 1-bit compressive sensing in synthetic aperture radar imaging. *IEEE Geoscience and Remote Sensing Letters*, 12(6):1237–1241, 2015. doi: 10.1109/LGRS.2015.2390623.
- [131] Intel Corporation. Intel xeon scalable platform. Technical report, Intel Corporation, 2017. URL <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/xeon-scalable-platform-brief.pdf>. Ref 0617/SC/HBD.
- [132] Inc Corsair Components. Rmx series. Technical report, 2017. URL [https://www.corsair.com/corsairmedia/sys\\_master/productcontent/RMx\\_2018\\_Manual.pdf](https://www.corsair.com/corsairmedia/sys_master/productcontent/RMx_2018_Manual.pdf).
- [133] Inwin. Iw-r400-01nopen-bay 4u server chassis. Technical report. URL <https://www.ipc.in-win.com/rackmount-chassis-iw-r400-01n>.
- [134] Gnu make. <https://www.gnu.org/software/make/manual/make.html>. Accessed: 23rd March 2021.
- [135] socket(2) — linux manual page. <https://man7.org/linux/man-pages/man2/socket.2.html>. Accessed: 23rd March 2021.
- [136] recvfrom(2) — linux manual page. <https://linux.die.net/man/2/recvfrom>. Accessed: 23rd March 2021.
- [137] pthreads(7) — linux manual page. <https://man7.org/linux/man-pages/man7/pthreads.7.html>. Accessed: 23rd March 2021.
- [138] Intel fpga products. URL <https://www.intel.co.uk/content/www/uk/en/products/programmable/fpga.html>. Accessed: 23rd March 2021.
- [139] Thevenin, Mathieu, Barbot, Loïc, Corre, Gwénolé, Woo, Romuald, Destouches, Christophe, and Normand, Stéphane. Digital real-time multiple channel multiple mode neutron flux estimation on fpga-based device. *EPJ Web of Conferences*, 106:05009, 2016. doi: 10.1051/epjconf/201610605009. URL <https://doi.org/10.1051/epjconf/201610605009>.
- [140] B. Lescop, P. Padeloup, J. Falaise, J. Lefevre, T. Nguyen, C. Sudres, and J. . Trama. Marine: a fast fully digitalized wide range neutron monitor system. In *2000 IEEE Nuclear Science Symposium. Conference Record (Cat. No.00CH37149)*, volume 1, pages 5/86–5/90 vol.1, 2000. doi: 10.1109/NSSMIC.2000.949089.

- [141] B. Bigot. Progress toward ITER's first plasma. *Nuclear Fusion*, 59(11): 112001, jun 2019. doi: 10.1088/1741-4326/ab0f84. URL <https://doi.org/10.1088/1741-4326/ab0f84>.
- [142] D. Sanz, M. Ruiz, R. Castro, J. Vega, M. Afif, M. Monroe, S. Simrock, T. Debelle, R. Marawar, and B. Glass. Advanced data acquisition system implementation for the iter neutron diagnostic use case using epics and flexrio technology on a pxie platform. *IEEE Transactions on Nuclear Science*, 63(2):1063–1069, 2016. doi: 10.1109/TNS.2016.2517402.
- [143] E. Barrera, M. Ruiz, A. Bustos, M. Afif, B. Radle, J. L. Fernández-Hernando, I. Prieto, R. Pedica, J. M. Barcala, J. C. Oller, and R. Castro. Implementation of iter fast plant interlock system using fpgas with compactrio. *IEEE Transactions on Nuclear Science*, 65(2):796–804, 2018. doi: 10.1109/TNS.2017.2783243.
- [144] A. Wojenski, K.T. Pozniak, D. Mazon, and M. Chernyshova. FPGA-based novel real-time evaluation and data quality monitoring system for tokamak high-performance GEM soft X-ray diagnostic. *JINST*, 13(12):P12024, 2018. doi: 10.1088/1748-0221/13/12/P12024.
- [145] Antonio J.N. Batista, Carlos Leong, Bruno Santos, Ana Fernandes, Ana Rita Ramos, Joana P. Santos, José G. Marques, Isabel C. Teixeira, João P. Teixeira, Jorge Sousa, and Bruno Gonçalves. Seu mitigation exploratory tests in a iter related fpga. *Fusion Engineering and Design*, 118:111 – 116, 2017. ISSN 0920-3796. doi: <https://doi.org/10.1016/j.fusengdes.2017.03.106>. URL <http://www.sciencedirect.com/science/article/pii/S0920379617303290>.
- [146] R.A. Charpie. The geneva conference. *Scientific American*, 193(4):27–33, 1955. URL <https://www.jstor.org/stable/10.2307/24943753>.
- [147] ITER Org. How the "zeta fiasco" pulled fusion out of secrecy. URL <https://www.iter.org/newsline/-/2905>. Accessed Nov. 17, 2020.
- [148] T. Donné and W. Morris. European research roadmap to the realisation of fusion energy. Technical report, 2018.
- [149] Nefeli Chrysanthopoulou, Panayiota Savva, M. Varvayanni, and N. Catsaros. Compilation of existing neutron screen technology. *Science and Technology of Nuclear Installations*, 2014:23 pages, 08 2014. doi: 10.1155/2014/395795.
- [150] R. A. Forster and T. N. K. Godfrey. Mcnp - a general monte carlo code for neutron and photon transport. In Raymond Alcouffe, Robert Dautray, Arthur

- Forster, Guy Ledanois, and B. Mercier, editors, *Monte-Carlo Methods and Applications in Neutronics, Photonics and Statistical Physics*, pages 33–55, Berlin, Heidelberg, 1985. Springer Berlin Heidelberg. ISBN 978-3-540-39750-2.
- [151] Ira B. Bernstein. Waves in a plasma in a magnetic field. *Phys. Rev.*, 109: 10–21, Jan 1958. doi: 10.1103/PhysRev.109.10. URL <https://link.aps.org/doi/10.1103/PhysRev.109.10>.
- [152] V F Shevchenko, R G L Vann, S J Freethy, and B K Huang. Synthetic aperture microwave imaging with active probing for fusion plasma diagnostics. *Journal of Instrumentation*, 7(10):P10016–P10016, oct 2012. doi: 10.1088/1748-0221/7/10/p10016. URL <https://doi.org/10.1088/1748-0221/7/10/p10016>.
- [153] Simon Freethy. *Synthetic aperture imaging of B-X-O mode conversion*. PhD thesis, York University, March 2012.
- [154] David A. Thomas. *Phased array imaging of two dimensional Doppler microwave backscattering from spherical tokamak edge plasmas*. PhD thesis, York University, July 2016.

## Colophon

This thesis is based on a template developed by Matthew Townson and Andrew Reeves. It was typeset with L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub>. It was created using the *memoir* package, maintained by Lars Madsen, with the *madsen* chapter style. The font used is Latin Modern, derived from fonts designed by Donald E. Kuniath.