

Durham E-Theses

FPGA-based High Performance Diagnostics For Fusion

KAI JAKOB BRUNNER

How to cite:

BRUNNER, KAI JAKOB (2017) FPGA-based High Performance Diagnostics For Fusion. Doctoral thesis, Durham University.

Use policy



This work is licensed under a [Creative Commons Attribution Non-commercial 3.0 \(CC BY-NC\)](https://creativecommons.org/licenses/by-nc/3.0/)

FPGA-based High Performance Diagnostics For Fusion

Kai Jakob Brunner

Thesis presented for the degree of
Doctor of Philosophy



Centre for Advanced Instrumentation

Department of Physics

University of Durham

England

2017-03-03

Dedicated to
our children

Abstract

High performance diagnostics are an important aspect of fusion research. Increasing shot-lengths paired with the requirement for higher accuracy and speed make it mandatory to employ new technology to cope with the increasing demands on digitization and data handling. Field programmable gate arrays (FPGAs) are well known in high performance applications. Their ability to handle multiple fast data streams whilst remaining programmable make them an ideal tool for diagnostic development. Both the improvement of old and the design of new diagnostics can benefit from FPGA-technology and increase the amount of accessible physics significantly. In this work the developments on two FPGA-based diagnostics are presented.

In the first part a new open-hardware low-cost FPGA-based digitizer is presented for the MAST-Upgrade (MAST-U) integral electron density interferometer. The system is shown to have an optically limited phase accuracy and a detection bandwidth of over 3.5 MHz. Data is acquired continuously at 20 MS/s and streamed to an acquisition PC via optical fiber. By employing a dual-FPGA approach real-time processing of the density signal can be achieved despite severely limited resources, thus providing a control signal for the MAST-U plasma control system with less than 8 μ s latency. Due to MAST-U being still inoperable, in-situ testing has been conducted on the ASDEX Upgrade, where fast wave physics up to 3.5 MHz could first be observed.

The second part presents developments to the Synthetic Aperture Microwave Imaging (SAMI) diagnostic. In addition to improving the utilization of long shot-lengths and enabling dual-polarized acquisition the system has been enhanced to continuously acquire active probing profiles for 2D Doppler back-scattering (DBS), a technique recently developed using SAMI. The aim is to measure pitch angle profiles to derive the edge current density. SAMI has been transferred to the NSTX-Upgrade and integrated into the experiment's infrastructure, where it has been acquiring data since May 2016. As part of this move an investigation into near-field effects on SAMI's image reconstruction algorithms was conducted.

Declaration

The work in this thesis is based on research carried out at the Culham Centre for Fusion Energy on behalf of the Centre for Advanced Instrumentation, Durham University, England. No part of this thesis has been submitted elsewhere for any other degree or qualification.

The work and results described in this thesis represent the combined efforts of the respective diagnostic teams and credit is given to the appropriate authors in the respective chapters. In all instances the main focus of my work was the development of the diagnostic hardware and FPGA firmware. Where not indicated to the contrary the presented work is entirely my own.

Copyright © 2016 by KAI JAKOB BRUNNER.

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

Acknowledgments

It is of course difficult to thank everyone responsible for making this thesis a success. Everyone who has contributed in even the slightest way, deserves an acknowledgement. However, there are a few people I wish to explicitly thank:

First and foremost my wife and children for enduring 4 years of video-conferences and 1000 km separation without a single complaint.

Next my supervisors Ray Sharples and Nigel Dipper for giving me this opportunity in the first place and trying their best to guide me over the past years.

I wish to thank Graham Naylor, who has given me a great deal of support and helped me immensely with my FPGA work. From the SAMI team I have to thank Roddy Vann, Simon Freethy and Vladimir Shevchenko, who have guided me with the microwave physics and technology of SAMI. Also Thomas O’Gorman and Rory Scannell, who have supported my every whim on the interferometer project.

I also want to explicitly thank my collaborators at AUG and NSTX-U. Most notably Gary Taylor from PPPL and Alexander Mlynek from IPP, Garching. Their support enabled two great collaborations as part of this work.

Last but not least I want to express my gratitude to electronics staff at CCFE and PPPL for their humongous help. Some of my circuit work would have certainly found a smoking grave without them.

Unless indicated to the contrary the analysis in this thesis was wholly conducted using Python 3 in combination with the Numpy and Scipy libraries[1]–[3]. Most of the plots were created using either the Matplotlib library or the open-source tools Inkscape or GIMP[4]–[6]. To all those unnamed programmers I am in debt.

The presented research is funded by Durham University and EPSRC grant EP/L01663X/1. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the EURATOM research and training programme 2014-2018 under grant agreement No. 633053 and from the RCUK Energy Programme [grant number EP/I501045]. Additional funding was received though the EUROfusion Enabling Research project grant CfP-WP15-ENR-01/CCFE-03.

Contents

Glossary	xx
1 General Introduction	1
1.1 Nuclear Fusion	1
1.2 Basic Physics	2
1.2.1 The Lawson Criterion	4
1.3 Magnetic Confinement Fusion Devices	5
1.3.1 Tokamaks	6
1.3.2 Spherical Tokamaks	7
1.3.3 Real-Time Fusion Diagnostics	9
1.4 Field Programmable Gate Arrays	10
1.4.1 The FPGA Hardware	10
1.4.2 The FPGA Design Flow	12
1.4.3 FPGA vs. Microprocessor/GPU	13
1.4.4 FPGAs In Nuclear Industry	15
1.5 Scope Of The Thesis	16
2 FPGA Digitization For The Integral Electron Density	18
2.1 Introduction	18
2.2 Basic Waves In Fusion Plasmas	22
2.2.1 The Cold Plasma Dispersion	22
2.2.2 Visible Light Approximation	25
2.3 Plasma Interferometry	26
2.3.1 Two-Color Interferometry	27
2.4 MAST-Upgrade Diagnostic Setup	29
2.4.1 Optical Setup	30
2.4.2 Digitization Hardware	32
2.5 SPEC Firmware	33
2.5.1 Top Level Layout	34

2.5.2	CO ₂ Processing	36
2.5.3	HeNe Processing	38
2.5.4	CORDIC Phase Difference	39
2.5.5	HeNe Hilbert Processing	41
2.5.6	PLL Generated Reference	45
2.5.7	Resource Usage	46
2.6	Aurora-PCIe Bridge	47
2.6.1	Diagnostic Control Software	51
2.7	Diagnostic Performance	51
2.7.1	Phase Noise	52
2.7.2	Phase Wrap Stability	56
2.7.3	Bandwidth Capabilities	59
2.7.4	Latency Results	62
2.8	ASDEX Upgrade In-Situ Testing	63
2.8.1	The AUG Experimental Setup	64
2.8.2	Firmware Modifications For AUG	65
2.8.3	MGI Induced Disruption Trace	67
2.8.4	High NTM Harmonic Detection	69
2.8.5	Indications For Fast Mode Physics On AUG	77
2.8.6	Operational Difficulties	80
2.9	Summary & Outlook	83
3	Continuous Pitch Angle Profile Measurements using SAMI	85
3.1	Introduction	85
3.2	Microwaves In Fusion Plasmas	89
3.2.1	Generalized Cut-Off & Resonance	89
3.2.2	Propagation Along B	91
3.2.3	Propagation Perpendicular To B	91
3.2.4	Backscattering Of Waves On Plasma Surfaces	92
3.3	Synthetic Apertures Image Formation	93
3.3.1	Fourier Inversion	94
3.3.2	Beamforming	96
3.3.3	Matrix Inversion	98
3.4	2D Doppler Back-Scattering	100
3.4.1	Doppler Reflectometry	100
3.4.2	Back-Scattering Efficiency	101
3.4.3	Deriving Magnetic Pitch Angle	103
3.5	The SAMI Diagnostic	105

3.5.1	The Analog Front-End	106
3.5.2	The Digitization Back-End	108
3.5.3	Diagnostic Limitations	109
3.6	Modifications To The SAMI FPGA Firmware	110
3.6.1	2D DBS Fork	112
3.6.2	Revised Clock Infrastructure	114
3.6.3	Aurora-based Data Transfer	115
3.6.4	Resource Usage	116
3.6.5	Aurora-PCIe Bridge	117
3.7	Installation On NSTX-U	119
3.7.1	Hardware Modifications	120
3.8	Far-Field To Near-Field Transition	122
3.8.1	Comparison Of Algorithms In Planar Geometry	123
3.8.2	Comparison Of Algorithms In Spherical Geometry	129
3.9	Initial NSTX-U Results	132
3.9.1	Firmware Benchmarking	136
3.10	Outlook	138
3.11	Summary	140
4	Concluding Remarks	142
4.1	Summary	142
4.2	Conclusions & Outlook	143

List of Figures

1.1	The binding energy as a function of atomic number. The maximum of the curve marks the transition from fusible to fissile atoms. The (currently) most relevant atoms for fusion are marked (adapted from [13]).	3
1.2	The fusion cross section for some of the lowest-Z elements. Deuterium-Tritium fusion has the highest cross section at the lowest energy (taken from [14])	4
1.3	The basic schematic of a tokamak (taken from [21]).	7
1.4	Basic elements contained in FPGA fabric. The logic elements in the top right of (a) can be replaced by LUTs in a CLB (marked in purple). The light purple elements mark configurable interconnects. Routing lanes are marked orange. The MUX element (purple in (b)) controls whether the CLB acts on a clock edge or is transition logic (adapted from [35]).	11
1.5	The basic source for a SET. A charged particle, e.g. generated by a neutron striking an atom, traverses the depletion region creating a cloud of electron-hole pairs in the doped silicon substrate thus potentially leading to an unwanted bit flip or a current surge (taken from [43]).	15
2.1	Measurement of the IED N_e during an ELM cycle. For size comparison : the MAST plasma has a major radius of 90 cm and a minor radius of 60 cm. The circle in the fast camera images on top mark the location of the interferometer measurement. The red line in the middle three camera images is a fitted field line. The density in the bottom plot shows a clear oscillation as the result of a filament moving past the measurement cord. Note that although a dip in density is clearly visible no increase is registered, which was expected at the position of the 'x' (taken from [46]).	20

2.2	Measurement of a 3.8 MHz CAE during MAST shot #18886. The left depicts the spectrogram of an OMAHA coil magnetic measurement. The right shows the major radius against the ion cyclotron frequency for that shot with R_0 denoting the location of the magnetic axis. The red line marks the mode frequency seen on the left (adapted from [55] with kind permission from author and publisher).	21
2.3	Basic schematic of a Michelson-Morley interferometer. The optical path length between the two arms differs according to the difference in propagation speed.	26
2.4	Example of two-color vibration interferometry data at the start of an AUG shot. The initial 30 ms no plasma is present and vibrations are the only source for the phase oscillations. After 40 ms the plasma density introduces a measurable phase offset from the null (black dotted line) as a result of the two individual phase-traces diverging. The residual oscillations in the compensated phase are the result of the non-co-linear beam paths. The compensated phase has been offset from 0 for better visualization (dotted line marks the null level taken as the average of the first 20 ms).	28
2.5	The optical path of the new IED interferometer on MAST-U. The optical bench sits in an isolated room and is attached to a wall. The retro-reflector is a roughly 80 kg granite sphere, which is isolated from the plasma vessel. The total optical beam path is roughly 30 m (adapted from [47] with kind permission from the author and AIP publishing).	30
2.6	A schematic of the diagnostic. The optical system is a superposition of a CO ₂ and a HeNe interferometer in Michelson-Morley setup. The digitization system is based on a tandem of two FPGAs, inter-linked by a fiber-optic cable.	31
2.7	Top level of the diagnostic firmware. The top level design with the communication and the deserialization core is written in VHDL. The processing core is designed using System Generator for MATLAB Simulink. Colors denote different clock regions as indicated in the figure. Orange components transition between clock domains.	34

2.8	A schematic of the data processing core using mixing I-Q generation on the HeNe branch. Signal TDM is required to maximize the use of DSP48 multipliers. The mixing frequency ω_{m1} for the CO ₂ signal is 36.36 MHz, the frequency ω_{m2} for the HeNe branch at 1.56 MHz. The delays between the different processing paths are matched by passing valid signals.	35
2.9	A schematic (not to scale) representation of the signals' spectral composition for different points in the CO ₂ processing pipeline as depicted in fig. 2.8. The underlying mathematical operation is described in equations 2.16 and 2.17.	37
2.10	A schematic (not to scale) representation of the signals' spectral composition for different points in the HeNe processing pipeline as depicted in fig. 2.8. The underlying mathematical operation is described in equations 2.16 and 2.17. Note the reference signal, which has higher harmonics of the fundamental wave due to its square shape. After down-mixing the lowest higher harmonic sits around 3.16 MHz.	39
2.11	The HeNe processing stage using Hilbert transformation filtering. The difference to fig. 2.8 is marked in orange. Because no additional low-pass filtering is available in this scheme, the digitized reference signal must be sinusoidal already. The initial decimation filter is identical to the one in the CO ₂ branch.	41
2.12	Hilbert filter attenuation compensation plotting the I against the Q component after the Hilbert transformation filter. The SPEC is fed with a 2.36 MHz signal from a signal generator. The left plot shows the Hilbert compensation factor measurement using a sinusoidal reference signal. Purple marks the uncompensated signal, orange the compensated one. An ideal I-Q signal should be circular in this plot (as indicate by the black dashed line). The middle plot shows the compensated Hilbert transformed and attenuation compensated reference signal using a square wave. The right plot shows the same signal with an analog filter added before the ADC.	42
2.13	The laser's temperature and resulting stability of the HeNe reference frequency over time. The measurement was conducted over a day in an temperature uncontrolled room, which cooled down over night. A strong temperature dependence of the reference frequency will require temperature control.	43

2.14	The short term stability of the HeNe reference frequency. The in-shot frequency stability is shown on the left and a histogram of the frequency distribution on the right. It can be seen that the frequency is on average stable although it shows large short spikes resulting in a broadening of the distribution.	44
2.15	An example for a frequency spike seen in figure 2.14. It can be seen that the frequency is highly variable on very short timescales.	44
2.16	The lock stability of 2 fast PID based PLLs attempting to lock on to the reference frequency. A perfect lock would manifest in a straight horizontal line indicating constant phase.	46
2.17	The protocol used on the MAST-U “Aurora network”. Data is transmitted in 4 B(yte) little endian packages. Only half of the bandwidth is used for data transmission. The MTS is used for transmission error checks, the META tag identifies the nature off the package (Read request, Write request etc.).	47
2.18	Schematic of the interferometer Aurora-PCIe bridge. Only if the data from Aurora qualifies as the correct data stream is the PCS data fork (orange) active. The connection between Aurora and PCIe only passes through a process polling the data for a command meant for the bridge. Anything else is forwarded without modification. Because the data has no relation to the clocks on the SP605 all processing happens at full rates and is controlled via valid flags.	48
2.19	Raw noise measurement for both processing channels separately (CO ₂ in purple, HeNe in orange)at full 20 MS/s sampling rate. The dashed lines indicate the standard deviation. The asymmetry of the HeNe signal around zero appears to be caused by residual higher harmonics of the square reference frequency. The processing filters’ amplitude response is visible in the noise spectrum. Some of the frequency spikes are the result of bad measurement conditions.	53
2.20	Comparison of digitization noise for down-conversion (top three in purple) and Hilbert transformation (bottom three in orange) I-Q-generation using an optical HeNe reference arm. The dashed lines indicate the standard deviation. The down-conversion design also uses the sinusoidal optical reference. The optical setup is as depicted in 2.3 and was placed entirely on an optical bench to minimize any external influence.	55

2.21	Schematic of the test setup employed to evaluate the phase wrap stability. The waveform of both signals is generated by a single clock source on the waveform generator ensuring a fixed phase relationship. Since only two outputs are available, the HeNe measurements had to be conducted after the CO ₂ measurements (indicated by the grayed signal). The unused inputs were floated.	57
2.22	The phase wrap stability of each laser color FPGA arm as a function of rate of density change (left) and mirror speed (right). The horizontal dashed lines indicate a 2π phase wrap, i.e. when the respectively colored line crosses it a phase wrap has been missed. The shaded area indicates an upper boundary estimate for the maximum rate-of-change that can be expected in a tokamak device.	58
2.23	Bandwidth detection limit of the CO ₂ phase measurement arm. The top plot shows the measured frequency spectrum over the modulation frequency (normalized to peak maximum). The width of the vertical bars indicates the spacing between tested probing frequencies and the overplotted dashed line indicates the expected peak location. The middle plot shows the PSD amplitude of the detected peak. The bottom plot shows the deviation from the set signal as a function of the modulation frequency.	61
2.24	The combined event timing diagram of the FPGA system responding to an <i>ARM</i> (purple) and a <i>TRIGGER</i> (orange) command. Both commands are handled by the same logic process in the firmware allowing the deduction of the system latency by simple algebra. . . .	63
2.25	The measurement cords of the two IED interferometers on AUG. The CO ₂ -HeNe interferometer optics are mechanically decoupled from the AUG vessel. The DCN laser's IF was too low to be compatible with the SPEC firmware. The V3 line was unavailable at the time of the measurements. Hence measurements were conducted on either the V1 or V2 line.	65
2.26	Digitizer setup at AUG. The analog signals from the interferometer's photo detectors arrive in a diagnostic rack in the control room, where the SPEC is placed (a). A loaned PC was fitted with the SP605 bridge and connected to the SPEC via an OM3 fiber under the floor (b). Setting everything up took about 4 hours.	67

2.27	IED trace using the V1 CO ₂ -HeNe interferometer during a MGI induced disruption of AUG shot #33658. The AUG digitizer is overplotted using the dashed line. The inlay plots the region of highest rate of density change, found to be $1.8 \times 10^{25} \text{ m}^{-2}\text{s}^{-1}$, with full detail for $\pm 40 \mu\text{s}$ around the time of maximum rate of change (time axes shifted for better visibility).	68
2.28	Total toroidal plasma current in the aftermath of a MGI induced disruption of AUG shot #33658. The slow current decay and flat periods in the aftermath of the disruption are a strong indicator for REs.	69
2.29	Time slice of shot #33618 with modulated heating power. The SPEC digitizer is measuring on cord V1. The top plots the IED measured using the SPEC-digitizer down-sampled by a factor of 10. Below is the PSD of the same signal up to 500 kHz.	70
2.30	Time slice of shot #33618 with modulated heating power. The top plots a spectrogram of the LFS midplane ballooning coil (B35-14). Below is the heating power and composition applied during the time period. The second plot from the bottom shows the D_α power and the bottom one the down-sampled SXR intensity.	71
2.31	ECE temperature spectra for ECE channels 39 (top), 47 (middle) and 56 (bottom). The decimated ECE temperature is over-plotted.	73
2.32	Smoothed measurements of the IED on the V1 line (top plots) and electron temperature profile (bottom plots). The left side shows a slice with 400 kHz oscillations, the right plot a slice about 100 ms later, where the oscillations have disappeared. The time offset of each plot is indicated by t_0 at the top.	74
2.33	Spectrogram of the SPEC based density measurement on the HFS V2 cord at the start of AUG shot #33624. The top plots the IED, with its spectrogram below. The second plot from the bottom prints the LOS averaged Alfvén velocity and the bottom plot the V2 cord length from the equilibrium reconstruction. The vertical line in each plot is the onset of NBIH.	78

2.34	Time slice of AUG shot #33624. The top right plot shows the raw density trace from the SPEC digitizer (note the logarithmic scale). The top left plot shows the corresponding spectrogram. The bottom left depicts the applied heating power, the bottom right is the normalized CORDIC amplitude. The vertical lines in each plot mark the start of the ICRH as set by the AUG control system.	81
3.1	Schematic of the plasma pressure in L-mode and H-mode. The formation of a strong edge pressure gradient and pedestal are the defining properties of H-mode.	85
3.2	Diagram of the ELM peeling-ballooning stability boundary, with the stable region marked in orange. The dynamics in the parameter-space of all three types of ELMs are indicated. The pressure or current density rises in direction of the instability boundary, before the fast crash occurs though the unstable region[102].	86
3.3	Basic geometry behind the FF synthetic aperture principle.	95
3.4	Basic principle behind Doppler back-scattering on a density corrugation of scale κ for monostatic and bistatic setups. The motion of the plasma corrugation results in the scattered signal being Doppler shifted.	101
3.5	Wavenumber component map at the cut-off surface of a 16 GHz microwave launched omni-directionally from (0,0) for MAST shot #27969 230 ms into the shot. The left plot depicts the magnitude of the component parallel to the local magnetic field. The right one shows the amplitude for the component perpendicular to the local magnetic field. The map was calculated using the TORBEAM code with measurements from TS. The shaded area at the top marks the position of a MAST poloidal field coil obscuring the view (adapted from [111] with kind permission from the author and AIP publishing).	102
3.6	The schematic of Doppler back-scattered spectral power around the probing frequency f_{probe} (shifted to 0 MHz) in two opposing directions. In a) the Doppler shift results in more red-shifted power, in b) (looking in the opposing direction) more blue shifted power is observable (adapted from [111] with kind permission from the author and AIP publishing).	103

3.7	2D red-blue imbalance map for MAST shot #27969. The colorbar is given in units of volts. The data is generated using the SAMI diagnostic for a 16 GHz omni-directional probing beam 300 ms into the shot. The overplotted curved lines indicate the magnetic field lines taken from MSE constrained EFIT. The EFIT location was matched against the SAMI probing frequency with the help of TORBEAM using TS data (taken from [111]).	104
3.8	The SAMI RF box. On the left is the synthetic aperture array using the linearly polarized sinuous antennae (only the antennae not covered in foil are used for imaging). On the right is the RF box containing the hybrid couplers and down-conversion electronics. To minimize RF losses the box is directly attached to the array and bolted on to the machine.	106
3.9	Top level schematic of the diagnostic hardware. The acquisition array is depicted in the inlay. The two shielded emitting antennae used for 2D DBS are marked. The FPGA based acquisition consists of two boards described in section 3.5.2.	107
3.10	Modifications to the SAMI FPGA firmware. New additions are marked in orange. Different shades of signal lines denote different sampling rates. The signal composition and the parts used for different parts of the design are used on top of the design (adapted from [116] with kind permission from AIP publishing).	111
3.11	Processing core for the 2D DBS pipeline. The initial filters select the orange part of the signal spectrum in fig. 3.10. Due to resource limitations only 4 of the 8 channels are interleaved for the down-conversion. The sine wave used for mixing is generated using a LUT with a synchronized address counter.	112
3.12	Revised clock infrastructure. Purple components are only synchronous to the board they are on. Orange components are synchronous across boards. The dashed signal line has a stable but unknown phase relationship between channels. Different shades of signal lines denote different sampling rates (adapted from [116] with kind permission from AIP publishing).	114

3.13	Layout of the SAMI firmware logic on the Virtex-6 FPGA generated using Xilinx PlanAhead. The serial data pins from the ADC are placed in the lower left of the FPGA. The pins related to streaming the data via fiber optic are in the top right. The logic in between is heavily constrained in placement due to the clock speeds it is operated at and the location of the physical pins it controls, most notably the RAM controller.	116
3.14	Aurora transmission protocol employed in SAMI. Each channel pair is identified by the combination of the top bits of each channel word (marked orange). The EOS can only be encoded by transmission of a decreasing top bit count.	118
3.15	The layout of the SAMI hardware on NSTX-U. The diagnostic is protected from Lithium conditioning via an EPICS controlled shutter. The diagnostic hardware has to be physically spread out requiring the implementation of optical interfaces (red) for all but the IF and RF signals (black).	121
3.16	Comparison of image reconstruction algorithms for a planar source emitting at 16 GHz. In each quadrant the simulated geometry is depicted at the top, with ρ given in meters; purple triangles mark the array geometry (scaled by a factor of 20 for visualization purposes). The polar plots in the middle use FF reconstruction with BF (left) and Fourier inversion (right). The bottom plots show the NF reconstruction using BF (left) and SVD inversion (right). Axis descriptions have been omitted for illustration purposes. Cartesian axis labels are in meters. The opening angle in the polar plots is in degree.	124
3.17	Image reconstruction algorithms compared for different distances of a planar source emitting at 26.5 GHz. The simulation geometry and plot-layout are identical to fig. 3.16.	126
3.18	Normalized distortedness δ over ρ for all four image reconstruction methods and two different probing frequencies. <i>delta</i> is normalized to the average over ρ . Quantitative comparison of the distortedness curves should be taken with great care.	128
3.19	Image reconstruction algorithms compared for different distances of an ellipsoid source emitting at 20.5 GHz. The plot-layout is identical to fig. 3.16.	130

3.20	Normalized distortedness δ over ρ for a probing frequency of 20.5 GHz in spherical geometry. δ is normalized to the average over ρ . Quantitative comparison of the distortedness curves should be taken with great care.	131
3.21	Normal incidence cut-off surfaces for an L-mode and an H-mode section of NSTX-U shot #204672. The R-wave cut-off ν_R (see sec. 3.2.2) is plotted in solid blue, the density cut-off ν_{pe} is plotted in magenta. The SAMI acquisition frequencies are marked with circles (O-mode) and crosses (X-mode). (Taken from [115]).	132
3.22	Comparison of 2D DBS power received for 10 GHz acquisition with the SAMI diagnostic on MAST and NSTX-U. Each measurement is compared against the null measurement without plasma (purple). The NSTX-U 2D DBS power is significantly higher above the noise floor (adapted from [115] with kind permission from the author and AIP publishing).	133
3.23	Pitch angle trace for NSTX-U shots #204620 and #204944 using 10 GHz and 16 GHz 2D DBS respectively. The SAMI measurement is in purple, the EFIT fit is in orange. The agreement is very good over most of the shot for the 16 GHz shot. A significant overestimation of the pitch angle (in comparison to EFIT) is found for shot #204620 probing at 10 GHz (adapted from [115] with kind permission from author and AIP publishing).	134
3.24	Selection of plasma parameters for the shots shown in figure 3.23. The top depicts the TS calculated IED. Below are the measured toroidal plasma current and the midplane Deuterium- α light intensity. The bottom is the applied NBIH power.	135
3.25	Simulated result of the 2D DBS down-conversion logic operating on SAMI data. The top plot depicts the spectrum before and after 2D DBS processing using data from NSTX-U shot #204620 240 ms into the shot. The acquisition frequency was 10 GHz. The spectrum is calculated from a 1 ms timeslice. The bottom plot shows the frequency response of each of the three filter stages as well as the combined response. The data-sets have all been shifted to have the respective active probing frequency peak at DC.	137

3.26 Image reconstruction algorithms compared for the array configurations using Vivaldi (left) and sinuous antennae (right). The source is emitting at 16 GHz is at a distance of roughly 1.3 m from the array. The plot-layout is identical to fig. 3.16. 139

Glossary

W7-X	Wendelstein 7-X
WRS	White Rabbit switch
ADC	analog-to-digital converter
AE	Alfvén Eigenmode
AL	asynchronous logic
AOM	acousto-optical modulator
ASDEX	Axially Symmetric Divertor EXperiment
ASIC	application specific integrated circuit
AUG	ASDEX Upgrade
BAE	beta-induced AE
BF	beamforming
BRAM	block RAM
CAE	compressional AE
CCFE	Culham Centre for Fusion Energy
CLB	configurable logic block
CORDIC	COordinate Rotation DIgital Computer
CPU	central processing unit
CXS	charge-exchange spectroscopy
DAC	digital-to-analog converter
DBS	Doppler back-scattering
DC	direct current
DCM	digital clock manager
DDR	double data-rate
DEMO	DEMOstration power plant
DIII-D	Doublet 3 D

DMA	direct memory access
DRO	dielectric resonator oscillator
DSP	digital signal processing
EBW	electron-Bernstein wave
ECCD	electron cyclotron current drive
ECE	electron cyclotron emission
ECRH	electron cyclotron resonance heating
ELM	edge localized mode
em	electro-magnetic
EOS	End-Of-Stream
EPICS	Experimental Physics and Industrial Control System
FF	far-field
FIFO	first-in-first-out
FIR	finite-impulse response
FM	frequency modulation
FMC	FPGA mezzanine card
FPGA	field programmable gate array
FSM	finite state machine
GAE	global AE
GPIO	general purpose IO
GPU	graphics processing unit
HDD	hard-disk-drive
HDL	hardware description language
HFS	high-field side
HLS	high level synthesis
HTSC	high temperature super-conductor
ICF	inertial confinement fusion
ICRH	ion cyclotron resonance heating

IED	line of sight (LOS)-integrated electron density
IF	intermediate frequency
IO	input/output
IPP	Max-Planck-Institute for Plasma Physics
ITER	International Thermonuclear Experimental Re- actor
JET	Joint European Torus
KAW	kinetic Alfvén wave
KSTAR	Korean Superconducting Tokamak Advanced Research
LAN	local area network
LCFS	last closed flux-surface
LFS	low-field side
LHD	Large Helical Device
LO	local oscillator
LOS	line of sight
LPC	low-pin-count
LSB	least significant bit
LUT	look-up table
MAST	Mega-Amp Spherical Tokamak
MAST-U	MAST-Upgrade
MCF	magnetic confinement fusion
MGI	massive gas injection
MHD	magneto hydro-dynamic
MSB	most significant bit
MSE	motional Stark-effect
MST	medium sized tokamak
MTS	mini-time-stamp
MUX	multiplexer

NBI	neutral beam injection
NBIH	NBI heating
NF	near-field
NSTX	Nearly Spherical Torus eXperiment
NSTX-U	NSTX-Upgrade
NTM	neo-classical tearing mode
OHWR	Open HardWare Repository
OS	operating system
PCB	printed circuit board
PCI	peripheral component interconnect
PCIe	PCI Express
PCS	plasma control system
PFC	plasma facing component
PID	proportional-integral-differential
PLL	phase-locked loop
PM	phase modulation
PPPL	Princeton Plasma Physics Laboratory
PROM	programmable ROM
PSD	power spectral density
PWM	pulse width modulation
RAM	random access memory
RE	run-away electron
RF	radio frequency
ROM	read-only memory
SAMI	Synthetic Aperture Microwave Imaging
SET	single-event-transient
SEU	single-event-upset
SFP	small form-factor pluggable
SL	synchronous logic
SNR	signal-to-noise ratio

SOC	system-on-chip
SODIMM	small-outline-dual-inline memory module
SOL	scrape-off layer
SPEC	simple PCIe carrier
SRAM	static RAM
SSD	solid-state drive
ST	spherical tokamak
START	Small Tight Aspect Ratio Tokamak
SVD	singular value decomposition
SXR	soft X-ray
TAE	toroidal AE
TDD	time-division de-multiplexing
TDM	time-division multiplexing
TFTP	Trivial File Transfer Protocol
TS	Thomson scattering
TTL	transistor-transistor logic
UART	Universal Asynchronous Receiver/Transmitter
UDP	user datagram protocol

Chapter 1

General Introduction

1.1 Nuclear Fusion

Harnessing nuclear fusion as a power source is probably one of the most important scientific challenges of our time. The world's demand for energy has more than doubled within the past 40 years with the majority of growth coming from third world countries such as China[7]. It is undisputed that the growth rates will increase as the less industrialized parts of the world are catching up.

The production of energy from fossil fuels, which currently contributes over 80% of the world's energy production, releases large amounts of CO₂ into the atmosphere[7]. It is accepted that man-made CO₂ as a greenhouse gas is a major contributor to global warming and the international community's consensus is that future emissions should be reduced[8]. In addition our heavy usage of fossil fuels as the traditional supplier of energy will start to grow increasingly scarce within the coming decades and start depleting within the coming century -oil being the most prominent example[9]. The currently well known alternatives are renewable energy sources, such as wind, hydro and solar, or nuclear fission. All of these sources however have some drawbacks.

Renewable energy sources are generally intermittent with a very low energy density[8]. Producing the vast amounts of energy we demand therefore requires large amounts of space that could otherwise be used for other tasks such as food production. In addition the intermittency requires the installation of large amounts of back-up capacity to cover for times when there is no electricity produced. Not all forms of renewable energy share this problem. However, for these forms the production potential is either locally limited and saturated, e.g. hydro, or the technology in its infancy, e.g. offshore wave. Future developments especially in the field of energy storage will most likely increase the potential for renewable energy

production. However, although the potential for electricity generation is sufficient to meet the world's energy needs, it is nonetheless unlikely that renewable energy will be able to supply all of it[10].

An already established alternative to renewable energy production is nuclear fission. Unlike renewable sources nuclear fission is much more suited for continuous gigawatt scale energy production, which is what the grid in most of the world is optimized to handle[8]. On the other hand there are a few well known disadvantages to the process, i.e. the production of long lived radioactive isotopes and the possibility for runaway chain reactions. This has long been the source for public criticism, although newer reactor designs attempt to reduce or mitigate these problems[11]. In addition the possibility to use the technology to produce nuclear weapons continues to prevent the implementation of more advanced reactor concepts. This in turn limits the currently economically accessible fuel reserves to about 130 years, although next generation reactor technology would increase this limit to well over 10,000 years[11], [12].

Nuclear fusion can promise the best of both worlds: the production of gigawatt scale power without the greenhouse gas emissions or long-lived nuclear waste products. The known fuel reserves on earth are already plentiful enough to supply the world for many millions of years.

1.2 Basic Physics

Nuclear fusion, like nuclear fission, harvests the binding energy freed when two atoms fuse - or in the case of fission split. The amount of energy liberated is determined by Einstein's most famous equation $E = \Delta mc^2$, where Δm is the mass defect, i.e. the difference in mass between the sources and products of the nuclear reaction, and c is the vacuum speed of light. This mass defect originates from the difference in total binding energy between source atoms and reaction products, which is highly dependent on the atomic number.

Figure 1.1 depicts the binding energy as a function of atomic number. As can be seen the maximum of the curve occurs at ^{56}Fe , which has the highest binding energy of the stable isotopes. As indicated by the graph, atoms with a lower atomic number will generally release energy when fusing. Although in principle all elements below iron can be used to generate electricity, the most relevant elements for the initial development of a power plant are the ones freeing the most energy per fused atom. In fig. 1.1 Deuterium and Tritium fusing to Helium-4 are marked as the most important candidates. Although one can find combinations of higher gain one has




Figure 1.1 – The binding energy as a function of atomic number. The maximum of the curve marks the transition from fusible to fissile atoms. The (currently) most relevant atoms for fusion are marked (adapted from [13]).

to consider another aspect namely the cross-section or “How easy is it to fuse two atoms?”.

In order to fuse two atoms together one has to surpass the Coulomb barrier, i.e. the repulsive force between two positively charged nuclei. This repulsive force increases exponentially the closer two of them get together until the strong binding force of the atomic nucleus starts to act at about 10^{-15} m. A particle would thus have to have a kinetic energy high enough to surpass this energy barrier if it was to fuse with another particle. This immediately leads to the conclusion, that low Z elements will be easiest to fuse due to the lower nucleic charge, where Z is the atomic number.

Nonetheless one would probably deem nuclear fusion power plants impossible if only classical forces were considered. Quantum physics however tells us that particles can tunnel through the Coulomb barrier if the distance between the two nuclei is small enough[15]. This effect significantly increases the fusion cross section making the concept feasible again. Since the cross section is such an important factor, the selection of a fuel can be based on this parameter. Figure 1.2 depicts the fusion cross section for a few low-Z elements. As can be seen the Deuterium-Tritium fusion reaction has the highest cross section at the lowest energy making it the best





Figure 1.2 – The fusion cross section for some of the lowest- Z elements. Deuterium-Tritium fusion has the highest cross section at the lowest energy (taken from [14])

candidate to build an initial fusion power plant.

Tritium and Deuterium fuse via the following reaction:



freeing a total of 17.6 MeV of energy[15]. This mostly goes into the kinetic energy of the neutrons. Extracting this energy is well established in nuclear fission where one uses water to slow them down. The heat generated this way can be turned into electricity via a traditional turbine[16].

1.2.1 The Lawson Criterion

Strong heating is necessary to get the particles to the required energies. In most cases this will result in the fuel to transition into the plasma state, i.e. a quasi-neutral collection of semi-freely moving charged particles. This will initially come from external sources, and even in a final plant it is likely that some auxiliary heating P_H is applied. However, the ultimate goal in a fusion power plant is for the heating to be supplied entirely by the fusion reaction in 1.1. The power with which a fusion plasma can heat itself is proportional to the reaction rate $\langle\sigma\nu\rangle$, which incorporates the previously discussed reaction cross section σ and relative particle velocity ν , which is a function of temperature[15]. In addition the density n and the energy freed per reaction ϵ are important in calculating the self-heating power, yielding

$$P_\alpha = \frac{n^2}{4} \langle\sigma\nu\rangle \epsilon. \quad (1.2)$$

Equation 1.2 has already been maximized with the ideal Deuterium-Tritium density ratio.

At the same time any energetic system not in equilibrium with the surroundings continuously loses energy. The rate at which energy is lost, can be simply described via

$$P_L = \frac{E}{\tau_C},$$

where τ_C is the energy confinement time of the “container”. For a magnetic confinement device (as described in the following sections) the energy contained in the plasma can be estimated with $E = 3\overline{nT}V$, where V is the plasma volume and \overline{nT} is the volume averaged product of density and temperature[15]. In steady state conditions the losses have to be balanced by the heating terms, i.e.

$$\frac{3\overline{nT}V}{\tau_C} = P_H + \frac{n^2}{4} \langle \sigma \nu \rangle \epsilon, \quad (1.3)$$

and ideally there will be a point at which the self heating term P_α is large enough, so that external heating is not required. This point is called ignition and can be roughly described by

$$n\tau_C > \frac{12}{\langle \sigma \nu \rangle} \frac{T}{\epsilon}. \quad (1.4)$$

The inequality 1.4 is called the “Lawson Criterion”[15]. It denotes a limit for the product of energy and confinement time for a given temperature in order to reach a “burning” condition. The criterion here is specific to magnetically confined devices, however the simplicity of the assumptions makes it possible to derive the same criterion for any power plant concept. All the research into nuclear fusion has been aimed at reaching this very point. Several concepts have been proposed to this end, e.g. inertial confinement fusion (ICF). In this concept the fuel is quickly compressed, e.g. by driving it radially inward using lasers[17]. The confinement time τ_C is therefore very small, but the density is very high, resulting in the criterion being met. The work presented in this thesis is relevant to the magnetic confinement fusion (MCF) concept, and more specifically tokamaks, to which we will confine the descriptions. For details on other concepts such as ICF the reader is kindly referred to literature as this would be beyond the scope of this thesis, e.g. Pfalzner[18].

1.3 Magnetic Confinement Fusion Devices

Magnetic confinement fusion is currently the most heavily researched concept aimed at developing a fusion power plant. It relies on magnetic field coils to keep the fully ionized fuel, or plasma, within a confined space, thus reducing the recombination

losses occurring when the plasma touches a solid wall. Again several subcategories can be established, e.g. by using magnetic cusps[19]. However, so far the toroidal confinement device has been the most successful.

The base features are a set of magnetic field coils arranged in a torus shape as depicted in figure 1.3. Since the particles inside the field are charged they will have to follow the magnetic field lines around the torus and thus are confined to it. However, this alone would not be enough to achieve high confinement due to plasma drifts induced by the curved magnetic fields, namely the curvature and grad-B drift. These two drifts cause charge separation in the plasma and result in a secondary drift ($E \times B$ -drift) due to the resulting electric field. This drift in turn causes the particles to move radially out of the confinement region. In order to improve confinement the axial field created by the torus coils needs to be twisted to average out the drifts.

The creation of this poloidal field component is where stellarators and tokamaks differ, which are the two major MCF concepts currently under development. The work presented here has been conducted on a tokamak type device and thus focus will be laid on this family of MCF devices. For a more thorough description of stellarators and heliotron devices the reader is again kindly referred to literature, e.g. Wakatani[20].

1.3.1 Tokamaks

Tokamak is the transliteration from Russian meaning “toroidal chamber with magnetic coils”. This concept uses the plasma itself to create the poloidal field component by inducing a toroidal current. This works in essence like using the plasma as the secondary circuit of a transformer. The plasma thus confines itself. In reality there is more to this than just the simple current drive as the shape of the plasma, e.g. triangularity and ellipticity of the cross-section, as well as the design of the plasma-vessel-interface, e.g. divertor and wall, and many other parameters play significant role in the confinement properties. All of these need to be optimized to achieve the plasma performance required for a power-plant-grade plasma[15].

The advantage of the tokamak setup over the stellarator, where the poloidal field component is being created externally without the plasma using a separate set of coils, is the much simpler and thus cheaper setup. However, this comes at the cost of a few disadvantages. One immediately notices that the tokamak (in its simplest form) can only operate in pulsed operation. This is simply due to the fact that the current in the plasma cannot be reversed as that would lead to a short loss of confinement. In consequence the power supplies driving the plasma current will eventually meet their voltage limit and will have to ramp back. This limit is well known and there




Figure 1.3 – The basic schematic of a tokamak (taken from [21]).

are efforts to develop scenarios for continuous operation, e.g. by driving the current via neutral beam or electro-magnetic (em) wave current drive[22]. Another known disadvantage of the tokamak setup is the fact that driving plasma currents can also drive instabilities in the plasma, increasing the performance demands needed for real-time feedback systems to prevent the plasma from becoming unstable.

Despite these disadvantages the tokamak is currently considered the fastest path towards a fusion power plant. Hence the international efforts currently focus on this family with the biggest experiment being ITER - previously named International Thermonuclear Experimental Reactor. This experiment is currently being built near Cadarache, France, and aims to be the first experiment to reach a Q factor of 10, which is the ratio of output fusion power over supplied heating power, thus generating 500 MW of fusion power[23].

1.3.2 Spherical Tokamaks

As has been pointed out in the previous section, the geometrical properties of the fusion plasma play a vital role in the confinement properties of the tokamak. One of the central geometrical quantities defining a tokamak plasma is the aspect ratio, i.e. the radius of the torus (known as major radius) over the radius of the plasma (known

as minor radius). Conventional aspect ratio tokamaks, such as the Joint European Torus (JET), commonly have an aspect ratio of 2 to 3[24]. A special case of the tokamak geometry is the so called spherical tokamak (ST), where the aspect ratio is brought close to 1 or even lower, leading to an almost spherical highly elongated plasma cross section as the name suggests[24]–[26]. The low aspect ratio and high elongation has shown to significantly increase the Troyon beta limit[24]. This hard magneto hydro-dynamic (MHD) limit describes the amount of plasma current that can be supported in a tokamak for a given magnetic field strength and plasma radius. It is found to be

$$\beta_L = 4l_i \frac{I_p}{B_\theta a}. \quad (1.5)$$

Here I_p is the plasma current, a is the minor plasma radius, i.e. the distance from the center of the plasma column to the plasma edge, and B_θ is the on-axis magnetic field[15]. The quantity $\beta = 2\mu_0 p/B^2 < \beta_L$ is the ratio of plasma pressure to magnetic pressure and l_i is the internal inductance of the plasma. If the plasma current is too high, low order MHD instabilities will start appearing and quickly lead to a disruption - a violent breakdown of plasma confinement. In conventional tokamaks the plasma beta generally cannot exceed a few tens of percent, although most conventional experiments nowadays will not surpass 5%. In a ST on the other hand the theoretical limit has been found to be 100%, where experiments like Small Tight Aspect Ratio Tokamak (START) have achieved record values of 40%[25]. There are multiple effects responsible for this improvement amongst which is the increased magnetic shear, which substantially reduces turbulent transport across field lines.

A higher plasma beta means that smaller coils can be used for the same amount of confinement, thus making a ST potentially cheaper, but definitely smaller than a conventional tokamak. This fact has resulted in several endeavours to design ST based power plants[25]. Such economic considerations are important for the entire field, since an eventual fusion power plant will have to be economically competitive with respect to the conventional alternatives. The very desirable properties of the ST have resulted in several experiments being built around the world. The two largest are the NSTX-Upgrade (NSTX-U) in the USA and the Mega-Amp Spherical Tokamak (MAST) in the UK, where most of this thesis' work was conducted[25], [26]. Both carry plasma currents of the order of 1 MA.

Although also considered an advantage, the ST's compactness also results in the disadvantage of having little space for neutron shielding. This increases the material demands, which are already problematic in normal aspect ratio devices.

Since ST power plants will likely rely on high temperature super-conductors (HTSCs), which are susceptible to neutron damage, this is of particular importance[25]. The compactness in fact is a problem not only for the center column. The smaller device also means that there is less space available to place peripheral components, most importantly the Tritium breeding blanket, which generates the Tritium fuel from Li and is a central component to any power plant concept[27]. It is also conceivable that the ST is more susceptible to fast particle losses, the most prominent of which are the fusion- α particles, since Alfvénic modes, which can enhance fast particle losses, are more easily excited in a low field, high density device[26].

Nonetheless there are currently industrial efforts underway to pursue the ST concept as a faster cheaper path to fusion than the current ITER baseline[28].

1.3.3 Real-Time Fusion Diagnostics

As has been pointed out before the tokamak concept is inherently susceptible to current driven instabilities, such as neo-classical tearing modes (NTMs). These instabilities can result in a violent loss of plasma confinement[15]. One of the worst case events that can happen is a so-called disruption. There are many causes that can lead to a disruption and hence it is not surprising that today’s experimental reactors quite often experience such events. In these relatively small machines they do not carry much energy and thus are often simply accepted byproducts. However, on a machine like ITER and the even bigger machines of the future the energy carried by a disruption can be large enough to cause significant damage to the first wall components or even the whole of the vessel[29]. Repairing such damage can take months and thus would not be acceptable for a commercial operator. In consequence disruptions need to be avoided, which is done by controlling the initiating instabilities.

The timescales of these are often machine size dependent, but can range from tens of microseconds to milliseconds. Since the underlying stability criteria are often marginal and their “eruption” is not always predictable, real-time feedback systems are required to control the plasma. These often rely on diagnostics, which, being a part of the protection system, have to supply their data robustly with low latency.

The diagnostic set on a fusion experiment is vast. ITER has at least 40 different diagnostics of which several deliver measurements of the same plasma parameter[30], [31]. This is to ensure that the critical parameters are adequately controlled. Many of the diagnostics will contribute to basic machine protection and control and thus have high requirements regarding latency as well as temporal and spatial resolution[30], [32].

The requirements in terms of latency and accuracy are not limited to the mere

signal digitization. On the contrary many stability conditions depend on 3 or more parameters, which often are acquired by separate diagnostics. In consequence real-time processing of large amounts of data for a reliable feedback control is required[33].

To tackle the latency and throughput margins modern technologies have to be employed. Most notably the use of graphics processing units (GPUs) and more recently field programmable gate arrays (FPGAs) as parallel data processors is of interest here. From a real-time high-performance low-latency diagnostic point of view, FPGAs are the technology of choice.

1.4 Field Programmable Gate Arrays

A FPGA is a semiconductor device that can be reconfigured to create any arbitrary logic circuit. It is thus very different from an application specific integrated circuit (ASIC), which is static hardware. FPGAs evolved out of programmable ROM (PROM) technology, which could implement logic circuits using its address decoder[34]. The technology is based on static RAM (SRAM) and thus greatly benefited from Moore's law applicable to the rest of the semiconductor industry. Despite these rapid improvements it was only in recent years that FPGAs became capable enough to be useful in the field of plasma diagnostics.

1.4.1 The FPGA Hardware

Any logic circuit can be described using only *NAND* or *NOR* gates[35]. In modern microprocessors from companies such as Intel *NAND* logic is preferred, since it can be implemented more efficiently and only requires a single type of logic element[36]. A FPGA creates logic circuits by replacing the gate logic using configurable look-up tables (LUTs), which can be freely interconnected. This is indicated in schematic 1.4a.

The central component of a FPGA is the configurable logic block (CLB) (see figure 1.4b)[35]. Although the specific hardware varies between vendors and chips the component is generally made up of a 4 to 6 input LUT as well as a register in the form of a flip-flop. The switch-matrix is the component allowing arbitrary connection of CLBs to design logic circuits (bright purple in fig. 1.4a). These logic circuits generally change state on the rising or falling edge of some global clock, which is distributed throughout the chip using dedicated routing lanes. This so called synchronous logic (SL) very flexible enabling the composition of logic ranging from a simple adder up to full blown system-on-chip (SOC) running an operating system. Its counter-part is called asynchronous logic (AL), where information is




Figure 1.4 – Basic elements contained in FPGA fabric. The logic elements in the top right of (a) can be replaced by LUTs in a CLB (marked in purple). The light purple elements mark configurable interconnects. Routing lanes are marked orange. The MUX element (purple in (b)) controls whether the CLB acts on a clock edge or is transition logic (adapted from [35]).

directly processed as it arrives. This has the disadvantage of requiring meticulous planning to ensure that information from different sources arrives at a processing element at the correct time. SL significantly eases this design difficulty by a “divide-and-conquer” approach, i.e. only the timing from one synchronous element to the next must be carefully planned out. The information travels between these elements at a well-known rate defined by the clock, which eases the combination of different logic devices. These advantages have resulted in SL to become the basis for most of modern digital logic.

In modern FPGAs the CLB is not the only component. To enhance the functionality of the FPGA fabric companies have embedded dedicated delay lines, memory elements, fast (de-)serialization components for data acquisition or streaming, dedicated multipliers and more[37]. Even though the functionality they provide can in principle be replicated by interconnecting CLBs the higher performance and density of these components have allowed a great deal of resource savings and performance increases in most FPGA-designs. The components are often spread throughout the actual chip to allow fast logic to meet timing.

Since FPGAs are based on SRAM technology they are volatile, i.e. the configuration of the chip is lost if it loses power. To program a FPGA a configuration file, also known as “bitstream”, is stored in a non-volatile memory on the carrier board. At power up this memory is read out and writes the configuration to the FPGA chip. Once programmed the logic circuit is permanent -unless manually reprogrammed- until the chip loses power again. On newer generation chips it is even possible to only partially reconfigure the chip during operation, thus allowing “hotswapping” or

“patching” of logic.

This feature is especially useful in applications where the outside interface to the FPGA must not change during operation, e.g. co-processing via peripheral component interconnect (PCI) or PCI Express (PCIe). Although the PCIe standard per-se is hot-pluggable, the implementation on motherboards often is not[38]. On the contrary the PCIe interface on a motherboard needs to be up and running within the first 100 ms of power up, so that the BIOS can poll and enumerate all PCIe devices correctly to make them available to the operating system (OS). Using partial reconfiguration the interface visible to the carrier system can be kept identical, whilst the functionality of the interfaced logic can change as needed.

1.4.2 The FPGA Design Flow

FPGAs are programmed using a hardware description language (HDL). The two most commonly used are VHDL and Verilog. Each of them has its own advantages and disadvantages and for most of the work conducted in this thesis VHDL was chosen. Programming using a HDL is different from the commonly known languages such as C. While most languages used to program standard central processing unit (CPU) programs are serial, i.e. the first instruction written down is executed first, then the second one written down etc., a HDL is used to describe synchronous or asynchronous logic operating on the chip. This means that the position of the code in the file says nothing about the order of execution. On the contrary, because one describes synchronous logic, every statement in the code is executed simultaneously. A FPGA design can often be pictured as a set of blocks that are interconnected via wires, which is the basis for most high level synthesis (HLS) tools. These try to abstract the HDL level design using a graphical approach.

The rate of instruction execution is determined by the clock frequency used in the synchronous logic and one can use multiple unrelated clocks in one design. A designer thus has to not only make sure that the logic is sound, but also that any data moving between clocks transitions correctly. Even within a single so-called “clock domain” meeting timing can be difficult as logic gets more complex and clocks get faster (clocks of several hundred megahertz are not uncommon). Timing problems within a single clock domain essentially originate from the fact that information requires a set amount of time to travel from point A to point B on the chip. If the clock is faster than the information can travel then the logic will behave erratically because information might get updated before it has been used. This way no well-defined logic circuit can be built. Another cause for issues might be the quality of the clock known as “jitter”, which describes the frequency and duty cycle stability of the clock.

To design working synchronous logic this quality factor needs to be well known. One also needs to consider that logic paths running on the same clock may differ too much in their respective paths-length.

To circumvent these issues logic is often designed to use many short paths, operations are conducted on more than one clock cycle or made less complex. Due to the flexibility of the FPGA hardware it is also possible to reduce the clock speed to the point that timing can be met. With modern FPGA design tools the task of meeting timing is much eased as many optimizations are employed automatically or can be iterated to find the optimum build-configuration. For average FPGA-designs the tools can meet timing even for suboptimally designed logic. However, for heavily constraint designs the tools quite often fail (even randomly). In such cases the solution to tools not meeting timing often requires redesigning of logic - sometimes in a non-obvious way involving educated guessing, trial and error. This is the result of the built process, which involves random numbers to determine build paths. Because building a FPGA bitstream can easily take over an hour, even on modern hardware, designing FPGA firmware can be lengthy and time consuming. For this reason simulation of the logic is often employed to reduce the time spent in the building chain. It is to be noted, that these problems generally tend to occur in heavily constraint designs only.

1.4.3 FPGA vs. Microprocessor/GPU

FPGAs are not omnipotent. When comparing it to other technologies such as microprocessors and GPUs, each is found to have very specific benefits and weaknesses[39]. Even though in principle it is possible to run a SOC, i.e. a full microprocessor system, on a FPGA, this is generally not the operation scenario FPGAs are good at. The clocks driving FPGA logic are generally an order of magnitude slower than those found in hard-wired microchips[39]. Nonetheless FPGAs have the advantage in certain operation scenarios. Most commonly a FPGA will excel at parallel pipelined processing, i.e. multiple streams of data samples having the same operation applied to them. FPGAs can naturally operate on each sample of a data stream without communication overhead, i.e. while the first value has the second operation applied the second value can already have the first operation applied simultaneously. As long as resources on the FPGA permit, any number of parallel pipelines can be set up, thus naturally increasing the parallel processing throughput. Although pipelining is also implemented in virtually all modern processors, their pipelining capabilities are usually limited by memory read/write operations[40]. For example a simple addition requires at least fetching two values and the instruction from memory followed

by applying the instruction. Finally the data has to then be written to memory again. Using multiple processors does not necessarily remedy this situation, since the communication with memory prevents linear scaling of the processing throughput. This limitation is not present on a FPGA, however the benefit in pipeline processing speed can of course be nullified by the speed of data supply to the FPGA. In fact it is often the case that the data interface bandwidth is the limiting factor[41].

Examples for applications FPGAs are not well suited for are floating point arithmetic and loop-type processing, unless the number of iterations is known and relatively small[39]. If the number of iterations is unknown, for example when a fitting procedure has a threshold value exit condition, the implementation on a FPGA is often difficult, if not impossible. In consequence these operations are often much better suited to be completed on a microprocessor. Another downfall, even in modern FPGAs, is the still limited amount of resources available on the chip. Even though this is becoming less and less of an issue there are still scenarios, where the microprocessor has the advantage due to the vast amounts of memory available to it. This is often the case in massively parallel codes, which so far have been firmly placed on CPUs and GPUs.

Because each technology has its own advantage the other cannot cover but most problems involve touching both regimes, the newest generation of FPGA devices actually attempt to bridge the gap by integrating the advantages of both into one chip. The Zynq architecture, which was introduced in 2013, combines a complete hard-wired ARM processor core including a memory controller with on-the-fly reconfigurable FPGA fabric[42]. The processor's logic is physically edged into the silicon, thus being capable of Gigahertz scale clock speeds, which -so far- is beyond the reach of the "soft" logic that can be programmed into the FPGA-fabric. However, the FPGA-fabric can directly interface with the ARM processor via a common bus and has the same capabilities as the standard modern FPGA. This type of architecture enables the use of FPGAs to speed up certain processing tasks, for which it is well suited, and leave the more processor suited tasks to the ARM processor.

Apart from the just discussed raw processing speed FPGAs have a much more important advantage when it comes to real-time control applications. Since the synchronous logic is all laid out on the chip and the operation speed is fixed and the latency of a data package traversing the processing pipeline is well-defined. As a consequence there is no uncertainty regarding the latency of data apart from clock jitter. This is a significant advantage over microprocessor systems, where the overhead to manage other parts of the system often results in instructions waiting for their execution. Even with interrupts, which are signals that can give absolute

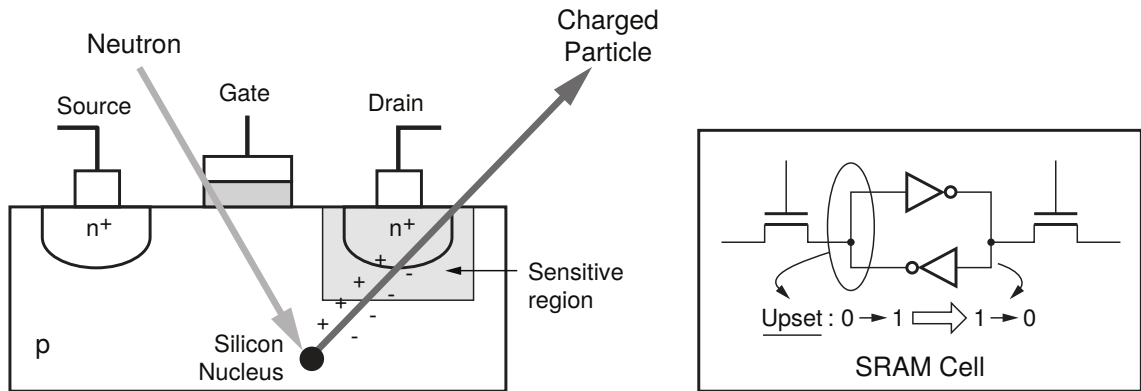


Figure 1.5 – The basic source for a SET. A charged particle, e.g. generated by a neutron striking an atom, traverses the depletion region creating a cloud of electron-hole pairs in the doped silicon substrate thus potentially leading to an unwanted bit flip or a current surge (taken from [43]).

priority to a timing critical task on a processor, there is a significant amount of uncertainty involved. In particular when it comes to real-time control and machine protection, these levels of uncertainty can be unacceptable, making FPGAs the technology of choice.

1.4.4 FPGAs In Nuclear Industry

Because more and more FPGAs are being used in today’s fusion experiments one should raise the question, how FPGAs as a volatile semi-conductor device will behave under radiation and in particular high energy neutron radiation.

Ionizing radiation can have significant impact on electrical circuitry. If a charged particle traverses doped material it will generally leave a cloud of electron-hole pairs as shown in fig. 1.5. In a transistor this can lead to a sudden increase in conductivity through the p-depletion region or a sudden current pulse emanating from the drain. The source of the charged particle can be natural alpha radiation from unstable isotopes or a neutron-neutral collision leading to the atom ionizing and traversing the substrates. These events are often called single-event-transient (SET), and usually last for only about 100 ps[43]. The effects on FPGA logic can be varying. A potentially dangerous effect is a latch-up, i.e. a permanent short circuit in a transistor. The uncontrolled currents running as a result can lead to severe overheating of the underlying silicon circuitry and cause damage. The manufacturers have in consequence built proprietary design tools that will place logic in a way that prevents these destructive latch-ups from occurring[43]. The second more common effect is a change of the value stored in a memory cell also known as bit-flip. These are generally caused by an SET, i.e. a current spike, coinciding with the rising edge

of the underlying logic memory cell, which will then store an incorrect value. In this case one speaks of a single-event-upset (SEU) as logic can behave erratically due to the change. Although a bit-flip is not dangerous to the FPGA fabric itself the resulting erratic behaviour can cause damage to secondary components, e.g. if the FPGA is part of a critical protection system.

The problem of SETs is commonly found in aerospace engineering, as satellites are continuously bombarded with high fluxes of radiation, as well as particle physics research[44]. In consequence manufacturers are offering radiation hardened FPGAs, which generally have a lower logic element density as these need to be made bigger. Other approaches check the currently operating logic against a secured copy and dynamically correct errors in the programming, if a SEU is detected[45]. Even though the conditions in current fusion experiments do not yield high neutron fluxes future devices will exhibit significantly higher reaction rates and therefore neutron fluxes. The issue must therefore be considered in the future. Even in today's machines the effect may be of importance, as will be discussed for the SAMI diagnostic in chapter 3.

1.5 Scope Of The Thesis

The main task of this thesis was the development of two fusion diagnostics using FPGA technology. Each of the diagnostics has different performance requirements and both make use of the capabilities discussed in section 1.4 to meet these.

Since each of the diagnostics are based on fundamentally different principles they will be explained separately, thus showing the flexibility inherent to the FPGA technology and its usefulness in the field of high-performance diagnostics. In chapter 2 the new digitization system developed for the MAST-Upgrade (MAST-U) electron density interferometer is presented. This system has strong latency requirements and requires real-time processing of data to have a reliable integration into the MAST-U plasma control systems. The underlying basic plasma wave physics diagnostic will be presented in section 2.2 and the hard- and firmware development detailed in sections 2.4, 2.5 and 2.6. Due to the MAST-U upgrade being currently underway the initial capabilities of the diagnostic are presented using synthetic data in section 2.7. The flexibility of the FPGA technology however enabled in-situ experiments on ASDEX Upgrade (AUG) to confirm its capabilities. Some of the results obtained at this experiment are presented in 2.8.

The second diagnostic addressed as part of this thesis is the Synthetic Aperture Microwave Imaging (SAMI) diagnostic, for which several modifications to the existing

hard- and firmware have been developed. It is described in chapter 3. Because the physics description used for the interferometer is too general, an expansion applicable to the diagnostic will first be derived in section 3.2. This will be followed by the necessary principles underlying the synthetic aperture approach in 3.3. After the theoretical basis has been established the diagnostic itself is explained in section 3.5, followed by a description of the modifications conducted to the FPGA firmware in section 3.6. The diagnostic has been moved to the NSTX-Upgrade (NSTX-U), which, being differently organized from MAST, yielded several further changes to the diagnostic and its behaviour. The necessary steps undertaken to conduct the integration are detailed in section 3.7 and an investigation into the effect of the changed geometry on the diagnostic behaviour is presented in sec. 3.8. The first results obtained by the SAMI group are presented in section 3.9 and an outlook to future developments is given in sec. 3.10.

The final chapter 4 will try to draw conclusions from the developments of these two diagnostics for the general field of high-performance and real-time diagnostics in fusion.

Chapter 2

FPGA Digitization For The Integral Electron Density

2.1 Introduction

As has been pointed out in 1.2.1 the plasma density is one of the fundamental plasma parameters in a fusion device. It directly determines the plasma performance and thus is continuously monitored and controlled. This is particularly important if operational limits leading to disruptions, such as the Greenwald limit, are considered, since they must be avoided. The most established technique to acquire this parameter is interferometry, which measures the line of sight (LOS)-integrated electron density (IED). This also acts as a reference for other diagnostics such as Thomson scattering (TS)[46]. The importance of the parameter requires, that IED is measured accurately and reliably with low latency.

The interferometer system used to measure the IED on the Mega-Amp Spherical Tokamak (MAST) was a traditional two-color setup using a CO₂-HeNe laser combination[47]. It was digitized using a commercial LeCroy JorWay model 1819 Phase Digitizer, which generally operated at 50 kSample/s (kS/s). The motivation for upgrading this system originated from the study of filaments using the MAST-Upgrade (MAST-U) fast TS diagnostic[46], [48]. Figure 2.1 shows a measurement of the MAST fast D_α-camera and IED during an edge localized mode (ELM) cycle. For size comparison : the MAST plasma has a major radius of 90 cm and a minor radius of 60 cm, i.e. the distance from the right edge of the picture to the visible plasma edge in the figure are about 120 cm. The IED cord is marked with a circle in the camera images. The red line in the middle three plots indicates a fitted magnetic field line. The density measurements in the bottom plot show a clear dip as the filament (bright lines in the camera images) is passing the measurement cord. Such

detail could only be obtained using a 1 MS/s digitization system, which however was not available continuously. The measurements show the advantage of having a high time resolution density measurement for this type of study and thus motivated the development of a continuous fast acquisition system. A point of note to figure 2.1: although the density dip is clearly visible, the increase in density above the background as the filament moves past (marked with an x) is only marginal. It was pointed out that the limited temporal resolution of the density measurement could be the source for this, although the integral nature of the measurement, i.e. the measurement cord passing through a hill and a trough at the same time, is a likely cause as well[46].

An additional motivation was the interest in studying Alfvén Eigenmodes (AEs) on MAST-U. AEs are weakly-damped magneto hydro-dynamic (MHD) modes resting in band gaps of the Alfvénic continuum[49]. The band gaps result from periodic perturbations of the mode’s refractive index along an energetic particle orbit around the torus. This can be caused by various different mechanisms, e.g. magnetic field ripple or density perturbations[49], [50]. The gap confines the mode to a particular flux surface and thus prevents the continuum damping. AEs can be driven unstable by energetic (Super-Alfvénic) particles, which occur in some cases in current machines using neutral beam injection (NBI), but will most definitely occur in a “burning” plasma with a fast α -particle population. They are of particular interest for high performance machines, since they can provide an exit mechanism for fast particles, e.g. fusion alpha particles, via magnetic ripple trapping, which in turn can cause fatal damage to the vacuum vessel[51]. At the same time there may be great benefits from AEs, since they may open up a method of efficient ion heating[52]. AEs come in various sub-classes and frequencies from tens of kilohertz (beta-induced AE) to several megahertz (compressional AE)[49], [52], [53]. Fast measurement of the IED has until recently not been commonly used to identify fast AEs. However, improvements of the diagnostic systems on experiments such as Doublet 3 D (DIII-D) have shown the benefit of fast density measurement for this field of studies[53]. AEs can also be used as a diagnostic tool via so called Alfvén spectroscopy. The gap structure and resulting limitations imposed on the evolution of so-called Alfvén cascades enables the deduction of information on the q-profile[54].

For MAST-U there is great interest in studying AEs. During the MAST campaigns a variety of different AEs could be identified using the fast OMAHA coil array[55]. Figure 2.2 depicts an example of a particularly high frequency AE during MAST shot #18886. The magnetic OMAHA coil spectrogram on the left shows a fast 3.8 MHz mode. The right plot compares the mode’s frequency (red line) to the




Figure 2.1 – Measurement of the IED N_e during an ELM cycle. For size comparison : the MAST plasma has a major radius of 90 cm and a minor radius of 60 cm. The circle in the fast camera images on top mark the location of the interferometer measurement. The red line in the middle three camera images is a fitted field line. The density in the bottom plot shows a clear oscillation as the result of a filament moving past the measurement cord. Note that although a dip in density is clearly visible no increase is registered, which was expected at the position of the 'x' (taken from [46]).

ion cyclotron frequency ω_{ci} (blue line). The frequency of the mode is above the ω_{ci} found on the low-field side (LFS), to which the mode has been localized using theoretical analysis taking the resonance condition into account. This suggest the compressional nature of the mode, i.e. a compressional AE (CAE). Currently only fast magnetic measurements are routinely used to investigate these modes, although fast reflectometry is occasionally employed, when available. An equally fast measurement of the electron density with megahertz bandwidth capabilities would be beneficial to study the mode evolution. The new digitization system was thus designed with a high bandwidth limit to study these high frequency modes.

In addition to the high demands on temporal resolution the system also has




Figure 2.2 – Measurement of a 3.8 MHz CAE during MAST shot #18886. The left depicts the spectrogram of an OMAHA coil magnetic measurement. The right shows the major radius against the ion cyclotron frequency for that shot with R_0 denoting the location of the magnetic axis. The red line marks the mode frequency seen on the left (adapted from [55] with kind permission from author and publisher).

to have the ability for real-time feedback control, i.e. for parameter control and machine safety. Generally the real-time control systems cannot respond much faster than $100 \mu\text{s}$ to several ms, e.g. due to the plasma inductance. The exact number is machine specific and depends on the operation scenario. The response systems in place on MAST had latencies of the order of a millisecond and it was estimated that the interferometer latency should be kept below $30 \mu\text{s}$. At the same time the upgrade of the experiment to MAST-U will enable the experiment to run 5 s long shots. In consequence the data acquisition and integration needed to be quasi-continuous.

The integration into the plasma control system (PCS) introduces an additional constraint. Changing parts in this system tends to require vast amounts of testing. While it was desirable to spend as little money as possible on the FPGA hardware it was also important to have good contingency in case of a hardware failure. The density measurement is so critical that the machine will not operate if the diagnostic is broken. In consequence high availability of both the diagnostic as well as the replacement parts are mandatory.

The new MAST-U IED interferometer system was designed with these considerations in mind. It consisted of a new optical setup, which was designed predominantly by Thomas O’Gorman as well as a field programmable gate array (FPGA)-based digitization back-end, which was designed as part of this thesis’ work[47], [56].

In this chapter the development of the new MAST-U digitization system will be described. First the basic theory underlying plasma interferometry is revised in section 2.2 and 2.3. Secondly in section 2.4, the optical setup including the digitization hardware is presented. In sections 2.5 and 2.6, a description of the

FPGAs' firmwares is given. This includes the interaction with the control PC and the design process, which determined a few design choices. This is followed by the results of the conditioning tests showing the diagnostic capabilities in section 2.7. Lastly the results from a short measurement campaign on the ASDEX Upgrade (AUG) experiment in Germany are presented in section 2.8, showcasing the FPGA back-end's capabilities on a working experiment.

Some of the work on the interferometer and its digitization system presented here has been published previously[47], [56].

2.2 Basic Waves In Fusion Plasmas

This section will give a brief overview of the physics involved in measuring the LOS-integrated electron density (IED) in a plasma, which involves deriving the cold plasma dispersion relation and relating it to the IED. This prepares the introduction of interferometry and the extension of two-color interferometry for vibration compensation, which will be introduced in the subsequent section.

2.2.1 The Cold Plasma Dispersion

A plasma is a quasi-neutral accumulation of free charged particles

$$n_e \approx Z_{\text{eff}} n_i,$$

where n_e and n_i are the electron and ion density respectively and Z_{eff} is the ion's effective charge number. Hence an electro-magnetic (em)-wave propagating through it is profoundly influenced by it. In general the electrons are the particles interacting with high frequency radio frequency (RF) waves, since the ions are too heavy to react to a fast changing em-field given their 1000 times higher mass.

To find a mathematical description for this process one starts off with Maxwell's equations:

$$\nabla \underline{E} = \frac{\rho}{\varepsilon_0} \tag{2.1a}$$

$$\nabla \underline{B} = 0 \tag{2.1b}$$

$$\nabla \times \underline{E} = -\dot{\underline{B}} \tag{2.1c}$$

$$\nabla \times \underline{B} = \mu_0 \underline{j} + \frac{1}{c^2} \dot{\underline{E}} \tag{2.1d}$$

Where the dot indicates the time derivative, E and B represent the electric and

magnetic field, \underline{j} denotes the current density, c the speed of light and μ_0 and ε_0 are the vacuum permeability and permittivity respectively. Taking the time derivative of Ampere's law (2.1d) and the curl of Faraday's law (2.1c) and setting the resulting equations equal to each other one arrives at the wave equation in its most basic form:

$$\nabla \times (\nabla \times \underline{E}) = -\mu_0 \dot{\underline{j}} + \frac{1}{c^2} \ddot{\underline{E}}. \quad (2.2)$$

By inserting the generalized Ohm's law $\underline{j} = \underline{\underline{\sigma}} \underline{E}$, where $\underline{\underline{\sigma}}$ is the electric conductivity tensor, we incorporate the plasma response into the equation. In addition we can linearize the equation assuming that $\underline{E} = \underline{E}_0 e^{i(kx + \omega t)}$ yielding the general dispersion relation for the propagation of a wave in a plasma:

$$(\underline{k} \otimes \underline{k} - k^2 \underline{I}) + \frac{\omega^2}{c_0^2} \underline{\underline{\epsilon}} \underline{E} = \underline{0}, \quad (2.3)$$

where we have used a vector identity to simplify the double cross-product. In 2.3 $\underline{\underline{\epsilon}}(\omega) = \underline{I} + i \frac{\underline{\underline{\sigma}}}{\omega \varepsilon_0}$ is the dielectric tensor, which physically describes the plasma response to the propagating wave and its effect on the wave propagation. It is a function of the propagating wave's frequency. The refractive index is often a more useful parameter so we further simplify equation 2.3 by using the definition for the refractive index ($N = \frac{\omega}{k}$):

$$(\underline{N} \otimes \underline{N} - N^2 \underline{I} - \underline{\underline{\epsilon}}) \underline{E} = \underline{0}, \quad (2.4)$$

To get a meaningful description we need to find a description for the dielectric tensor. As mentioned above the electrons are the most important particle in the plasma-wave interaction and we will therefore neglect the ion motion for now. To describe the plasma response we start with the electron's equation of motion in the em-field $m_e \underline{\dot{v}}_e = e_0 (\underline{E} + \underline{v}_e \times \underline{B})$. This single particle description is valid for cold plasmas, where kinetic effects, such as pressure gradients, can be neglected. As before we linearize the equation and without loss of generality assume that the magnetic field points in the direction of the z-axis. This leads to a set of coupled equations describing the electron motion perpendicular to the magnetic field, which can be solved to receive a matrix equation describing the electron motion:

$$\begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = \frac{ie_0}{\omega m_e} \begin{pmatrix} \frac{\omega^2}{\omega^2 - \omega_{ce}^2} & -i \frac{\omega \omega_{ce}}{\omega^2 - \omega_{ce}^2} & 0 \\ i \frac{\omega \omega_{ce}}{\omega^2 - \omega_{ce}^2} & \frac{\omega^2}{\omega^2 - \omega_{ce}^2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix}. \quad (2.5)$$

In this $\omega_{ce} = \frac{e_0 B_0}{m_e}$ is the electron cyclotron frequency. By comparison with the generalized Ohm's law we can get an expression for the electric conductivity

tensor, which we can insert back into eqn. 2.4. Because the dielectric tensor is critical in describing the plasma-wave interaction is useful to rewrite it using so-called Stix-parameters:

$$\underline{\underline{\epsilon}} = \begin{pmatrix} S & -iD & 0 \\ iD & S & 0 \\ 0 & 0 & P \end{pmatrix} \quad (2.6)$$

with

$$S = 1 - \frac{\omega_{pe}^2}{\omega^2 - \omega_{ce}^2}$$

$$D = -\frac{\omega_{pe}^2}{\omega^2 - \omega_{ce}^2} \frac{\omega_{ce}}{\omega}$$

$$P = 1 - \frac{\omega_{pe}^2}{\omega^2}$$

In the above equation set $\omega_{pe} = \frac{e_0^2 n_e}{\epsilon_0 m_e}$ is the plasma frequency. Because the magnetic field introduces an anisotropy it is important to differentiate between different directions of wave propagation. In consequence we will make the additional assumption that \underline{k} is of the form

$$\underline{k} = (k \sin \psi, 0, k \cos \psi),$$

i.e. without loss of generality the wave propagates either along the magnetic field or along the x-axis perpendicular to it.

Bringing everything together, we can now write the final version of the wave equation:

$$\begin{pmatrix} S - N^2 \cos^2 \psi & -iD & N^2 \cos \psi \sin \psi \\ iD & S - N^2 & 0 \\ N^2 \cos \psi \sin \psi & 0 & P - N^2 \sin^2 \psi \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix} = \underline{0}. \quad (2.7)$$

Equation 2.7 only has non-trivial solutions if its determinant is zero. This yields a quadratic equation of the form

$$AN^4 - BN^2 + C = 0, \quad (2.8)$$

with

$$A = S \sin^2 \psi + P \cos^2 \psi,$$

$$B = RL \sin^2 \psi + PS(1 + \cos^2 \psi),$$

$$C = PRL,$$

where $R = S+D$ and $L = S-D$. This can be solved for N^2 and is equivalent to

$$\tan^2 \psi = \frac{P(N^2 - R)(N^2 - L)}{(SN^2 - RL)(N^2 - P)}, \quad (2.9)$$

We thus have acquired a generalized form for an arbitrary wave's dispersion relation. One has to keep in mind that this formula is based on a cold plasma with stationary ions, which is valid for the high wave frequencies discussed here. In chapter 3 the description is expanded for lower frequency microwaves in magnetized plasmas.

2.2.2 Visible Light Approximation

Because visible or near infra-red light is very common in interferometry this case should be looked at in more detail. This spectrum corresponds to frequencies in the 400 THz range. The cyclotron frequency for current and future fusion devices is less than 200 GHz so it may be assumed that $\omega \gg \omega_{ce}$. This allows us to re-evaluate the Stix parameters as $S \approx D \approx P$. Given this approximation we can solve eqn. 2.8 for N^2 and get

$$N^2 = P = 1 - \frac{\omega_{pe}^2}{\omega^2} \quad (2.10)$$

showing us that in this approximation the plasma density is the only factor influencing the wave propagation. This is very convenient as it allows this parameter to be directly measured, for which the most established technique is interferometry, discussed in the next section.

As a point of note the high frequency approximation is particularly important in magnetized plasmas. Here the magnetic field induces birefringence between the ordinary and the extraordinary mode, which will be explained in more detail in chapter 3. The birefringence adds ellipticity to the wave, known as Cotton-Muton-Effect, which induces an error on the phase measurement of $\sigma = \omega_{ce}/\omega \cos \alpha$ [57]. Here α is the angle between the propagation wave vector and the magnetic field vector. However, while for normal phase measurement introduced here this poses an error, the effect can be used to measure the magnetic field along the line of sight, using polarization sensitive detection schemes[58]. This is known as polarimetry, which will not be discussed here.




Figure 2.3 – Basic schematic of a Michelson-Morley interferometer. The optical path length between the two arms differs according to the difference in propagation speed.

2.3 Plasma Interferometry

Interferometry is an established technique to measure the variation in optical path length between two measurement arms and has many different applications both in and outside of science.

Figure 2.3 shows a basic interferometric setup. A coherent light beam, usually sourced by a laser, is split into two different beams using a beam-splitter. One beam arm, usually referred to as the scene arm, passes through a dispersive medium with a refractive index N_1 , another, usually referred to as the reference, passes through a second arm without a perturbed refractive index. The optical path length between the two arms differs according to their spatial length difference as well as the difference in refractive index. The beams are recombined at a detector and the phase difference between the two is measured against some previously defined null-value using the interference pattern.

In case of a plasma interferometer the aim is to track the change of refractive index along the LOS of the probing laser. The optical path in a double pass interferometer as shown in fig 2.3 can be calculated as :

$$\Delta s = 2 \int_0^L [N_{\text{air}} - N_{\text{pl}}(l)] dl \approx 2 \int_0^L \left[1 - \sqrt{1 - \frac{\omega_{pe}^2}{\omega^2}} \right] dl. \quad (2.11)$$

Here N_{pl} denotes the plasma's refractive index. The visible light approximation from eqn. 2.10 is already employed. It can be translated to a phase difference between the two paths by $\Delta\phi = \frac{2\pi}{\lambda}$. The square root can be expanded in a Taylor series again assuming $\omega_{pe} \ll \omega$. Given this we can write an equation for the phase difference in a double pass plasma interferometer:

$$\Delta\phi_{\text{pl}} \approx \frac{\lambda e_0^2}{2\pi c_0^2 \varepsilon_0 m_e} \int_0^L n_{\text{pl}}(l) dl \quad (2.12)$$

$$\int_0^L n_{\text{pl}}(l) dl = \frac{2\pi c_0^2 \varepsilon_0 m_e}{e_0^2 \lambda} \Delta\phi_{\text{pl}} = \frac{\lambda n_c}{2\pi} \Delta\phi_{\text{pl}} \quad (2.13)$$

In here the plasma density is denoted as n_{pl} , which for ideal hydrogen plasmas is equivalent to the electron density n_{pe} used before. Equation 2.13 was simplified using the cut-off density $n_c = \frac{\lambda^2 \epsilon_0^2}{4\pi^2 c_0^2 \epsilon_0 m_e}$, i.e. the plasma density at which the refractive index for a wave of a given wavelength approaches zero thus preventing it from propagating. Using 2.13 the IED can be calculated. The laser wavelength can often be assumed constant or very slowly varying so that slow control measurements are often sufficient. The linear dependence on λ also indicates that higher wavelengths are beneficial for the measurements, since the phase shift per length-integral is bigger thus making the measurement more sensitive to small density changes. However, a balance needs to be struck, since a higher wavelength places the beam closer to its cut-off, which means the visible light approximation breaks down. In addition the beam starts diverging significantly from a straight path in plasma regions with large density gradients due to the accompanying gradients in the refractive index. In a curved plasma these gradients are unavoidable and particularly disturbing during fast plasma events such as disruptions. There is also the additional problem of 2π wraps, which will be addressed later in section 2.7.2.

2.3.1 Two-Color Interferometry

As has been initially noted the optical path difference in an interferometer can be caused by a difference in refractive index as well as a difference in spatial dimension. A complex experiment such as MAST-U with long optical paths is not a mechanically rigid system. In consequence the setup and therefore measurement will be significantly affected by vibrations. The measured phase difference will be a superposition of the plasma dispersion and the optical components vibrating:

$$\Delta\phi = \Delta\phi_{\text{pl}} + \Delta\phi_{\text{vib}} = \frac{\epsilon_0^2 \lambda}{2\pi c_0^2 \epsilon_0 m_e} \int_0^L n_{\text{pl}}(l) dl + \frac{2\pi}{\lambda} \Delta l \quad (2.14)$$

The phase change due to the mechanical displacement of all optical components will add up linearly, so that it is possible to describe it as a single component being displaced by a length Δl . One can see, that unlike the phase shift due to the plasma density the vibrations have an inversely proportional dependence on the wavelength. This allows us to devise a compensation scheme to extract only the density dependence.

First we solve eqn. 2.14 for Δl . If one now takes the same measurement with two different wavelength, but identical optical paths, we can set the two measurements equal to each other and solve for $\int n_e dl$, yielding




Figure 2.4 – Example of two-color vibration interferometry data at the start of an AUG shot. The initial 30 ms no plasma is present and vibrations are the only source for the phase oscillations. After 40 ms the plasma density introduces a measurable phase offset from the null (black dotted line) as a result of the two individual phase-traces diverging. The residual oscillations in the compensated phase are the result of the non-co-linear beam paths. The compensated phase has been offset from 0 for better visualization (dotted line marks the null level taken as the average of the first 20 ms).

$$\int n_e dl = \frac{2r_{ce}}{\lambda_1} \frac{\Delta\phi_1 - \frac{\lambda_2}{\lambda_1} \Delta\phi_2}{\lambda_2^2 - \lambda_1^2}. \quad (2.15)$$

In this equation $r_{ce} = \frac{e_o^2}{2\pi c^2 \epsilon_0 m_e}$ is the classical electron radius. The equation allows a measurement of the IED, without the unwanted contribution from mechanical vibrations by using co-linearly propagating laser beams of different colors, which are measured simultaneously.

Figure 2.4 shows an example of two color vibration compensation based on real measurements at the AUG tokamak. The HeNe phase has been scaled by $\frac{\lambda_1}{\lambda_2}$ for visualization. The compensated phase is calculated using eqn. 2.15 without the constant density scaling factor. As can be seen the phases of both colors show significant oscillations, which however are roughly identical and thus the compensated phase stays flat. By 40 ms the plasma density has risen enough to cause a significant contribution and the traces start diverging. The result is the compensated phase rising above the level found at the start of the trace. In fig. 2.4 the phase was offset from 0 for better visualization. A point of note: the real phase is measured against an arbitrary non-zero reference phase taken at some point in time and traced from then on. The null value (in fusion experiments often an average over the first milliseconds without a plasma) is subtracted after the measurement. In fig. 2.4 the

first 20 ms were averaged to define the null phase (dotted black line).

Two-color electron density interferometers have been used on other machines such as DIII-D[59], AUG[60] and the JET[61]. The combination of colors is often a matter of balancing the sensitivity towards plasma density changes and the sensitivity towards vibrations and beam refraction. The optical setups are generally simple Michelson-Morley or Mach-Zehnder setups, which combine and split the light of two lasers using beam splitters and filters.

The analytically perfect vibration compensation in eqn. 2.15 is based on the optical path of the two colors being identical. However in a two-color setup with two laser sources there are parts of the lasers' paths, that are not co-linear, e.g. the sections to the beam combiner and the section after the color separator. The vibrations in these parts of the optical setup are not compensated, which is also the cause of the residual vibrations seen in fig. 2.4. This is a well known problem and can only be reduced by careful design of the optical layout.

For completeness it shall be mentioned here that the problem of the traditional two-color vibration compensation can be mitigated by building a more sophisticated setup, which nonetheless is based on the same principle. This is called dispersion interferometry, which only relies on a single laser source[62]. The laser is passed through a frequency doubling crystal, generating the first harmonic of the laser wavelength. The superimposed beams are then passed through the plasma and by means of an acousto-optical modulator (AOM) and another frequency doubling crystal superimposed onto each other. The benefit of this scheme is that both colors have passed through identical optical components and thus see the identical vibration phase shift. This enables near perfect vibration compensation. A down-side is the convoluted measurement signal, which is very difficult to evaluate. Systems like this have been implemented on several machines such as Wendelstein 7-X (W7-X)[63] and the Large Helical Device (LHD)[64].

2.4 MAST-Upgrade Diagnostic Setup

The IED interferometer for MAST-U, just like its predecessor, is designed as a conventional two-color setup to reuse some of the already existing hardware. The development of a new digitization back-end was intended to ensure that the diagnostic meets the previously stated criteria :

- High digitization accuracy, i.e. ensure an optically limited diagnostic
- High digitization bandwidth




Figure 2.5 – The optical path of the new IED interferometer on MAST-U. The optical bench sits in an isolated room and is attached to a wall. The retro-reflector is a roughly 80 kg granite sphere, which is isolated from the plasma vessel. The total optical beam path is roughly 30 m (adapted from [47] with kind permission from the author and AIP publishing).

- Continuous acquisition
- Low latency integration into the PCS
- Low cost, with readily available replaceable parts

Meeting these criteria led to a set of design choices, which are described here. Although not part of this work the description of the optical setup is important for the design of the digitization system and will be described first before describing the FPGA system.

2.4.1 Optical Setup

The interferometer for MAST-U is set up as a Michelson-Morley interferometer. This is a necessity, since the scene arm is passing through the center of the vessel close to the central column and exits the vessel almost on the opposite side from the entrance port as shown in figure 2.5. Apart from a Mach-Zehnder setup being prevented by the geometry of the experiment, using a Michelson-Morley setup has additional benefits, such as the compensation of density gradient induced beam divergence and a larger phase shift induced by the plasma density leading to an increased signal-to-noise ratio (SNR).

Compared to the system operated on MAST the optical layout of the interferometer was changed as part of the upgrade. To reduce the effect of vibrations the system is fully decoupled from the MAST-U vessel. A separate isolated room has been built




Figure 2.6 – A schematic of the diagnostic. The optical system is a superposition of a CO₂ and a HeNe interferometer in Michelson-Morley setup. The digitization system is based on a tandem of two FPGAs, inter-linked by a fiber-optic cable.

in the North Annex to reduce the influence of disturbing effects such as thermal convection in the beam path. In addition the reflector needed for the laser path is separate from the experiment and made from an 80 kg granite sphere suspended by air vibration-isolation mounts (see inlay in the figure)[47]. The overall beam path as shown in fig. 2.5 is about 30 m.

Figure 2.6 depicts a schematic of the optical setup and it's relation to the FPGA digitizer. The system is essentially made of two independent interferometers. The major one most sensitive to the density is based on a water cooled 8W Merit SL CO₂ laser from Access Lasers at 10.6 μm. The beam is split and frequency modulated using an AOM operating at 40 MHz. This modulation is necessary for the heterodyne detection of the phase, since it is otherwise impossible to discern the direction of the phase change from the (stationary) interference pattern. The reference arm is then folded on the optical bench to match the 30 m scene arm length and recombined with the scene arm on a VIGO PVM-10.6 photo-voltaic detector. The detector then passes through a 35 MHz to 45 MHz band-pass filter and is digitized using the FPGA digitizer. The reference frequency from the AOM is directly digitized by the FPGA on a separate input channel.

The second color for the interferometer is based on a Agilent 5517B HeNe interferometer operating at around 633 nm. The laser cavity in this commercial interferometer produces laser light with two orthogonal polarizations, which are

frequency offset by about 2.236 MHz. The polarization components are separated using a polarizer and then recombined using a half-waveplate and the same polarizer as depicted in 2.6. A reference path is not required for the HeNe due to its long coherence length. If two wave packages from the same source differ in propagation path by more than this distance a stable phase relationship is no longer present between them and no well-defined interference pattern is formed, which would make the density measurement impossible. The coherence length of a HeNe laser however generally exceeds 100 m and thus no folded reference arm is required for the HeNe interferometer as the maximum path difference is 30 m. The recombined light is digitized using a Thorlabs PDA8A silicon amplified photo detector, which is directly digitized using no filter. This choice is necessitated by the stability of the Agilent interferometer modulation frequency, which is explained in more detail in section 2.5.5.

The modulation frequency reference is supplied by the 5517B interferometer as an analog square wave, which is directly digitized by the FPGAs as well. The fact that the reference is supplied as a square significantly impacts the FPGA design, which will be discussed further in the following sections.

As has been discussed previously the vibration compensation scheme using two colors only works well given that the beam paths are fully identical. As can be seen from figures 2.5 and 2.6 there are actually significant parts of the optical setup, which have no common path between the two colors. Most notably the CO₂ reference arm and the HeNe reference path are completely separate system, so any vibration on these optical paths will be uncompensated. For this reason all uncompensated parts of the optical setup are placed on an optical bench, so the contribution of vibrations on the measurements will be dominated by the retro-reflector RR3 and the 90° reflector, which are not on the optical bench (see fig. 2.5). Nonetheless this scheme will reduce the accuracy of the final system.

2.4.2 Digitization Hardware

The analog signals described in the previous section are digitized using a FPGA based digitization back-end. This is based on the simple PCIe carrier (SPEC) board designed by the Open HardWare Repository (OHWR), which is maintained by CERN[65]. The SPEC board carries a Xilinx Spartan-6 XC6SLX45T-3FGG484C FPGA and has a FPGA mezzanine card (FMC) low-pin-count (LPC) connector for peripheral expansion. This expansion slot carries the FmcAdc100M14b4cha FMC expansion board, which is a 4 channel 100 MS/s 14 bit analog-to-digital converter (ADC) board designed by the OHWR as well[66]. The ADC also has an additional

input line to receive a trigger from the MAST-U systems. All hardware specifications and circuit diagrams by the OHWR are open source, which means they are available for use by anyone. Thus in case of a failure the part can always be replaced, even if no seller for the board can be found. This meets the proliferation requirement given earlier. In addition customization of the hardware is more easily implemented. In this case the ADC was slightly modified to increase the input bandwidth from 40 MHz to 45 MHz by removing a notching coil used as anti-aliasing filter.

The SPEC board's main responsibility is to act as a digitizer, pre-processor and data streamer. No data is stored on the FPGA itself. Instead the data is reduced enough to be able to stream it in real-time using a fiber-optic link, which has been implemented using the on-board small form-factor pluggable (SFP) module. Because all of this is implemented digitally very few analog components are required, significantly reducing the cost of the system. Communication with the board for control, e.g. arm/disarm/status, is also conducted using the optical link. The details of the firmware are presented in section 2.5.

As can be seen in figure 2.6 the SPEC board does not connect directly to a PC. Instead it connects to a second FPGA. In this case a Xilinx SP605 carrier board is used, which holds the same chip as the SPEC board[67]. It has two responsibilities. Firstly it takes the data received by the fiber-optic link and sends them via PCI Express (PCIe) to the PC memory. Secondly it forks the data from the optical link, processes it and then sends out a control signal via a fiber to the MAST-U control systems. The details of this firmware are presented in section 2.6.

2.5 SPEC Firmware

In this section the firmware for the SPEC board as the primary digitizer and pre-processor is described. The firmware of the SP605 based Aurora-PCIe bridge is detailed in section 2.6, at which point the interaction with the diagnostic PC is also explained. Unless specified differently, all firmware design has been conducted using Xilinx ISE v14.7 and System Generator for MATLAB Simulink. Most of the top level logic is written in VHDL, however certain standard components, such as first-in-first-outs (FIFOs) buffers and digital clock managers (DCMs) were generated using Xilinx Coregen[68]. This includes the employed filters. In addition components from the OHWR General-Cores library were used for certain sub-components[69].




Figure 2.7 – Top level of the diagnostic firmware. The top level design with the communication and the deserialization core is written in VHDL. The processing core is designed using System Generator for MATLAB Simulink. Colors denote different clock regions as indicated in the figure. Orange components transition between clock domains.

2.5.1 Top Level Layout

Figure 2.7 depicts the top level schematic of the SPEC board’s firmware. The deserialization core has been modified from an early version of the SPEC reference design. The Whishbone bus interface logic, which is of no use in this design, has been removed. In addition the core was enhanced with a finite state machine (FSM) to control the ADC’s state transitions, i.e. arming, wait-for-trigger and disarming, and logic to handle triggers correctly. The top level module in particular handles the communication with the outside world via a fiber-optic link.

This link is implemented using the open Xilinx Aurora 8b/10b protocol[70]. The link operates at 2.0 Gbit/s in duplex mode transferring 32 bit packages. It streams data with low latency to the Aurora-PCIe bridge described in section 2.6. The data received via this link is handled by a communication core. Its purpose is to respond to request and control signals received via the fiber-optic link and set all relevant FSMs to the appropriate states. It also handles sending of confirmation signals back via the link.

Another very crucial role is the selection, tagging and multiplexing of the processed data to send it via the optic link as well as ensuring proper clock domain crossing. A clock domain is a group of synchronous logic operating at the same frequency. In particular if these clocks originate from a different physical source they have no well-defined phase relationship. The different clock domains in the top-level module are indicated by different colors in fig. 2.7. Hence to transition data from one clock

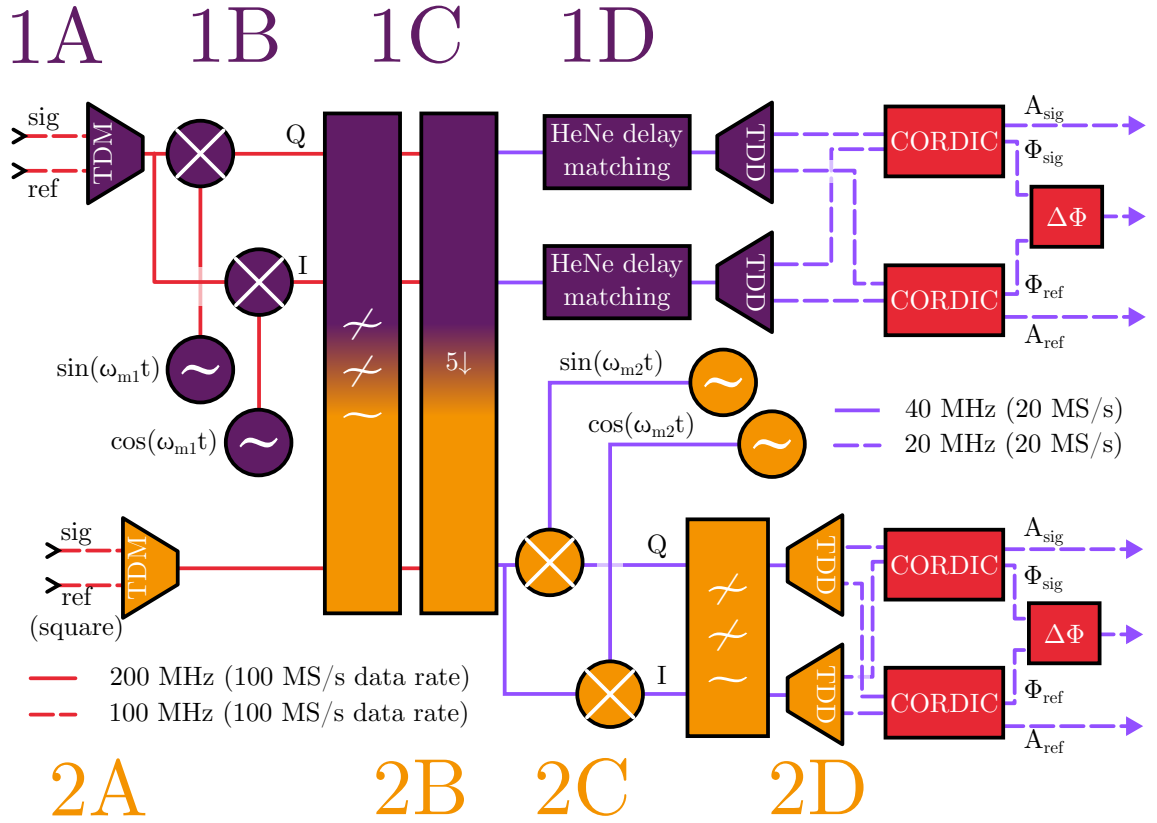


Figure 2.8 – A schematic of the data processing core using mixing I-Q generation on the HeNe branch. Signal TDM is required to maximize the use of DSP48 multipliers. The mixing frequency ω_{m1} for the CO₂ signal is 36.36 MHz, the frequency ω_{m2} for the HeNe branch at 1.56 MHz. The delays between the different processing paths are matched by passing valid signals.

region to another great care needs to be taken to prevent data loss or corruption. The communication core transitions the data coming from the data digitization and processing core (6 channels at 20 MS/s at the output) to an intermediate state machine for channel selection, multiplexing/down-sampling (operated at 150 MHz) and the handing over to the Aurora streaming core (operating at 50 MHz). Each of these clock domains is sourced on a different physical oscillator and therefore have no phase relationship. The transitions are done using asynchronous FIFOs (indicated in orange in fig. 2.7). The FPGA chip employed is only capable of moderate clock frequencies and is limited in its logic routing speed. For this reason the FIFOs are mostly implemented using low-depth distributed logic, which is beneficial for finding a possible layout of the logic on the FPGA fabric that meets the logic’s timing constraints.

The heart of the design is the processing core in the top of fig. 2.7. The core evaluates the digitized signals for both lasers ensuring equal delay. To prevent glitches from appearing when transitioning between the ADC and the FPGA the

entire processing core is clocked from the ADC digitization oscillator. The core itself traverses several slower clock domains which are driven by clocks derived off of the ADC's digitization clock. Only after the processing is complete does the signal get transferred to a FPGA-oscillator based clock domain.

Figure 2.8 shows a schematic of the initial down-conversion processing core implemented on the SPEC board. It acts on both laser signals in the same way to calculate the phase. In the following sections each subcomponent will be described in detail.

2.5.2 CO₂ Processing

The processing path for the CO₂ branch is marked purple in figure 2.8. Since the modulation frequency of the CO₂ interferometer is 40 MHz the design directly heterodyne-down-converts the CO₂ signal using a 36.36 MHz sine and cosine generated on the FPGA using a read-only memory (ROM) element programmed into the logic, which is read out with a wrapping address counter. This will yield the reference and scene signal's in-phase (I) and quadrature (Q) component, where the quadrature component has all spectral components shifted by a quarter of the wavelength with respect to the I-component:

$$\begin{aligned} \sin \omega_{\text{mod}} t \cdot \sin \omega_{\text{mix}} t &= \sin [(\omega_{\text{mod}} - \omega_{\text{mix}})t] + \cancel{\sin [(\omega_{\text{mod}} + \omega_{\text{mix}})t]} \\ &= \sin \omega_{\text{IF}} t \end{aligned} \quad (2.16)$$

$$\begin{aligned} \sin \omega_{\text{mod}} t \cdot \cos \omega_{\text{mix}} t &= \sin \omega_{\text{mod}} t \cdot \sin \left(\omega_{\text{mix}} t + \frac{\pi}{2} \right) \\ &= \sin \left[(\omega_{\text{mod}} - \omega_{\text{mix}})t - \frac{\pi}{2} \right] + \cancel{\sin \left[(\omega_{\text{mod}} + \omega_{\text{mix}})t + \frac{\pi}{2} \right]} \\ &= \cos \omega_{\text{IF}} t \end{aligned} \quad (2.17)$$

Equation 2.16 describes the generation of the I component and eqn. 2.17 that of the Q component. The two together locate the signal within the 4 quadrants of the unit circle enabling a phase sensitive measurement. In both equations the crossed out contributions are high frequency components that are filtered out. Figure 2.9 shows a schematic not to scale representation of the generation in (real) frequency space. It must be noted that ω_{mod} can be spectrally broad. The resulting low intermediate frequency (IF) for both signals is at 3.64 MHz. This frequency directly determines the available measurement bandwidth given by the distance to direct current (DC) and the final low-pass filter's cut-off frequency. No digital anti-aliasing

filter is applied before the initial down-conversion, since the analog band-pass filter before the digitizer acts in the same way and pick-up in the short distance from the filter to the digitizer can be assumed negligible.

The down-converted signal is then 5-fold decimated, i.e. filtered and down-sampled by a factor of 5, using a 115 tap 9 MHz low-pass decimation filter, which also removes the high frequency component from the down-mixing process. This reduces the Nyquist frequency to 10 MHz and is necessary to reduce the amount of data to a level that is transmittable via the optical link. The Spartan-6 chip on the SPEC board has a total of 58 fast multiplication elements (DSP48). Since filtering requires many multipliers and scene and reference require identical processing, time-division multiplexing (TDM) has been employed to double the number of available multipliers for processing. Because all information is in the phase of the signal the filters have been implemented as finite-impulse response (FIR) filters, which ensures a linear phase delay across the signal spectrum. Any decimation filters are implemented as multi-phase filters, i.e. they use the same multiplier element for multiple samples thus significantly reducing the number of DSP48s required.

The resource and timing constraints on the FPGA determine the number of possible multipliers available for filtering, thus setting the 115 coefficients available for the CO₂-branch's filter. It was implemented using a Butterworth coefficient structure, which has a maximally flat pass-band amplitude response. The coefficients were calculated using the MATLAB FDA tool. This reduces the effect of attenuation ripple in the pass-band on the measurements[71]. The final filter has a stop band suppression of 60 dB. After the filter the scene and reference signal are passed to a time-division de-multiplexing (TDD) element before being passed on to the phase measurement core.




Figure 2.9 – A schematic (not to scale) representation of the signals' spectral composition for different points in the CO₂ processing pipeline as depicted in fig. 2.8. The underlying mathematical operation is described in equations 2.16 and 2.17.

2.5.3 HeNe Processing

The HeNe signals have to be handled different from the CO₂ branch. The biggest challenge is introduced by the square reference signal, which is supplied by the Agilent interferometer and cannot be changed. The phase measurement core relies on sinusoidal signals for its calculation so that the square wave cannot be used directly. In addition the IF of the HeNe interferometer only sits at about 2.236 MHz. To have the best vibration compensation with the CO₂ signal this frequency should be kept as high as possible - ideally the same as the CO₂ IF. During the course of system development several proposals to design this stage have been made. The initial implementation approached the square wave by means of a digital low pass filter basically forcing the square wave into a sinusoid. This scheme is depicted in fig. 2.8 in orange and a schematic of the signal processing procedure in real frequency space shown in fig. 2.10. Scene and reference are multiplexed and then first passed through the same decimation filter the CO₂ signal passes through after down-conversion. This is possible because of the already low IF. Next, the signal is down-converted with a digital 1.56 MHz wave similar to the CO₂ branch, which results in a roughly 800 kHz IF. Subsequently the signal is passed through a 60 tap 1 MHz low-pass filter. This filter has a stop band suppression of 70 dB and will filter out the higher harmonics of the square wave, the lowest of which sits at 3.16 MHz. This way the down-converted signal is sinusoidal after mixing and ready for the phase measurement core. Although the HeNe final IF sits at only 800 kHz the HeNe signal is *not* decimated a second time. This way the sampling rate is the same as in the CO₂ branch and no interpolation is required to compensate the vibrations - or in other words, the filter has already conducted all interpolation necessary.

After this processing the CO₂ and HeNe signals have experienced different delays with the HeNe having been delayed more due to traversing 2 filter stages. In consequence a matching logic is inserted after the decimation stage adaptively delaying the CO₂ signal until it has the same delay as the HeNe signals. This is controlled via two valid flags, passed in parallel through both processing stages experiencing identical delay and decimation to the respective signals. In principle the delay is fixed after the firmware development has been finished, however implementing an adaptive matching logic allows the design to be transferred to different FPGAs and designs with more complex filters in the future without having to rematch the circuits. After the delay matching the signals can be demultiplexed and passed to the phase calculation stage.

There are several down-sides to this way of processing the data. Number one is the additional significant filtering required to create a sinusoidal reference signal at




Figure 2.10 – A schematic (not to scale) representation of the signals’ spectral composition for different points in the HeNe processing pipeline as depicted in fig. 2.8. The underlying mathematical operation is described in equations 2.16 and 2.17. Note the reference signal, which has higher harmonics of the fundamental wave due to its square shape. After down-mixing the lowest higher harmonic sits around 3.16 MHz.

the cost of stop-band suppression in the initial decimation filter. This increases the overall noise. Secondly the generation of the I and Q component by down-mixing significantly reduces the IF, which in turn reduces the available bandwidth on the HeNe branch. Nonetheless the down-conversion processing core has been implemented successfully, since it is the lowest-cost solution[56]. The lower band-widths on the HeNe branch can be accepted under the assumption that vibrations will not be found in this frequency range and thus no mechanical vibrations go uncompensated. Since the main sensitivity with respect to plasma density is found on the CO₂ branch, the HeNe branch contributes little towards the ability to detect high frequency modes in the first place. However, if they are detectable on the CO₂-arm there will be a 10% error on the high-frequency density fluctuation amplitude δn_e , which will have to be accepted. Since most of the high-frequency analysis is concerned with the mode evolution in frequency space, this is likely an acceptable trade-off. The results of the down-conversion design’s conditioning tests will be presented in section 2.7.

2.5.4 CORDIC Phase Difference

The phase calculation is conducted using the COordinate Rotation DIgital Computer (CORDIC) algorithm. This method is very resource efficient and well established in digital signal processing (DSP) applications[72]. The cordic algorithm calculates the angle of a signal by rotating the 2D vector (x,y) over successively smaller angles. For efficiency the angle steps are chosen to be $\arctan(2^{-i})$ in digital circuits, where i indicates the i^{th} rotation. Each iteration of the next vector is thus calculated by

$$\begin{aligned}
x_{i+1} &= x_i - \alpha_i y_i 2^{-i}, \\
y_{i+1} &= y_i + \alpha_i x_i 2^{-i}, \\
\theta_{i+1} &= \theta_i + \alpha_i \arctan(2^{-i}).
\end{aligned} \tag{2.18}$$

Here θ_i is the angle against the x-axis, whilst $\alpha_i = \pm 1$ indicates the direction of the next rotation[73]. To calculate the arcus-tangent of the input vector (x,y) the rotation can be conducted such that the y component is 0 and the vector is directed parallel to the x-axis. This enables the calculation of both amplitude and phase via

$$\begin{aligned}
x_N &\approx S_N \sqrt{x_0^2 + y_0^2}, \\
\theta_N &\approx \arctan\left(\frac{y_0}{x_0}\right),
\end{aligned}$$

with

$$S_N = \prod_{i=0}^{N-1} \frac{1}{\arccos(\arctan(2^{-i}))}.$$

Since the arcus-tangent is only defined on the interval $(-\frac{\pi}{2}, \frac{\pi}{2})$ the signal is first coarsely rotated into the first quadrant before the actual iteration. After the iteration has been completed the signal is rotated back by the same amount, thus delivering the value anywhere on the unit-circle. The S_N is the CORDIC scale factor required to keep the vector at equal length after each iteration[73]. The operations of eqn. 2.18 can be implemented very efficiently using bit-shift and add operations and thus require very little resources. The measurement accuracy increases with each step and thus can be taken all the way to the single bit accuracy. Using this approach the accuracy of the end-result is mainly determined by the internal accuracy of the operations[73].

The phase calculation is identical for both signal arm, as shown in red in fig. 2.8. Four identical Xilinx CORDIC cores are used to calculate the phase of each signal. Following the phase calculation the difference between respective scene and reference arm signals are calculated in a two-step process. First the scene and reference arm's phase wraps are synchronized making use of the fact that the CORDIC phase will progress in one direction around the unit circle. This improves the wrap detection of the subsequent unwrapping algorithm. The two-step approach increases the robustness of the system towards 2π phase wraps. This matter will be discussed in more detail in section 2.7.2. After the phase calculation the data is passed on to the communication core.

An additional feature of the processing core is that it passes a single trigger line




Figure 2.11 – The HeNe processing stage using Hilbert transformation filtering. The difference to fig. 2.8 is marked in orange. Because no additional low-pass filtering is available in this scheme, the digitized reference signal must be sinusoidal already. The initial decimation filter is identical to the one in the CO₂ branch.

signal though the core with identical delay to the laser signals. This signal can then be used by the communication core to insert a data package into the data stream allowing the design to synchronize the data to the MAST-U data acquisition clock via a set of well known triggers.

2.5.5 HeNe Hilbert Processing

A HeNe processing scheme using a Hilbert transformation filter can be used as an alternative to the down-conversion design. A Hilbert filter performs a phase transformation across the signal spectrum, thus not changing the spectral composition of the signal. In this case we are interested in generating the quadrature component and consequently will employ a Hilbert filter that shifts the signal by 90°. The IF, which is being reduced in the down-conversion design, is not changed in this scheme. Such a filter can be efficiently constructed using FIR filters and will only introduce an amplitude mismatch[74].

The overall processing scheme is almost identical to what is depicted in figure 2.8. The difference is the replacement of the down-mixing stage with the bespoke Hilbert transformation filter, as shown in orange in figure 2.11. The output of the filter will be the Q component of the required signal. Hence a Z-transform, i.e. a time delay equal to the filter’s delay, is necessary to propagate the I component past the filter. The signals’ spectral composition is in essence as depicted in figure 2.10 on page 39 for points 2A and 2B. Since the Hilbert filter will attenuate the signal passing it, the I component is multiplied by a constant to match the amplitude of the filtered signal. This constant had to be measured on chip by supplying a 2.236 MHz clean synthetic reference sinusoid from a signal generator to the SPEC and plotting the Q against the unmatched I component after the Hilbert filter using Xilinx ChipScope. The measurement is shown in figure 2.12 in the left plot. The ratio of the semi-axes

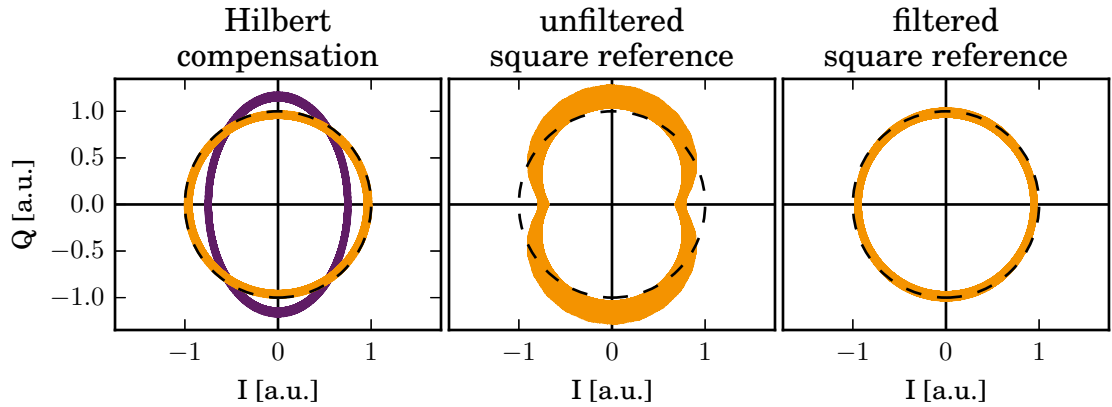


Figure 2.12 – Hilbert filter attenuation compensation plotting the I against the Q component after the Hilbert transformation filter. The SPEC is fed with a 2.36 MHz signal from a signal generator. The left plot shows the Hilbert compensation factor measurement using a sinusoidal reference signal. Purple marks the uncompensated signal, orange the compensated one. An ideal I-Q signal should be circular in this plot (as indicate by the black dashed line). The middle plot shows the compensated Hilbert transformed and attenuation compensated reference signal using a square wave. The right plot shows the same signal with an analog filter added before the ADC.

of the uncompensated elliptical signal (drawn in purple) is the attenuation factor at the nominal HeNe reference frequency. A perfect compensation should result in the compensated signal (drawn in orange) to be perfectly circular (as indicated by the black dashed line). As can be seen the signal is adequately attenuation compensated for sinusoidal reference frequencies. One has to note that the compensation will be less accurate as the scene and/or reference frequency deviate from 2.236 MHz due to the attenuation ripple in the frequency domain that is inherent to the Hilbert filter.

As has been mentioned before the reference signal for the HeNe is a square wave. The middle plot in fig. 2.12 shows the previously gained compensation factor applied to a square wave on the reference input. As can be seen the compensated signal is significantly distorted. This is caused by the higher harmonics of the square wave, which are not filtered out in this scheme. They will propagate to the phase measurement stage and cause oscillations at the harmonics making the final signal harder to interpret and the compensation worse.

The problem of a square reference can be tackled in several ways, the most obvious of which is the use of an analog low-pass filter with a steep transition band at 2.5 MHz. This suppresses every harmonic of the fundamental wave. The rightmost plot in figure 2.12 shows the I-Q plot for a square wave that has been passed through a Mini-Circuits BLP-2.5+ low pass filter before entering the ADC. As can be seen the signal is still not perfectly circular, which again is due to the




Figure 2.13 – The laser’s temperature and resulting stability of the HeNe reference frequency over time. The measurement was conducted over a day in an temperature uncontrolled room, which cooled down over night. A strong temperature dependence of the reference frequency will require temperature control.

residual higher harmonics. In addition the analog filter introduces an additional problem: analog low-pass filters have a non-linear group delay across the spectrum, which is also true for the ones employed in this system[75]. The group delay is around the reference frequency is around 450 ns and particularly non-linear, varying between 350 ns to 500 ns[75]. It will therefore introduce an unknown distortion in the phase measurement for fast changes in the reference frequency. This is particularly true if a high acquisition bandwidth is desired. Since this delay cannot be accurately compensated in logic it will lead to an unknown error on the phase measurement depending on the stability of the reference frequency.

To estimate the possible error the laser’s frequency stability was measured. For this the laser was operated on a desk and it’s square reference frequency put into the scene arm input of the interferometer down-conversion design. The reference was fed with a 2.236 MHz sinusoid from an Agilent 33600A Series Waveform Generator with a frequency resolution of 1 μ Hz[76]. The slope of the phase measurement from the digitization is thus proportional to the difference in frequency between the two signals. In the process both the ambient room temperature and the laser case temperature were measured for a long term stability dependence check. Figure 2.13 shows the long term stability with respect to the laser and ambient temperature. As can be seen there is a strong frequency dependence even after the laser has been operating for almost 24 hours and should be thermalized. The variations in room temperature are the result of the day and night cycle. The results show that temperature control will be required in the laser room.

Figure 2.14 shows the results of the short term frequency stability check. The




Figure 2.14 – The short term stability of the HeNe reference frequency. The in-shot frequency stability is shown on the left and a histogram of the frequency distribution on the right. It can be seen that the frequency is on average stable although it shows large short spikes resulting in a broadening of the distribution.

long term stability (from the evaluation perspective) is seen on the left and shows the frequency to be relatively stable, although the spectrum appears to be very broad. A histogram of the measured frequencies is seen on the right of the figure showing long “tails” on both ends of the distribution function. As can be seen in figure 2.15, which depicts a zoom on one of the spikes, the reference signal shows significant variations on very short time scales, which appear to be caused by a beating frequency of the order of 50 kHz. This is a typical frequency for the noise ripple of HeNe power supplies. On random occasions massive spikes in the reference frequency of the order of 100 kHz on very short timescales appear, which seem to be caused by a quasi-random constructive superposition of the beating oscillation’s




Figure 2.15 – An example for a frequency spike seen in figure 2.14. It can be seen that the frequency is highly variable on very short timescales.

spectral components. Similar features are well known in phenomena such as freak waves[77]. As has been explained before, this can lead to significant errors on the phase measurement because of the high bandwidth this system is designed to achieve.

The simplest solution to the reference signal issue would be the use of an optically generated reference by splitting a portion of the HeNe signal via a beam-splitter and not passing it through the plasma. The thus generated signal would be sinusoid and allow a phase calculation without the need to consider the filter dispersion. However, as the economic situation on the MAST-U experiment is one of the driving factors for attempting a low-cost solution another digital approach was attempted first.

2.5.6 PLL Generated Reference

Instead of using a filter a phase-locked loop (PLL) on the FPGA can be implemented, which locks on to the reference signal and generates a digital ideal sine. The digital sinusoid produced by this would mimic the reference frequency changes and because the delay of the system is well defined it could be compensated on-board. A down-fall of digital PLLs is that they can never perfectly mimic a signal due to bit granularity and will always oscillate around the true input frequency. In addition certain PLL designs have difficulties locking on to fast varying signals and can “run off”.

However, because little experience could be found with such a scenario, a set of different PLL designs, each with different capabilities in terms of stability and response time, have been developed to test the scheme. The PLLs have been designed and simulated in System Generator. To test, whether the PLLs could lock on to the laser signal, the measurements from fig. 2.14 were fed into a logic simulation and the resultant output frequency measured.

Figure 2.16 shows the simulation results. Only two of the four designs attempted actually managed to lock on and are plotted here. These were based on proportional-integral-differential (PID) controllers. This type of feedback controller is well known and can be applied to a variety of different operating scenarios. The left plot in figure 2.16 has the integral feedback disabled and handles the residual frequency offset by means of resets. This introduces spurious modes in the reference frequency, but allows a very fast locking and tracing of the reference signal. The plot on the right is from a design using a full PID implementation, which allows for a true match of the reference frequency without the spurious modes and residual error. However, it requires more time to lock on to the signal and respond to fast changes in frequency.

The test simulates the down-conversion ADC design with the reference frequency supplied by the PLLs and the scene signal directly supplied from the measurements. The phase measurement at the output should ideally be a flat straight line indicating




Figure 2.16 – The lock stability of 2 fast PID based PLLs attempting to lock on to the reference frequency. A perfect lock would manifest in a straight horizontal line indicating constant phase.

no mismatch between reference and scene signal. However, neither PLL design is free of oscillations. In both cases the error is of the order of 100 mrad, which is two orders of magnitude larger than the bit noise error achieved with the down-conversion design (see section 2.7.1). This renders the digital PLL scheme useless for implementation in this case. Nonetheless both PLLs perform very well given the rapidly changing reference frequency and maintain a stable lock. Hence the designs may be useful in future FPGA projects.

Since neither filters nor PLL based attempts to rectify the square reference frequency were satisfactory the final recommendation can only be an optical reference path. However, the final decision on this matter is to be made on an economic basis. An optical reference is comparatively expensive and thus might not be justifiable for the improvement in measurement accuracy gained over an analog filter approach. The final implementation can be chosen freely by reprogramming the SPEC board with the corresponding bitstreams, which have been prepared. No modifications to the code are necessary as part of this process and the time it takes for the modifications is mere minutes.

2.5.7 Resource Usage

Both designs presented occupy over 91 % of the slices available on the FPGA, using all of the fast DSP48 multiplier elements and more than 75 % of the available look-up table (LUT) logic. This heavy usage results in the placing and routing tools to have significant problems finding logic placement layouts that meet timing and fit on the chip. Therefore the tools had to be adjusted and carefully optimized in their




Figure 2.17 – The protocol used on the MAST-U “Aurora network”. Data is transmitted in 4B(yte) little endian packages. Only half of the bandwidth is used for data transmission. The MTS is used for transmission error checks, the META tag identifies the nature off the package (Read request, Write request etc.).

configuration parameters to be able to build the design. Even then the design was did not always meet timing, since the build process uses a random seed to determine the initial location to start the place and route process. The Xilinx SmartXplorer tool was employed to iterate and find random seeds that meet timing. The seed changes every time the firmware is modified and even changes, when using a different PC for the building process. With a working seed the entire process of generating a bitstream takes between 35 and 50 minutes. If the seed is incorrect the SmartXplorer tool would on average have to try around 20 seeds before finding another solution, which depended on the significance of the modification made.

A future project would likely use a recent FPGA generation and thus a more capable logic layout with more multipliers. Even low-end series 7 FPGAs would significantly ease the resource problems of this design. However, at present the SPEC board is the only open hardware FMC platform for sale. Since the proliferation benefit of open-hardware was one of the factors intended for this design, no alternative is available at the moment.

2.6 Aurora-PCIe Bridge

As has been mentioned in the previous sections, the diagnostic actually uses two FPGAs. One of the reasons for not having the SPEC directly stream to the PC using PCIe is that the PC cannot be in the optics room. The thermal convection introduced by the PC could significantly disturb the interferometer measurements. Hence it will be kept in a separate room. At the same time one wants to keep the analog cables from the photo diodes short, so that the digitizer is to be kept in the optics room. For this reason the SPEC sends its data out via a fiber-optic link using the Aurora protocol. Since the FPGA-chip used is comparatively small, its thermal influence is negligible.

The data stream is formatted as shown in figure 2.17. This formatting is specific to the MAST-U infrastructure which utilizes it in its White Rabbit switch

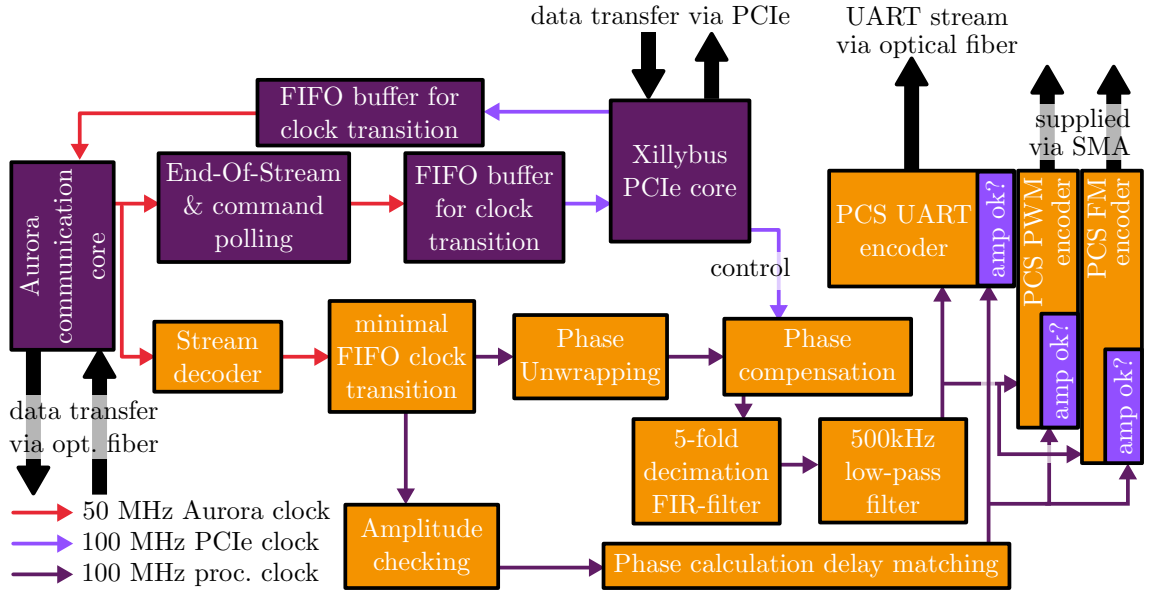


Figure 2.18 – Schematic of the interferometer Aurora-PCIe bridge. Only if the data from Aurora qualifies as the correct data stream is the PCS data fork (orange) active. The connection between Aurora and PCIe only passes through a process polling the data for a command meant for the bridge. Anything else is forwarded without modification. Because the data has no relation to the clocks on the SP605 all processing happens at full rates and is controlled via valid flags.

(WRS) infrastructure. The WRS is based on the equally named CERN open-source project and routes packages received from a FPGA to another one on the “Aurora network” [78]. This infrastructure is an effort to organise the increasing number of FPGA-based diagnostics on MAST-U. The initial implementation of the diagnostic however will be without the WRS.

The data packages hold different sets of information. Only half of the link-bandwidth holds diagnostic data (16 bit). The rest is a source/destination address (9 bit), a META tag stating the package purpose (3 bit) and a mini-time-stamp (MTS), which can be used as a sub-command tag or a transmission check.

The data sent out by the SPEC board is received by a Xilinx SP605 board housing the same FPGA chip as the SPEC[67]. This board acts as a bridge between the fiber-optic Aurora link and a PC serving two functions:

- 1 : send the optical stream from the FPGA to PC memory via the PCIe interface
- 2 : fork and process the data from the stream for the PCS

Figure 2.18 outlines the top level logic on the bridge FPGA. The SP605 is equipped with a generation 2 single lane PCIe chip (PCIe 2.0 x1). This chip is being controlled by the FPGA firmware using a Xilinx PCIe controller core, which is wrapped by Xillybus[79]. This wrapper core abstracts the communication of the

firmware with the PCIe interface and significantly reduces the amount of control logic that needs to be written in the firmware. The data from the fiber-optic link is directly passed into two separate FIFOs. The first passes the data from the Aurora to the SP605's local oscillator (LO) clock domain. In there the stream is scanned for an End-Of-Stream (EOS) command, which is used to signal the PC that the stream has ended and the device file can be closed. If this particular frame is not found the data is passed on to another FIFO transferring the stream to the PCIe clock domain and on to the PCIe interface.

In the other direction the data is first scanned to check whether the package received is for the SP605 as part of the Aurora network. If it is not, it is simply passed on to the fiber-optic link. This operation scheme in essence allows both the link and the PC to act as part of the optic Aurora network.

The data fork for the PCS (orange in fig. 2.18) is taken from the second FIFO and directly calculates the phase compensation using formula 2.15. Since the PCS cannot handle data coming in faster than about 1 MS/s the data is then passed through two 5-fold decimators to reduce the high frequency oscillations. Filtering is mandatory, since the aliasing noise induced by the high frequency oscillations can result in unwanted feedback oscillations in the PCS. Even though no decimation would result in lower latency it is still necessary to prevent aliasing. The decimation reduces the data rate to 800 kS/s and thus a 400 kHz Nyquist-bandwidth. This is still well above the 40 kHz bandwidth available on the control hardware - a limitation imposed by the digital-to-analog converters (DACs).

Since at the time of writing the final interface for the interferometer has not been decided the data for the PCS is passed on to three different data encoders. The first is a Universal Asynchronous Receiver/Transmitter (UART) encoder, sending samples at 250 kHz with 12 bit accuracy. The encoder always samples the most recent sample from the processing pipeline depicted in fig. 2.18 and drops any samples that arrive in between. This scheme provides the lowest possible latency. The small amount of aliasing noise introduced by the under-sampling does not introduce additional noise, as it is below the 12 bit resolution available for transmission. For a robust integration into the control systems the supplied signal must have some form of error check. For this reason the laser amplitude signal, which is also encoded in the data-stream coming from the SPEC, is monitored and checked against some configurable level. If the laser amplitude drops out it is detected in the stream and no data is supplied to PCS, which in turn flags an error on the receiving end. A second encoder supplies the density as a pulse width modulation (PWM) signal with the same 12 bit accuracy. This signal can be directly converted into an analog

voltage signal using a photo-diode. A third encoder provides the signal as a frequency modulation (FM) signal with 9 bit accuracy, with the frequency ranging from 150 kHz to 400 kHz denoting the lowest and highest density of the set dynamic range. This signal is compatible with one of the old PCS control boards. The benefit of the last two encoded signals is, that they can be evaluated analog and can potentially provide a lower latency control feedback. However, they are also more susceptible to transmission noise.

The limitation of the PCS to accept no more than 12 bit data imposes a limitation on the available dynamic range of the density signal. Because the signal supplied to PCS has little relevance from a physics analysis point of view it is not necessary to supply the signal at the highest possible accuracy. The maximum achievable density is expected to be the Greenwald density, which describes the density for a given plasma current and plasma cross-section at which disruptions tend to appear. For a circular plasma cross-section the Greenwald density is found to be

$$n_G = \frac{I_p}{\pi a^2}, \quad (2.19)$$

where I_p is the toroidal plasma current and a is the plasma's minor radius. Generally a tokamak cannot exceed the density limit in normal operation without triggering disruptions. However, strong shaping and during transient events, e.g. an L-H mode transition, the density limit can be exceeded. On MAST the density was just short of $1 \times 10^{20} \text{ m}^{-3}$ at 1 MA plasma current and 60 cm minor radius, which is slightly above the value for a circular cross-section plasma, due to the spherical shaping. Although MAST generally operates below the Greenwald density it has shown to exceed the limit by a factor of 1.4[80].

Since upgrades to available heating power are already planned, the granularity and dynamic range required by PCS may change over the lifetime of the diagnostic. The system has thus been designed to change the granularity (on the fly if necessary), by setting registers on the Aurora-PCIe bridge. This is done via the PCIe interface and does not interfere with the diagnostic stream. This enables maximum freedom in implementing the interferometer back-end into the MAST-U PCS.

The supplied density signals have been tested using the respective interface cards, which supply an analog signal to the PCS. Both the FM and the UART card responded as expected. However, care needed to be taken with the granularity setting, which resulted in unwanted oscillations in the control waveform, when set too low. An issue encountered during testing with the UART interface card was that the optical receivers on the card would saturate due to the strong light levels supplied by the optical transmitters used. This resulted in the interface card being

unable to lock on to the transmitted signal stream correctly. This problem can be mitigated by using longer fiber-optic cables and bulk-heads to reduce the signal level. Since the final installation of the system will employ these anyway, the issue is not expected to be a problem. Nonetheless, the levels must be checked after installation.

2.6.1 Diagnostic Control Software

The Aurora-PCI bridge described in the previous section is placed in a standard PC running CentOS-7 Linux. The communication with the Aurora bridge is conducted via device files, which are created by the Xillybus driver running on the PC and enable sending and receiving data via the PCI interface. All communication is fully controlled using Python 3 as well as the Numpy and Scipy libraries[1]–[3]. The scripts automate acquisition of shots and enable the configuration of both the SPEC as well as the SP605. They are sufficiently modular to be used for communication with any FPGA on the previously mentioned Aurora network, as long as they adhere to the defined standard.

The data stream is stored on the hard-disk-drive (HDD) in raw encoded form. Hence the scripts are also responsible for decoding the streamed data and calculating the vibration compensated density, which can then be supplied to the MAST-U central data system. The running of the scripts will eventually be controlled using the “mastda” control program, which has been in operation on MAST already.

2.7 Diagnostic Performance

The diagnostic performance was initially measured using artificial data. For this an Agilent 33600A Series Waveform Generator was used to generate clean data and to simulate experimental data[76]. The conditioning was originally conducted for the down-conversion scheme only, however where available and necessary comparisons to the Hilbert design will be made. At the time the conditioning measurements for the down-conversion design were made the HeNe reference frequency was assumed stable. In consequence the measurements were conducted using a square reference frequency passed through a 2.5 MHz low-pass filter, which would not introduce a measurement error with a stable reference frequency. The phase is always measured against some semi-arbitrary reference point anyway and the filter delay is of no concern (if constant). Only the design of the Hilbert scheme and the subsequent investigation into the frequency stability revealed the problems with using an analog filter. Nonetheless the measurements presented here indicate the capabilities of the digitization system.

2.7.1 Phase Noise

The noise introduced by the FPGA-digitizer due to the digitization process as well as the algebraic truncation as a result of the limited resources was measured. For this the ADC was supplied with a clean 40 MHz sine wave for the CO₂ scene and reference input and a 2.36 MHz signal for the HeNe inputs (sinusoid on the scene and square on the reference). All signals were chosen to be at 900 mV peak-to-peak amplitude.

Figure 2.19 on page 53 shows the raw phase difference measured on CO₂ (top 3 plots in purple) and HeNe (bottom 3 plots in orange) pipelines. As can be seen the phase noise is under 2 mrad standard deviation. The main source for this error is bit error induced noise from the processing stage, although the measurement conditions were not ideal and introduced frequency components which could not be identified. This is particularly obvious in the CO₂ spectrum, where the frequency spikes seen (apart from the one at 3.56 MHz) cannot be introduced by the processing stage, since nothing is operating at these frequencies (or harmonics thereof). Similarly the HeNe section shows many frequency components that cannot be the result of the firmware processing, although a significant portion will be the result of residual harmonics of the square reference frequency, which cannot be sufficiently filtered out. Both the CO₂ and HeNe noise spectra show a particular amplitude response, which is the result of the convolution of the filters employed in the processing stage.

A point of note is that the noise appears to be asymmetrical around zero. This may be the result of algebraic error due to truncation. This naturally introduces a bias. Nonetheless for the purpose of error estimation we will neglect the asymmetry, as it is marginal. To translate the phase error to a density error one can conduct a quadratic error propagation on equation 2.15 assuming all errors to be normally distributed. This assumption may not be valid given the many parasitic frequency components, but gives a good estimate. To approximate the error for the wavelength ratio a Taylor expansion is used and the measurement error for each individual wavelength assumed small compared to the ratio[58]. One then arrives at the following error propagation formula:

$$\begin{aligned} \sigma_{\int n_e dx}^2 &= \sum_{p=1}^n \left(\sigma_p \frac{\partial}{\partial x_p} \left(\int n_e dl \right) (x_1 \dots x_n) \right)^2 \\ &\approx \frac{K^2}{\lambda_{\text{CO}_2}^2} \left[\left(\frac{\sigma_{\Phi_{\text{CO}_2}}}{1 - \alpha^2} \right)^2 + \left(\frac{\alpha \sigma_{\Phi_{\text{HeNe}}}}{1 - \alpha^2} \right)^2 + (\Phi_{\text{HeNe}} \sigma_\alpha)^2 \right], \end{aligned} \quad (2.20)$$

where $\alpha = \lambda_{\text{CO}_2} / \lambda_{\text{HeNe}}$ is the wavelength ratio. The lasers' wavelength stabilities

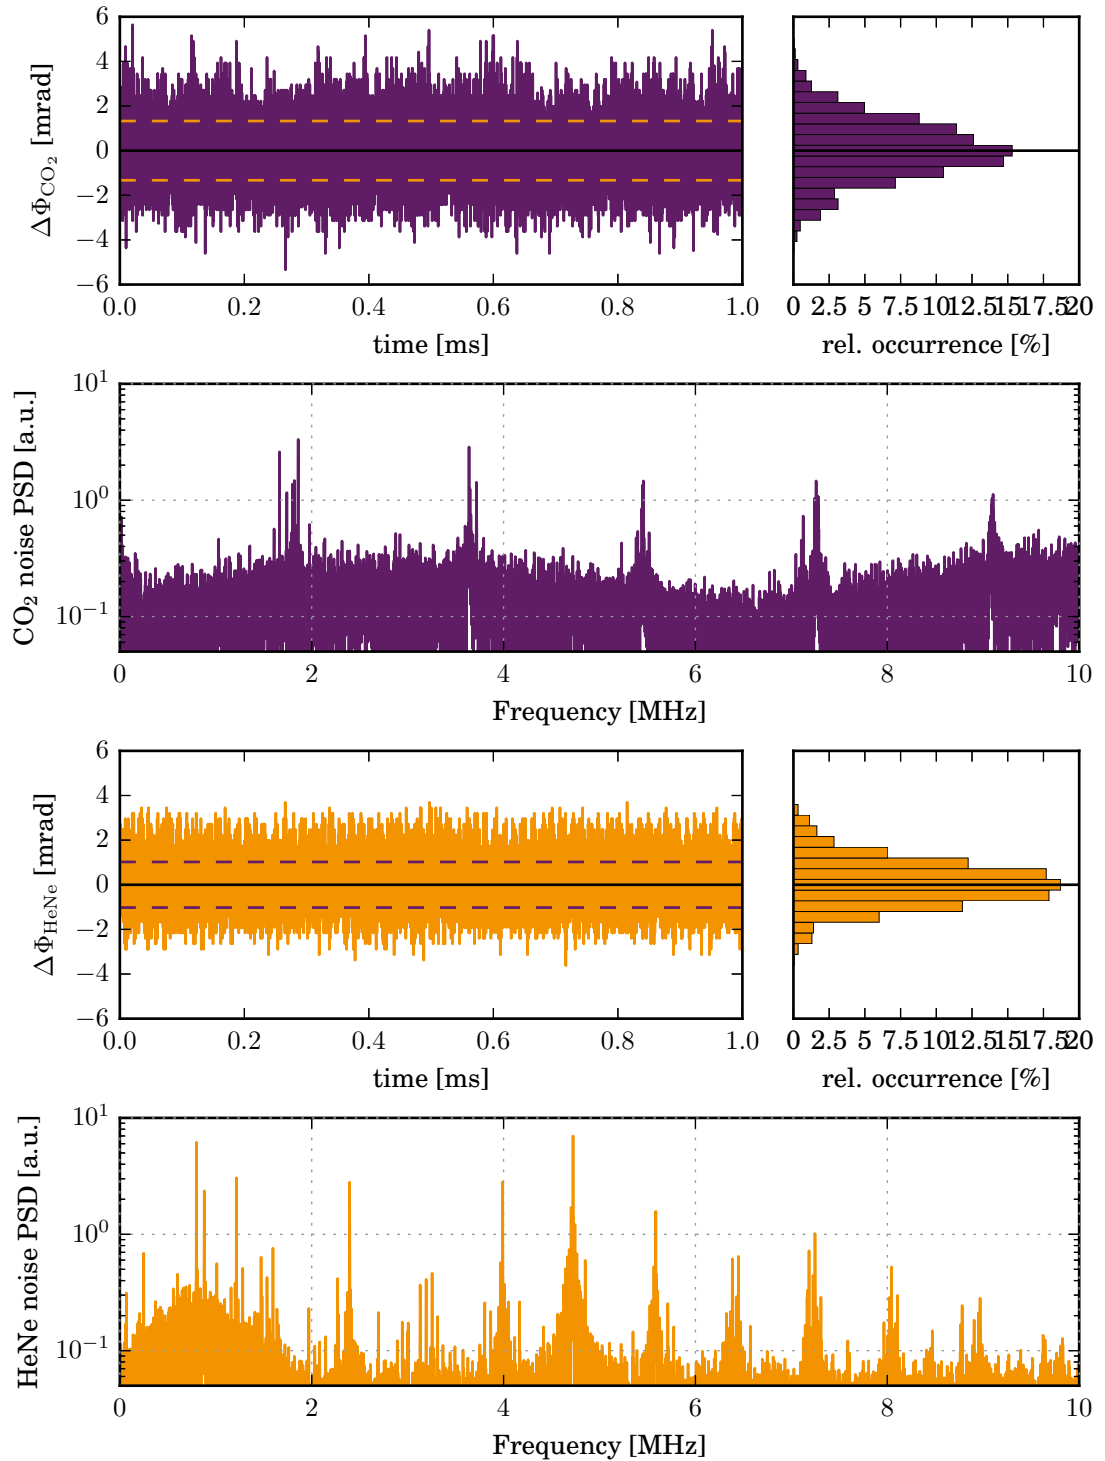


Figure 2.19 – Raw noise measurement for both processing channels separately (CO₂ in purple, HeNe in orange) at full 20 MS/s sampling rate. The dashed lines indicate the standard deviation. The asymmetry of the HeNe signal around zero appears to be caused by residual higher harmonics of the square reference frequency. The processing filters' amplitude response is visible in the noise spectrum. Some of the frequency spikes are the result of bad measurement conditions.

are given by the manufacturers[47]. Using error propagation this results in a wavelength ratio error of under 3×10^{-6} dominated by the error of the CO₂ laser, which is 3 orders of magnitude higher than the HeNe one's. Hence the error σ_n induced on the vibration compensated electron density measurement due to the digitization system is found to be better than $1.4 \times 10^{17} \text{ m}^{-2}$ at full acquisition bandwidth. While the accuracy is of the same order as other current interferometers (see for example the Joint European Torus (JET) system with $3 \times 10^{17} \text{ m}^{-2}$ [81]) the achievement of this accuracy at the full multi-megahertz bandwidth is novel.

For comparison with the Hilbert transformation design the HeNe arm was additionally tested using an optical setup. A simple Michelson-Morley interferometer as shown in figure 2.3 on page 26 was setup on an optical bench using a HeNe reference arm. The reference was then digitized the same way the signal arm was. Figure 2.20 on page 55 shows the noise measured for the down-conversion design in purple (three plots on top) and the Hilbert design in orange (bottom three plots). As can be seen the noise level is slightly higher for the Hilbert transformation design than for the down-conversion design and generally two orders of magnitude higher than the pure digitizer noise previously presented. In particular the spectra in figure 2.20 show that the difference in noise level between the down-conversion and the Hilbert transformation design is the result of the employed filtering. Since the Hilbert only low-pass filters the signal once for the initial decimation a large fraction of the high-frequency noise is kept in the signal and propagated to the CORDIC-stage. This raises the overall noise level. However, this should not be seen as a problem, as it also provides better results in the high-frequency domain, as discussed before. In the down-conversion design the last anti-aliasing filter in particular removes a large portion of this high frequency contribution.

The order of magnitude higher noise in the optical measurement compared to the normal digitization with a clean simulated signals, which is seen when comparing fig. 2.20 to fig. 2.19 must be introduced by the measured signal. The main contribution to the measurement error in this case is thus the noise generated by the optical setup as well as the photo diode electronics. This shows that the final interferometer system will be accuracy limited by its optical and analog properties, not by the digitization system's bit accuracy.

A point of note is that the disturbing frequency components in the HeNe data, which could be seen in figure 2.19 on page 53, are not present or at least far below the pass-band signal level. This indicates that the frequency components seen in the down-conversion conditioning test were the result of sub-optimal measurement conditions as well as the square reference frequency, giving another argument for an

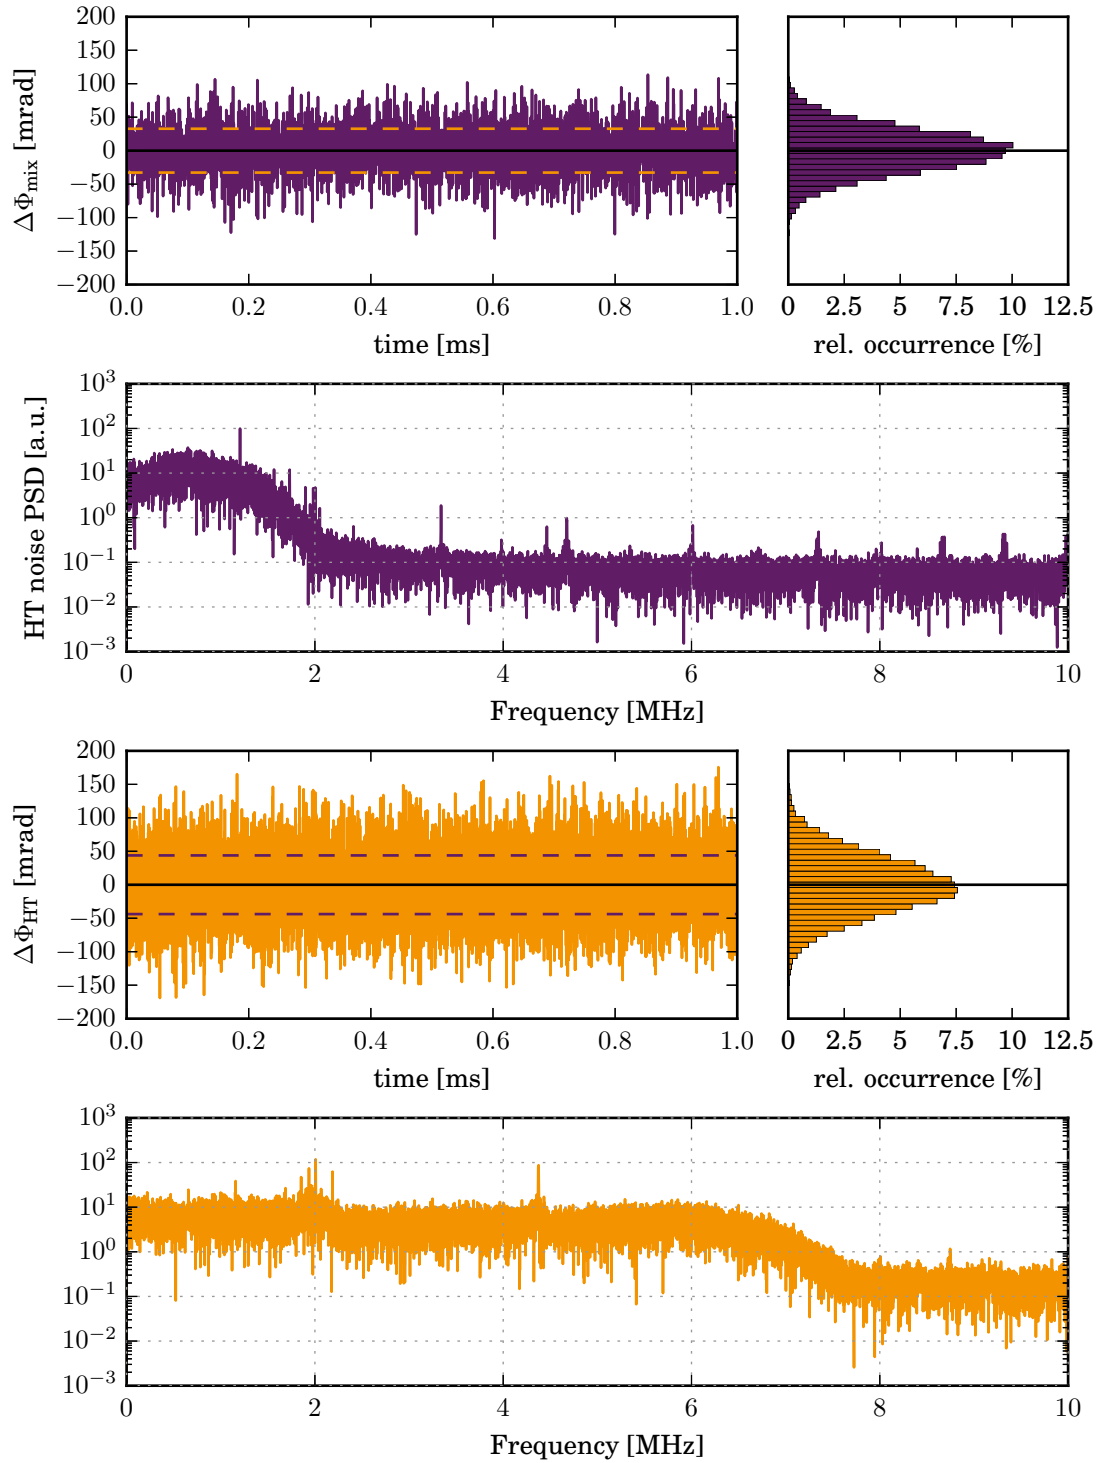


Figure 2.20 – Comparison of digitization noise for down-conversion (top three in purple) and Hilbert transformation (bottom three in orange) I-Q-generation using an optical HeNe reference arm. The dashed lines indicate the standard deviation. The down-conversion design also uses the sinusoidal optical reference. The optical setup is as depicted in 2.3 and was placed entirely on an optical bench to minimize any external influence.

optical reference generation for the HeNe arm.

2.7.2 Phase Wrap Stability

Phase wrap stability is a major concern for many two-color interferometer setups and an important quantity to be aware of in a real-time-protection system. As the phase delay between the reference and the scene arm increases the phase difference eventually becomes greater than π . However, the CORDIC algorithm will only understand values that lay on a unit circle and will wrap around to $-\pi$. This phase wrap needs to be captured accurately as a miss could lead to an error in the density larger than 10^{19} m^{-2} (using eqn. 2.15). In a feedback control scenario, where no post-processing correction can take place, such an error can result in dangerous conditions, such as the low density limit for safe NBI operation (i.e. shine-through threshold) or the approach of electron cyclotron resonance heating (ECRH) cut-off limits. Phase wraps are common to all interferometer measurements. On the CO_2 branch the wraps are generally induced by fast plasma density changes and are more common in the Michelson-Morley interferometer setup employed due to the beam passing the plasma twice. On the HeNe branch the wraps are usually caused by mechanical vibrations to which the signal is very susceptible due to its short wavelength. Generally phase wraps will occur much more frequently on HeNe signals.

The phase unwrapping algorithm is a simple comparator based approach that checks the previous value against the current one. A phase difference of bigger π is evaluated as a change across the cyclic 2π boundary, i.e.

$$\Delta x|_{|x|>\pi} = -\text{sgn}(\Delta x) \cdot 2\pi + \Delta x.$$

Phase wrap misses occur if the phase difference between the scene and the reference signal changes so fast, that the difference between two successive samples is more than π . In this case an algorithm will not be able to distinguish between a wrap and a change of phase by less than π in the opposite direction due to the cyclic boundary. From this it can be seen that the phase wrap stability is directly linked to the sampling rate.

The stability against phase wraps of the present system was evaluated by using synthetic signals generated by the Agilent 33600A Series Waveform Generator. A schematic setup is depicted in figure 2.21. The waveform generator can generate two signals with a fixed phase dependence due to the fact that the internal DACs are driven by a common clock source. However, since only two signals can be generated the CO_2 and HeNe tests had to be conducted separately (floating the unused inputs)

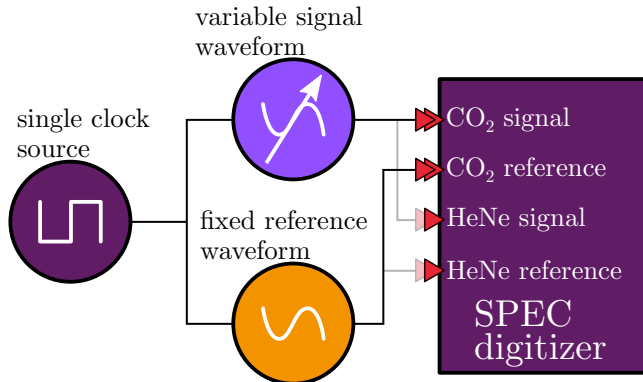


Figure 2.21 – Schematic of the test setup employed to evaluate the phase wrap stability. The waveform of both signals is generated by a single clock source on the waveform generator ensuring a fixed phase relationship. Since only two outputs are available, the HeNe measurements had to be conducted after the CO₂ measurements (indicated by the grayed signal). The unused inputs were floated.

This is indicated by the gray signal connection.

For the wrap stability test the reference frequency was kept at 40 MHz (CO₂) and 2.36 MHz (HeNe) respectively. The scene signal was varied with respect to the reference frequency and the phase difference measured. A different signal and reference input frequency translates directly to a continuously increasing or decreasing phase difference thus resulting in periodic phase-wraps with a period determined by the frequency difference. The wrapping phase was then subjected to the phase unwrapping algorithm. The slope of the resultant phase measurement must be a linear function of the frequency difference over the measured time interval:

$$\Delta\phi_{\text{per-sample}} = 2\pi\Delta f_{\text{set}}/f_{\text{sample}}.$$

Figure 2.22 depicts the deviation from the ideal slope as a function of Δf_{set} . The left figure has the frequency difference translated into rate of change of the IED using equation 2.13, which is the primary cause for wraps on the CO₂ measurements. The right figure has the frequency difference calculated as integrated mirror movement speed - i.e. using the second vibration term in eqn. 2.14 - to which the HeNe is mainly susceptible. The dashed horizontal lines indicate, where the slope deviation measured on the respective signal evaluates to a missed phase wrap over the measured time interval. The tests were conducted with 1 ms time-slices.

The phase wrap stability of the CO₂ signal can be evaluated from the left plot in figure 2.22 and is found to be more than $9 \times 10^{20} \text{ m}^{-2}/\mu\text{s}$. The shaded area depicts a gross estimate for an upper limit on the rate of density change that could potentially be expected in current and future tokamaks. It was calculated by assuming that one loses the entire core plasma density of a high density ITER shot (ca. 10^{21} m^{-2}) within a microsecond. Looking at figure 2.1 on page 58 this is 2 times the IED depicted, which is assumed to drop to null 100 times as fast. This extreme scenario




Figure 2.22 – The phase wrap stability of each laser color FPGA arm as a function of rate of density change (left) and mirror speed (right). The horizontal dashed lines indicate a 2π phase wrap, i.e. when the respectively colored line crosses it a phase wrap has been missed. The shaded area indicates an upper boundary estimate for the maximum rate-of-change that can be expected in a tokamak device.

assumes a loss rate an order of magnitude higher than the expected rates during a disruption on JET and will likely cover all MAST-U scenarios[82]. The measurements show that the diagnostic is wrap stable against fast rates of density changes on the CO_2 line. The HeNe phase is - expectedly - influenced less by the plasma density change and thus wrap stability is of little concern.

The right hand side of fig. 2.22 shows the same data evaluated as rates of integrated mirror speed, i.e. all movements of mirror elements in the optical path summed up. This is the equivalent to a single mirror moving at the summed rate due to the linear dependency on Δl in eqn. 2.14. As expected the HeNe measurement is the most affected by mechanical motion due to the smaller wavelength. One can see that the maximum rate at which the mirrors may move in sum is only a few 100 mm/s, before the HeNe signal arm will loose track of some of the phase wraps. This would correspond to a 100 Hz sound wave of 1 mm amplitude. Curiously the limit is asymmetrical towards a positive frequency difference. The first wrap misses occur once the difference frequency is higher than 600 kHz. This would initially suggest a relation to the IF, which of roughly the same magnitude. However, as could be confirmed in simulations, the reason for phase wraps appearing at this frequency is actually related to the final anti-aliasing filter seen in figure 2.8. This filter's transition band is located between 1.5 MHz and 1.8 MHz. Modulating the 800 kHz

HeNe IF with a 700 kHz signal thus results in a 1.5 MHz frequency component, which is already significantly attenuated by the filter. The reduction in amplitude results in the CORDIC core to measure the phase unreliably, probably due to the limited accuracy of the internal calculations, which in turn results in the phase wraps.

The Hilbert design improves this stability by not requiring a decimation filter, thus keeping the full bandwidth of the signal. In consequence the limitation imposed by the filter is lifted bringing the phase wrap stability up to the 2.36 MHz IF as a difference frequency, in essence expanding the positive arm of the right plot in figure 2.22 to the transition band of the CO₂ decimation filter. This was confirmed in logic simulations.

On the negative end in the figure the limitation is given by the laser's IF, which appears to contradict the fact that the signals get down-converted to a 800 kHz IF. A frequency difference higher than this should result in the mixed signal being reflected on DC. However, one has to keep in mind that the down-converted signal is complex, since we calculate both the I and Q components. Thus negative (imaginary) frequency components are defined.

The most prominent candidate for vibrations is the RR3 retro reflector as shown in figs. 2.5 and 2.6. For this reason it has been isolated against mechanical vibrations as much as possible. The reflector is made of a 80 kg granite sphere suspended by shock absorbing feet with an extremely low Eigenfrequency[47]. In addition the sphere is encased in a acrylic glass casing to prevent air circulation from creating resonant oscillations or inducing fluctuations in refractive index. It is hard to estimate the extent of the vibrations that will occur in advance. Hence further tests will have to be done once MAST-U is in the conditioning phase with the interferometer installed and the surrounding infrastructure operating.

A point of note is that in both plots the line is not perfectly parallel to the x-axis. This is the result of the digitization frequency and the reference frequency not being perfectly identical. The effect can however be easily calibrated out on a shot by shot basis using the calibration trigger setup described.

2.7.3 Bandwidth Capabilities

Since one of the main targets of this diagnostic was to measure high frequency perturbations of the plasma density, the design's maximum detectable bandwidth was probed. Plasma modulations, and in particular high frequency oscillations with small amplitude, will mostly affect the CO₂ measurements the test was only conducted for this branch. However, since the root cause of the bandwidth limitation is the same for both lasers, the results obtained for the CO₂ branch can be used to

infer the limits of the HeNe branch.

The measurements were again conducted using the Agilent 33600A Series Waveform Generator, i.e. using the same setup used in the previous section. This was done by putting two 40 MHz sinusoids into the CO₂ scene and reference arm. The scene arm signal was phase modulated at varying frequencies using the waveform generator. The HeNe was supplied with a 2.36 MHz fixed sinusoidal signal on both inputs. The amplitude of the phase modulations (PMs) was chosen to be below a 2π wrap to prevent phase wraps from influencing the measurements.

Figure 2.23 shows the results of this measurement. Each vertical bar in the top plot represents the power spectral density (PSD) of one measurement with a given modulation frequency f_{mod} . This value is overplotted with a dashed line as expected value. The bars are not of equal width because the spacing between probing frequencies was varied. Each measurement has been normalized to the PSD peak value of the respective measurement.

The middle plot shows the detected PSD peak value normalized to the largest PSD peak value found - the modulation amplitude was not modified in between measurements. The bottom plot depicts the deviation of the measured PSD peak frequency from the set frequency.

As can be seen from the figure measured Fourier peak stays at a high level up to a frequency of over 3.6 MHz - the CO₂ IF. After this point significantly more peaks appear in the Fourier spectrum, which are the result of aliasing from residual high-frequency components and reflections on DC. Even before this point the spectrum shows unwanted higher frequency components, which are likely the result of the higher harmonics seen at lower modulation frequencies. Remarkably the peaks of the modulation frequency can still be detected correctly up to a PM frequency of 7 MHz. Only after this point does the automatic detection routine fail, as indicated by the drop in the bottom plot. However, although the peaks can be detected up to this point a physical interpretation of data at this frequency range might be difficult due to the artifacts introduced into the spectrum. The true bandwidth of the system should thus be denoted as the CO₂ IF of 3.6 MHz.

Notably the amplitude response in fig. 2.23 is not flat, but instead peaks at around 1.8 MHz, and drops off at lower frequencies. This is an effect of the decimation filter's amplitude response, which itself is not flat due to the limited number of DSP48 elements available. Since the measurements presented here have been taken the filter coefficients have been redesigned to use a Butterworth coefficient structure with a flatter pass-band amplitude response at the cost of slight stop-band suppression. This has not significantly increased the noise level, but reduced the effect of the filter




Figure 2.23 – Bandwidth detection limit of the CO₂ phase measurement arm. The top plot shows the measured frequency spectrum over the modulation frequency (normalized to peak maximum). The width of the vertical bars indicates the spacing between tested probing frequencies and the overplotted dashed line indicates the expected peak location. The middle plot shows the PSD amplitude of the detected peak. The bottom plot shows the deviation from the set signal as a function of the modulation frequency.

response ripple.

The bottom plot in figure 2.23 shows the deviation of the measured PSD peak’s frequency from the set frequency. As can be seen the values appears to linearly increase, which is very likely a result of a frequency mismatch between the digitizer and the signal generator’s oscillator. This mismatch is unknown and cannot be compensated. The roughness in the line is a result of the window function applied to the data for measurement, which distorts the spectrum.

The measurements shown here show that the system is capable of multi-megahertz bandwidth detection.

2.7.4 Latency Results

Since the diagnostic will be integrated into the real time PCS on MAST-U the diagnostic latency is an important quantity to know. In the case of the two-FPGA interferometer tandem the latency actually involves the latency across two FPGA designs conducting processing on the data and including the latency involved in physically transmitting the data from one carrier board to the next. As has been mentioned in the introduction, the target latency for the diagnostic was 30 μs . This was a rough estimate based on the reaction time of the subsequent control systems, which are of the order of milliseconds.

One of the large benefits of pipelined digital processing using FPGAs is the deterministic nature of the processing with frequency independent latency. This in principle allows the calculation of the system latency by simply adding up the number of steps involved in the calculation. However, this is often impractical and in many cases not possible, since the exact number of steps might be unknown for proprietary components. In consequence the latency was measured on the chip whilst acquiring with the diagnostic. This was done by adding a Xilinx Chipscope core to the Aurora-PCIe bridge. Chipscope works by adding block RAM (BRAM) elements to the firmware and storing the values at the inputs on each clock edge until the buffer is full. The time of each sample is thus directly determined by its index in the buffer and the speed of the sampling clock. Additional logic responsible for reading the buffer via UART is also inserted. For the SPEC the overall latency of the digitizer system can thus be derived from the response to a normal command and the time it takes for the first sample to arrive at the output of the PCS-processor on the bridge FPGA.

Figure 2.24 shows the system latency calculated from an event diagram of the system responding to commands sent via the Aurora link. First an *ARM* command is sent, which is acknowledged within 1 μs (purple). Secondly a *TRIGGER* command is sent, which is not acknowledged by the FPGA, however it sets the next - and all subsequent - digitized samples valid enabling their processing and flagging them to be streamed. The *TRIGGER* command package will pass through the same logic as the previous *ARM* command package before reaching the acknowledgment logic. Hence the latency of the system including Aurora streaming logic is the time from half of the *ARM* acknowledge time to the first received data sample totalling less than 4 μs . The measurement was conducted with a 1 m OM3 optical fiber so the




Figure 2.24 – The combined event timing diagram of the FPGA system responding to an *ARM* (purple) and a *TRIGGER* (orange) command. Both commands are handled by the same logic process in the firmware allowing the deduction of the system latency by simple algebra.

latency of the optical transmission is negligible[83]. This logic delay is thus the latency of the SPEC design plus a small amount of bridge logic.

As depicted in fig. 2.24 it takes an additional 3 μs after the first valid stream sample arrives, before the PCS output’s valid flag is raised, which is the total processing latency on the SP605. Hence the total latency for PCS purposes of the IED interferometer is better than 7 μs - a factor of 4 better than the target.

The main contributors to the latency are the filter stages. The delay introduced by an FIR filter is generally half the number of filter coefficients in clock cycles plus any pipe-lining registers introduced to meet timing. A signal coming from the ADC to the PCS output traverses roughly 4 filters with 60 to 130 coefficients. This can account for 99 % of the latency.

Low latency processing of data is a significant advantage FPGAs have over other digital processing schemes, which generally tend to take significant amounts of time for communication. A digital approach can naturally not compete with analog “processing” in terms of latency, but the advantage in cost and sometimes space as well as the ability to directly supply the data in almost any digital or analog format can often justify the increased latency. The design presented here is a showcase-example of this advantage.

2.8 ASDEX Upgrade In-Situ Testing

Since the MAST-U upgrades are still under way at the time of writing in situ-measurements could not be conducted with the target system. However, a significant advantage of the design implemented here is its versatility. Since interferometers are all based on the same fundamental principle, the design presented in chapter 2 can be applied to almost any interferometer as long as the IF is compatible. Due to the FPGA-based architecture small modifications can be made to the firmware

to make it more compatible to other setups, e.g. by changing the digital mixing frequency in the down-conversion processing core to be compatible with a different AOM frequency.

The experiment found to be the most compatible and accessible was AUG, which is equipped with a CO₂-HeNe interferometer of compatible IF, as shown in section 2.8.1. To test the digitization system the FPGAs were moved to AUG to take a few days of data, providing a test under real conditions. The necessary modifications will be explained in sec. 2.8.2. It was also an opportunity to compare the system to already established systems. Since the digitization unit was employed parasitically, i.e. forking signal-lines without impeding normal operations, on very short notice no influence could be had on the shot-scenarios run during the 2 days of measurements. However, since this was in essence a test case for the MAST-U system a selection of shots was taken that would show the target capabilities for MAST-U. The results are presented in sections 2.8.3 to 2.8.5. The issues faced at AUG and lessons learned as a result are discussed in 2.8.6.

2.8.1 The AUG Experimental Setup

The AUG is a medium sized tokamak (MST) with a conventional aspect ratio of about 3. It is equipped with 2 different interferometer systems to measure the IED. The first is a Deuterium Cyanide (DCN) laser based interferometer with 5 different horizontal lines perpendicular to the torus axis as shown in figure 2.25a[84]. The IF of this particular interferometer is of the order of 10 kHz[84], which is too low to be compatible with the SPEC design. Firstly the low IF in itself makes fast digitization unnecessary, since the available information is already severely limited and no benefit can be gained from the high bandwidth system developed here. Secondly the generation of the I-Q components via mixing or Hilbert transforming is undesirable, since it would require too many filter coefficients to achieve the necessary amplitude response steepness in the transition band.

The second interferometer is a CO₂-HeNe interferometer with 3 vertical lines as depicted in figure 2.25b. Similar to the MAST-U system the IF of the AUG system is 40 MHz, however unlike the MAST-U system it identical for both lasers. Although this is not identical to the MAST-U system it is a minor difference and for proof-of-principle measurements sufficiently close.

The optical setup of the interferometer is mechanically decoupled from the main AUG vessel. The 40 MHz IF signal from the photo detectors as well as the sinusoidal AOM reference frequency is transmitted via long single ended shielded copper lines (over 30 m) to the read out electronics, which are positioned in the main control




Figure 2.25 – The measurement cords of the two IED interferometers on AUG. The CO₂-HeNe interferometer optics are mechanically decoupled from the AUG vessel. The DCN laser’s IF was too low to be compatible with the SPEC firmware. The V3 line was unavailable at the time of the measurements. Hence measurements were conducted on either the V1 or V2 line.

room. This is also where the signal was forked for the measurements, i.e. just before the established readout electronics. The signal received by the SPEC is identical to the signal received by the fast readout scheme temporarily used at AUG[60]. The V3 line in figure 2.25b was unavailable at the time of the measurement, since it is being refitted to hold a CO₂ dispersion interferometer in the future. In addition the V2 line was not available for some of the shots at AUG due to the line’s signal strength decaying. This was induced by a cold weather front changing the temperature in the laser room and could not be recalibrated on-the-fly.

2.8.2 Firmware Modifications For AUG

To make the firmware design compatible with the AUG interferometer the diagnostic required only a change in the processing and the trigger acquisition. The trigger up to this point was edge sensitive and had to be changed to a non-retriggerable level sensitive trigger. This is done by using a “monostable”-core from the general-cores repository which samples the input level[69]. It provides a 1 μ s long pulse on the first *HIGH* level with a second pulse only allowed to occur, once the level has passed


LOW again.

The changes to the processing core were only necessary to the initial I/Q generation and filtering section. Since the 40 MHz modulation of the HeNe signal makes it identical to the CO₂ signal from a processing point of view the HeNe processing section was quickly replaced by a copy of the CO₂ section. Multiplexing of all 4 input signals was not an option, since the Spartan-6 architecture can only be operated up to 390 MHz, where 400 MHz would be needed to free up additional resources for filtering[85, p.52]. Because this resulted in the CO₂ and the HeNe branch passing identical processing circuits and thus experiencing identical delays, the delay matching circuit shown in figure 2.8 could be removed, thus freeing up resources and slightly reducing the processing latency by about 1 μ s.


After these initial modifications the resulting design had unused DSP48 elements, which could be used to increase the number of filter coefficients for the 5-fold decimation filter. This reduced the amplitude ripple in the filter's pass-band and increased its stop band suppression leading to a potentially lower analog noise contribution. However, the influence of additional bit noise, due to the increased number of multiplication operations on the CO₂ signals, could not be evaluated. The results presented in section 2.7 are still valid, albeit the entire issue of the differing HeNe reference, its stability or waveform don't apply here.

Applying these changes to the firmware took no longer than 1 day, where only 3 hours went into code work and the rest of the time into running the SmartXplorer routine to find a working implementation seed for the tools so that they would meet timing. The Aurora-PCIe bridge did not require any changes. This proves the versatility of both the FPGA architecture as well as the interferometer firmware developed for it.

The integration into the AUG systems went as smoothly as can be expected. The electrical compatibility was ensured by setting the FPGAs to the correct input impedance. The cable runs were quickly put together and the only equipment the AUG team had to supply was a Linux PC with a moderately recent kernel. The available PC had 12 GB DDR3 random access memory (RAM) with a dual core central processing unit (CPU) running Lubuntu 14.04. This Linux is organized very differently from the one mentioned in section 2.6.1, however the integration went without issues and was completed within 4 hours of the hardware arriving at the lab. This is also due to the fact that all processing is conducted using Python, which is readily available in most Linux environments. The ease of integration demonstrates that the FPGA design is almost a "plug-n-play" diagnostic digitization system, which could be employed in other scenarios as well. An example is the 1 mm RF



(a) The SPEC (marked red) placed near the analog data ports.



(b) The loaned diagnostic PC carrying the SP605 Aurora-PCIe bridge used at AUG.

Figure 2.26 – Digitizer setup at AUG. The analog signals from the interferometer’s photo detectors arrive in a diagnostic rack in the control room, where the SPEC is placed (a). A loaned PC was fitted with the SP605 bridge and connected to the SPEC via an OM3 fiber under the floor (b). Setting everything up took about 4 hours.

interferometer currently at Culham Centre for Fusion Energy (CCFE), which also has a 40 MHz modulation frequency, but uses microwaves instead of visible light.

2.8.3 MGI Induced Disruption Trace

The most critical test for the digitization system is the wrap stability in the presence of quick density changes. Some of the fastest rates of changes in density can be observed during massive gas injection (MGI) induced disruptions, which result in low order MHD instabilities. During a MGI a large amount of a gas with high atomic numbers (typically Neon or Nitrogen) is quickly puffed into the plasma using a set of valves distributed around the torus[86]. The cold gas interacting with the plasma at the outer plasma edge is cooling the latter, while at the same time strongly raising the edge density due to the neutrals being ionized. This “shock-front” is diffusing from the scrape-off layer (SOL) towards the plasma core.

As a result of the plasma cooling the resistivity rises leading to a flattening of the current profile[15]. Eventually the shock-front reaches the $q=2$ surface, resulting in a sudden drop of core temperature and the development of strong tearing modes. This is commonly referred to as thermal quenching. With the thermal quench the core resistivity rises rapidly leading to a decay of the current profile and therefore

confinement. This decay of current profile is referred to as current quench[86].

Disruptions can be fatal for large fusion machines such as ITER since the energy released into the vessel walls can lead to large mechanical stresses due to vertical displacements of the plasma and high thermal stresses due to the plasma touching the walls. In addition highly energetic electrons can be generated as a result of reduced collisionality at their respective energies, which leads to rapid acceleration due to the loop voltage. These are known as run-away electrons (REs) and can be a source of very localized heat loading, even resulting in punctures of the vacuum vessel. MGIs are actually a means to help mitigate the issues just described[86], [87].

The severity of the issue in particular for ITER has resulted in enhanced research effort into disruptions and REs, requiring fast diagnostics for the investigation. IED measurements often have difficulties tracing these events accurately, since the fast evolution of the density results in fringe jumps, i.e. the misinterpretation of fast density changes as 2π -wrap.

Figure 2.27 shows a IED measurement during a MGI induced disruption for AUG shot #33658. The measurements were conducted using the V1 line. The results of the old measurement electronics are overplotted in orange. As can be seen the edge density shows the well known sudden increase as the neutral gas is ionized. The large fluctuations present during this phase are the result of MHD instabilities. As the energy is further dispersed the recombination rate eventually surpasses the ionization rate at which point the electron density starts to drop. The long slowly-decaying




Figure 2.27 – IED trace using the V1 CO₂-HeNe interferometer during a MGI induced disruption of AUG shot #33658. The AUG digitizer is overplotted using the dashed line. The inlay plots the region of highest rate of density change, found to be $1.8 \times 10^{25} \text{ m}^{-2} \text{ s}^{-1}$, with full detail for $\pm 40 \mu\text{s}$ around the time of maximum rate of change (time axes shifted for better visibility).




Figure 2.28 – Total toroidal plasma current in the aftermath of a MGI induced disruption of AUG shot #33658. The slow current decay and flat periods in the aftermath of the disruption are a strong indicator for REs.

tail at the end of the density drop is a sign of REs, which in this case are detectable for almost 180 ms. The total plasma current similarly exhibits a strongly retarded decay with flat periods as be seen in figure 2.28. This is a good indicator for the presence of REs, which this particular experiment was aiming to investigate.

As can be seen the old AUG measurement electronics and the FPGA digitization trace each other well. However, the AUG electronics had to be corrected during the disruption due to fringe jumps occurring. This is usually done by hand using the null measurement at the end of the plasma. The FPGA electronics on the other hand trace the density without fringe jumps. The maximum rate of density change measured during this disruption was $1.8 \times 10^{25} \text{ m}^{-2}\text{s}^{-1}$, of which a more detailed plot is in the inset figure. This is 2 orders of magnitude below the wrap stability measured in section 2.7.2.

This result is proof that the measurement electronics are capable of investigating extreme density changes without risk of missing phase wraps, thus providing a means to study the fine structure in fast violent plasma events such as disruptions.

2.8.4 High NTM Harmonic Detection

Although the measurements shown in section 2.7.3 already prove the ability to resolve high-frequency modes, confirmation of this statement using a real system with optical noise is an important verification step. This will show whether high frequency plasma modes can be detected on an optically limited system.

Figure 2.29 shows the IED measurement using the SPEC digitizer on cord V1




Figure 2.29 – Time slice of shot #33618 with modulated heating power. The SPEC digitizer is measuring on cord V1. The top plots the IED measured using the SPEC-digitizer down-sampled by a factor of 10. Below is the PSD of the same signal up to 500 kHz.

for AUG shot #33618 between 2.8s and 3.2s. The top plot shows the raw density trace and the spectral composition is plotted below. Figure 2.30 on page 71 depicts additional data from the same shot. At the top is a spectrogram of the B31-14 ballooning coil. This diagnostic sits at the LFS midplane. The second plot is the heating power and composition applied during this section of the shot for NBI heating (NBIH) (dark purple), ECRH (orange) and ion cyclotron resonance heating (ICRH) (light purple). The levels are recorded using a 1 kS/s ADC. The bottom two plots show the Deuterium- α (D_α) and soft X-ray (SXR) intensity. The former is sampled at 40 kS/s and the latter at 2 MS, albeit the plot has been down-sampled by a factor of 50. The time slice here is entirely in low-confinement mode (L-mode).

Several fast modes can be seen in this shot. The first most prominent feature seen in the spectrograms of both figures is the evolution of a neo-classical tearing mode (NTM). This form of plasma instability is common in tokamak plasmas and the result of small magnetic islands in the presence of otherwise strong pressure gradients.




Figure 2.30 – Time slice of shot #33618 with modulated heating power. The top plots a spectrogram of the LFS midplane ballooning coil (B35-14). Below is the heating power and composition applied during the time period. The second plot from the bottom shows the D_α power and the bottom one the down-sampled SXR intensity.

The resultant flattening of the local pressure gradient results in a loss of bootstrap current, which in turn can lead to the island to grow and become unstable [88], [89]. NTMs and the accompanying magnetic islands pose a problem for tokamaks, as they induce a soft β -limit, i.e. limiting the maximum achievable plasma performance. In addition NTMs can grow to a point that they cause low-order MHD instabilities, which often result in disruptions. The research into NTMs has greatly improved the ability to control them, e.g. by using localized electron cyclotron current drive (ECCD) and ECRH [88], [90]. Magnetic islands can even be useful when controlled. One can for example introduce magnetic perturbations using special coils to control

the edge transport and as a result the edge pressure gradients. This may yield an efficient method for ELM control or mitigation[91]. The initial seeds for NTMs are usually found near rational flux surfaces, which always exist in tokamak plasmas. The NTM seen in figure 2.30 has a 3/2 periodicity, which has been confirmed using the AUG MTR code. This IDL-based code is commonly used on AUG to identify and analyze MHD-modes using several different diagnostics, most importantly the Mirnov coils, but also electron cyclotron emission (ECE) and more. To find the periodicity in the present case, the signals from several Mirnov coils distributed poloidally and toroidally around the torus were compared. The phase difference between the otherwise identical signals allows the derivation of the periodicity. The higher modes above the NTM's fundamental (between 10 kHz and 40 kHz) are harmonics. The frequency of the fundamental is constantly changing, which is partially caused by a modulation of the heating power resulting in the NTM's rotation velocity and size changing. This is indicated in the second plot of figure 2.30. A second contribution is probably the constantly changing pressure profile as indicated by ECE measurements.

In the time from 2.3 s to 2.7 s the plasma appears to experience sawtooth instabilities. These are identified by the repetitive rise of density, core SXR and temperature followed by a sudden drop in all three[92]. These MHD instabilities are found on almost all tokamaks and have a detrimental effect on the plasma profile. During the crash phase hot electrons are expelled from the plasma core thus flattening the pressure profile. Figure 2.31 on page 73 shows the spectra of 3 out of 60 different ECE channels up to the Nyquist frequency. The raw down-sampled ECE temperature is overplotted in each case. The selected channels contain the maximum amplitude of one of the observed fast modes visible in the density and magnetic spectra. As can be seen from the figure the sawteeth's characteristic ramping is visible on both channel 39 and 47. At the same time channel 56 appears to increase in temperature until all three channels settle around 1.5 keV. The sawtooth activity appears to increase up to about 2.55 s at which the SXR starts to strongly oscillate making the ramping indiscernible. This is accompanied by a rapid decline in the temperature profile (as indicated by the overplotted ECE temperatures in fig. 2.31). The decline is ended with a broadband spike in the magnetics spectrum leading into a period of very low SXR and flat temperature profile peaking at roughly 1 keV. The IED is largely unaffected by this.

The next prominent feature can be seen between 2.3 s and 2.4 s around 100 kHz to 150 kHz. This mode is visible on both the density and the magnetic spectrograms and, using the MTR code's mode analysis, can be confirmed as a toroidal AE (TAE). These modes are a low frequency extension to the plasma dispersion relation and belong to




Figure 2.31 – ECE temperature spectra for ECE channels 39 (top), 47 (middle) and 56 (bottom). The decimated ECE temperature is over-plotted.

the family of R-waves covered in 3.2. A TAE is a MHD wave propagating along the magnetic field and in this case has a $n=4$ toroidal periodicity. The maximum of the TAE can be observed on the ECE channel 47, but it is visible down to ECE channel 40. This places the temperature perturbation mainly on the high-field side (HFS), albeit several of the equivalent LFS ECE channels on the same flux surfaces were broken.

The most interesting feature in the density measurement is the appearance of several intermittent high frequency modes in the 400 kHz range between 2.8 s and 3.1 s. As can be seen the B35-14 ballooning coil does not see any of these oscillations and indeed none of the magnetic coils could pick this signal up. Nonetheless it is clearly visible on the density and ECE spectograms. Both of these modes appear at first sight to be harmonic to the NTM oscillation at the bottom, this is however only true for the mode occurring between 3.05 s and 3.15 s.




Figure 2.32 – Smoothed measurements of the IED on the V1 line (top plots) and electron temperature profile (bottom plots). The left side shows a slice with 400 kHz oscillations, the right plot a slice about 100 ms later, where the oscillations have disappeared. The time offset of each plot is indicated by t_0 at the top.

First one must consider, whether this harmonic mode is a mathematical artefact

induced by a highly localized island moving very quickly past the measurement cord or ECE radius. The Fourier-envelope of a function with very steep gradients in the temporal domain would have a sinc-like shape in the frequency domain and therefore very high frequency components. This could result in the high harmonics with a spectral gap in between, as seen in figures 2.29, 2.30 and 2.31. Figure 2.32 shows the raw density trace for two time slices. The left side shows a time slice, where the high harmonic is present, the right shows a time slice without the fast harmonic. The top plot on each side is the IED trace from the SPEC digitizer, the bottom is the midplane electron temperature profile over the major radius derived from ECE. In the density trace the vertical line marks the point in time the ECE profile was generated at. The radial position of the ECE channels plotted in figure 2.31 are marked. The reason why the higher ECE channel is positioned at a lower radius is due to the fact that the midplane radial location was calculated using an equilibrium reconstruction. This calculates the ECE location using flux coordinates, which are assumed quasi-axisymmetric (in the flux coordinate frame). As a consequence the ECE channel location gets reflected at the magnetic axis.

As can be seen from the density trace in figure 2.32 the oscillation is not a mathematical artefact induced by a steep rise in the density. The density trace clearly exhibits periodic fluctuations on top of the NTM perturbation. The fact that the oscillation is only visible on IED and ECE measurements may therefore be a result of the integral nature of the magnetic measurement. The actual source for these oscillations is difficult to ascertain. Comparing to figure 2.29 it can be seen the temperature profile exhibits a very strong flattening during the time of the oscillations (see the purple marked region in the bottom plots of figure 2.32). The fact that the fast oscillation at this point in time is harmonic to the NTM means that the effect seen in the frequency domain cannot be an independently propagating wave or oscillation. Such a wave would have its own evolution in the frequency domain according to the local dispersion relation for that wave. A modulation by the NTM would be symmetric around the wave's fundamental frequency, which would nonetheless evolve quasi-independently of the NTM; even assuming that the dispersion relation of such a wave would be governed by the local plasma parameters, which may be largely defined by the NTM island. An example for such a feature is the ion acoustic wave described later in this section. The only conclusion can therefore be that the effect seen is a result of the physical substructure of the NTM, e.g. small-scale field line bending or micro tearing. As such a structure moves past a measurement location with the main NTM island, the perturbation (such as the one seen in the top left of figure 2.32) would look like a periodic high frequency oscillation

in phase with the motion of the main island. The fact that the amplitude for this harmonic is found on ECE channel 39, which from figure 2.32 can be located in the core of the plasma, indicates, that the structure is located further away from the core of the main NTM's island. Using a fast spatially resolved density measurement, e.g. a fast reflectometer, could potentially confirm this, however such a measurement was not available for this analysis.

A feature to note in figure 2.32 on the right hand plot, which does not exhibit the just described oscillations, is the small rise of the IED coinciding with the NTM density dip. This hints at small bands of increased density located at the center of the magnetic island. This phenomenon has previously been observed and a theoretical description presented [91], [93]. The density spike seen here is much less pronounced than the previously recorded observations, which may be the result of the NTM location with respect to the measurement cord.

The second fast wave seen in fig. 2.29 is appearing around 2.8s. This wave is not harmonic to the NTM at the bottom. It chirps up in frequency, but is evidently modulated by the NTM rotation. This is the behaviour expected from an independent wave propagating according to the local plasma dispersion. A classification of the wave is again difficult. It can be seen, that the occurrence of the wave coincides with a rapid increase in the electron temperature, which is particularly obvious in ECE channel 39. This indicates that the wave is of ion acoustic character. The dispersion relation of an ion acoustic wave for a deuterium plasma can be found as

$$\begin{aligned}\omega &= v_S k \\ \text{with} \\ v_S &= \sqrt{\frac{\gamma_e k_B T_e + \gamma_i k_B T_i}{M_i}}.\end{aligned}\tag{2.21}$$

The derivation of this equation requires the warm plasma dispersion. In this equation k_B is Boltzman's constant, $\gamma_{e,i}$ are the adiabatic coefficient for the electrons and ions and $T_{e,i}$ are their respective temperatures [94, ch. 4]. The ion mass is indicated by M_i . The increase in frequency in combination with an increase in plasma temperature strongly suggests this class of waves. The density perturbation observable on the interferometer can be explained by the compressional nature of the wave propagation.

Although the phenomena seen in this section could not be fully explained, it is obvious that the fast measurement of the IED using the SPEC digitizer greatly helps in the investigation. In particular the combination with fast magnetic and

ECE diagnostics provide a powerful tool to understand high frequency activity in a tokamak plasma.

2.8.5 Indications For Fast Mode Physics On AUG

The high bandwidth capabilities of the SPEC digitizer provide access to high frequency physics regimes that are currently inaccessible using the diagnostic set on AUG. Although AUG has the ability for short 100 ms bursts of 10 MHz magnetic acquisition, the limited acquisition time requires the advance knowledge of the interesting time window, which is not always given. The continuous measurement at a fast rate using the SPEC digitizer is free of this limitation.

Figure 2.33 on page 78 shows a measurement of AUG shot #33624. The time window is at the start of the shot with the first application of NBIH marked with a vertical line. The NBIH is applied at a level of 2.4 MW with 59.1 keV Deuterium atoms. The top plot shows the raw density trace measured with the SPEC digitizer. Below that is a spectrogram of the same data. The spectrogram clearly shows a high frequency mode starting at 1.5 MHz several milliseconds after the NBIH is applied and reaches its maximum power level. In the initial phase the mode shows a burst of different “sub-modes” which all appear to be emerging roughly at the same time, quickly chirping up in frequency. In the second phase, after the initial chirp’s maximum frequency is reached at around 1.8 MHz, the chirp reverses and goes down in frequency, which also results in most of the modes dying off. During this entire time the IED increases.

After the downwards-chirp has stopped (around 400 ms) a single mode around 1.6 MHz is emerging as the dominant mode, which starts to chirp upwards. During this phase the IED appears to be stagnating and even decreasing. The plot below the spectrogram plots the LOS averaged Alfvén velocity $\langle v_A \rangle_l$. The initial up-down chirp phase in the spectrum appears to be following the same trend as the Alfvén velocity. One may therefore assume that the nature of the mode is Alfvénic.

After the upwards-chirp phase the mode appears to be disappearing around 500 ms, but reappears around 600 ms. The cause for this cannot be determined just from this measurement, but may be the result of the localized mode moving away from the measurement cord, e.g. by moving radially inward. The Alfvén velocity is decreasing during this period, although an increase would be expected from the increase in frequency at its reappearance. This supports the previous thesis about the mode moving with respect to the measurement cord. At this point a sudden rapid upwards chirp all the way to 3.5 MHz is observed, with the maximum observed amplitude around 600 ms. The mode’s frequency trend again appears to be consistent




Figure 2.33 – Spectrogram of the SPEC based density measurement on the HFS V2 cord at the start of AUG shot #33624. The top plots the IED, with its spectrogram below. The second plot from the bottom prints the LOS averaged Alfvén velocity and the bottom plot the V2 cord length from the equilibrium reconstruction. The vertical line in each plot is the onset of NBIH.

with the Alfvén velocity, since $\langle v_A \rangle_l$ also appears to rapidly increase. After 670 ms the mode cannot be seen any more.

The rapid upward chirp and amplitude maximum at 600 ms can be explained from the plasma evolution. The bottom plot shows the length of the V2 measurement cord l_{V2} and the inner most boundary of the plasma, both of which are derived from equilibrium calculations. At the start of the time window depicted in figure 2.33 the plasma is not yet fully formed and only has a circular cross section. As the beams are applied the cross section increases in diameter and triangularity. The X-point configuration is formed around 500 ms, marking the transition from a limited to a diverted plasma. Between 600 ms and 700 ms the plasma's geometric axis - and therefore its entire cross section - is being moved radially towards the LFS, thus rapidly decreasing l_{V2} as indicated in the bottom plot. The result is that the measurement cord transitions from a high density region to a low density region. Alfvénic modes are generally localized to a specific flux surface, due to the required band-gap. The increase and subsequent decrease in mode amplitude is likely a result of the poloidal geometric structure of the mode. Using the equilibrium flux surface reconstruction and the position of the measurement cord one can localize the mode to a flux surface encompassing 2/3 of the magnetic flux. The increase in observable frequency during this phase can be attributed to the fact that the Alfvén velocity is proportional to $n_e^{-1/2}$. Hence as the measurement cord is progressing through ever decreasing density layers the frequency of the Alfvénic mode will increase in these regions. This is in essence a representation of the Alfvén continuum.

The classification of this very high frequency mode is difficult, since no routine diagnostic at AUG has the ability to resolve it. The magnetic and ECE diagnostics are not able to investigate this, since they are sampling at a rate of 1 MS/s and are protected against aliasing using low-pass filters. In consequence not even the reflection around the Nyquist frequency is visible.

The Alfvénic character and high frequency of this particular mode suggests a CAE. This high frequency AE lives in the Alfvén continuum band gap induced by density fluctuations in the core of the plasma[50]. The excitation of CAEs is often from energetic ions during high power NBIH, which transfer energy via Landau damping. This is consistent with the observation here. In addition the frequency splitting observed during the emergence of the wave appears to be consistent with CAE poloidal and toroidal mode splitting[95]. However, several inconsistencies can be found with the CAE argument. Firstly CAEs are generally excited on the LFS, due to the requirement for a confining potential well[96]. The V2 cord however measures on the HFS even during the short start-up period. A mechanism found in simulations

is the conversion of a CAE into a kinetic Alfvén wave (KAW) via finite gyro radius effects[96]. KAWs are localized to the HFS of the plasma cross section and therefore could be visible in the V2 measurement cord. A second discrepancy is the fact that the Alfvén velocity of the NBI is roughly 2400 km/s, while the core plasma v_A is at least 6200 km/s. For Landau damping to occur the beam velocity generally must be 2 to 3 times the plasma’s Alfvén velocity, which is the reason why AEs are much more common on spherical tokamaks (STs), where a low field is present. This condition is not met with this AUG measurement, thus posing the question of how the initial excitation and energy transfer is conducted. A possibility is the excitation via a Doppler shifted harmonic cyclotron resonance condition $\omega = k_{\parallel}v_b - \omega_{ci}$, where v_b is the beam velocity and k_{\parallel} is the wave number of the excited wave. The ion cyclotron frequency at the start of the window is roughly 18 MHz on the magnetic axis so only about 10 times that of the emerging frequencies, which is within reasonable limits[96].

Instead of a CAE the emerging mode could also be a global AE (GAE). These modes have frequencies similar to a CAE, but show different evolution behaviour[95]. They are localized more to the core of the plasma, but can under certain condition move radially outwards. A possible hint that the observed mode is a GAE could be found in the initial phase with multiple simultaneous modes. At the highest chirp frequency some of the modes appear to be crossing in frequency space. This behaviour identifies GAEs[97]. However, since the modes die off at the same time and do not quite make it to the other mode this indicator is not very strong. Without further knowledge of the mode structure of the emerging wave, further analysis is difficult.

The measurement presented here motivates the necessity for the AUG experiment to improve its diagnostic set as there is a significant amount of interesting fast-wave physics that currently cannot be investigated. The just presented data should be taken with caution, in particular since the equilibrium reconstruction used for the analysis is likely unreliable at this stage of the plasma formation. Nonetheless, the digitization using the fast SPEC digitizer showed the ability to find fast wave physics that was previously inaccessible.

2.8.6 Operational Difficulties

The SPEC digitization system performed nominally under most conditions found during the time at AUG. However, during the measurements at AUG the digitizer appeared to have several shots with unreliable phase tracing.

Figure 2.34 shows a time slice from shot #33624 at the point of which the ICRH




Figure 2.34 – Time slice of AUG shot #33624. The top right plot shows the raw density trace from the SPEC digitizer (note the logarithmic scale). The top left plot shows the corresponding spectrogram. The bottom left depicts the applied heating power, the bottom right is the normalized CORDIC amplitude. The vertical lines in each plot mark the start of the ICRH as set by the AUG control system.

sets in. The top left shows a spectrogram of the SPEC density trace. The bottom left is the ICRH heating power. The top right plots the raw SPEC density trace and the bottom right the CORDIC amplitude, which is transmitted at a down-sampled rate with the data stream. As can be seen the density trace is normal until 1.399s. At this point the ICRH is applied, which results in the CORDIC amplitude for the CO₂ scene signal to become very noisy and rapidly changing with the average moving towards zero. In consequence the phase trace start to rise rapidly due to phase wraps occurring. Surprisingly, the HeNe scene signal is almost unaffected. The higher relative noise level on the HeNe scene signal is caused by the normalization combined with the fact that the HeNe signal level is only a fifth if the CO₂'s. Although an increase in noise can be detected with the onset of the ICRH it is not at the level of the CO₂ signal's and should not pose a problem. The reference signals in both cases do not get disturbed by the ICRH.

The cause for this is related to the AUG diagnostic setup. The ICRH antenna at AUG operates between 30 MHz and 80 MHz[98]. In this particular shot the frequency was set to 36.5 MHz, which is a common frequency used for many scenarios. As described in section 2.8.2 the signal from the interferometer detectors is carried to the digitization electronics via long copper cables. Although these cables have been heavily shielded they still act as a very good antenna for the ICRH signal, which has a wavelength of roughly 8 m. The ICRH signal is picked up and mixes with the 40 MHz AOM frequency. The resultant IF modulates all signals relevant to the phase measurement and causes significant interference resulting in unreliable phase measurements. The IF resulting from ICRH and AOM frequency is clearly visible as strong 3.5 MHz peak on the spectrogram in figure 2.34. The measurements presented here show, that the problem is actually mostly localized to the CO₂ scene signal. Although the shown effect is representative to ICRH shots on AUG, the severity of the problem depends on the power level employed for heating. As can be seen from the measurements in section 2.8.4, low powers of ICRH do not appear to be leading to phase wraps.

The problem described here is well known to the AUG staff and efforts have been undertaken to improve the shielding on the relevant cables, albeit to no avail. The SPEC digitizer could however provide the solution, since the problem appears to be caused by the long copper cables. The fiber-optic transmission implemented with the SPEC digitizer should mitigate this problem. A very simple and easily implemented solution would be to place the SPEC digitization board directly next to the optical detectors using only very short (≈ 10 cm) copper cables to connect the detector with the ADC. The ICRH wavelength being two orders of magnitude larger, should not be resonant to such short cables. The dimensions of the SPEC board itself are less than 30 cm, so no pick-up should be expected here either. Since the data from the digitizer is transmitted digitally via optical fiber, no further interference from the ICRH should be expected. This solution has been proposed to the AUG team.

The ICRH problem encountered at AUG is not relevant for MAST-U, since no ICRH is employed here. However, it is generally noteworthy, that pick-up from any source can be problematic and significantly impede the density measurement. The digitizer can minimize the pick-up by keeping leads short and making use of the fiber-optic transmission.

2.9 Summary & Outlook

In this chapter we have presented an FPGA based digitization back-end for the new MAST-U two-color IED interferometer. The digitization hardware is fully designed using open source hardware and has a total cost of under £1000 in 2016. The data acquired by the digitizer is transferred via fiber-optic link to a second FPGA, which has the benefit of mitigating electrical pick-up. This has been found to cause significant problems with the measurement electronics on machines such as AUG. The second FPGA has the purpose of storing the data into PC memory as well as processing it in real time to supply a control signal for the MAST-U PCS. The system can be operated continuously for an arbitrary time, where the limit in data acquisition (but not PCS control) is given by the hard drive capacity of the acquisition PC.

The digitization hardware has been extensively tested using synthetic data and shown to be phase wrap stable for all normal operation scenarios to be expected in fusion machines. Plasma density perturbations have been shown to be detectable up to 4 MHz with an accuracy of 10^{17} m^{-2} at full 20 MS/s digitization bandwidth. Using a second FPGA real-time processing of the acquired signals is achieved at a well defined latency of less than 8 μs . This is well below a 30 μs limit set by the MAST-U PCS response time.

The digitization hardware has been tested in-situ at the ASDEX Upgrade tokamak using the vertical CO_2 -HeNe interferometer. This has shown the versatility of the FPGA design and firmware, being able to integrate the system into a new experimental environment within a few hours. The measurements at AUG have confirmed the abilities of the system regarding phase wrap stability and bandwidth detection limits. During the short campaign very high frequency modes could be detected, which the diagnostics generally available at AUG at the time of the measurements could not. Most notably this includes the fast magnetic or ECE diagnostic, which could not pick-up due to their limited bandwidth. This discovery has been discussed with the AUG staff and may lead to diagnostic upgrades in the intermediate future. For MAST-U the discovery confirms that the investigation of fast Alfvén Eigenmodes will in the future be supported by the two-color interferometer. Some of the fast mode observations obtained during the AUG campaign could be worth investigating further.

Since at the time of writing the MAST-U upgrade is still under way, the final integration into its systems has not yet been finished. This will be the next step, which should not yield any major issues. Before the first shots of MAST-U start, an integrated systems test will have to be conducted including the optical system

and all of the final infrastructure. At the time of writing the optical and electrical installations are still not finished, which is why the tests were not conducted already.

The system designed here was strongly constrained from a FPGA hardware resource perspective, in particular regarding timing and fast multiplier usage, which can be attributed to the FPGA chip used. In the future using a more complex chip will greatly lift these constraints. However, at the time of writing no bigger open-hardware FMC FPGA board is available.

The experiences gained during the development of the system helped in designing the SAMI diagnostic, as will be seen in the next chapter. In particular the heterodyne down-conversion scheme employed in the SPEC digitizer was also employed (although adapted) and Aurora-PCIe bridge firmware could be reused with some minor modifications. This will be described in the next chapter.

Chapter 3

Continuous Pitch Angle Profile Measurements using SAMI

3.1 Introduction

The edge of a magnetically confined plasma is the interface between the “hot” core plasma that is oblivious to the vessel it is contained in and the “cold” vessel wall essentially acting like a perfect heat sink. In fusion plasmas the edge is believed to be one of the primary factors defining the overall plasma performance - most prominently in the high confinement mode (H-mode). This operation regime was first found on the Axially Symmetric Divertor EXperiment (ASDEX) in 1982 and later on various other fusion experiments - both tokamaks and Stellarators[99], [100]. The improved confinement in H-mode is the result of a strong suppression of turbulence in the edge, which is related to a strong increase of edge flow shear in tokamaks[100]. The reduced edge transport results in a rapid increase in plasma pressure over a short distance and the formation of a pedestal, i.e. a knee-like edge in the pressure profile, which is a feature not found in the low confinement mode (L-mode). This is schematically shown in figure 3.1. Although a very good empirical understanding of the conditions leading to the improved confinement has been established, the theory behind the H-mode onset is not yet well understood[100]. For this reason the plasma edge is subject to intensive studies, and edge diagnostics, which can provide detailed




Figure 3.1 – Schematic of the plasma pressure in L-mode and H-mode. The formation of a strong edge pressure gradient and pedestal are the defining properties of H-mode.




Figure 3.2 – Diagram of the ELM peeling-ballooning stability boundary, with the stable region marked in orange. The dynamics in the parameter-space of all three types of ELMs are indicated. The pressure or current density rises in direction of the instability boundary, before the fast crash occurs though the unstable region[102].

measurements of the plasma parameters in the edge, are of great interest. One of the most investigated edge-phenomena is the edge localized mode (ELM). This type of instability is found in most H-mode plasmas and poses a serious threat to ITER and even more so the DEMOnstration power plant (DEMO). An ELM can expel a significant fraction of the energy stored in the plasma in a very short time period, which can lead to significant erosion and damage on the divertor tiles due to the small area wetted by the plasma[101].

The stability of ELMs is a field of intensive study. There are several subclasses, however the major classification has three types, which separate the ELMs according to the underlying instability causing the “crash”. This is schematically shown in figure 3.2, where the stable region is marked in orange. The two main parameters of interest are the edge pressure gradient and the edge current density. If the former is too large the plasma becomes unstable against ballooning modes. These result in relatively small ELMs of type II. If the latter becomes too large peeling instabilities form resulting in type III ELMs, which are also of moderate size. The biggest ELMs, which also pose the most risk to future high performance machines, are of type I. In this case both the pressure gradient as well as the current density are beyond both the peeling and the ballooning instability boundary[102]. The resulting instability can act to reduce the pressure gradients, which induces large edge currents and further pushes the mode in direction of the peeling unstable region. This coupled mode is thus called Peeling-ballooning mode and results in the largest crashes.

The edge pressure is a relatively well diagnosed parameter. Diagnostics such as Thomson scattering (TS), which is found on almost all larger fusion experiments, provide measurements of the electron temperature and density with good temporal and spatial resolution. Other diagnostics include retracting probes for parameters in the far edge of the plasma and charge-exchange spectroscopy (CXs) for the ion temperature[103]. In consequence the ballooning stability limit (right hand side of fig. 3.2) is relatively well understood. The edge current density however is a difficult parameter to measure. In most cases it has to be derived from other parameters, such

as the magnetic pitch angle gradient in the plasma edge (via Ampère’s law) or using the pressure gradient (using the force balance equation). Diagnostics that have been able to measure edge current density either directly or indirectly use for example Zeeman polarimetry on either Lithium-pellets as they ablate or Lithium beams[104], [105]. Another example is the motional Stark-effect (MSE) diagnostic, which uses the Stark broadening instead[106]. Lithium is however a material that needs to be handled with care and is therefore not commonly found on fusion experiments. In consequence MSE is the only pitch angle diagnostic found on most fusion machines. MSE itself relies on neutral beam injection (NBI), which is undesirable in certain operation scenarios. Although not a concern for ELM-studies, NBI tends to induce the transition from L- to H-mode in divertor tokamaks. In consequence the pitch angle in L-mode can often not be investigated using this technique. The limited available diagnostic set for the edge current density results in a poor understanding of the peeling stability limit and thus has long been a highly desired parameter.

The Synthetic Aperture Microwave Imaging (SAMI) diagnostic is a first-of-its-kind diagnostic imaging the edge of a fusion plasma in the microwave range using a synthetic aperture approach. This type of imaging is commonly found in radio astronomy, e.g. the Very Large Array[107]. It has several advantages over optical approaches, most notably a superior field of view given a limited array size. In addition its plasma facing radio frequency (RF) components, i.e. the antennae, can be manufactured out of metal, making them resilient to the effects of fusion neutron bombardment. Although not of concern in current fusion experiments, machines like ITER and its successor, the DEMO, will have to deal with significant neutron loads, which are already predicted to cause problems for the traditional fusion diagnostic set and in particular optical diagnostics such as MSE[108].

SAMI’s original purpose was the study of electron Bernstein waves and the derivation of magnetic pitch angle by looking at B(ernstein)-X-O mode conversion windows[109], [110]. In addition the plasma can be probed actively by sending out a modulated RF-beam. The beam scatters off the plasma’s iso-density surfaces and is picked up by the measurement array[109]. This has recently been shown to enable the measurement of magnetic pitch angle via 2D Doppler back-scattering (DBS)[111]. However, magnetic pitch angle profiles, which would be necessary to derive the edge current density, could not yet be acquired with this technique due to limitations of the original diagnostic design.

The SAMI diagnostic relies heavily on field programmable gate array (FPGA) technology. The analog-to-digital converters (ADCs) are operated by a pair of Xilinx ML605 FPGA carriers[112]. These run a soft-processor based Linux to control the

data acquisition in the FPGA logic[113]. The original firmware was designed for the Mega-Amp Spherical Tokamak (MAST) experiment as a diagnostic controller and plain digitizer streaming data directly into on-board double data-rate (DDR) random access memory (RAM). The amount of RAM available on each FPGA limits the available acquisition time to just over 500 ms. Modern fusion experiments run shots on much longer time scales. The upgraded NSTX and MAST experiments for example are anticipated to run for up to 5 s. Experiments such as ITER will run for tens of minutes. Hence the original acquisition scheme in SAMI is insufficient.

The work presented here is part of the ongoing development of the SAMI diagnostic. The experience gained over the past years using the diagnostic on MAST showed several improvements that could be made to its design. Most notably, the 2D DBS technique developed by David Thomas *et al.* has yielded interesting possibilities to develop the SAMI diagnostic beyond its original intent[111]. To make use of long shot-lengths the diagnostic has been modified to enable continuous 2D DBS measurements of magnetic pitch angle profiles. In addition the passive raw data acquisition has been enhanced to allow windowed data acquisition. These modifications have been made as part of an integration into NSTX-Upgrade (NSTX-U), where the SAMI diagnostic has been operating since May 2016.

One of the differences between the MAST experiment, for which SAMI was originally developed, and the NSTX-U system is the substantially closer plasma. As will be discussed in this chapter, the far-field approximation, which has so far been used to evaluate the SAMI data, does not necessarily apply in the NSTX-U geometry. This motivated a preliminary investigation into near-field effects as well as the use of alternative image reconstruction methods using matrix inversion.

First the relevant theoretical background required to understand the SAMI diagnostic and the 2D DBS technique will be outlined. The wave physics will be based on the theory detailed in the previous chapter. In section 3.3 a thorough description of the various potential near and far field image inversion techniques that SAMI can employ is given. This is followed by a description of the 2D DBS technique and its application to derive the magnetic pitch angle in section 3.4. Both the diagnostic's firmware and hardware are introduced in 3.5 followed by a presentations of the improvements to the FPGA firmware in 3.6. Some of the firmware modifications were made in preparation of SAMI's installation on NSTX-U. The integration process is detailed in section 3.7 including the necessary hardware modifications. Since the move to NSTX-U makes near-field effects potentially relevant a preliminary investigation into the applicability of SAMI's image reconstruction algorithms is presented in section 3.8, which will also investigate near-field inversion routines.

Finally the initial results from NSTX-U are presented in section 3.9 followed by an extensive outlook into the future of SAMI in section 3.10.

The work presented here was a team effort and some of it was previously published[111], [114]–[116]. In particular the development of the 2D Doppler back-scattering technique was mainly conducted by David Thomas as part of his PhD project. The main focus of this thesis was the development of the diagnostic’s hardware and firmware, the investigation of near-field inversion routines as well as the integration into the NSTX-U systems.

3.2 Microwaves In Fusion Plasmas

In this section the fundamentals underlying the SAMI diagnostic are detailed. First the dispersion relation derived in section 2.2 is expanded for the case of microwaves. This will cover the principal solutions of propagation perpendicular and parallel to the magnetic field. Subsequently the back-scattering of waves on plasma surfaces is covered.

It shall be noted, that, although the SAMI diagnostic was originally designed to investigate electron-Bernstein wave (EBW) physics, the focus of this thesis was only marginally related to this field. Consequently the theory will be limited to the cold plasma approximation, which is sufficient for 2D DBS. For a further detail including the kinetic description of a hot plasma the reader is kindly referred to literature[110], [117], [118].

3.2.1 Generalized Cut-Off & Resonance

The visible light approximation used for the interferometer in section 2.2.2 is not generally applicable for microwaves. In most cases the wave frequency ω is close to -if not less than- the electron cyclotron frequency ω_{ce} . In consequence the propagation of waves is direction and polarization dependent. The basis for a description of the wave propagation are the generalized equations derived in section 2.2.1. However, the formalism in chapter 2 neglected the ions’ response to the plasma. This was acceptable in the interferometer’s high frequency approximation, however for the microwave case the ions’ contribution must be recognized. For this we first solve the equation of motion for a general charge carrier, thus rewriting equation 2.5 as

$$\begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = \frac{iq_s}{\omega m_s} \begin{pmatrix} \frac{\omega^2}{\omega^2 - \omega_{cs}^2} & -i \frac{\omega_{cs}\omega}{\omega^2 - \omega_{cs}^2} \text{sgn}(q_s) & 0 \\ i \frac{\omega_{cs}\omega}{\omega^2 - \omega_{cs}^2} \text{sgn}(q_s) & \frac{\omega^2}{\omega^2 - \omega_{cs}^2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix}. \quad (3.1)$$

The subscript 's' identifies any free charge-carrying particle species in the plasma and 'sgn' the signum function. As before we can compare equation 3.1 with the generalized Ohm's law, which contains the sum over all species' contribution, and arrive at a description for the conductivity tensor. The subsequent derivation of the generalized dispersion relation is identical to section 2.2.1. The only difference is found in the Stix parameters, which now read

$$\begin{aligned}
S &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2 - \omega_{cs}^2} \\
D &= \sum_s \operatorname{sgn}(q_s) \frac{\omega_{ps}^2}{\omega^2 - \omega_{cs}^2} \frac{\omega_{cs}}{\omega} \\
P &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2} \\
R = S + D &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2} \left(\frac{\omega}{\omega + \omega_s} \right) \\
L = S - D &= 1 - \sum_s \frac{\omega_{ps}^2}{\omega^2} \left(\frac{\omega}{\omega - \omega_s} \right)
\end{aligned}$$

In the above description $\omega_{ps} = \frac{q_s^2 n_s}{\epsilon_0 m_s}$ and $\omega_{cs} = \frac{|q_s| B}{m_s}$ are the species' plasma and cyclotron frequency respectively. Equation 2.8 has not changed and for simplicity we shall assume from here on that only electrons and a single ion species is present in the plasma.

To describe the behaviour of microwaves in a plasma the generalized dispersion relation can first be used to formulate a generalized resonance condition. For a wave resonating with the medium it propagates, in the refractive index goes to infinity, i.e. $N^2 \rightarrow \infty$, resulting in energy being transferred from the wave to the medium or vice versa. In eqn. 2.8 this forces A to approach zero at a quadratic rate and hence

$$\begin{aligned}
S \sin^2 \psi + P \cos^2 \psi &= 0 \\
\tan^2 \psi &= -\frac{P}{S}.
\end{aligned} \tag{3.2}$$

A similarly generalized criterion can be found for the cut-off condition. Here the refractive index approaches zero, i.e. $N \rightarrow 0$, resulting in the wave being unable to propagate. In this case C is the only relevant term in 2.8, and thus

$$PRL = 0. \tag{3.3}$$

A wave approaching the cut-off surface will generally be reflected, although effects such as mode conversion and tunnelling can occur[110], [117]. As will be seen later beam refraction is an important factor in 2D DBS.

3.2.2 Propagation Along B

The first class of waves we consider propagate along the magnetic field. In this case $\psi = 0$ and thus the solution to the wave equation is

$$P(N^2 - R)(N^2 - L) = 0.$$

This yields 3 solutions. For $P = 0$ the plasma oscillation at the plasma frequency ω_{pl} is found. It essentially describes the plasma's lattice vibration. For $N^2 = R$ ($N^2 = L$) the dispersion for right (left) hand circularly polarized waves is obtained.

Their resonances are found using 3.2 and are the electron cyclotron frequency ω_{ce} and the ion cyclotron frequency ω_{ci} for the respective co-rotating wave. Since the resonance generally involves energy being transferred between the propagating wave and the respective particle, waves at the cyclotron frequency or harmonics thereof are very efficient at heating the plasma and hence have been employed in many fusion devices[15]. The respective heating schemes are ion cyclotron resonance heating (ICRH) and electron cyclotron resonance heating (ECRH).

The waves' cut-off can be found from equation 3.3. For a cold plasma the solution is generally split up into a high and a low density approximation. For the R-wave this is found to be

$$\omega_{cR} = \left\{ \begin{array}{ll} \omega_{ce} \left(1 + \frac{\omega_{pe}^2}{\omega_{ce}^2} \right), & : \omega_{pe} \ll \omega_{ce} \text{ (low density),} \\ \omega_{pe} + \frac{\omega_{ce}}{2}, & : \omega_{pe} \gg \omega_{ce} \text{ (high density)} \end{array} \right\}.$$

Similarly the L-wave has its cutoff at the respective ion frequency, i.e.

$$\omega_{cL} = \left\{ \begin{array}{ll} \omega_{ci} \left(1 + \frac{\omega_{pi}^2}{\omega_{ci}^2} \right), & : \omega_{pi} \ll \omega_{ci} \text{ (low density),} \\ \omega_{pi} + \frac{\omega_{ci}}{2}, & : \omega_{pi} \gg \omega_{ci} \text{ (high density)} \end{array} \right\}.$$

3.2.3 Propagation Perpendicular To B

From SAMI's perspective the wave propagation perpendicular to the magnetic field is the more interesting one, since it is the direction in which the SAMI array points towards the plasma. With $\psi = \frac{\pi}{2}$ the solution to the wave equation 2.9 yields

$$(SN^2 - RL)(N^2 - P) = 0,$$

and thus gives 2 solutions:

$N^2 = P$ describes the so called ordinary wave (O-mode) with the wave polarized parallel to the magnetic field. Since the electrons following the electric field vector are unhindered by the magnetic field they propagate according to the high-frequency approximation derived for the interferometer. As already described in ch. 2 the density cut-off is the (electron) plasma frequency ω_{pe} and no resonance exists.

The second solution ($N^2 = \frac{RL}{S}$) is the extraordinary wave (X-mode). This wave has an elliptical polarization, which is partially longitudinal and partially perpendicular to the magnetic field[117]. It has two resonances at the upper and the lower hybrid frequency :

$$\omega_{UH}^2 = \omega_{pe}^2 + \omega_{ce}^2 \quad (3.4)$$

$$\omega_{LH}^2 = \omega_{ce}\omega_{ci} \left(\frac{\omega_{pe}^2 + \omega_{ce}\omega_{ci}}{\omega_{pe}^2 + \omega_{ce}^2} \right) \quad (3.5)$$

The X-mode resonance is of particular importance for SAMI, since the O-mode can convert into a X-mode and vice versa at the upper hybrid frequency[110]. The mode conversion process has been studied by the SAMI team with the aim of using the inherent angular anisotropy in the mode conversion process to derive the magnetic pitch angle[109], [110].

Since $N_{X\text{-mode}}^2$ is directly proportional to both R and L the X-mode experiences the cut-off of both the R and the L wave derived in the previous section.

3.2.4 Backscattering Of Waves On Plasma Surfaces

A wave propagating in the direction of increasing density will eventually arrive at the cut-off density at which point it will generally refract in direction of lower density and/or get reflected, neglecting finite scale effects. However, a plasma's iso-density surface is not a flat hard surface. The density corrugation along the path of propagation is equivalent to a modulation of the refractive index, which results in scattering.

The process of microwave scattering on a plasma surface can be described by energy and momentum conservation[119], i.e.

$$\omega = \omega_i - \omega_s \quad \text{and} \quad \underline{\kappa} = \underline{k}_i - \underline{k}_s.$$

In this description ω and κ hold the temporal and spatial information of the corrugation the wave scatters off. $k_{i,s}$ and $\omega_{i,s}$ describe the incident and scattered microwave's wavenumber and frequency. The process of scattering can be assumed

to have no influence on the wave's wavelength and therefore $|k_i| = |k_s| = K$ [119]. This enables the reduction of the problem into a single plane via

$$o\kappa = K (\sin \theta_i - \sin \theta_s),$$

with θ denoting the waves' respective angles measured from the scattering surface's normal and $o \in \mathbb{N}$ indicating the order of the scattered maximum for either forward ($o > 0$) or backward ($o < 0$) scattering. For back-scattering the scattered wave has to return in the direction of the incident wave, or $\theta_s = -\theta_i$, and thus

$$o\kappa \approx 2K_{\text{bs}} \sin \theta_i. \quad (3.6)$$

If the surface being scattered off is moving the scattered wave will experience a lab-frame Doppler shift of $\Delta\omega$ with

$$\Delta\omega = \underline{v} \cdot \underline{K}, \quad (3.7)$$

where \underline{v} is the velocity of the corrugation. The wave vector \underline{K} is taken at the scattering location and thus incorporates the local plasma dispersion.

This effect can be used to investigate the velocity distribution of turbulent structures at the cut-off layer and is known as Doppler back-scattering (DBS). DBS has been employed on various machines such as ASDEX Upgrade (AUG)[119], W7-AS[120] and MAST[121] as a tool to study edge turbulence. This is done by aligning an antenna so that it emits radio waves perpendicular to the magnetic field. The back- and forward-scattered radiation can be measured using the same or a second antenna. The spectral composition of the received signal can then be mapped to a velocity spectrum of the turbulence.

The 2D DBS technique developed for SAMI is an extension on that very principle using the inherent 2D capabilities of the SAMI diagnostic. However, to understand it one has to first understand how SAMI images the plasma.

3.3 Synthetic Apertures Image Formation

The SAMI diagnostic employs the synthetic aperture principle. This type of imaging uses the spatial information contained in the phase of the incoming radiation to form a 2D image. The phase is measured using an array of phase-sensitive detectors. Assuming that the source is spatially incoherent, i.e. that constructive interference can only occur from a single location, the source's intensity distribution can be recovered by summation with a relative phase offset.

The recovery of this spatial information can be conducted in multiple ways. The SAMI diagnostic has so far implemented the far-field Fourier inversion and beam-forming technique, both of which will be detailed first. The underlying assumption in both is a plane-parallel incident wave. In addition more complex algorithms can be employed, which incorporate near-field effects. Each of the methods relevant to SAMI will be described here.

3.3.1 Fourier Inversion

The most common synthetic aperture inversion algorithm is based on Fourier inversion due to its robustness and relatively low computational expense. Considering a setup such as depicted in figure 3.3 the signal at an antenna i coming from a distributed source at distance r can be described as

$$s_i(\underline{r}_i, \underline{k}, t) = \int \frac{1}{|\underline{r}_i|} \left(\underline{E}_s(\underline{r}_i, t) \cdot \underline{g}_i(\underline{r}_i, \underline{k}) \right) e^{ik \cdot \underline{r}_i} d\underline{r}_i^3. \quad (3.8)$$

Here \underline{E}_s is the electric field oscillation at the source, \underline{k} is the incident wave's wave vector and g_i is the antenna's voltage response function, which depends on frequency and polarization of the incident wave. The distance to the source can be defined with respect to some predetermined reference, e.g. the center of the antenna array, or $\underline{r}_i = \underline{r} - \underline{a}_i$. Additionally one assumes that the wave is plane-parallel, i.e. $\underline{r}_i \cdot \underline{k} \approx \hat{\alpha} \cdot \underline{a}_i$, where $\hat{\alpha}$ is the unit vector in direction of the source, and \underline{a}_i the location vector of antenna i . This far-field (FF) approximation yields

$$\begin{aligned} s_i(\hat{\alpha}, \underline{a}_i, t) &= \int \left(\underline{E}_s(\hat{\alpha}, t) \underline{g}_i(\hat{\alpha}) \right) e^{ik(\hat{\alpha} \cdot \underline{a}_i)} d\hat{\alpha} \\ &= \int E(\eta, \zeta, t) \hat{g}_i(\eta, \zeta) e^{ik(a_{i,y}\eta + a_{i,z}\zeta)} d\eta d\zeta, \end{aligned} \quad (3.9)$$

where the problem was transferred to a more convenient coordinate system using $\eta = \sin \phi$ and $\zeta = \cos \phi \sin \theta$. It is assumed that the 2D antenna array is positioned in the y - z -plane, which is in accordance with the SAMI setup. The \hat{g} denotes the response function including the polarization mismatch in the projection of \underline{E} onto antenna i . Because of the FF approximation the $1/r$ -amplitude fall-off can be assumed identical for all antennae and is therefore neglected.

Equation 3.9 can be identified as a Fourier transformation, i.e.

$$s_i \left(\frac{a_{y,i}}{\lambda}, \frac{a_{z,i}}{\lambda}, t \right) = \mathcal{FT} \{ E(\eta, \zeta, t) \hat{g}_i(\eta, \zeta) \},$$

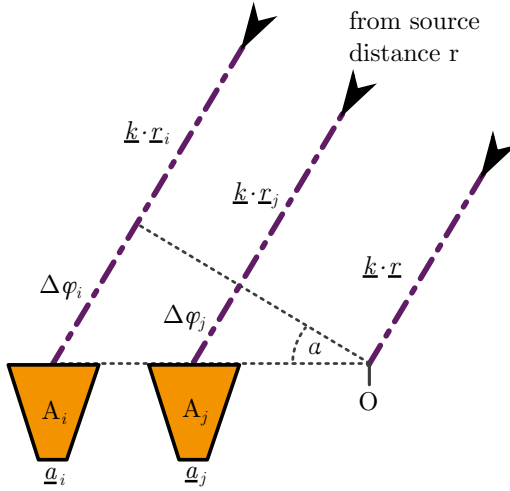


Figure 3.3 – Basic geometry behind the FF synthetic aperture principle.

and consequently the inversion can be calculated with

$$E(\eta, \zeta, t) \hat{g}_i(\eta, \zeta) \approx \sum_{i=0}^{N_a} s_i e^{-ik(a_{i,y}\eta + a_{i,z}\zeta)}. \quad (3.10)$$

Equation 3.10 has been discretized assuming an array of N_a antennae and hence is an estimate of the complex field at the source given a collection of point samples with the entirety of the information contained in the phase. A handy tool for analyzing phase relationships is the cross correlation

$$K_{ij}(t) = \int_{-\infty}^{+\infty} s_i(t) s_j^*(t - \tau) d\tau = \langle s_i s_j \rangle_t, \quad (3.11)$$

where * denotes the complex conjugate. Applying equation 3.9 enables us to find a description for the intensity distribution in real space:

$$\begin{aligned} K_{ij} &= \left\langle \int E(\eta, \zeta) \hat{g}_i(\eta, \zeta) e^{ik(a_{i,y}\eta + a_{i,z}\zeta)} d\eta d\zeta \right. \\ &\quad \left. \int E^*(\eta', \zeta') \hat{g}_j^*(\eta', \zeta') e^{-ik(a_{j,y}\eta' + a_{j,z}\zeta')} d\eta' d\zeta' \right\rangle_t \\ &= \int \int \langle E(\eta, \zeta) E^*(\eta', \zeta') \rangle_t \hat{g}_i(\eta, \zeta) \hat{g}_j^*(\eta', \zeta') e^{ik(a_{i,y}\eta + a_{i,z}\zeta - a_{j,y}\eta' - a_{j,z}\zeta')} d\eta d\zeta d\eta' d\zeta'. \end{aligned} \quad (3.12)$$

In the above equation the FF approximation has already been applied. In addition one may now assume the source to be spatially incoherent[122]. The integral may thus be solved assuming that the cross-correlation vanishes for every point - or in this case direction - except for the one under observation, or

$$\langle E(\eta, \zeta, t)E^*(\eta', \zeta', t) \rangle_t = I(\eta', \zeta')\delta(\eta - \eta')\delta(\zeta - \zeta').$$

This yields the van Cittert-Zernicke Theorem[122]:

$$K_{ij}(t) = \int I(\eta, \zeta, t)G_{ij}(\eta, \zeta)e^{i\frac{2\pi}{\lambda}(\sigma_{ij,x}\eta + \sigma_{ij,z}\zeta)}d\eta d\zeta. \quad (3.13)$$

Here we have defined $\underline{\sigma}_{ij} = \underline{a}_i - \underline{a}_j$ as the baseline vector running between antenna i and j . As before eqn. 3.13 can be identified as a two-dimensional Fourier transform, this time of the intensity distribution I multiplied by the antenna baseline gain pattern G_{ij} . Via inversion one can thus finally acquire an estimate for the desired quantity :

$$I_{\text{FI}}(\eta, \zeta, t)G_{ij}(\eta, \zeta) \approx \sum_i \sum_j K_{ij}(t)e^{-i\frac{2\pi}{\lambda}(\sigma_{y,ij}\eta + \sigma_{z,ij}\zeta)}. \quad (3.14)$$

Here we have already discretized the equation. Since a real system would ultimately have a limited set of antennae one only acquires an estimate for the true intensity distribution. The Fourier inversion algorithm is heavily used in radio astronomy, where the FF approximation is given naturally[123]. With SAMI this technique has been predominantly used for the passive EBW studies[110].

3.3.2 Beamforming

Beamforming (BF) was originally developed for ultra-sound body scanning. Nowadays the technique is heavily employed in many fields, such as satellite radar imaging[124], [125]. The idea is similar to that of Fourier-inversion, i.e. assuming spatial incoherence of the source results in destructive interference in every direction but the desired one. The selection of direction is conducted by applying a phase shift to each antenna signal corresponding to the constructive interference condition for that direction:

$$s_{\text{B}}(\underline{\alpha}, f) = \sum_{j=1}^{N_a} W_j(\underline{\alpha}, f)s_j e^{i\frac{2\pi f}{c_0}\underline{\sigma}_{1j}\underline{\alpha}}. \quad (3.15)$$

As can be seen the reference is taken to be antenna 1 and we have again assumed the FF approximation as per the Fourier inversion. W_j is a scaling factor, which will generally be related to the response function used before. The function of equation 3.15 is to suppress every signal not coming from the selected direction and thus the intensity at the given point can be calculated using

$$I_{\text{BF}}(\underline{\alpha}) = \int_{-f_{\text{Nyq}}}^{f_{\text{Nyq}}} |\mathcal{FT}_{\eta\psi}^{-1} \{s_{\text{B}}(\underline{\alpha}, f)\}|^2 df, \quad (3.16)$$

where the digitization discretization has been taken into account by integration within the Nyquist bounds. The above equation can be easily evaluated for any direction $\underline{\alpha}$. The BF technique is remarkably simple, but at the same time quite computationally expensive. For SAMI this technique has been found to perform better under certain conditions than the Fourier inversion. Most notably the 2D DBS technique entirely relies on this image formation algorithm[111].

The FF approximation employed so far may not necessarily be applicable in fusion applications. The transition between the FF and the near-field (NF) is commonly assumed to be the point at which the wave curvature can no longer be disregarded and the source distribution has to be described using Fresnel diffraction[123]. In Gaussian optics this boundary is described by the Rayleigh distance

$$R = \frac{2L^2}{\lambda}, \quad (3.17)$$

with L being the array's defining scale length (usually the largest baseline) and λ the measured wave's wavelength. Using Gaussian optics in this context is only marginally justified. Using BF as an active measure to create a Gaussian beam is possible and thus the same must be true in reverse via reciprocity, i.e. one can assume that the beam selected via superposition of the beams has the properties of a Gaussian beam. However, the emission pattern of the Vivaldi antennae used in the SAMI array are non-Gaussian, which translates to the formed image via G_{ij} and W_j in equations 3.14 and 3.16. However, for the cases considered here this measure is the best guess.

In section 3.8 it is shown that the NF regime will definitely have to be considered for SAMI. So far the image was in essence formed for the entire hemisphere, being only a function of viewing angle. In the NF the wave curvature requires the amplitude fall-off to be considered for each antenna separately and is not negligible. Also, since the antenna response is anisotropic it may not be disregarded. This requires that the source's location is known, which in itself is a very complex question not discussed in detail here. Fortunately the simplicity of the BF technique makes it very easy to implement a NF technique by using the original expression for the antenna signal (eqn. 3.8). The spatial incoherence assumption still holds in the NF, so that one only has to revise the signal beam:

$$s_B(\underline{r}_{\text{pt}}, f) = \sum_{j=1}^{N_a} \frac{W_j(\underline{r}_j - \underline{r}_{\text{pt}})}{|\underline{r}_j|} s_j e^{i \frac{2\pi f}{c_0} |\underline{r}_j - \underline{r}_{\text{pt}}|}. \quad (3.18)$$

Here $\underline{r}_{\text{pt}}$ is the location of the point under observation, which has been assumed to emit spherically. The rest of the evaluation is identical to the FF case. NF BF is a common problem in acoustic applications, where the exact source is not previously known and unwanted noise is to be suppressed[126]. In such situations the weighting factor W_j will often have to be adapted to account for the unknowns. In the case of a fusion plasma the source location can be located using equilibrium reconstruction codes such as EFIT[127], [128]. These do however have an inherent level of uncertainty around the applicability of the equilibrium assumption, which is enhanced by the generally “soft” nature of the emission surface.

3.3.3 Matrix Inversion

As has been pointed out, the BF algorithm, although very simple and effective, is very computationally expensive. This was the original motivation to use the Fourier inversion on SAMI. Another alternative method of forming an image is the inversion of a forward model. For this one first goes back to eqn. 3.8, which is the complete description of the antenna signal. As before the interesting quantity is the intensity distribution, for which we will again calculate the cross-correlation :

$$\begin{aligned} K_{ij}(t) &= \int \frac{1}{|\underline{r}_i(\underline{r}_{\text{pt}})|} \frac{\Omega_i(\underline{r}_{\text{pt}})}{2\pi} \left(\underline{E}(\underline{r}_{\text{pt}}) \underline{g}_i(\underline{r}_{\text{pt}}) \right) e^{ik(\underline{r}_{\text{pt}}) \cdot \underline{r}_i(\underline{r}_{\text{pt}})} \\ &\quad \frac{1}{|\underline{r}_j(\underline{r}_{\text{pt}})|} \frac{\Omega_j(\underline{r}_{\text{pt}})}{2\pi} \left(\underline{E}^*(\underline{r}_{\text{pt}}) \underline{g}_j^*(\underline{r}_{\text{pt}}) \right) e^{-ik(\underline{r}_{\text{pt}}) \cdot \underline{r}_j(\underline{r}_{\text{pt}})} d\underline{r}_{\text{pt}}^3 \\ &= \int \frac{d\underline{r}_{\text{pt}}^3}{r_i r_j} \frac{\Omega_i \Omega_j}{4\pi^2} \hat{G}_i \hat{G}_j e^{ik(r_i - r_j)}. \end{aligned} \quad (3.19)$$

This is a full description of the wave propagation from source to array baseline with Ω_i describing the solid angle the antenna occupies in the emission hemisphere of the source. As before, the spatial incoherence condition has been applied. It is to be noted, that all quantities in eqn. 3.19 are functions of the source location $\underline{r}_{\text{pt}}$. The integral in the equation can be discretized by assuming that the emission surface is a 2D structure, which is reasonable when imaging a plasma surface:

$$K_{ij}(t) = \sum_{n=0}^{N_{\text{pts}}} \frac{I_n}{r_{in}r_{jn}} \frac{\Omega_{in}\Omega_{jn}}{4\pi^2} \hat{G}_{in}\hat{G}_{jn} e^{ik(r_{in}-r_{jn})} \Delta A_n. \quad (3.20)$$

Here we have assumed that the grid used on the 2D surface has $N_{\text{pts}} = X$ points and each point represents a surface of area ΔA_n . There are $N_{\text{bl}} = \frac{N_a}{2}(N_a - 1) = Y$ unique baselines holding non-redundant information in an array of N_a antennae. If an index m is assigned to each unique baseline (i, j) , we can describe equation 3.20 using a matrix equation :

$$K_m(t) = \sum_{n=0}^X I_n \Gamma_{mn} \hat{G}_{mn} e^{ik\Delta r_{mn}} \Delta A_n$$

$$\begin{pmatrix} K_0(t) \\ \vdots \\ K_Y(t) \end{pmatrix} = \begin{pmatrix} \Gamma_{00}\hat{G}_{00}e^{ik\sigma_{00}}\Delta A_0 & \dots & \Gamma_{0X}\hat{G}_{0X}e^{ik\Delta r_{0X}}\Delta A_X \\ \vdots & \ddots & \vdots \\ \Gamma_{Y0}\hat{G}_{Y0}e^{ik\sigma_{Y0}}\Delta A_0 & \dots & \Gamma_{YX}\hat{G}_{YX}e^{ik\Delta r_{YX}}\Delta A_X \end{pmatrix} \cdot \begin{pmatrix} I_0 \\ \vdots \\ I_X \end{pmatrix}$$

$$K_m(t) = P_{mn} I_n \quad (3.21)$$

Inside the matrix elements we have substituted $\Gamma_{mn} = \frac{\Omega_{in}\Omega_{jn}}{r_{in}r_{jn}4\pi^2}$. The image on the model surface is described by the vector \underline{I} . To calculate it we have to apply the inverse of the propagation matrix \underline{P} to the cross-correlation vector \underline{K} . However, the propagation matrix \underline{P} will generally be highly non-square and thus a unique inverse cannot be calculated.

However, a least squares approximation of the inverse can be obtained by reformulating the inversion as a minimization problem[129], [130]. In the present case we minimize the intensity distribution across the visible model space given the propagation constraints, i.e.

$$\min \|\underline{I} \cdot \underline{I}^{\text{H}}\| \quad \text{subject to} \quad \underline{P} \cdot \underline{I} = \underline{K},$$

where superscript-H denotes the conjugate transpose of the matrix. This kind of minimization problem is well known and can be solved using Lagrangian multipliers. The result is the Moore-Penrose Pseudo Inverse[130], [131]:

$$\underline{I} = \underline{P}^{\text{H}} (\underline{P} \cdot \underline{P}^{\text{H}})^{-1} \underline{K}. \quad (3.22)$$

This generalized inverse is the least-squares solution of the ill-posed problem

defined by 3.21. It always exists and thus will always give a solution[131]. To calculate the inverse methods such as singular value decomposition (SVD) can be used, for which most programming languages have readily available solutions - e.g. the Numpy library provides the “pinv” routine via the linear algebra sub-module (numpy.linalg.pinv).

The matrix inversion method, apart from being a NF solution, can potentially improve the computational effort of reconstructing an image. Although the calculation of the propagation matrix and subsequent inversion is an order of magnitude slower than the Fourier-inversion, it is still an order of magnitude faster than the beam-forming algorithm. An additional benefit is that the matrix pseudo-inverse only needs to be calculated once for a given geometry. Therefore only the matrix multiplication has to be calculated as long as the emission surface’s geometry does not change significantly. Matrix multiplications have been implemented very efficiently on many architectures, including FPGAs.

Having established the methods of forming an image using synthetic aperture data, we can now describe how 2D DBS is conducted to acquire magnetic pitch angle.

3.4 2D Doppler Back-Scattering

In section 3.3.2 we have established the basic theory behind DBS. Using a synthetic aperture approach this technique can be used to derive 2D information about the back-scattered signal. In the paragraphs leading up to 3.4.3 this section will show how this enables the derivation of the magnetic pitch angle on the back-scattering surface. To describe this method we will first have to explain normal Doppler reflectometry before making use of the 2D imaging capabilities to derive magnetic pitch angle.

3.4.1 Doppler Reflectometry

Reflectometry is the process of sending a microwave signal, which is often chosen to sweep in frequency, in a direction normal to the plasma’s cut-off surface. A parameter that can be derived with this is the density profile; usually by measuring the time-of-flight of the cut-off reflected signal. Doppler reflectometry was developed by accident from normal reflectometry while investigating an effect called “phase runaway”, which was attributed to the misalignment of the antennae and the resulting oblique incidence[119]. It was later found that this could be utilized as a tool to measure plasma rotation, which resulted in deliberate misalignment at fixed frequencies.

As shown in figure 3.4 Doppler reflectometers can be monostatic (1 antenna

to transmit & receive) or bi-static (separate transmit & receive antenna), often depending on whether forward- or back-scattering is to be investigated and which peak order is desired. As will be seen in section 3.5, the SAMI diagnostic is a bi-static setup looking at the 0th order back-scattered signal.

3.4.2 Back-Scattering Efficiency

As has been pointed out before, a wave propagating in direction of the cut-off will generally get reflected neglecting any finite scale effects occurring close to the cut-off. At an oblique angle the wave propagating towards the cut-off surface will see a density corrugation along its propagation path, which, as described in sec. 3.2.4, will result in back-scattering to occur due to the modulation of the refractive index.

In a fusion plasma a source of density corrugations comes in the form of turbulent eddies, which are omni-present even in a H-mode edge plasma. Generally turbulent structures are elongated along magnetic field lines, since charged particles can essentially travel freely along them whilst being severely limited in their motion perpendicular to them[119]. For back-scattering of $\underline{K} = (k_{\perp}, k_{\parallel})$ on the corrugation of $\underline{\kappa}$ only the parallel components contribute. As a consequence a back-scattered signal can only be expected from the wave component k_{\perp} perpendicular to the magnetic field, since $\underline{\kappa}$ is aligned as such.

Along the path of propagation the back-scattering location is highly localized. As has been established, back-scattering occurs when the propagation vector and the density corrugation vector are situated in one plane. The amplitude of the turbulence causing the corrugation scales with its scale length as κ^{-3} [132]. In addition the back-scattering process has an efficiency that scales with κ^{-2} [133]. As the wave propagates towards the cut-off and approaches the density surface K decreases until the back-scattering condition 3.6 is met. This high order dependence strongly localizes the back-scattering location to the point of minimum K , which is roughly the density cut-off surface. Generally beam divergence becomes more pronounced as the beam's angle of incidence becomes more oblique, which moves the scattering




Figure 3.4 – Basic principle behind Doppler back-scattering on a density corrugation of scale κ for monostatic and bistatic setups. The motion of the plasma corrugation results in the scattered signal being Doppler shifted.

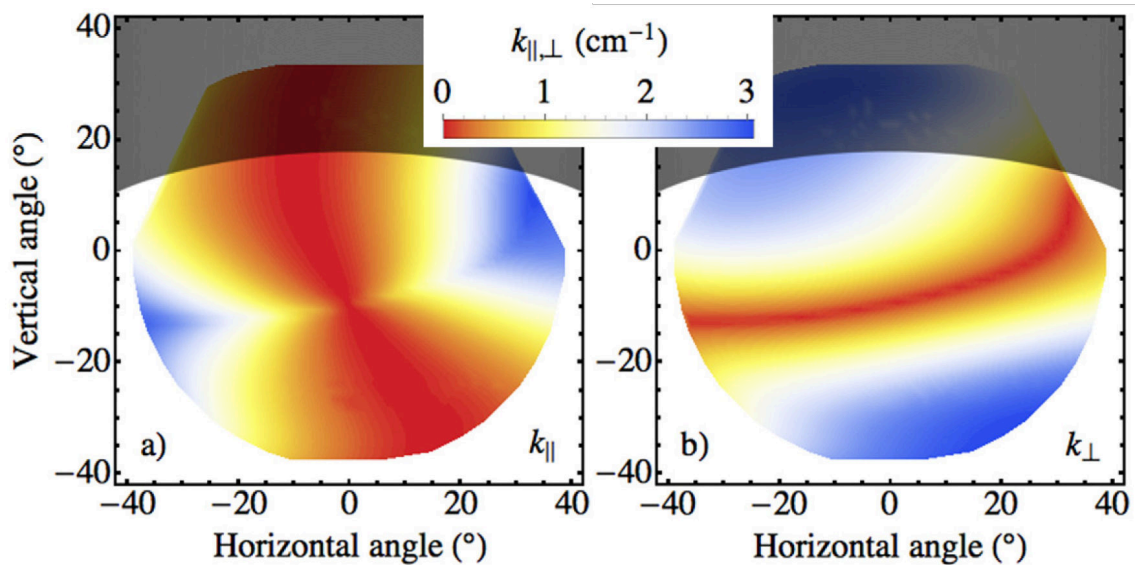


Figure 3.5 – Wavenumber component map at the cut-off surface of a 16 GHz microwave launched omni-directionally from (0,0) for MAST shot #27969 230 ms into the shot. The left plot depicts the magnitude of the component parallel to the local magnetic field. The right one shows the amplitude for the component perpendicular to the local magnetic field. The map was calculated using the TORBEAM code with measurements from TS. The shaded area at the top marks the position of a MAST poloidal field coil obscuring the view (adapted from [111] with kind permission from the author and AIP publishing).

location further away from the normal-incidence cut-off.

Figure 3.5 on page 102 shows the wave vector components along the magnetic field ($k_{||}$) and parallel to it (k_{\perp}) at the density cut-off of an omni-directional 16 GHz microwave beam launched from (0,0). The map was calculated by the TORBEAM code with the help of TS for MAST shot #27969[111], [134]. This weakly-relativistic beam-tracing code launches a beam in direction of the plasma and propagates it making use of the toroidal symmetry. The beam refraction is calculated using temperature and density profiles from TS. Unlike a ray-tracer a beam includes diffraction effects as it incorporates the spatial extent of the beam. The sample for each point in figure 3.5 is taken at the point of minimum K along the calculated path of propagation. The back-scattering condition requires that the k_{\perp} component is maximized, i.e. that the microwave is propagating mostly perpendicular to the magnetic field (blue regions in fig. 3.5b). At the same time the back-scattering efficiency argument requires that the absolute wavenumber $K = \sqrt{k_{\perp}^2 + k_{||}^2}$ is minimized. As can be seen, the efficiency will thus be highest along a line perpendicular to the magnetic field for a launch angle of about 20° from the normal. While this result is calculated for the SAMI diagnostic, the result is nonetheless applicable to any monostatic DBS setup.




Figure 3.6 – The schematic of Doppler back-scattered spectral power around the probing frequency f_{probe} (shifted to 0 MHz) in two opposing directions. In a) the Doppler shift results in more red-shifted power, in b) (looking in the opposing direction) more blue shifted power is observable (adapted from [111] with kind permission from the author and AIP publishing).

3.4.3 Deriving Magnetic Pitch Angle

A normal DBS system will use a highly directional beam only illuminating a single point in the plasma[119]. To receive a 2D image therefore requires mechanical or quasi-optical steering of the probing beam, which is very slow, requiring seconds to generate a single 2D image. Using the synthetic aperture approach one can acquire a 2D DBS image in one go by illuminating the entire plasma using an omni-directional probing beam. The back-scattered signal can be spatially resolved using the image formation techniques discussed in section 3.3. It is also possible to get a spatially resolved measurement of the spectrum by omitting the integral over frequency space in equation 3.16.

As has been pointed out, the back-scattered signal is Doppler shifted according to the local motion of the turbulent structures. This Doppler shift has two contributions[121]:

$$\underline{v} = \underline{V}_0 + \tilde{v}.$$

The turbulent phase velocity \tilde{v} results in a simple broadening of the back-scattered signal, since the phase at the time of scattering is random. However, the second contribution due to the movement of the underlying structure \underline{V}_0 is directional[119]. Depending on the operating scenario this can be mainly $\underline{E} \times \underline{B}$ -drift or pressure-force driven. Due to its direction-dependence it is used by normal DBS diagnostics to derive rotation speed via the induced bump in the back-scattered spectrum.

SAMI cannot discern this Doppler peak directly. This is due to the incomplete directional separation as a result of the number of antennae in the acquisition array so that the peak is “hidden” in the Doppler-broadened direct reflection signal[111]. However, even with no clear bump from the Doppler shifted peak visible, the contribution of the peak can still be seen in the received Doppler-broadened signal as




Figure 3.7 – 2D red-blue imbalance map for MAST shot #27969. The color-bar is given in units of volts. The data is generated using the SAMI diagnostic for a 16 GHz omni-directional probing beam 300 ms into the shot. The over-plotted curved lines indicate the magnetic field lines taken from MSE constrained EFIT. The EFIT location was matched against the SAMI probing frequency with the help of TORBEAM using TS data (taken from [111]).

shown in figure 3.6. In this image the spectral power received by a synthetic aperture array around the probing frequency f_{probe} is plotted for two directions. In a) the back-scattered signal is selected from a location, where \underline{V}_0 is expected to result in red shift, in b) the direction is chosen to result in blue shift. In both cases the effect is a subtle but obvious asymmetry around the peak. If one takes the difference of integrated spectral power above and below the probing frequency, i.e.

$$\delta P_{\text{RB}}(\underline{r}) = \int_{-\infty}^{f_{\text{probe}}} P(\underline{r}, f) df - \int_{f_{\text{probe}}}^{\infty} P(\underline{r}, f) df, \quad (3.23)$$

the results have opposing signs for a) and b).

This effect can now be combined with the 2D back-scattering efficiency detailed in the previous section. The surface of a toroidally confined plasma is always rotating. One can therefore expect the Doppler shift resulting from a beam illuminating the plasma to be inverse axisymmetric with only blue-shifted power coming from one direction and only red-shifted power coming from the other. At the same time the back-scattering efficiency has been shown to only occur efficiently perpendicular to the magnetic field.

Figure 3.7 shows a 2D map of δP_{RB} for MAST shot #27969 300 ms into the shot. The probing frequency for this shot was 16 GHz, for which the cut-off is roughly at a normalized radius of .95 on a standard L-mode shot on MAST[111]. The over-plotted curved black lines indicate the magnetic field of the MAST plasma. They have been generated using MSE constrained EFIT. The EFIT location was matched against the SAMI probing frequency using TORBEAM[111]. TS data was used to find the radial location corresponding to the probing frequency again by looking for the minimum K along the path of propagation, at which the EFIT pitch-angle was then calculated. As can be seen the combination of back-scattering efficiency and Doppler shift power

imbalance results in two distinct peaks of back-scattered power imbalance. These are aligned perpendicular to the magnetic field due to the aforementioned co-alignment of the turbulence. By drawing a line between the two maxima an average magnetic pitch angle can be measured.

The difficulty of the technique is in separating the mixed polarizations, which are received by the array. Since the dispersion of O and X mode differs and the cut-off frequencies, i.e. the upper-hybrid frequency ω_{UH} and the plasma frequency ω_p , can be radially separated by 1 cm in a MAST plasma, a superposition of the back-scattering from both polarizations introduces significant errors. In addition the maximum back-scattering may not actually come from the cut-off surface in particular for more oblique directions, where the beam refraction is more severe. It is thus difficult to match the measured pitch angle to the correct flux surface. Lastly an inherent problem is the fact that the technique relies on turbulence in the edge to work. This is naturally given in L-mode. In H-mode however the edge turbulence is suppressed and thus measurements are less reliable. This can be compensated with higher power however.

A problem currently more related to the hardware implementation is the fact that long integration times are required for decent pitch angle measurements. It is assumed that the pitch-angle and the plasma surface itself does not change during that time, which is not necessarily given. The integration time can potentially be reduced by increasing the side-lobe suppression, which requires more antennae in the synthetic aperture and a more demanding digitization system.

The SAMI diagnostic is the first diagnostic to show the feasibility of this technique, which is not limited to spherical tokamaks, but can in theory be applied to any toroidally confined magnetic confinement fusion (MCF) device. The implementation will be described in the next section.

3.5 The SAMI Diagnostic

The Synthetic Aperture Microwave Imaging diagnostic is a first-of-its kind diagnostic that images the plasma using a synthetic aperture. It has been operating on MAST since the beginning of 2011 and has conducted several novel measurements, which were only possible due to its 2D capabilities. In this section the diagnostic will be described including some of the limitations identified in the original design before elaborating on the development work in the subsequent sections.




Figure 3.8 – The SAMI RF box. On the left is the synthetic aperture array using the linearly polarized sinuous antennae (only the antennae not covered in foil are used for imaging). On the right is the RF box containing the hybrid couplers and down-conversion electronics. To minimize RF losses the box is directly attached to the array and bolted on to the machine.

3.5.1 The Analog Front-End

Figure 3.9 shows a schematic of the acquisition hardware. The actual realization of the RF hardware, i.e. the synthetic aperture array and the down-conversion electronics, can be seen in figure 3.8. The SAMI diagnostic uses an array made of 8 antipodal Vivaldi antennae also seen in figure 3.8. This type of antenna is based on printed circuit board (PCB) technology, which makes it very cheap to manufacture. The antennae themselves have a very wide field of view and are extremely broadband[135]. In addition their 3D geometry ensures very good polarization separation, which however can come at the cost of distorting cross-talk effects[110]. The array is arranged in an optimized geometry, which reduces the amplitude of the nearest sidelobe[110]. The two antennae marked at the top of the inlay in figure 3.9 are used for active probing and emit a broadband omni-directional modulated signal at the imaging frequency to illuminate the entire plasma surface for 2D DBS. They are shielded by aluminium foil to reduce the direct cross-talk to the passive acquisition array nearby. Since the cross-talk would be correlated on all antennae it would introduce unacceptable distortions in the image reconstruction. The signals measured by the acquisition array are amplified and passed through a set of hybrid couplers. These act as a 90° phase shifter and thus supply both the I and the Q component for each antenna (16 signals in total), similar to the digital Hilbert filter approach described in chapter 2. This is important as the phase information could otherwise

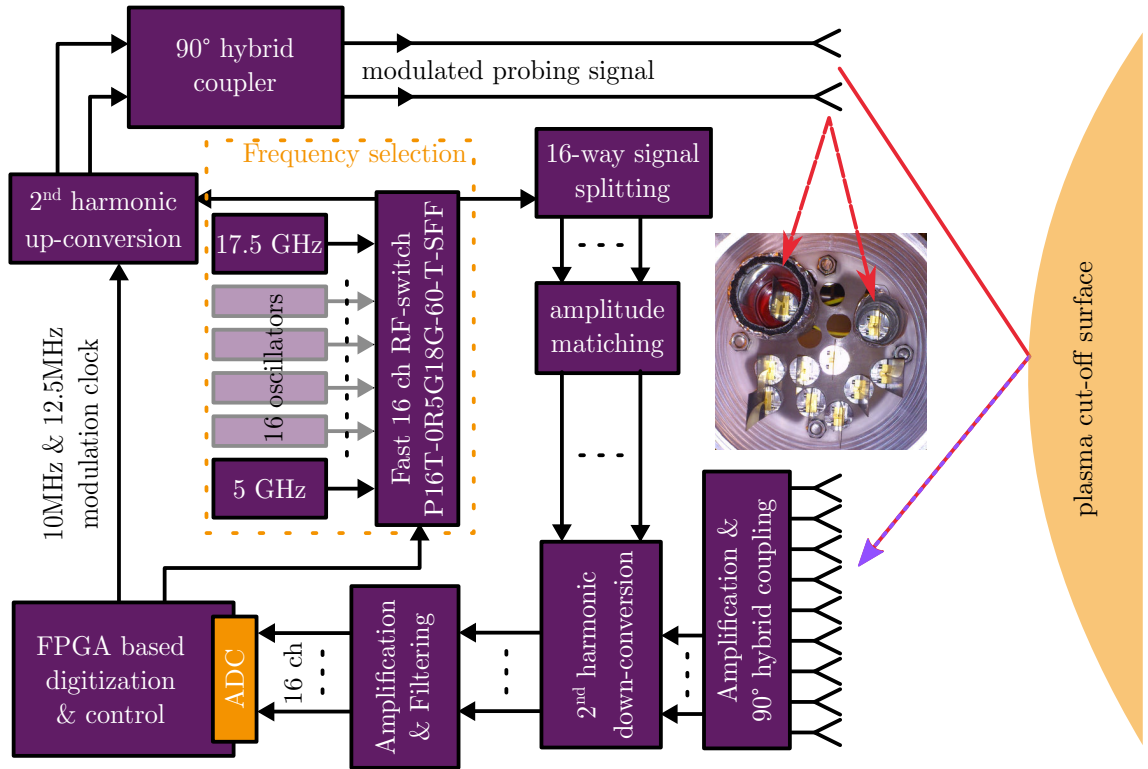


Figure 3.9 – Top level schematic of the diagnostic hardware. The acquisition array is depicted in the inlay. The two shielded emitting antennae used for 2D DBS are marked. The FPGA based acquisition consists of two boards described in section 3.5.2.

not be recovered after the subsequent down-conversion[110].

SAMI images the plasma at 16 discrete frequencies from 10 GHz to 35.5 GHz. The imaging frequency is supplied by a set of 16 dielectric resonator oscillators (DROs), which range from 5 GHz to 17.25 GHz. A particular imaging frequency is selected using a fast RF switch with a very low settling time of 100 ns. The chosen frequency signal is amplified and split into 16 separate signals with identical amplitude. These are used to mix the antennae’s signals via second-harmonic down-converters to direct current (DC). The result is passed on to be amplified and filtered with a 90 MHz low-pass filter to prevent aliasing and a DC-block to prevent biasing. The final digitization is conducted by a FPGA digitizer. The probing signal for 2D DBS cannot be at the imaging frequency due to the necessary DC-block. For this reason two modulation signals are supplied by the FPGA digitizer at 10 MHz and 12.5 MHz. The active probing and down-conversion signal are supplied by the same physical source via a power splitter. The modulation is conducted using the same type of second-harmonic mixer the down-conversion stage uses. Another set of hybrid couplers ensures that the signal sits only in the upper half of the complex acquisition spectrum.

3.5.2 The Digitization Back-End

The SAMI diagnostic relies heavily on FPGA technology. The signals are digitized using 2 Xilinx Virtex-6 ML605 boards equipped with a FMC108 ADC made by 4DSP acquiring 8 channels at 250 MS/s with 14 bit accuracy[109], [112], [113], [136]. The high digitization speed is critical in supplying a high accuracy in the image reconstruction, but requires the handling of a large amount of data. Since the data has to be stored byte-aligned to be useful to the processing PC, this amounts to 8 GB/s for all the antennae's I and Q components. The ML605 board used for the SAMI diagnostic uses DDR3 memory with a 64 bit wide interface and is specified up to 800 MTransfers/s[112]. In consequence a single board could not handle the amount of data generated by two FMC108 cards, which is why two boards are used for digitization. The two FPGAs are stacked and each acquires the I and Q component of 4 out of the 8 antennae resulting in 8 channels per board, for which one FMC108 per board is sufficient.

Because a stable phase relationship across the channels is critical, the data acquisition on both boards must be synchronized. This is ensured via several means:

- The FPGA firmware is identical on each board and uses a MASTER-SLAVE approach. i.e. one board generates all timing relevant signals and supplies them to all digitization boards including itself. The difference in behaviour between MASTER and SLAVE board was originally ensured via switch controlled multiplexer (MUX) elements.
- To make sure the digitization happens at an equal rate the digitization clock is generated on the MASTER board and supplied via on-board SMA connectors to both boards through the FM108s' external clock input.
- The acquisition trigger for both boards is interpreted by the MASTER board, which, via similar mechanisms used for the digitization clock, supplies a synchronized trigger to both boards.

Any remaining timing mismatch between the two boards' acquisition can be compensated by the calibration.

The acquired data is directly streamed into the FPGA's small-outline-dual-inline memory module (SODIMM). This results in a data rate of 4 GB/s. Since the ML605 is capable of supporting SODIMMs up to 2 GB, the acquisition time is limited to a total of just over 500 ms of raw data acquisition. The stored data was downloaded in between shots using a user datagram protocol (UDP) via a private local area network (LAN) connection. The protocol was specifically designed for SAMI[113].

The ethernet communication is controlled via a MicroBlaze central processing unit (CPU) running a Linux specifically built for SAMI[137], [138]. The kernel has been compiled to only use the lowest 120 MB of the FPGA’s memory-space and enables full control of the connected peripherals, including the digitization core, without modifying any data above the 120 MB boundary. Since the data is streamed at such a high rate, the MicroBlaze CPU cannot be given access to the memory interface during data acquisition. The FPGAs thus switch the memory pins from the CPU bus to the digitization core during this stage. The communication required by the CPU to maintain the Linux’s operation is buffered using first-in-first-outs (FIFOs)[137]. In general the MicroBlaze system and the digitization core are two separate components only interacting with each other via general purpose IO (GPIO) and dual-port block RAM (BRAM) elements. Shared components like the RAM use MUX elements to enable/disable access interfaces when necessary.

3.5.3 Diagnostic Limitations

The original diagnostic had several limitations, which were identified during the course of this work. The most prominent one was the fact that acquisition had to be conducted in a single continuous data stream limited to only 500 ms. Already tokamaks such as the Korean Superconducting Tokamak Advanced Research (KSTAR) experiment operate in the region of several minutes[139]. MAST-Upgrade (MAST-U) and NSTX-U, the experiments SAMI will operate on, will run for several seconds. Eventually fusion machines will have to run continuously and thus diagnostics like SAMI will accordingly have to acquire continuously.

A second limitation of unknown influence is SAMI’s inability to separate polarization components. This makes exact localization of the emission surface difficult and results in signals from both polarizations mixing in the image. Some of the measurement error in the MAST data can be attributed to this[111].

The 2D DBS technique developed from MAST data currently requires about 10 ms of averaging[111]. For profile acquisition, i.e. at multiple interleaved frequencies, the FPGA firmware had a minimum switching rate of 2.8 kHz. As a consequence the maximum time available for the spectral analysis at one probing frequency while switching between frequencies was 357 ms. Using the SAMI diagnostic as it was operated on MAST, i.e. with 8 antennae and single polarization acquisition, this was not enough to measure pitch angle, since the red-blue imbalance described in chapter 3.4.3 cannot be discerned. To circumvent the time limitation, spectral stitching across multiple acquisition intervals was attempted. However, this failed due to the fact that the active probing modulation signal had no stable phase relationship with

the digitization clock. Hence the 2D DBS results acquired so far have all been in single frequency acquisition mode, where no frequency switching is conducted.

During calibration on NSTX-U it was found that the ADC calibration routine and underlying firmware had a flaw, which could render the data useless. In essence the calibration of the data lines transferring the data from the ADC to the FPGA has to be phase matched with the sampling clock, which is also transferred. The calibration routine only calibrated 8 least significant bits (LSBs), but not the most significant bits (MSBs). Fortunately and purely by chance the problem appeared to not be present during the campaign on MAST, probably because the lines did not require calibration with the logic layout of the employed bitstream. However, the change in the firmware components resulted in a change of the physical layout of the logic circuits on the chip, which caused the problem with the calibration to emerge. Ageing and a different climate at Princeton Plasma Physics Laboratory (PPPL) likely added to the problem and required the calibration to be rewritten.

Lastly the ethernet solution for transferring the acquired data at a fast rate from the FPGAs to the acquisition PC was unreliable, therefore occasionally resulting in significantly increased transfer times (up to 2 h). It also depended on a specifically configured router, which could not be replaced in case of a hardware failure and was thus a matter of great concern, as will be seen in section 3.7.

The limitations described here are the result of the novelty of this diagnostic. During the initial development it was not known that these design choices would pose a problem. A major part of the work conducted as part of this project went into mitigating or minimizing these limitations.

3.6 Modifications To The SAMI FPGA Firmware

The limitations pointed out in section 3.5.3 resulted in a set of modifications to the SAMI diagnostic, most notably the SAMI FPGA firmware. To make more use of long shot lengths the acquisition of multiple arbitrarily sized and spaced windows was enabled for the acquisition at the full sampling rate. As shown in figure 3.10 the corresponding control logic had to be inserted before the RAM interface as the internal structure of the RAM controller would not allow the introduction of a control mechanism without impeding the rate of data transfer. The final solution employs a write-controller to the data acquisition FIFO. This FIFO acts as a buffer. It is read out faster by the RAM controller, than data is streamed to it. During the RAM's refresh cycles, i.e. when no data can be written to the SODIMM due to the data in the memory cells being rewritten, the FIFO fills up, however it is chosen big




Figure 3.10 – Modifications to the SAMI FPGA firmware. New additions are marked in orange. Different shades of signal lines denote different sampling rates. The signal composition and the parts used for different parts of the design are used on top of the design (adapted from [116] with kind permission from AIP publishing).

enough to not overflow. Writing to the FIFO was previously only controlled by the trigger. Now a write-controller core can interrupt the acquisition after the system has been triggered. The acquisition times used by the controller core are selected from a BRAM element holding the respective acquisition window and pause lengths (bottom left of fig. 3.10). Since the write cycles must be synchronous across boards the write-controller core operates on a cross-board synchronized clock (see sec. 3.6.2). The contents of the BRAM element can be set using the embedded Linux on the FPGAs between shots. The software ensures that both boards' BRAMs hold the same values. In addition to controlling the data acquisition the controller core is also responsible for controlling the frequency switching. Due to the synchronized clock it was possible to increase the switching time to a quasi-arbitrary length (the previous limitation was a result of the clocking infrastructure). An additional modification to the core was the ability to supply a control signal for a set of dual-polarization antennae, detailed in section 3.10. Since it is not yet known what the exact switching pattern will be, it can be arbitrarily controlled from the MicroBlaze Linux before

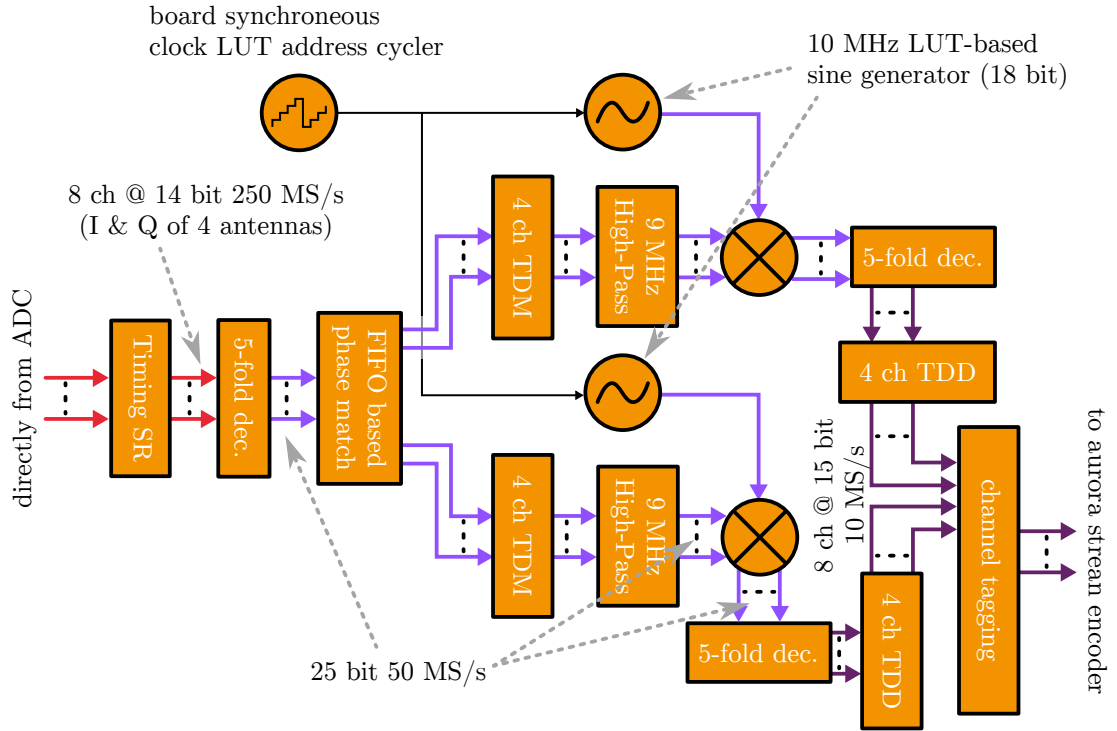


Figure 3.11 – Processing core for the 2D DBS pipeline. The initial filters select the orange part of the signal spectrum in fig. 3.10. Due to resource limitations only 4 of the 8 channels are interleaved for the down-conversion. The sine wave used for mixing is generated using a LUT with a synchronized address counter.

the shot by writing to the same BRAM elements used for frequency switch control.

3.6.1 2D DBS Fork

The 2D DBS signal occupies only a small portion of the overall SAMI bandwidth, as indicated in the top of figure 3.10. The actual information for the 2D DBS image reconstruction is contained within ± 1.5 MHz of the probing peak, as will be seen in section 3.9. It is therefore feasible to reduce the signal using phase linear operations without destroying the information contained therein. The ML605 is equipped with a small form-factor pluggable (SFP) cage, which can be equipped with a module for gigabit communication using optical fibers, just like the simple PCIe carrier (SPEC) board in the interferometer. The maximum speed at which the optical link can be operated at is 2.5 Gbit/s, which was limited by the partner board’s hardware. To make use of this link a data reduction core was implemented.

The schematic of this core is seen in figure 3.11. It forks the data received from the ADC using timing registers. The active probing signal at 12.5 MHz is first reduced in data rate using a 100 tap 5-fold decimator and then isolated via a 130 tap 9 MHz high-pass filter. For the initial decimation filter the coefficients were designed

using the SciPy.signal library. All other filters were designed using the MATLAB FDA tool. As with the interferometer design of chapter 2, the use of finite-impulse response (FIR) filters ensures phase linearity and thus the preservation of the phase information. The signal is then down-converted using a digitally generated 10 MHz sinusoidal waveform, similar to the line of sight (LOS)-integrated electron density (IED) digitizer. During this process the preservation of phase information across FPGAs is ensured via three means:

- The sinusoid is generated from BRAM elements acting as a look-up table (LUT). These hold identical bit patterns on each FPGA. Hence any phase distortion introduced due to limited bit resolution is identical for each signal and will thus cancel during image formation.
- The LUT element generating the sinusoid is read out using a cyclic address counter, which is incremented using a clock derived from the sampling clock. This ensures a stable phase relationship.
- The generation of the sinusoid starts with the first valid sample arriving at the mixer and the LUT address is only updated with a new sample. This way each sample from each channel will see exactly the same phase of the sinusoid during mixing.

The process of mixing brings the 12.5 MHz probing signal down to exactly 2.5 MHz (in the reference frame of the sampling clock). This is now analytically exact since the modulation clock is derived off the sampling clock as detailed in section 3.6.2. After down-conversion the signal passes through final 5-fold decimator using 150 taps. The final sampling rate of the data channels is thus 10 MS/s.

All calculations within the filters are conducted using full precision, however the output of each filter is truncated to 25 bit (see grey arrows in fig. 3.11). This is motivated by the ability of the Virtex-6 fast multipliers (DSP48E1) to multiply one 18 bit with a 25 bit value[140]. The coefficients were chosen with 18 bit so that the signal can be maximized to 25 bit. This eases logic placement to meet timing and uses the minimum number of fast multipliers. The processor's output signals are chosen to be 15 bit wide, with an additional bit for channel identification (tagging is done in the top bit by the processor). The data rate can hence be calculated as 160 MB/s, which is slow enough to be transferred via a 2.0 Gbit/s Xilinx Aurora Stream link. The transfer is encoded using a separate Aurora control and encoder core as shown in figure 3.10. This addition enables continuous 2D DBS measurements.

3.6.2 Revised Clock Infrastructure

The addition of the 2D DBS fork ability for synchronized windowed acquisition required revising of the digitization clock infrastructure, which had to be completely redesigned. An overview is depicted in figure 3.12. The new infrastructure ensures synchronization of logic via several clocks derived from the main sampling clock. Some of these are generated using counters, which in turn are started via the synchronization pulses. As a diagnostic trigger is received the phase-locked loops (PLLs) generating the sampling clock are reset to ensure that the physical components can adequately acquire a lock on the reference clock. Once the propagation of a stable sampling clock is ensured, the PLLs and counters responsible for cross-board synchronized logic are reset by a second MASTER board generated pulse. This ensures that the clock edges appear at the same time on each board. Only after stability of the derived clocks is ensured does the actual acquisition start via a last pulse. The time between the received trigger and final acquisition start pulse is well

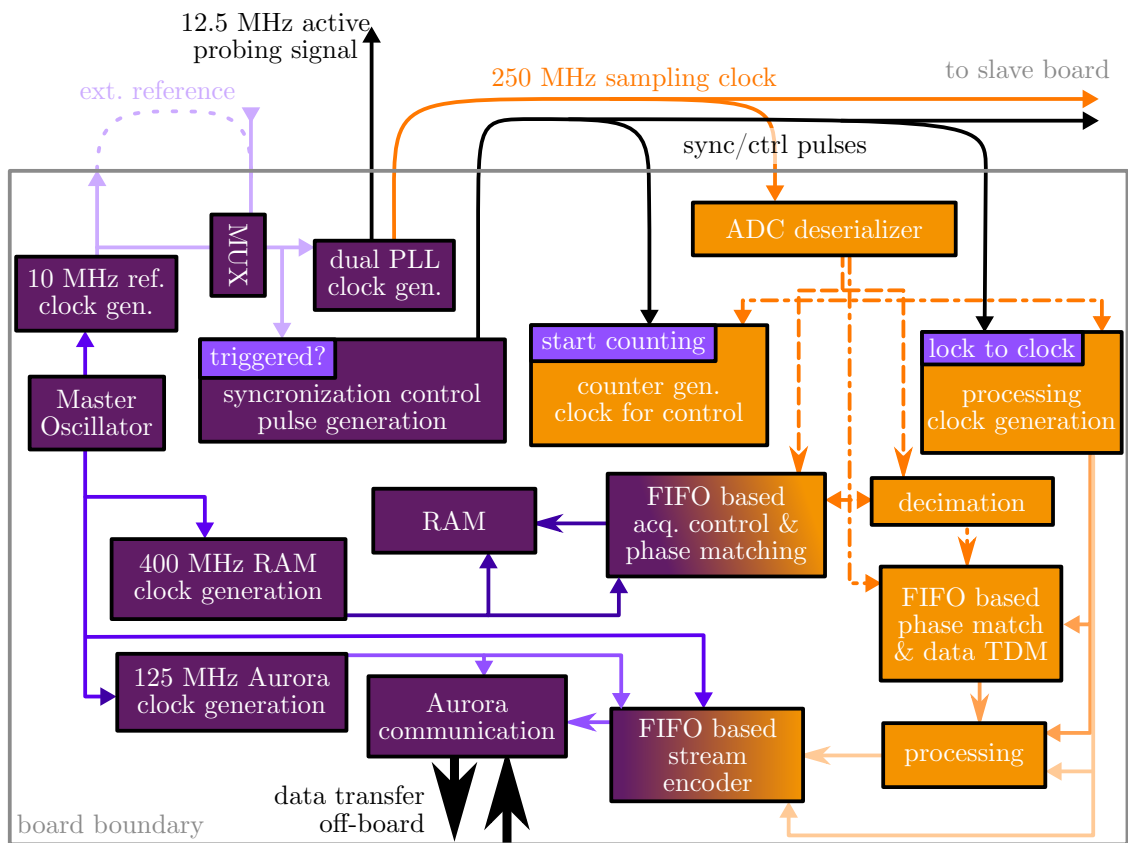


Figure 3.12 – Revised clock infrastructure. Purple components are only synchronous to the board they are on. Orange components are synchronous across boards. The dashed signal line has a stable but unknown phase relationship between channels. Different shades of signal lines denote different sampling rates (adapted from [116] with kind permission from AIP publishing).

defined and can be controlled. There is a level of uncertainty in the synchronicity of the clock edges across boards as a result of manufacturing accuracy and material ageing. However, an upper limit can be placed on this uncertainty based on setup and hold-time requirements of the underlying logic. Given the 250 MHz maximum sampling rate one can assume the uncertainty at less than 1 ns. Since the period of the fastest cross board-synchronous clock is 4 ns no significant difference is to be expected.

As shown in schematic 3.12 the 2D DBS data forked in figure 3.10 is initially out of phase. This is related to the physical position of the pins bringing the digital data onto the FPGA chip. Due to the speed of the sampling clock the initial decimation has to be conducted on a per-channel basis before FIFO elements can be used to match the phase of each sampling clock domain. This phase matching only synchronizes clock edges, which is necessary for the subsequent processing. It does not eliminate the digitization phase mismatch. This is handled by both the ADC calibration prior to acquiring data and the RF calibration in post-processing.

3.6.3 Aurora-based Data Transfer

The link used for continuous 2D DBS can also be used to transfer data in between shots as indicated in figure 3.10. To achieve this the digitization core's Aurora interface is connected to an external AXI Stream interface in between shots[141]. Similar to the RAM controller this switch-over is conducted using MUX elements. The switch-over connects the Aurora core to a direct memory access (DMA) peripheral connected to the MicroBlaze system. This core can quickly read chunks of data from RAM and feed them to an AXI streaming interface. The chunks are predefined by writing instructions into memory (in this case the 8 MB between the Linux's memory space and the region holding the acquired data).

The benefit of a DMA core is that it is entirely logic driven without the need to involve the relatively slow processor. It can therefore supply data at a much faster rate being only limited by the driving clock's frequency and the throughput of the memory interface. Since there is no feedback involved in the streaming interface and no communication overhead due to the transmission of headers when using an Aurora Stream interface, the transfer speed is only determined by the link speed. This in itself is already 2 times faster than the MicroBlaze's gigabit ethernet interface, even without considering the speed reduction due to the CPU's speed.




Figure 3.13 – Layout of the SAMI firmware logic on the Virtex-6 FPGA generated using Xilinx PlanAhead. The serial data pins from the ADC are placed in the lower left of the FPGA. The pins related to streaming the data via fiber optic are in the top right. The logic in between is heavily constrained in placement due to the clock speeds it is operated at and the location of the physical pins it controls, most notably the RAM controller.

3.6.4 Resource Usage

The ML605 board carries a Virtex-6 XC6VLX240T-1FF1156 FPGA. This chip is substantially larger than the Spartan-6 chip used for the interferometer in chapter 2. Consequently resources are more plentiful and less of a concern. Even though the SAMI firmware logic is a lot more complex than in the interferometer’s the design itself occupies only 45 % of the slices available with an average resource usage between 10 % and 50 % (strongly component dependent). For comparison: the utilization of DSP48 elements, which were heavily constrained in the interferometer design, is only 24 % (190 out of 768) and thus of no concern. However, a strongly constrained resource on the chip is the number of global clocking elements (BUFG). These elements are necessary to enable PLL-generated clocks to clock logic globally. Due to the complicated and expanded clocking infrastructure described in the previous section, many of these elements are required to generate all the necessary clocks, of which many operate across multiple FPGA domains. The utilization here is 100 %, making use of several substitutions using counters and local clocking components where possible.

Despite the moderate resource utilization the design is nonetheless very difficult to fit onto the FPGA. The reason can be seen in figure 3.13. The schematic shows an outline of the FPGA chip with the position of relevant logic marked. The MicroBlaze

system, the RAM controller and the digitization & control logic were mostly present before the modifications implemented here. As can be seen the RAM controller, which is the component operating at the fastest clock rate and thus the least flexible when it comes to placement, is positioned at the upper center of the chip, where the respective pins are located. This enables the shortest distance to the source of the fast ADC data as well as the MicroBlaze logic, which itself is predominantly placed in the top right of the system. The MicroBlaze also needs access via a high speed bus to the fast ethernet-controller, which is situated in the lower right of the chip, where its fast transceiver pins are located. The ethernet core in turn needs fast access to the DDR controller in the top left. Due to this three way lock the ADC fork logic, which itself requires many multipliers operating at high frequencies, is “pushing” the logic into the top of the chip, since it needs to stay close to the ADC’s input pins. These pins bringing the digitized data onto the FPGA are located in the lower left of the chip as marked by the shaded area.

The difficulty now arises when the position of the SFP transceiver pins is considered, which are located in the top right of the chip (shaded in fig. 3.13). The data coming in on the ADC has to be processed using a variety of multiplication stages. Although the multipliers are plentiful, they are scattered throughout the chip in strips, which are shorter than the number of taps the processing filters demand. Hence multiple strips are necessary, which requires a lot of routing resources. This can lead to congestion. As a result the processing logic has to be confined to specific regions of the chip to prevent interference with the MicroBlaze logic. This in turn is problematic due to the many different global clocks regions required for the processing, since there is a limit as to how many of these global clocks can coexist in one region of the chip.

To reduce the required number of slices necessary for the DBS fork and ensure timing closure the processing pipeline consequently utilizes partial time-division multiplexing (TDM) as indicated in figure 3.10 (4 channels are interleaved in two parallel pipelines). The number of filter taps was also chosen to result in a routable design, although more multipliers were available and latency is not a concern for SAMI. Finally a large number of pipeline registers was employed.

3.6.5 Aurora-PCIe Bridge

To receive the fiber optic data stream the Aurora-PCIe bridge designed for the interferometer (see section 2.6) was adapted for SAMI. The plasma control system (PCS) fork was removed and additional logic added to enable remote power control of the digitizer box. This is necessary as the MicroBlaze Linux running on the ML605

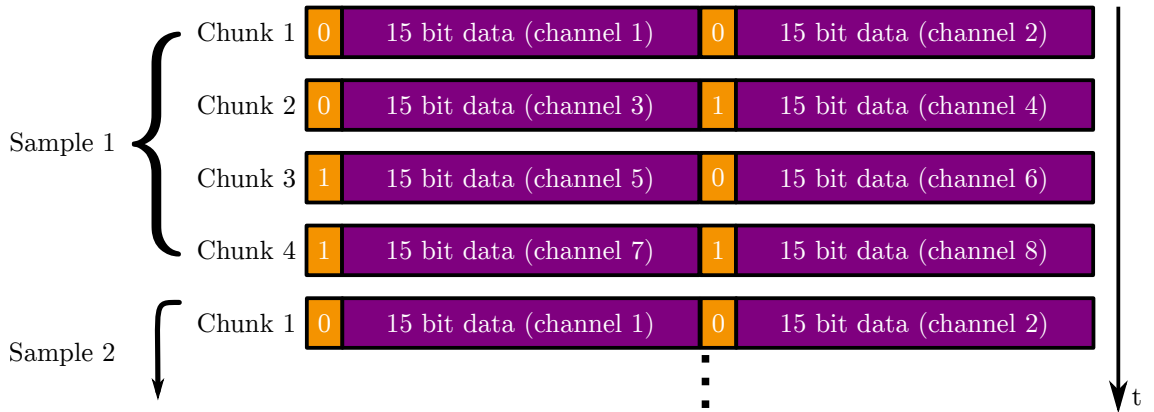


Figure 3.14 – Aurora transmission protocol employed in SAMI. Each channel pair is identified by the combination of the top bits of each channel word (marked orange). The EOS can only be encoded by transmission of a decreasing top bit count.

occasionally freezes during acquisition, requiring a power cycle of the FPGA rack. The cause for the lock-up is unknown. Since it predominantly occurs during the acquisition phase a likely source is a buffer overflow on the RAM FIFOs. If the MicroBlaze has to conduct a lot of work requiring memory access right before the acquisition starts, e.g. by a set of network requests, the FIFO buffering the CPU's RAM-access requests might overflow during the acquisition, when RAM access is only given to the digitizer core. However, the diagnosis of this issue is difficult and no definite source could thus far be pin-pointed. A previously employed watchdog implemented to automatically power cycle the FPGAs could be removed, since its function can be fulfilled in software and the hardware required repair. This also freed resources on the CPU. Since the diagnostic has to be operated remotely and acquire data autonomously, a software controlled power control was also necessary in case the FPGA is powered off, e.g. due to a power outage, which occasionally do occur at PPPL.

In addition the logic to detect the End-Of-Stream (EOS) signal had to be modified. Unlike the interferometer, the SAMI stream cannot use the MAST Aurora protocol due to the limited available bandwidth. The SAMI link is configured to transmit 4 Byte chunks, where each contains the data of 2 channels as shown in figure 3.14. Each sample is therefore transmitted in 4 chunks. Each channel pair is identified by the combination of the channels' top bits (marked in orange). The available 2 bit combinations are therefore not sufficient to identify 4 channels and an EOS chunk. Hence the EOS is encoded via the temporal evolution of the top bits. During a normal stream these will count up (as shown in fig. 3.14). To encode a EOS the top bits are transmitted in a reverse sequence. It is to be emphasized that no EOS *frame* can be used, because there is a chance that the random data contained in the

bottom 15 bit might accidentally reproduce the same bit pattern. The detection of the stream is conducted via a FIFO-based demultiplex-scan-multiplex stage. This has the additional purpose of acting as a buffer and a “ready-to-receive” control for the otherwise unchecked Aurora link. As a result additional reset logic had to be implemented to ensure reliable detection.

Two of these bridge cards with quasi-identical firmware are operating in parallel in a Linux PC. Due to the internals of the Xillybus core employed, an identical firmware is not possible. However, since the data is fully digital at this stage and not being modified by the bridge, this is not a necessity. The link is only used for data transfer. The diagnostic control is still conducted using a web-interface as per the original design[113]. However, the link is configured to operate both ways and could be used to transmit data to the ML605 boards, if future developments so require.

3.7 Installation On NSTX-U

The upgrade to the MAST experiment provided time for extensive diagnostic development. However, due to construction delays the down-time of the experiment was longer than originally anticipated preventing testing of the modifications on MAST-U. As has been pointed out in section 1.3.2, the Nearly Spherical Torus eXperiment (NSTX) is the second major spherical tokamak (ST) next to MAST. Similar to MAST the experiment went into an upgrade phase in 2012, out of which it emerged as NSTX-Upgrade (NSTX-U) at the end of 2015. The two experiments share some defining parameters such as similar aspect ratios (disregarding the MAST-U Super-X geometry) and a similar on-axis toroidal field strength. At the same time there are crucial differences such as the available heating power.

Since SAMI’s initial aim was to investigate B-X-O mode conversion windows an interesting aspect of the NSTX is the capability to condition the NSTX-U wall using Lithium. Wall conditioning, i.e. covering the plasma facing components (PFCs) via evaporation for example, with low-Z materials, most notably Li, has been employed on many machines[142]. The aim is to improve the plasma performance by reducing the recycling of high-Z materials from the walls back into the plasma. Wall conditioning using Li was first employed on NSTX in 2006 and has since been expanded significantly to cover the entire lower divertor. The Li also showed to have a significant impact on the frequency of ELMs, which is thought to be due to the impact on the edge transport and thus pressure gradients[142]. Unrelated to this primary purpose it was found that the B-X-O mode conversion efficiency is significantly influenced by Li-conditioning, since it moves the conversion layer towards

the hotter core, thus reducing collisional damping[143]. Previous investigations into this effect had to be conducted using lengthy shot-per-shot steering of a directional radiometer. SAMI's 2D capabilities make it ideally suited to conduct further studies into this effect and acquire a more extensive dataset.

A second benefit of moving the diagnostic to NSTX-U is the ability to test the previously described changes to the diagnostic firmware, which should result in the first pitch angle profile measurements using synthetic aperture imaging.

3.7.1 Hardware Modifications

NSTX-U is setup differently from MAST and hence modifications to the diagnostic were required to integrate it into its systems. One of the main concerns raised during the initial survey was that computer components close to the NSTX-U vessel regularly fail. Although this normally included failed hard-disk-drives (HDDs) there were also reports of mainboards and solid-state drives (SSDs) failing. The exact cause for this is unknown, but both NBI neutrons as well as stray magnetic fields were speculated. Due to this uncertainty and since FPGAs and SSDs are based on the same technology, the decision was made to try and place the FPGA-digitizer box as far away from the vessel as possible.

Figure 3.15 shows a schematic of the installation on NSTX-U. The down-conversion RF electronics have to be kept as close to the acquisition array as possible to minimize the cable attenuation across the 10 GHz to 40 GHz signal spectrum. These components are therefore bolted directly to the machine on Bay I using a re-entrant midplane port, as indicated by the inlay. Since the signals used for mixing are below 18 GHz the frequency selection circuit can be kept in a separate rack. However, the total cable length still has to stay below 3 m to minimize attenuation. As can be seen in the schematic the amplifier power control for the down-conversion unit is placed next to the frequency selection circuit. The digitizers are positioned on a higher floor several meters away from the experiment. The distance between the FPGA rack, which generates the control signals, and the RF-switch necessitated the use of optical control signals. The main reason for this was the prevention of harmful ground loops. In addition the voltage drop of the transistor-transistor logic (TTL) control signals for the RF switch over the cable distance was thus eliminated, which could have resulted in unreliable frequency selection. The control circuitry was changed accordingly to employ optical links (marked in red in figure 3.15).

An additional difficulty arose due to the Lithium conditioning. To protect against the build-up of metal on the wall a pneumatically controlled shutter is placed in front of the window, which slides in front of the window in between shots. Under normal




Figure 3.15 – The layout of the SAMI hardware on NSTX-U. The diagnostic is protected from Lithium conditioning via an EPICS controlled shutter. The diagnostic hardware has to be physically spread out requiring the implementation of optical interfaces (red) for all but the IF and RF signals (black).

circumstances this shutter is scheduled to open 2 minutes before the start of a shot. However, the pneumatic control mechanics for these shutters have been known to fail regularly and thus there is a substantial chance that it might not open before a shot. On MAST no control for the active probing signal was required and hence the active probing antennae were always sending out their probing signal. A closed shutter could in such a scenario cause damage to the SAMI hardware due to the fact that SAMI uses a re-entrant port on NSTX-U. The projected power on the antenna array could easily exceed the maximum average power allowed into the initial low noise amplifiers. Since a NSTX-U shot can last up to 5 s this could have damaged these components. Consequently a power control circuit using a high-side switch was implemented to cut the power to the active probing signal in such an event.

The shutter itself is controlled via the NSTX-U Experimental Physics and Industrial Control System (EPICS)[144]. This control framework is widely used in many laboratories around the world. However, since MAST uses a different control system, a new control routine had to be developed to ensure the automation of SAMI at NSTX-U. This incorporated the introduction of a shutter status check to prevent the diagnostic from starting the acquisition cycle if there is a fault. Due to the

complexity of the new control and calibration routine, a lot of the work previously conducted by the MicroBlaze Linux on the MASTER FPGA was offloaded to the acquisition PC, which is now solely responsible for controlling the FPGAs. The MASTER-SLAVE architecture was thus confined to the digitization logic. This also enabled the automated calibration of the ADC logic, which can be scheduled at regular intervals, e.g. once a month on a weekend.

3.8 Far-Field To Near-Field Transition

The NSTX-U geometry results in the SAMI acquisition array being much closer to the plasma's density cut-off surface. As a rough estimate the cut-off surface can have a distance from the array as little as 30 cm. The SAMI array's longest baseline is roughly 8.2 cm.

In consequence the far-field approximation based on the Rayleigh distance (eqn. 3.17) is not necessarily valid any more. To investigate the influence of the plasma's proximity on the image reconstruction a dedicated Python based simulation has been written, generating synthetic data based on the emission from a surface comprised of independent point sources. The validity of this approximation is not necessarily given, when considering real plasmas when looking at the optical thickness of plasmas in the microwave range. However, since the purpose of this simulation was not the evaluation of plasma physics on the image formation it is appropriate.

Each point is emitting with a realistic spectrum around the imaging frequency, which is generated separate for each point, thus ensuring the spatial incoherence requirement. Since the sole purpose was the investigation of the image reconstruction, no plasma dispersion or plasma physics is involved in the model, i.e. all waves emitted from the emission surface are considered to propagate in free-space. The simulation can incorporate all relevant diagnostic parameters including antenna response, which however has been disabled for the initial work shown here. This was deemed acceptable, since both the FF as well as the NF inversion algorithms can incorporate the response function and neither should benefit over the other by its inclusion. Future work may incorporate this calculation, however for a more complete picture on the fundamental distortions introduced by the response on the image reconstruction. The simulation can also be used to investigate other array geometries as will be seen in section 3.10.

3.8.1 Comparison Of Algorithms In Planar Geometry

First a basic comparison of the reconstruction algorithms introduced in section 3.3 was performed. For simplicity the emission surface was chosen to be a plane of 3 m by 3 m emitting from 3 distinct “blobs” of 20 cm diameter. They are positioned in a circle of 1 m radius and 120° apart (see top picture in figure 3.16 on page 124). This way distortions introduced by the image formations can be recognized. The data has a 10% uniform noise-floor. All four FF and NF algorithms from section 3.3 are compared. For the NF reconstruction the emission surface is naturally known, which would not be the case for a real surface. Hence additional noise is added to the free-space propagation to incorporate this fact.

In figure 3.16 on page 124 the image reconstruction is shown for several different distances between the emission source emitting at 16 GHz and the imaging array. In each of the figure’s quadrants the top shows the geometry of the model. The purple triangles indicate the array geometry as depicted in figure 3.8 on page 106 scaled up by a factor of 20 for better visualization. The middle plots show the BF image reconstruction (left) and image formed via Fourier inversion (right) in far-field approximation. The bottom plots show near-field inversions using BF (left) and SVD inversion (right). The difference between Rayleigh distance R and average distance to the array \hat{d} has been given in meters on top as

$$\rho = \hat{d} - R.$$

The first most obvious feature seen in the images is that the BF naturally exhibits a significantly higher level of base noise. This is due to the fact that BF in essence includes the self-correlation information, which Fourier and SVD inversion omit. This holds the information on the total received power. Comparing the reconstruction accuracy the plots show that the already employed BF and Fourier inversion algorithms are the most robust. In the very far-field (top left) the SVD inversion is only marginally reconstructing the correct image. Most of the power appears to be localized to the border of the reconstructed surface. The NF-BF algorithm reconstructs the location of maximum emission correctly, but assumes it to have a significantly higher spatial extent. In both cases the source for the problem is likely the limitation on the maximum resolution of the array geometry. The largest baseline determines the smallest resolvable structure. The further the emission surface goes into the distance, the smaller the phase difference between the points in the reconstruction surface, which leads to washing out of the true intensity distribution due to spatial undersampling. Using an angular difference based grid, as for the far-field reconstruction, may be beneficial, but increases computational effort.




Figure 3.16 – Comparison of image reconstruction algorithms for a planar source emitting at 16 GHz. In each quadrant the simulated geometry is depicted at the top, with ρ given in meters; purple triangles mark the array geometry (scaled by a factor of 20 for visualization purposes). The polar plots in the middle use FF reconstruction with BF (left) and Fourier inversion (right). The bottom plots show the NF reconstruction using BF (left) and SVD inversion (right). Axis descriptions have been omitted for illustration purposes. Cartesian axis labels are in meters. The opening angle in the polar plots is in degree.

As the emission surface moves closer to the array, each reconstruction method shows the correct behaviour. It can be seen that the SVD inversion appears to be robust in the intermediate far-field.

In the bottom right quadrant the emission surface is approaching the NF regime. However the FF image reconstruction methods appear to still be functioning. This is surprising as the fundamental assumption behind these methods is broken. Nonetheless ρ is still positive with 0.553 m in this geometry and thus the result is not impossible. The NF inversion algorithms show significant distortions on the reconstruction surface. Again the BF algorithm appears to outperform the SVD inversion as the washed out maxima align better with the expected location. However, the strong distortions in line with the high overall noise level present in the BF reconstruction make the interpretation of the image difficult. The SVD inversion's reconstruction appears to exhibit significant distortions at the center of the reconstructed image, which appear to resemble the three source blobs scaled and mirrored. This suggests a mathematical problem in the inversion, e.g. a local minimum convergence. The subjectively better performance of the FF inversion algorithms when compared to the NF must be viewed with care. Since the FF inversion uses an angular grid with even spacing, a reconstruction onto a true geometry will cause the same distortions the BF algorithm exhibits.

Since the array geometry and imaging frequency are strongly related, the same investigation was conducted for a source emitting at 26.5 GHz (see fig. 3.17 on page 126). On first sight the SVD reconstruction performs significantly better in the far field. However, it has to be noted that the furthest distance at this imaging frequency corresponds to a ρ of 1.58 m, for which the matrix inversion yielded adequate results in the 16 GHz case as well.

For this higher frequency case the image artefacts are significantly more pronounced in all reconstructions. Based on the case for a ρ of 1.58 m this is not directly a result of near-field effects, since the FF inversions were adequate for lower values of ρ in the 16 GHz case. The effect is therefore likely caused by the array geometry, i.e. an effect of sub-optimal sidelobe suppression at this frequency. The SAMI array geometry has been optimized to minimize the nearest sidelobe level given the Fourier reconstruction[110]. This optimization has been conducted for 20 GHz, which is roughly in the middle of SAMI's range of frequencies. It is conceivable that the suppression is less adequate for higher frequencies. The observations of the simulation match experiences from SAMI's real data, in which higher frequency channels exhibit more artefacts as well.

Due to the increasing number of artefacts in the FF inversion images the recon-

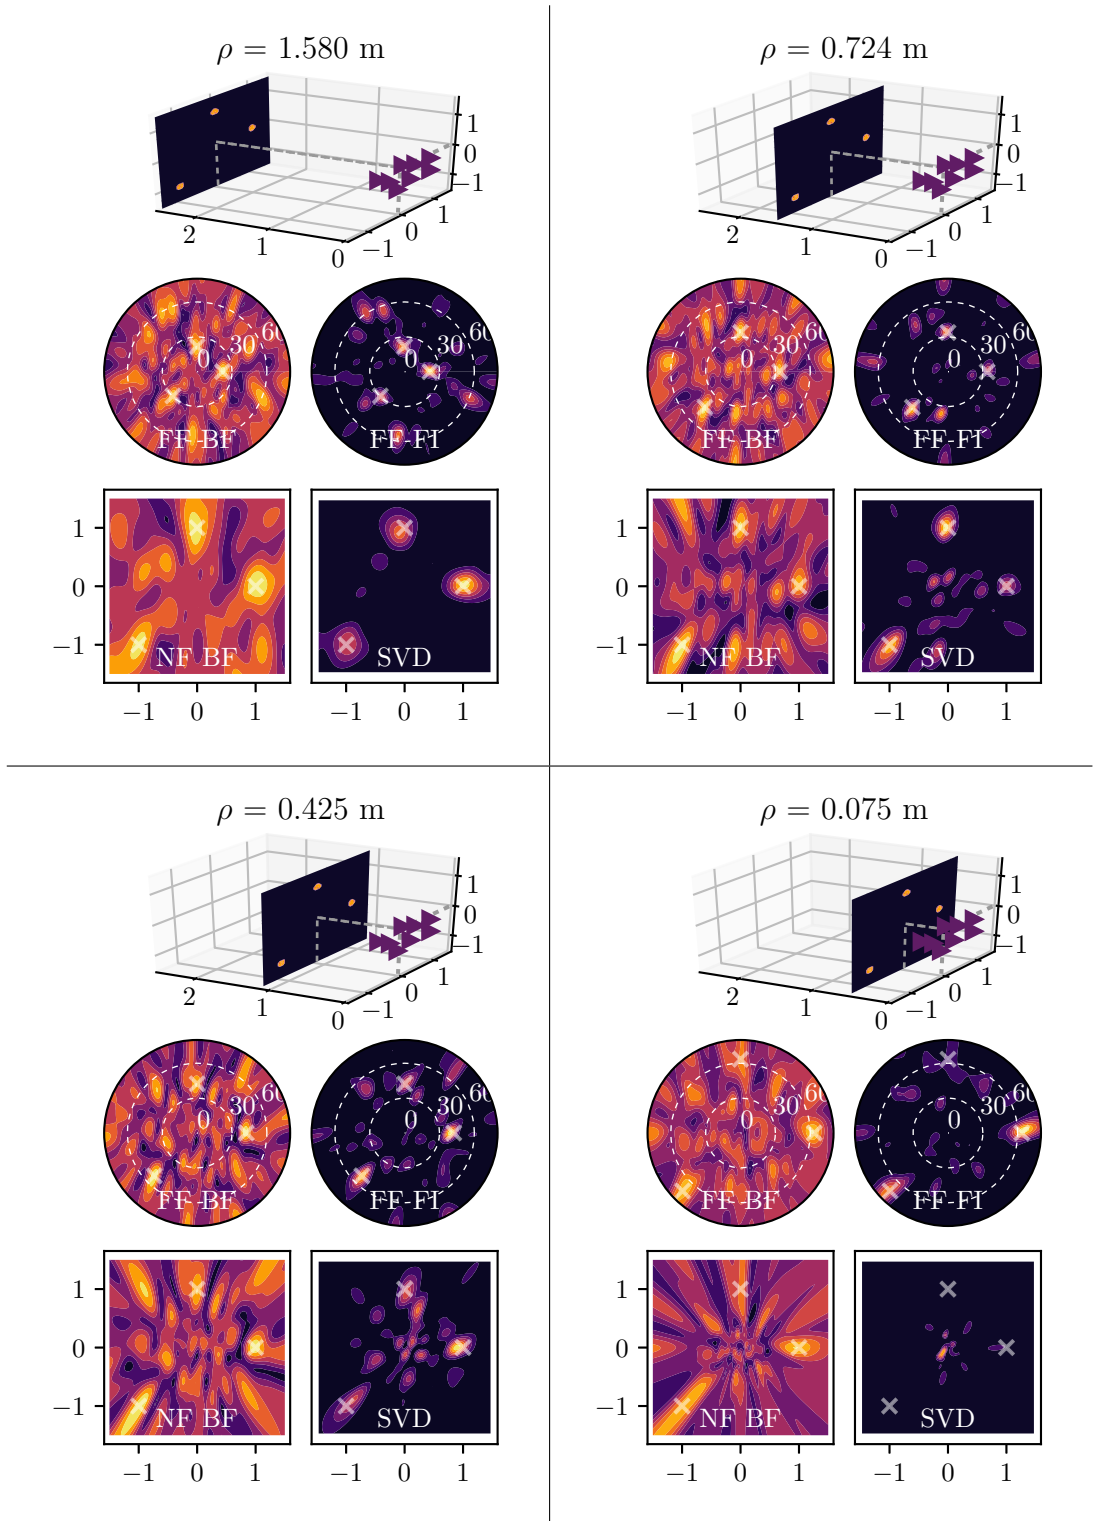


Figure 3.17 – Image reconstruction algorithms compared for different distances of a planar source emitting at 26.5 GHz. The simulation geometry and plot-layout are identical to fig. 3.16.

struction becomes increasingly hard to interpret as ρ approaches 0 m and identification of the correct maxima is most difficult for FF BF at 1 m distance from the array due to the generally higher noise level. The Fourier inversion appears to provide a more accurate reconstruction due to the lower base noise, but here as well the image reconstruction is distorted for lower ρ . However, some of the distortions may be the result of the spatial sampling, i.e. the choice of the grid.

The near field inversion algorithms are a lot more useful in this scenario. The NF BF inversion, just like its FF counter-part, exhibits a natural level of artefacts. The three emission peaks are clearly distinguishable all the way to the lowest ρ case, although the distortions would make the identification difficult without prior knowledge of the source. The SVD inversion subjectively appears to be delivering easy to interpret images up to a ρ of 0.425 m. In the very near-field regime however the reconstruction appears to have broken down, localizing the emitted power in the center of the reconstruction surface. This behaviour of assuming too much power near the center as the NF regime is approached appears to be inherent to this inversion technique. Even for the other ρ cases distortions always start appearing at the center as the source gets closer to the array. This also appeared in figure 3.16. The reverse appears to be true as well, since it was visible in fig. 3.16 for the high ρ case that too much power was assumed at the edge of the source. The cause for this could be an incorrect weighting of the points in the reconstruction surface, i.e. by assuming an incorrect area that each point represents leading to a numerical bias towards certain sections of the reconstruction surface. Again the artefacts appearing at the center of the reconstructed image appear to be mimicking the original source, but are scaled and rotated.

So far the quality of the reconstruction was only evaluated “by eye”. A quantitative analysis is difficult, as it is hard to find a quantitative measure for the “distortedness” of the reconstruction. An attempt was made to calculate a distortion factor δ via

$$\delta = \left\| \frac{\text{PD}_{\text{rec}}(x, y, z)}{\|\text{PD}_{\text{rec}}\|} - \frac{\text{PD}_{\text{true}}(x, y, z)}{\|\text{PD}_{\text{true}}\|} \right\|. \quad (3.24)$$

In this equation PD is the reconstructed and true power distribution on the emission surface. For NF algorithms this is a function of x , y and z as indicated, for FF inversion algorithms this is only a function of θ and ϕ , to which the true probability distribution has been translated. The chosen norm is the Frobenius norm. The distortedness δ is higher, the more power is reconstructed away from the true power distribution. However, it says nothing about the distribution of that incorrectly placed power and thus is no measure of how well the defining features

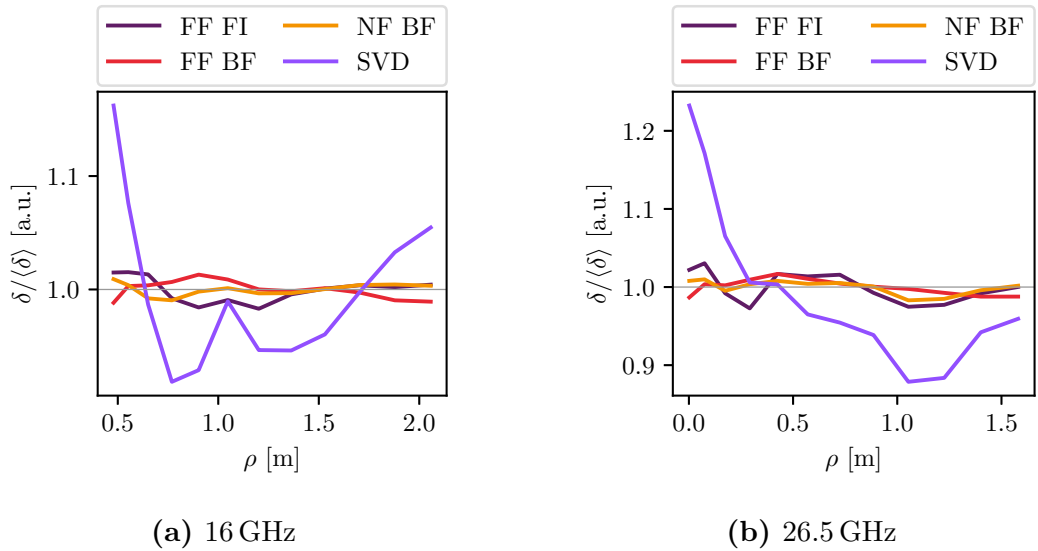


Figure 3.18 – Normalized distortedness δ over ρ for all four image reconstruction methods and two different probing frequencies. δ is normalized to the average over ρ . Quantitative comparison of the distortedness curves should be taken with great care.

can then still be identified.

Figure 3.18 shows the distortedness δ for the planar emission surface simulation over ρ . δ is normalized to the average over ρ . The absolute values of the distortedness should only qualitatively be compared to each other. As can be seen the inversion routines exhibit very different distortion levels. As expected both FF inversion routines become increasingly inaccurate for lower ρ , i.e. closer to the NF-regime. The SVD inversion on the other exhibits strong variability with respect to ρ , which was already evident from figures 3.16 and 3.17. In the intermediate regime around $\rho = 1$ the inversion appears to do significantly better, when compared to the extreme ρ cases. The NF BF inversion shows the least amount of variability and exhibits roughly the same distortion level for all ρ . However, it must be reminded that this may well be due to the generally higher level of distortedness on the inversion, which may hide the actual variability due to the generally higher δ . Nonetheless it appears that this inversion routine exhibits the most predictable level of inversion accuracy.

One may note in figure 3.18b that the FF Fourier inversion shows a small dip in distortedness at a ρ of roughly 0.25 m. Similarly there is a drop in distortedness for the FF BF routine in the NF. This hints at a discretization effect, i.e. that the choice of the reconstruction grid influences the variability and accuracy of the reconstruction, since we are spatially undersampling in the reconstruction.

3.8.2 Comparison Of Algorithms In Spherical Geometry

The planar geometry, although instructive, is unphysical. A real plasma emission surface more closely resembles that of an ellipsoid. Due to the shaping of the emission surface parts of it may well be in the NF regime while others are still outside. Since the effects were unknown an ellipsoid emission surface was simulated and its distance to the array changed. This also helps investigate the expected change in image distortion due to the transition from MAST to NSTX-U.

Figure 3.19 on page 130 shows the simulation results. The plots are laid out as in the previous examples. The emission source was simulated to emit at 20.5 GHz. As can be seen, the reconstruction is adequate for all of the image reconstruction algorithms in the case of a distant plasma. The FF inversions show more distortion than the NF ones, but none that make interpretation difficult. The scenario depicted in the bottom left is a setup that resembles the MAST distance. Evidently the FF algorithms start showing notable distortions, but interpretation is still obvious. The NF BF already shows substantial distortion in this geometry leading to difficult interpretation. This, as has been pointed out before, is likely a result of the array geometry in combination with the reconstruction surface's geometry. The SVD inversion on the other hand appears to be correctly reconstructing the image with acceptable distortions, which is even true for the very NF case. With a very close emission surface the correct interpretation of the reconstructed image is impossible for the FF algorithms. The NF BF inversion could with difficulties still be interpreted correctly. Nonetheless it requires interpretation of the immense washing out effect. Although the results presented in fig. 3.19 indicate that the SVD inversion outperforms the other algorithms, the quality of the inversion is still strongly frequency dependent and occasional SVD inversions have failed completely, probably due to the randomness in the emission source.

Figure 3.20 depicts the normalized distortedness δ over ρ for this simulation scenario. As can be seen the NF BF again exhibits the least amount of variability in δ when it comes to the reconstruction. The SVD inversion is again highly variable in its distortedness and becomes significantly worse for low ρ . Surprisingly the distortedness shows the opposite behaviour for FF inversion to what would be expected when looking at figure 3.19. For lower ρ the distortedness appears to decrease again, although the plotted power distribution suggests otherwise. This shows that the norm chosen for δ is not a reliable measure to quantitatively compare the algorithms.

In conclusion the FF algorithms, in particular the Fourier inversion, appear to be working very robustly, even close to the NF regime and most certainly for




Figure 3.19 – Image reconstruction algorithms compared for different distances of a ellipsoid source emitting at 20.5 GHz. The plot-layout is identical to fig. 3.16.




Figure 3.20 – Normalized distortedness δ over ρ for a probing frequency of 20.5 GHz in spherical geometry. δ is normalized to the average over ρ . Quantitative comparison of the distortedness curves should be taken with great care.

MAST. However, distortions are to be expected in the NSTX-U geometry that could make interpretation using FF inversions difficult. If possible, near-field inversion algorithms should be used, of which the NF beam-forming algorithm, although susceptible to distortions, is the most robust at this stage. This is a subjective assessment based on the fact that the distortedness of the reconstruction is similar in both the NF and the FF regime. In one would expect that an automated evaluation routine, which is necessary to obtain a pitch angle time trace, would very likely have difficulties in coping with the distortions introduced by the SVD inversion. Based on the reconstructed images' power distribution, the distortions introduced by the NF BF algorithm, would subjectively only increase the variability of the pitch angle measurement, since there are still maxima in the correct location. This assessment however is hypothetical until such an automation routine has been written.

The SVD based matrix inversion has shown, that it can in some cases outperform the BF reconstruction. As such, it may be that future improvement towards constructing the inversion matrix may yield better results. An initial starting point would be the introduction of a regularization, e.g. Thikonov regularization, to improve the stability of the inversion and prevent the apparent convergence problems[131]. However, the difficulty of such methods is that the regularization parameter is problem dependent and may vary strongly as the plasma properties change even slightly. Another approach would be the use of a maximum entropy inversion instead of the least squares approach employed here. This type of inversion distributes the intensity more smoothly on the reconstruction surface. In particular the accumulation of reconstructed intensity at the borders or the center of the reconstruction surface may be prevented this way. However, the smoothing effect also results in more washed

out features in the reconstruction.

3.9 Initial NSTX-U Results

The initial measurements on NSTX-U were conducted with a configuration that closely resembles the MAST setup, i.e. the same array configuration using the Vivaldi antenna array and a single frequency acquisition shot. This was done as a sanity check to compare the results to the ones obtained on MAST.

Figure 3.21 shows the location of normal-incidence cut-off surfaces for NSTX-U shot #204672 for two timeslices in L-mode and H-mode. In this initial measurement the lower SAMI imaging frequencies are located at the very edge of the plasma; much more than on MAST[111], [115]. The lowest frequency of 10 GHz, which was used for the first measurements, was probably already in the scrape-off layer (SOL). This may influence the level of back-scattered signal for 2D DBS as well as the fluctuation level present in the signal.

Figure 3.22 compares the 2D DBS power present on MAST and NSTX-U for the 10 GHz acquisition. The plot is normalized to the maximum received power. The time points in both shots exhibit roughly the same edge density and temperature




Figure 3.21 – Normal incidence cut-off surfaces for an L-mode and an H-mode section of NSTX-U shot #204672. The R-wave cut-off ν_R (see sec. 3.2.2) is plotted in solid blue, the density cut-off ν_{pe} is plotted in magenta. The SAMI acquisition frequencies are marked with circles (O-mode) and crosses (X-mode). (Taken from [115]).




Figure 3.22 – Comparison of 2D DBS power received for 10 GHz acquisition with the SAMI diagnostic on MAST and NSTX-U. Each measurement is compared against the null measurement without plasma (purple). The NSTX-U 2D DBS power is significantly higher above the noise floor (adapted from [115] with kind permission from the author and AIP publishing).

gradients, thus it can be expected that turbulence levels are roughly equal and the optical thickness of the plasma is roughly the same. In both cases the plasma density was peaking around $2 \times 10^{19} \text{ m}^{-3}$ and the maximum electron temperature read roughly 700 eV. Turbulence levels could not be compared due to the reflectometry diagnostic being unavailable on NSTX-U. It is evident that the back scattered power on NSTX-U is much higher above the general noise floor than on MAST. Here the width of the Doppler shift previously stretched less than 1 MHz in either direction of the direct reflection peak. The measurements on NSTX-U show information up to 1.5 MHz in either direction. This may be very beneficial, as it reduces the uncertainty on the imbalance measurement. This result led to the decision to disable the 10 MHz active probing signal, since it would have resulted in an overlap of the two DBS spectra for 10 MHz and 12.5 MHz, leading to difficult interpretation of the imbalance spectrum. The developments of the active probing stream presented in section 3.6 already incorporate this. The cause for this increased power is likely the result of the acquisition array being closer to the emission surface.

In figure 3.23 a pitch angle trace compared with EFIT is plotted for the NSTX-U shots #204620 (imaging at 10 GHz) and #204944 (imaging at 16 GHz). EFIT’s pitch angles have again been matched to the probing location using TORBEAM. The acquisition was conducted using the conventional raw digitization into RAM with a single frequency for one continuous 500 ms acquisition window and a far-field BF reconstruction. This is also roughly the shots’ length. As can be seen from the left plot the diagnostic can trace the pitch angle adequately with an accuracy comparable to the results obtained on MAST[111]. A comparison to MSE was not available for this measurement campaign. Figure 3.24 shows plasma parameters relevant to the pitch angle measurements. The top plots the IED derived from TS. Below is the total measured toroidal plasma current and the midplane Deuterium- α intensity. The bottom shows the applied NBI heating (NBIH) power.





Figure 3.23 – Pitch angle trace for NSTX-U shots #204620 and #204944 using 10 GHz and 16 GHz 2D DBS respectively. The SAMI measurement is in purple, the EFIT fit is in orange. The agreement is very good over most of the shot for the 16 GHz shot. A significant overestimation of the pitch angle (in comparison to EFIT) is found for shot #204620 probing at 10 GHz (adapted from [115] with kind permission from author and AIP publishing).

In shot number #204620 SAMI probed the plasma using 10 GHz, which is in the very edge of the EFIT reconstruction domain and potentially beyond the last closed flux-surface (LCFS) in the NSTX-U plasma, i.e. within the residual plasma present in the area with the first open field lines that connect to the divertor, also known as SOL (figure 3.21 on page 132 can be used for a rough localization of the probing locations). The shot managed to trigger the transition from L-mode to H-mode 210 ms into the shot, which is indicated by the sudden drop in $D\text{-}\alpha$ emission as well as the subsequent rise in density due to the improved confinement. The H-mode transitions back into an L-mode around 430 ms, where the $D\text{-}\alpha$ trace suddenly increases. Surprisingly the SAMI measurements do not respond to the L-H-mode transition with an increased variability, which would be expected from the reduced level of edge turbulence resulting in a higher measurement error. This was a typical feature seen on MAST. NSTX-U shot #204944 on the other hand stays in L-mode throughout the shot. SAMI probed the plasma at 16 GHz, which is in the outer edge of the plasma .

In both plots the pitch angle increases throughout the shot, which is a result of the plasma current increasing, since this also increases the poloidal field in the edge. In NSTX-U shot #204944 the pitch angle increases continuously at the start as the current is increased (see fig. 3.24a). Around 250 ms it can be seen that the flattening of the current ramp coincides with a flattening of the EFIT pitch angle ramp in figure 3.23, showing that the toroidal plasma current is the primary factor defining the edge magnetic pitch. Surprisingly the SAMI pitch angle measurement starts to deviate from the EFIT trace at this point, since it appears to continue rising. The magnitude of the deviation is small and within the discrepancy found between unconstrained EFIT and SAMI[111]. However, NSTX-U shot #204620



(a) NSTX-U Shot #204944

(b) NSTX-U Shot #204620

Figure 3.24 – Selection of plasma parameters for the shots shown in figure 3.23. The top depicts the TS calculated IED. Below are the measured toroidal plasma current and the midplane Deuterium- α light intensity. The bottom is the applied NBIH power.

SAMI overestimates the pitch-angle by a factor of roughly 2 throughout the shot when comparing to the EFIT measurements. This error appears to be a systematic scaling mismatch, however in the light of the deviation seen in shot #204620 it is uncertain, whether this factor just appears to be fixed.

The cause for the discrepancy is unknown. One is initially tempted to say that EFIT is less accurate at the plasma boundary, in particular the SOL. However, the EFIT equilibrium reconstruction tends to be more reliable at the edge than at the core, although this strongly depends on the available diagnostic set to constrain the code[128]. Judging from a comparison of the plasma current seen in figure 3.24 it is unlikely that the SAMI-measured pitch angle is correct, even under the assumption

that the sampling location of the EFIT pitch is not properly aligned with the probing frequency's location in the plasma.

As has been seen in the investigation of near-field effects in section 3.8, the near-field can lead to significant distortions in the SAMI image reconstruction, which likely also applies to the 2D spectral distribution. This could end in a flattening of the profile, as the relevant maxima get smeared out. It would also be interesting to see, whether a large poloidal plasma rotation, beyond what has been observed with SAMI on MAST so far, influences the automated 2D DBS routine. Although the assumption on the minimum k back-scattering efficiency is not flawed, a large poloidal rotation may result in the SAMI 2D DBS routine finding the maxima of red-blue shifted power-difference at a different point (when compared to a lower poloidal rotation), since the LOS seeing the largest amount of Doppler shift, does not coincide with the LOS seeing the maximum back-scattering power.

Unfortunately no diagnostic measuring the poloidal rotation was available for the above two shots so such a correlation could not be tested. This should be done in the future.

3.9.1 Firmware Benchmarking

During the short measurement campaign the raw acquisition data was automatically downloaded using the fiber optic link. On average the data download took about 12s per board for the entire 2 GB of raw data acquired. The download procedure can be conducted from both boards at the same time using a RAM-disk. However, since the resultant data rate exceeds the write speed of the HDD, the downloads are normally conducted in sequence. The link's reliability was verified by comparing the transfer against a Trivial File Transfer Protocol (TFTP) transfer, which employs error checking and can thus be used as a benchmark. So far no transfer errors could be detected. This download speed is an order of magnitude improvement over the previously implemented ethernet approach[137]. In combination with the recently developed code employing a graphics processing unit (GPU) to calculate the FF Fourier inversion, results of the SAMI diagnostic will be available within one minute of the end of a shot[114]. This will greatly benefit future physics campaigns.

The continuous 2D DBS could not be tested during the measurement campaign due to several issues with the diagnostic followed by a hardware fault in the NSTX-U central solenoid. A replacement test could however be conducted due to the unique way SAMI acquires data. Since the raw acquisition data that is stored directly into RAM is identical to the data being used for 2D DBS streaming (see section 3.6) the down-conversion logic can be tested in simulation with real data, as it would be




Figure 3.25 – Simulated result of the 2D DBS down-conversion logic operating on SAMI data. The top plot depicts the spectrum before and after 2D DBS processing using data from NSTX-U shot #204620 240 ms into the shot. The acquisition frequency was 10 GHz. The spectrum is calculated from a 1 ms timeslice. The bottom plot shows the frequency response of each of the three filter stages as well as the combined response. The data-sets have all been shifted to have the respective active probing frequency peak at DC.

during a real acquisition. A corresponding simulation has been written and supplied with the raw data acquisition of one of the SAMI antenna’s data. Simulating all 8 antennae was not possible due to the computational complexity of the simulation.

Figure 3.25 shows the results of the simulation. The top plot shows the normalized spectra of the raw data as well as the simulation-processed data. Evidently the noise floor around the main active probing peak appears to have increased in the processed data with respect to the raw data. This is likely the result of bit-accuracy, mixing attenuation and aliasing noise. As a rough guideline all calculations on the in the 2D DBS down-conversion stage have to be conducted with roughly 25 bit fix point accuracy. Although the FPGA could in principle calculate with any accuracy, the resource and timing limitations explained in section 3.6.4 prevent this. In addition the final accuracy of the transmission cannot exceed 15 bit due to the limited Aurora transmission bandwidth. The necessary truncation operations raise the overall noise floor slightly. In addition the mixing process introduces a level of attenuation into the mixing due to the fact that the down-conversion waveform is at a different signal level than the mixed signal, which reduces the signal-to-noise ratio (SNR) of the

down-converted signal. Lastly, the noise floor is probably slightly raised due aliasing noise. Although the stop-band suppression has been maximized given the number of available filter coefficients, there will naturally be residuals after the filtering, which translate to aliasing noise due to the down-sampling employed.

The most important aspect for the 2D DBS technique is a distortion-less down-conversion of the signal spectrum, since the information is contained in the shape of the back-scattering peak. A slight distortion can be seen in particular on the high frequency side. This is likely the result of filter attenuation. For a comparison the bottom plot shows the spectral filter amplitude response for all three filter stages employed in the down-conversion calculation shifted to have the respective probing frequency at DC. The attenuation ripple can only be minimized with more filter coefficients, but can never be prevented. As can be seen the slight increase in the spectrum coincides with a peak in the attenuation. This distortion would lead to an overestimation of blue shifted power. However, since the filter response is well defined it can be accounted for via deconvolution. Since no other major distortions are visible the 2D DBS technique will work with the stream-processing.

3.10 Outlook

The start of the next measurement campaign will test the new 2D DBS streaming implementation to measure the pitch angle continuously. In addition frequency switching will be employed to acquire the first pitch angle profile measurements using 2D DBS.

All of the measurements will initially have been taken using the Vivaldi antipodal antenna array. However it has been pointed out, that there is a significant amount of uncertainty in the localization of the cut-off surfaces due to SAMI's inability to separate polarization[111]. In addition the 3D geometry of the Vivaldi array can lead to antenna cross-talk, which has been speculated to introduce artefacts in the image reconstruction[110]. The SAMI array has thus been prepared to employ a new set of dual-polarized sinuous antennae. This type of logarithmic antenna has no 3D geometry and can be designed for large bandwidth applications[145]. Just like the Vivaldi antenna, the sinuous antenna is a PCB based element and thus cheap to manufacture.

The sinuous antenna array will be arranged in a more regularly spaced circular pattern. This is done to reduce the deformation of observed structures. The deformation effect can be seen in the reconstruction of figure 3.16 on page 124. In particular the far field reconstruction for ρ above 1 m exhibits an elongation of the




Figure 3.26 – Image reconstruction algorithms compared for the array configurations using Vivaldi (left) and sinuous antennae (right). The source is emitting at 16 GHz is at a distance of roughly 1.3 m from the array. The plot-layout is identical to fig. 3.16.

reconstructed point sources along the vertical axis. This is an effect of the array geometry, which, as depicted in the same figure, is elongated along the horizontal axis. The counter-intuitive behaviour is the result of the higher spatial resolution along the horizontal, since the longer baselines are placed in this dimension. Hence the point source is more washed out along the vertical.

The new sinuous antenna array, although circular, is nonetheless placed in a pattern ensuring that none of the baselines are redundant. This makes most use of the information obtainable with 8 antennae. Figure 3.26 directly compares the point source reconstruction for the old Vivaldi and the new sinuous antenna configuration. The array geometry is again plotted with purple triangles scaled by a factor of 20. As can be seen the point source reconstructed with the old array is elongated along the vertical, while the sinuous array reconstructs a circular image. In particular automated evaluation routines, such as the one used for pitch angle traces in figure 3.23, will benefit from this, since the smear-out effect can result in incorrect maximum recognition.

A second reason for using a more regular circular array is the tighter packing of the antennae. The information SAMI is interested in is mainly contained at relatively large opening angles of over 20° from the normal. As pointed out in the previous

section the current Vivaldi array geometry is optimized for 20 GHz signals. However, as has been shown, the higher acquisition frequencies appear to exhibit significant amounts of distortion. By using a tighter packed more circular array, the higher frequencies may become less distorted, albeit at the cost of the lower frequencies' image quality. On average it is believed that this will yield easier to interpret images.

An effect that was not investigated with the image reconstruction was the effect of the antenna response directionality, which was assumed perfect in all directions for the simulation. However, the Vivaldi response in particular is highly directional, which likely leads to additional distortions on top of the ones related to the array geometry. The sinuous antennae have a significantly flatter and more uniform antenna response, which will hopefully improve the quality of the image reconstruction and an easier interpretation.

The dual polarization measurement using the new sinuous antennae will not be simultaneous. This would require twice the amount of RF down-conversion and digitization hardware. A partial polarization separation can however be achieved by measuring the cross-polarization on a subset of the available baselines, e.g. by measuring half of the antennae in the opposite polarization direction. With four different combinations the majority of the polarization information can be reconstructed. The changes to the clocking infrastructure presented in section 3.6 were mandatory for this to work, since the evaluation of the resultant data will require spectral stitching. The FPGA firmware is now set up to control the diagnostic with dual polarized antennae. On the hardware side an additional set of RF switches have to be installed directly behind each antenna, which will then be controlled by the FPGAs. This change will be conducted once the Vivaldi measurements are finished.

3.11 Summary

Significant development has been undertaken on the Synthetic Aperture Microwave Imaging diagnostic. The FPGA firmware as well as the diagnostic hardware have been modified to fix several issues discovered during the initial phase of experiments conducted on MAST. Most importantly the diagnostic has been modified to enable continuous 2D Doppler back-scattering measurements with the aim of acquiring magnetic pitch angle profiles.

The diagnostic has also been transferred to NSTX-U and integrated into the local systems, where it has been running since May 2016. The initial measurements showed improvements in the back-scattered power and a discrepancy of the pitch

angle trace when compared to EFIT, which could so far not be explained. These measurements could not be compared to MSE, as it was not available.

The influence of near-field effects on the image reconstruction was evaluated as a result of the move to NSTX-U showing that near-field effects should be considered and may potentially have a detrimental effect on the ability to locate the maximum of a point source adequately. This may be a cause for the current discrepancy between EFIT and the SAMI pitch angle trace.

SAMI is now ready to attempt the first measurements of pitch angle profiles using 2D DBS, which will commence after NSTX-U restarts. Following these measurements the current Vivaldi array will be swapped for a dual-polarized sinuous antenna array with a circular array arrangement, which will allow the partial separation of O mode and X mode, thus improving the measurement and localization error.

Chapter 4

Concluding Remarks

The development of diagnostics is a key element of our efforts towards developing a fusion power plant. The complexity and scale of fusion experiments results in a wide variety of different diagnostics employed to study them. In addition to the increased demands regarding diagnostic accuracy, modern experiments themselves have an ever increasing need for reliable real-time protection systems, since their larger size increases the danger to the hardware due to violent plasma instabilities. In consequence an increasing amount of data must be processed with very low latency. Meeting these demands requires the implementation of new technology.

Field programmable gate arrays (FPGAs) deliver excellent real-time performance. Their ability to handle and process large amounts of streamed data, as it is generated by fast analog-to-digital converters (ADCs), enables the development of high-performance diagnostics that can deliver real-time information on plasma properties. In particular fusion diagnostics can benefit greatly from utilizing this technology by improving old diagnostics and enabling new ones.

4.1 Summary

The developments presented in this thesis concerned two different diagnostics, both of which depend strongly on FPGA technology. The first is the new MAST-Upgrade (MAST-U) integral electron interferometer described in chapter 2. A FPGA based flexible real-time digitization system has been designed as a back-end to the optical front-end. The system delivers very high bandwidth and accuracy, whilst at the same time delivering excellent real-time performance. The flexibility of the system has been demonstrated by transferring it to the ASDEX Upgrade (AUG), where integration into the local systems could be conducted within a few hours. The measurements from AUG showed the potential for discovering fast plasma dynamics one might have

previously been oblivious to. For AUG the system may solve the issue of ion cyclotron resonance heating (ICRH) cross-talk, which pollutes the local interferometer signal to the point of being unusable. The digitization system presented will be integrated into the MAST-U plasma control system, to which it will deliver processed density data in 7 μ s, whilst maintaining the full 3.5 MHz bandwidth for diagnostic purposes. This level of data handling was only possible using a tandem FPGA approach.

The second diagnostic that was worked on as part of this thesis was the Synthetic Aperture Microwave Imaging (SAMI) diagnostic presented in chapter 3. This system could only be built using the fast data-processing capabilities FPGAs offer. This diagnostic is a first of its kind and development of it is still ongoing. The experience gained with the diagnostic during the initial campaigns revealed several aspects of the diagnostic system requiring revising. The flexibility to reconfigure the firmware of a FPGA, something not possible on an application specific integrated circuits (ASICs), allowed the relevant parts of the firmware to be redesigned. The diagnostic could in addition be enhanced to supply continuous 2D Doppler back-scattering (DBS) data, showing the ability to process 4 GB/s of data on chip. SAMI has been transferred to the NSTX-Upgrade (NSTX-U), where it has been acquiring data since May 2016. The initial results from NSTX-U, which had the primary purpose of benchmarking the system in the new environment, showed promising results. The 2D DBS power is significantly higher, which hopefully increases the accuracy of the pitch angle trace as well as reducing the required amount of averaging. In some shots the pitch angle traces exhibited a systematic scaling offset when compared to EFIT for which the cause could so far not be identified. An analysis conducted into image distortion effects due to the experiment's geometry on the image reconstruction indicates that this may well be a result of near-field effects. The diagnostic is aiming to deliver the first measurements of magnetic pitch angle profiles using the developments here. In the near future, the Vivaldi array will be replaced by a dual-polarized sinuous array to mitigate some of the array geometry induced image distortions.

4.2 Conclusions & Outlook

Both of the diagnostics presented here use the strengths of the FPGA technology to its fullest. The flexibility to modify a FPGA-based diagnostic with relative ease can free up development resources. In particular the ability to adapt an existing system from one experiment to make it compatible with another reduces the amount of redundant development work, thus saving money and time. Both diagnostic projects have shown this here. In addition FPGAs-based diagnostics have the ability

to grow with the experiment by enabling further firmware development even after the diagnostic has been finalized. In particular the SAMI diagnostic has greatly benefited from this capability. Lastly FPGAs can significantly reduce the cost of diagnostic hardware by implementing analog circuitry in high performance digital circuits. Both diagnostics would have been significantly more expensive if filtering and down-conversion were to be conducted using analog circuitry. The techniques developed as part of this thesis to handle large amounts of data can be employed in these future diagnostic designs, even ones unrelated to interferometry or microwave imaging.

The difficulties encountered during the firmware design phase were mainly a result of limited resources available on the FPGAs. However, it has to be noted that the designs presented here all employ Series 6 Xilinx FPGAs. This technology is now succeeded by the Series 7 devices, which hold significantly more resources at a lower price, while at the same time being capable of faster clock speeds. The continued improvement of the technology will therefore eliminate the resource bottleneck in the future.

A lesson learned from the development in particular on the interferometer project is the order in which FPGA-based diagnostics should be developed. The benefit of FPGA technology is that it is easily simulated, although the computational expense can be significant if large designs are to be modelled. This allows a developer to design and even test the firmware for a FPGA before actually purchasing the device. In many current diagnostic projects the designing is often conducted the other way around, i.e. the hardware is ordered before the firmware is developed. The experience here shows that this practise should be revised as it may impose unnecessary tight or even prohibitive constraints on the development process as new knowledge is gained along the way. In particular in the field of experimental physics this is more the rule than the exception.

It can be expected that novel fusion diagnostics will increasingly employ FPGA technology as the tools make development easier and the chips become more and more capable. Companies such as PI already heavily rely on it for their products, but tend to abstract much of the underlying development by using high level synthesis (HLS) tools. While this eases development, it can lead to suboptimal performance and higher cost. The development of firmware free from these proprietary constraints will benefit the fusion community and as such the world as a whole. Newer classes of FPGAs, especially architectures like the Zynq, will enable even more sophisticated diagnostics employing the power of both the FPGA and the CPU world. The tools used to develop firmware for these chips is much more sophisticated making the

development of FPGA firmware more accessible to a wider audience. In fact there are fusion diagnostics projects already employing this architecture[146].

The flexibility and cost advantage over ASICs give FPGAs a significant advantage in many fields such as high-performance computing, real-time feedback control and diagnostic development. In an ever more demanding world of fusion research we should harness its capabilities to bring us faster to the common goal of fusion power.

Bibliography

- [1] *Python*, version 3.5, Python Software Foundation, 2001-2016. [Online]. Available: <https://www.python.org> (visited on 09/01/2016).
- [2] S. van der Walt, S. C. Colbert, and G. Varoquaux, “The numpy array: A structure for efficient numerical computation,” *Computing in Science Engineering*, vol. 13, no. 2, pp. 22–30, Mar. 2011, ISSN: 1521-9615. DOI: [10.1109/MCSE.2011.37](https://doi.org/10.1109/MCSE.2011.37).
- [3] E. Jones, T. Oliphant, P. Peterson, *et al.*, *SciPy: Open source scientific tools for Python*, version 0.17, 2016. [Online]. Available: <http://www.scipy.org> (visited on 07/01/2016).
- [4] J. D. Hunter, “Matplotlib: A 2d graphics environment,” *Computing in Science Engineering*, vol. 9, no. 3, pp. 90–95, May 2007, ISSN: 1521-9615. DOI: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55).
- [5] Inkscape Project, *Inkscape*, version 0.91, Jan. 28, 2015. [Online]. Available: <https://inkscape.org> (visited on 09/01/2016).
- [6] The GIMP team, *Gnu image manipulation program*, version 2.8, Sep. 1, 2016. [Online]. Available: <http://gimp.org> (visited on 09/04/2016).
- [7] “Key world energy statistics 2015,” International Energy Agency, research rep., 2015. [Online]. Available: https://www.iea.org/publications/freepublications/publication/KeyWorld_Statistics_2015.pdf.
- [8] “Electricity: The key to a sustainable and climate-compatible energy system,” Deutsche Physikalische Gesellschaft e.V., research rep., 2010. [Online]. Available: https://www.dpg-physik.de/veroeffentlichung/broschueren/studien/energy_2011.pdf.
- [9] (Jun. 2015). BP statistical review of world energy, BP p.l.c., [Online]. Available: <https://www.bp.com/content/dam/bp/pdf/energy-economics/statistical-review-2015/bp-statistical-review-of-world-energy-2015-full-report.pdf>.
- [10] Intergovernmental Panel on Climate Change - United Nations, *Renewable energy sources and climate change mitigation: Special report of the*. CAMBRIDGE UNIV PR, Nov. 11, 2011, 1076 pp., ISBN: 1107607108. [Online]. Available: <http://www.ipcc.ch/pdf/special-reports/srren/Chapter%201%20Renewable%20Energy%20and%20Climate%20Change.pdf> (visited on 01/11/2017).

- [11] (Jul. 2, 2016). Gen iv reactor design, Generation 4 International Forum, [Online]. Available: https://www.gen-4.org/gif/jcms/c_40275/gen-iv-reactor-design.
- [12] “Uranium 2014: Resources, production and demand,” OECD Nuclear Energy Agency, and International Atomic Energy Agency, research rep., 2014. [Online]. Available: <https://www.oecd-nea.org/ndd/pubs/2014/7209-uranium-2014.pdf>.
- [13] FastFission. (). Binding energy curve - common isotopes, Wikimedia Commons, [Online]. Available: https://commons.wikimedia.org/wiki/File:Binding_energy_curve_-_common_isotopes.svg.
- [14] (Jul. 1, 2016). Inertial electrostatic confinement fusion, University of Wisconsin–Madison, [Online]. Available: <http://iec.neep.wisc.edu/operation.php>.
- [15] J. Wesson, *Tokamaks*, 3rd ed. Oxford University Press, 2004, ISBN: 0198562934.
- [16] E. Bogusch. (2004). Balance of plant requirements and concepts for tokamak reactors, 9th course on technology of fusion tokamak reactors. Erice, 26 July to 1 August 2004, EFET / Framatome ANP GmbH, [Online]. Available: <http://www-ferp.ucsd.edu/LIB/MEETINGS/0407-Erice/Bogusch.pdf> (visited on 09/01/2016).
- [17] J. Nuckolls, L. Wood, A. Thiessen, and G. Zimmerman, “Laser compression of matter to super-high densities: Thermonuclear (CTR) applications,” *Nature*, vol. 239, no. 5368, pp. 139–142, Sep. 1972. DOI: [10.1038/239139a0](https://doi.org/10.1038/239139a0). [Online]. Available: <http://www.nature.com/nature/journal/v239/n5368/pdf/239139a0.pdf>.
- [18] S. Pfalzner, *An introduction to inertial confinement fusion*, Series in Plasma Physics. CRC Press, 2006, ISBN: 0-7503-0701-3.
- [19] T. McGuire, *The lockheed martin compact fusion reactor*, PPPL Thursday Colloquium, Princeton Plasma Physics Laboratory (PPPL), Aug. 6, 2015.
- [20] M. Wakatani, *Stellarator and heliotron devices*. Oxford University Press on Demand, 1998, vol. 95, ISBN: 0-1950-7831-4.
- [21] D. Mazon, C. Fenzi, and R. Sabot, “As hot as it gets,” *Nat Phys*, vol. 12, no. 1, pp. 14–17, Jan. 2016. DOI: [10.1038/nphys3625](https://doi.org/10.1038/nphys3625).
- [22] M. Ushigome, S. Ide, S. Itoh, E. Jotaki, O. Mitarai, S. Shiraiwa, T. Suzuki, Y. Takase, S. Tanaka, T. Fujita, P. Gohil, Y. Kamada, L. Lao, T. Luce, Y. Miura, O. Naito, T. Ozeki, P. Politzer, Y. Sakamoto, and the JT-60 Team, “Development of completely solenoidless tokamak operation in jt-60u,” *Nuclear Fusion*, vol. 46, no. 2, pp. 207–213, Jan. 2006. DOI: [10.1088/0029-5515/46/2/003](https://doi.org/10.1088/0029-5515/46/2/003). [Online]. Available: <http://stacks.iop.org/0029-5515/46/i=2/a=003>.
- [23] B. J. Green, I. I. Team, and P. Teams, “Iter: Burning plasma physics experiment,” *Plasma Physics and Controlled Fusion*, vol. 45, no. 5, p. 687, 2003. DOI: [10.1088/0741-3335/45/5/312](https://doi.org/10.1088/0741-3335/45/5/312). [Online]. Available: <http://stacks.iop.org/0741-3335/45/i=5/a=312>.

- [24] A. Sykes, “Progress on spherical tokamaks,” *Plasma Physics and Controlled Fusion*, vol. 36, no. 12B, B93, 1994. DOI: [10.1088/0741-3335/36/12b/007](https://doi.org/10.1088/0741-3335/36/12b/007). [Online]. Available: <http://stacks.iop.org/0741-3335/36/i=12B/a=007>.
- [25] A. Sykes, M. P. Gryaznevich, D. Kingham, A. E. Costley, J. Hugill, G. Smith, P. Buxton, S. Ball, S. Chappell, and Z. Melhem, “Recent advances on the spherical tokamak route to fusion power,” *IEEE Trans. Plasma Sci.*, vol. 42, no. 3, pp. 482–488, Mar. 2014. DOI: [10.1109/tps.2014.2304569](https://doi.org/10.1109/tps.2014.2304569).
- [26] M. Ono and R. Kaita, “Recent progress on spherical torus research,” *Physics of Plasmas*, vol. 22, no. 4, p. 040 501, Apr. 2015. DOI: [10.1063/1.4915073](https://doi.org/10.1063/1.4915073).
- [27] A. Li Puma, J. Berton, B. Brañas, L. Bühler, J. Doncel, U. Fischer, W. Farabolini, L. Giancarli, D. Maisonnier, P. Pereslavytsev, S. Raboin, J.-F. Salavy, P. Sardain, J. Szczepanski, and D. Ward, “Breeding blanket design and systems integration for a helium-cooled lithium–lead fusion power plant,” *Fusion Engineering and Design*, vol. 81, no. 1-7, pp. 469–476, Feb. 2006. DOI: [10.1016/j.fusengdes.2005.08.041](https://doi.org/10.1016/j.fusengdes.2005.08.041).
- [28] (Aug. 23, 2016). Tokamak energy - a faster way to fusion, Tokamak Energy, [Online]. Available: <http://www.tokamakenergy.co.uk/about-us/>.
- [29] M. Lehnen, K. Aleynikova, P. Aleynikov, D. Campbell, P. Drewelow, N. Eidiētis, Y. Gasparyan, R. Granetz, Y. Gribov, N. Hartmann, E. Hollmann, V. Izzo, S. Jachmich, S.-H. Kim, M. Kočan, H. Koslowski, D. Kovalenko, U. Kruezi, A. Loarte, S. Maruyama, G. Matthews, P. Parks, G. Pautasso, R. Pitts, C. Reux, V. Riccardo, R. Roccella, J. Snipes, A. Thornton, and P. de Vries, “Disruptions in ITER and strategies for their control and mitigation,” *Journal of Nuclear Materials*, vol. 463, pp. 39–48, 2015, PLASMA-SURFACE INTERACTIONS 21 Proceedings of the 21st International Conference on Plasma-Surface Interactions in Controlled Fusion Devices Kanazawa, Japan May 26-30, 2014, ISSN: 0022-3115. DOI: [10.1016/j.jnucmat.2014.10.075](https://doi.org/10.1016/j.jnucmat.2014.10.075). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0022311514007594>.
- [30] A. Costley, T. Sugie, G. Vayakis, and C. Walker, “Technological challenges of ITER diagnostics,” *Fusion Engineering and Design*, vol. 74, no. 1-4, pp. 109–119, Nov. 2005. DOI: [10.1016/j.fusengdes.2005.08.026](https://doi.org/10.1016/j.fusengdes.2005.08.026). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S092037960500390X>.
- [31] A. Donné, A. Costley, R. Barnsley, H. Bindslev, R. Boivin, G. Conway, R. Fisher, R. Giannella, H. Hartfuss, M. von Hellermann, E. Hodgson, L. Ingesson, K. Itami, D. Johnson, Y. Kawano, T. Kondoh, A. Krasilnikov, Y. Kusama, A. Litnovsky, P. Lotte, P. Nielsen, T. Nishitani, F. Orsitto, B. Peterson, G. Razdobarin, J. Sanchez, M. Sasao, T. Sugie, G. Vayakis, V. Voitsenya, K. Vukolov, C. Walker, K. Young, and the ITPA Topical Group on Diagnostics, “Chapter 7: Diagnostics,” *Nucl. Fusion*, vol. 47, no. 6, S337–S384, Jun. 2007. DOI: [10.1088/0029-5515/47/6/s07](https://doi.org/10.1088/0029-5515/47/6/s07). [Online]. Available: <http://iopscience.iop.org/article/10.1088/0029-5515/47/6/S07/meta>.

- [32] J. Snipes, D. Beltran, T. Casper, Y. Gribov, A. Isayama, J. Lister, S. Simrock, G. Vayakis, A. Winter, Y. Yang, and L. Zabeo, “Actuator and diagnostic requirements of the ITER plasma control system,” *Fusion Engineering and Design*, vol. 87, no. 12, pp. 1900–1906, Dec. 2012. DOI: [10.1016/j.fusengdes.2012.04.002](https://doi.org/10.1016/j.fusengdes.2012.04.002). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0920379612002530>.
- [33] D. Moreau and I. Voitsekhovitch, “Plasma control issues for an advanced steady state tokamak reactor,” *Nuclear Fusion*, vol. 39, no. 5, p. 685, 1999. [Online]. Available: <http://stacks.iop.org/0029-5515/39/i=5/a=308>.
- [34] S. Brown and J. Rose. (2000). Architecture of fpgas and cplds: A tutorial. D. of Electrical and C. E. U. of Toronto, Eds., [Online]. Available: <http://www.eecg.toronto.edu/~jayar/pubs/brown/survey.pdf> (visited on 09/01/2016).
- [35] C. Maxfield, *The design warrior’s guide to fpgas, devices, tools and flows, 2004*. Elsevier, 2004, ISBN: 0-7506-7604-3,
- [36] PC Tech. (Sep. 11, 2011). Principles of cpu architecture – logic gates, mosfets and voltage, [Online]. Available: <https://www.pctechguide.com/cpu-architecture/principles-of-cpu-architecture-logic-gates-mosfets-and-voltage> (visited on 09/18/2016).
- [37] (Sep. 1, 2016). Free fpga and embedded design training, Xilinx Inc., [Online]. Available: <http://www.xilinx.com/training/free-video-courses.htm> (visited on 09/01/2016).
- [38] PCI-SIG. (2016). Pci express base specification, [Online]. Available: <http://pcisig.com/specifications> (visited on 09/01/2016).
- [39] R. Bittner, E. Ruf, and A. Forin, “Direct GPU/FPGA communication via PCI express,” *Cluster Computing*, vol. 17, no. 2, pp. 339–348, May 2013, This document is the authors’ preliminary version of an article to appear in Cluster Computing, a publication of Springer-Verlag, which holds copyright to the final work. The official publication is available at <http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s10586-013-0280-9..> DOI: [10.1007/s10586-013-0280-9](https://doi.org/10.1007/s10586-013-0280-9). [Online]. Available: <http://research.microsoft.com/apps/pubs/default.aspx?id=192884>.
- [40] Wikibooks. (Sep. 2, 2016). Microprocessor design, Wikimedia Foundation, [Online]. Available: https://en.wikibooks.org/wiki/Microprocessor_Design/Pipelined_Processors (visited on 09/02/2016).
- [41] R. Bittner. (2012). Speedy bus mastering pci express, Microsoft Research, [Online]. Available: <http://research.microsoft.com/pubs/172728/20120628%20Speedy%20PCIe%20FPL%20Final.pdf>.
- [42] L. H. Crockett, R. A. Elliot, M. A. Enderwitz, and R. W. Stewart, *The zynq book, Embedded processing with the arm® cortex®-a9 on the xilinx® zynq®-7000 all programmable soc*, 1st ed. Jul. 1, 2014. [Online]. Available: www.zynqbook.com (visited on 09/01/2016).

- [43] J. Hussein and G. Swift. (2015). Mitigating single event upsets. X. Inc., Ed., [Online]. Available: http://www.xilinx.com/support/documentation/white_papers/wp395-Mitigating-SEUs.pdf.
- [44] J. J. Wang. (2003). Radiation effects in fpgas, Actel Corporation, [Online]. Available: <http://lhc-electronics-workshop.web.cern.ch/LHC-electronics-workshop/2003/PLENARYT/WANG.PDF>.
- [45] M. Gokhale, P. Graham, E. Johnson, N. Rollins, and M. Wirthlin, “Dynamic reconfiguration for management of radiation-induced faults in FPGAs,” in *18th International Parallel and Distributed Processing Symposium, 2004. Proceedings.*, Institute of Electrical and Electronics Engineers (IEEE), 2004. DOI: [10.1109/ipdps.2004.1303127](https://doi.org/10.1109/ipdps.2004.1303127). [Online]. Available: <http://ieeexplore.ieee.org/document/1303127/>.
- [46] R. Scannell, A. Kirk, N. B. Ayed, P. G. Carolan, G. Cunningham, J. McCone, S. L. Prunty, and M. J. Walsh, “Experimental investigation into elm filament formation on mast,” *Plasma Physics and Controlled Fusion*, vol. 49, no. 9, p. 1431, Jul. 2007. DOI: [10.1088/0741-3335/49/9/006](https://doi.org/10.1088/0741-3335/49/9/006).
- [47] T. O’Gorman, G. Naylor, R. Scannell, G. Cunningham, K. J. Brunner, R. Martin, and D. Croft, “Design of a real-time two-color interferometer for MAST upgrade,” *Review of Scientific Instruments*, vol. 85, no. 11, 11D861, p. 11D861, Nov. 2014. DOI: [10.1063/1.4894394](https://doi.org/10.1063/1.4894394).
- [48] R. Scannell, M. J. Walsh, P. G. Carolan, N. J. Conway, A. C. Darke, M. R. Dunstan, D. Hare, and S. L. Prunty, “Enhanced edge thomson scattering on mast,” *Review of Scientific Instruments*, vol. 77, no. 10, 10E510, pp. -, 2006. DOI: [10.1063/1.2237488](https://doi.org/10.1063/1.2237488).
- [49] W. W. Heidbrink, “Basic physics of alfvén instabilities driven by energetic particles in toroidally confined plasmas,” *Physics of Plasmas*, vol. 15, no. 5, 055501, p. 055501, 2008. DOI: [10.1063/1.2838239](https://doi.org/10.1063/1.2838239). [Online]. Available: <http://dx.doi.org/10.1063/1.2838239>.
- [50] S. M. Mahajan, “Spectrum of alfvén waves, a brief review,” *Physica Scripta*, vol. 1995, no. T60, p. 160, 1995. [Online]. Available: <http://stacks.iop.org/1402-4896/1995/i=T60/a=020>.
- [51] R. B. White, E. Fredrickson, D. Darrow, M. Zarnstorff, R. Wilson, S. Zweben, K. Hill, Y. Chen, and G. Fu, “Toroidal alfvén eigenmode-induced ripple trapping,” *Physics of Plasmas*, vol. 2, no. 8, pp. 2871–2873, 1995. DOI: [10.1063/1.871452](https://doi.org/10.1063/1.871452). [Online]. Available: <http://dx.doi.org/10.1063/1.871452>.
- [52] E. D. Fredrickson, N. Gorelenkov, C. Z. Cheng, R. Bell, D. Darrow, D. Johnson, S. Kaye, B. LeBlanc, J. Menard, S. Kubota, and W. Peebles, “Observation of compressional alfvén modes during neutral-beam heating on the national spherical torus experiment,” *Phys. Rev. Lett.*, vol. 87, no. 14, Sep. 2001. DOI: [10.1103/physrevlett.87.145001](https://doi.org/10.1103/physrevlett.87.145001).
- [53] M. A. Van Zeeland and G. J. Kramer, “Measurements of alfvén eigenmodes in fusion plasmas,” *AIP Conference Proceedings*, vol. 988, no. 1, pp. 103–109, 2008. DOI: [10.1063/1.2905049](https://doi.org/10.1063/1.2905049).

- [54] S. E. Sharapov, B. Alper, H. L. Berk, D. N. Borba, B. N. Breizman, C. D. Challis, A. Fasoli, N. C. Hawkes, T. C. Hender, J. Mailloux, S. D. Pinches, D. Testa, and E. work programme, “Alfvén wave cascades in a tokamak,” *Physics of Plasmas*, vol. 9, no. 5, p. 2027, 2002. DOI: [10.1063/1.1448346](https://doi.org/10.1063/1.1448346).
- [55] M. Gryaznevich, S. Sharapov, M. Lilley, S. Pinches, A. Field, D. Howell, D. Keeling, R. Martin, H. Meyer, H. Smith, R. Vann, P. Denner, and E. Verwichte, “Recent experiments on alfvén eigenmodes in MAST,” *Nucl. Fusion*, vol. 48, no. 8, p. 084003, Jun. 2008. DOI: [10.1088/0029-5515/48/8/084003](https://doi.org/10.1088/0029-5515/48/8/084003).
- [56] K. J. Brunner, T. O’Gorman, G. Naylor, R. Scannell, G. Cunningham, R. Sharples, and N. A. Dipper, “Fpga-based high bandwidth integral electron density interferometer for mast-u,” in *Proceedings of the 1st EPS conference on Plasma Diagnostics (ECPD2015). 14-17 April 2015. Frascati, Italy*, Proceedings of Science, 2015, p. 138. [Online]. Available: http://pos.sissa.it/archive/conferences/240/138/ECPD2015_138.pdf.
- [57] I. Hutchinson, *Principles of plasma diagnostics*, 2nd ed. The Pitt Building, Trumpington Street, Cambridge: Cambridge University Press, 2002, ISBN: 978-0-521-67574-1.
- [58] M. A. Van Zeeland, R. L. Boivin, D. L. Brower, T. N. Carlstrom, J. A. Chavez, W. X. Ding, R. Feder, D. Johnson, L. Lin, R. C. O’Neill, and C. Watts, “Conceptual design of the tangentially viewing combined interferometer-polarimeter for ITER density measurements,” *Review of Scientific Instruments*, vol. 84, no. 4, 043501, pp. -, 2013. DOI: [10.1063/1.4798602](https://doi.org/10.1063/1.4798602).
- [59] M. A. V. Zeeland and T. N. Carlstrom, “Phase error correction method for a vibration compensated interferometer,” *Review of Scientific Instruments*, vol. 75, no. 10, pp. 3423–3425, Oct. 2004. DOI: [10.1063/1.1786641](https://doi.org/10.1063/1.1786641). [Online]. Available: <http://dx.doi.org/10.1063/1.1786641>.
- [60] A. Mlynek, G. Pautasso, M. Marashek, H. Eixenberger, and A. U. team, “Infrared interferometry with submicrosecond time resolution in massive gas injection experiments on asdex upgrade,” *Fusion Science and Technology*, vol. 61, no. 4, pp. 290–300, 2012.
- [61] A. Murari, L. Zabeo, A. Boboc, D. Mazon, and M. Riva, “Real-time recovery of the electron density from interferometric measurements affected by fringe jumps,” *Review of Scientific Instruments*, vol. 77, no. 7, 073505, 2006. DOI: [10.1063/1.2219731](https://doi.org/10.1063/1.2219731).
- [62] T. Akiyama, R. Yasuhara, K. Kawahata, K. Nakayama, S. Okajima, K. Urabe, K. Terashima, and N. Shirai, “Development of dispersion interferometer for magnetic confinement plasmas and high-pressure plasmas,” *Journal of Instrumentation*, vol. 10, no. 09, P09022, 2015. [Online]. Available: <http://stacks.iop.org/1748-0221/10/i=09/a=P09022>.
- [63] R. König, J. Baldzuhn, W. Biel, C. Biedermann, H. S. Bosch, S. Bozhenkov, T. Bräuer, B. B. de Carvalho, R. Burhenn, B. Buttenschön, G. Cseh, A. Czarnecka, M. Endler, V. Erckmann, T. Estrada, J. Geiger, O. Grulke, D. Hartmann, D. Hathiramani, M. Hirsch, S. Jabłonski, M. Jakubowski, J.

- Kaczmarczyk, T. Klinger, S. Klose, G. Kocsis, P. Kornejew, A. Krämer-Flecken, T. Kremeyer, M. Krychowiak, M. Kubkowska, A. Langenberg, H. P. Laqua, M. Laux, Y. Liang, A. Lorenz, A. O. Marchuk, V. Moncada, O. Neubauer, U. Neuner, J. W. Oosterbeek, M. Otte, N. Pablant, E. Pasch, T. S. Pedersen, K. Rahbarnia, L. Ryc, O. Schmitz, W. Schneider, H. Schuhmacher, B. Schweer, T. Stange, H. Thomsen, J.-M. Travers, T. Szepesi, U. Wenzel, A. Werner, B. Wiegel, T. Windisch, R. Wolf, G. A. Wurden, D. Zhang, A. Zimbal, and S. Zoletnik, “The set of diagnostics for the first operation campaign of the wendelstein 7-x stellarator,” *Journal of Instrumentation*, vol. 10, no. 10, P10002–P10002, Oct. 2015. DOI: [10.1088/1748-0221/10/10/p10002](https://doi.org/10.1088/1748-0221/10/10/p10002). [Online]. Available: <http://stacks.iop.org/1748-0221/10/i=10/a=P10002>.
- [64] T. Akiyama, R. Yasuhara, K. Kawahata, S. Okajima, and K. Nakayama, “Dispersion interferometer using modulation amplitudes on LHD (invited)a),” *Review of Scientific Instruments*, vol. 85, no. 11, p. 11D301, Nov. 2014. DOI: [10.1063/1.4886777](https://doi.org/10.1063/1.4886777).
- [65] O. H. R. (OHWR). (Jul. 2, 2016). Simple PCIe carrier v4 (SPEC), CERN, [Online]. Available: <http://www.ohwr.org/projects/spec> (visited on 07/01/2016).
- [66] —, (Jul. 1, 2016). FMC ADC 100m 14b 4cha, CERN, [Online]. Available: <http://www.ohwr.org/projects/fmc-adc-100m14b4cha> (visited on 07/01/2016).
- [67] (Jul. 23, 2016). SP605 hardware user guide, Ug526 (v1.8), Xilinx Inc., [Online]. Available: http://www.xilinx.com/support/documentation/boards_and_kits/ug526.pdf (visited on 07/23/2016).
- [68] (Jul. 1, 2016). Xilinx core generator system, Xilinx Inc., [Online]. Available: <https://www.xilinx.com/products/design-tools/coregen.html> (visited on 07/01/2016).
- [69] —, (Jul. 2, 2016). Platform-independent core collection, CERN, [Online]. Available: <http://www.ohwr.org/projects/general-cores> (visited on 07/01/2016).
- [70] (Jul. 1, 2016). Aurora 8B/10B, Xilinx Inc., [Online]. Available: <http://www.xilinx.com/products/intellectual-property/aurora8b10b.html> (visited on 07/01/2016).
- [71] (2013). Introduction to the filter design and analysis tool (fdatool), The MathWorks Inc., [Online]. Available: <http://uk.mathworks.com/help/signal/examples/introduction-to-the-filter-design-and-analysis-tool-fdatool.html> (visited on 09/12/2016).
- [72] B. Lakshmi and A. S. Dhar, “CORDIC architectures: A survey,” *VLSI Design*, vol. 2010, 2010. DOI: [10.1155/2010/794891](https://doi.org/10.1155/2010/794891).
- [73] (Sep. 4, 2016). Logicore ip cordic v5.0 - product specification, Xilinx Inc., [Online]. Available: http://www.xilinx.com/support/documentation/ip_documentation/cordic/v5_0/ds858_cordic.pdf (visited on 09/04/2016).

- [74] D. E. T. Romero and G. Jovanovic, “Digital FIR Hilbert transformers: Fundamentals and efficient design methods,” in *MATLAB - A Fundamental Tool for Scientific Computing and Engineering Applications - Volume 1*, InTech, Sep. 2012. DOI: [10.5772/46451](https://doi.org/10.5772/46451). [Online]. Available: <http://dx.doi.org/10.5772/46451>.
- [75] Mini-Circuits. (Jul. 10, 2016). Blp-2.5+ coaxial low-pass filter performance data, [Online]. Available: <http://www.minicircuits.com/pdfs/BLP-2.5+.pdf>.
- [76] 2.-2. Keysight Technologies, *33600a series trueform waveform generators - data sheet*, Aug. 4, 2014, 15ff. [Online]. Available: <http://literature.cdn.keysight.com/litweb/pdf/5991-3272EN.pdf?id=2431617> (visited on 07/10/2016).
- [77] A. Chabchoub and M. Fink, “Time-reversal generation of rogue waves,” *Phys. Rev. Lett.*, vol. 112, no. 12, Jul. 11, 2016. DOI: [10.1103/physrevlett.112.124101](https://doi.org/10.1103/physrevlett.112.124101). [Online]. Available: <http://dx.doi.org/10.1103/PhysRevLett.112.124101>.
- [78] O. H. R. (OHWR). (Jul. 1, 2016). White rabbit switch, CERN, [Online]. Available: <http://www.ohwr.org/projects/white-rabbit/wiki/Switch> (visited on 07/01/2016).
- [79] (Jul. 1, 2016). An FPGA IP core for easy dma over PCIe with windows and linux, Xillybus Ltd., [Online]. Available: <http://xillybus.com/> (visited on 03/06/2015).
- [80] A. Sykes, R. Akers, L. Appel, E. Arends, P. Carolan, N. Conway, G. Counsell, G. Cunningham, A. Dnestrovskij, Y. Dnestrovskij, A. Field, S. Fielding, M. Gryaznevich, S. Korsholm, E. Laird, R. Martin, M. Nightingale, C. Roach, M. Tournianski, M. Walsh, C. Warrick, H. Wilson, S. You, M. Team, and N. Team, “First results from MAST,” *Nuclear Fusion*, vol. 41, no. 10, pp. 1423–1433, Oct. 2001. DOI: [10.1088/0029-5515/41/10/310](https://doi.org/10.1088/0029-5515/41/10/310). [Online]. Available: http://www-pub.iaea.org/MTCD/Publications/PDF/csp_008c/pdf/ov4_1.pdf (visited on 01/24/2017).
- [81] M. Brombin, A. Boboc, L. Zabeo, and A. M. and, “Real-time electron density measurements from cotton–mouton effect in JET machine,” 10, Oct. 2008, 10E718. DOI: [10.1063/1.2956831](https://doi.org/10.1063/1.2956831).
- [82] J. Wesson, R. Gill, M. Hugon, F. Schüller, J. Snipes, D. Ward, D. Bartlett, D. Campbell, P. Duperrex, A. Edwards, R. Granetz, N. Gottardi, T. Hender, E. Lazzaro, P. Lomas, N. L. Cardozo, K. Mast, M. Nave, N. Salmon, P. Smeulders, P. Thomas, B. Tubbing, M. Turner, and A. Weller, “Disruptions in jet,” *Nuclear Fusion*, vol. 29, no. 4, p. 641, 1989. [Online]. Available: <http://stacks.iop.org/0029-5515/29/i=4/a=009>.
- [83] Cisco Systems, Inc. (2008). Fiber types in gigabit optical communications, Cisco Systems, Inc., [Online]. Available: https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/transceiver-modules/white_paper_c11-463661.html (visited on 02/04/2017).

- [84] A. Mlynek, G. Schramm, H. Eixenberger, G. Sips, K. McCormick, M. Zilker, K. Behler, J. Eheberg, and A. U. Team, “Design of a digital multiradian phase detector and its application in fusion plasma interferometry,” *Rev. Sci. Instrum.*, vol. 81, no. 3, p. 033 507, 2010. DOI: [10.1063/1.3340944](https://doi.org/10.1063/1.3340944).
- [85] (Jan. 30, 2015). Spartan-6 FPGA electrical characteristics, Xilinx Inc., [Online]. Available: xilinx.com/support/documentation/data_sheets/ds162.pdf (visited on 07/17/2016).
- [86] G. Pautasso, D. Coster, T. Eich, J. C. Fuchs, O. Gruber, A. Gude, A. Herrmann, V. Igochine, C. Konz, B. Kurzan, K. Lackner, T. Lunt, M. Maraschek, A. Mlynek, B. Reiter, V. Rohde, Y. Zhang, X. Bonnin, M. Beck, and G. Prausner, “Disruption studies in ASDEX upgrade in view of ITER,” *Plasma Phys. Control. Fusion*, vol. 51, no. 12, p. 124 056, Nov. 2009. DOI: [10.1088/0741-3335/51/12/124056](https://doi.org/10.1088/0741-3335/51/12/124056). [Online]. Available: stacks.iop.org/PPCF/51/124056.
- [87] EQUIPE TFR, “Decoupled electrons and abnormal diffusion of electrons trapped in local mirrors of the tfr tokamak,” *Nuclear Fusion*, vol. 16, no. 3, p. 473, 1976. [Online]. Available: <http://stacks.iop.org/0029-5515/16/i=3/a=012>.
- [88] M. Maraschek, “Control of neoclassical tearing modes,” *Nucl. Fusion*, vol. 52, no. 7, p. 074 007, Jul. 2012. DOI: [10.1088/0029-5515/52/7/074007](https://doi.org/10.1088/0029-5515/52/7/074007). [Online]. Available: <http://stacks.iop.org/NF/52/074007>.
- [89] H. Zohm, “Stabilization of neoclassical tearing modes by electron cyclotron current drive,” *Physics of Plasmas*, vol. 4, no. 9, p. 3433, 1997. DOI: [10.1063/1.872487](https://doi.org/10.1063/1.872487).
- [90] H. Zohm, G. Gantenbein, F. Leuterer, A. Manini, M. Maraschek, Q. Yu, and the ASDEX Upgrade Team, “Control of mhd instabilities by eccd: Asdex upgrade results and implications for iter,” *Nuclear Fusion*, vol. 47, no. 3, pp. 228–232, Mar. 2007. DOI: [10.1088/0029-5515/47/3/010](https://doi.org/10.1088/0029-5515/47/3/010). [Online]. Available: <http://stacks.iop.org/0029-5515/47/i=3/a=010>.
- [91] F. Waelbroeck, “Theory and observations of magnetic islands,” *Nuclear Fusion*, vol. 49, no. 10, p. 104 025, Sep. 2009. DOI: [10.1088/0029-5515/49/10/104025](https://doi.org/10.1088/0029-5515/49/10/104025). [Online]. Available: <http://stacks.iop.org/0029-5515/49/i=10/a=104025>.
- [92] I. T. Chapman, “Controlling sawtooth oscillations in tokamak plasmas,” *Plasma Physics and Controlled Fusion*, vol. 53, no. 1, p. 013 001, Nov. 2010. DOI: [10.1088/0741-3335/53/1/013001](https://doi.org/10.1088/0741-3335/53/1/013001). [Online]. Available: <http://stacks.iop.org/0741-3335/53/i=1/a=013001>.
- [93] P. de Vries, A. Donné, S. Heijnen, C. Hugenholtz, A. Kramer-Flecken, F. Schuller, and G. Waidmann, “Density profile peaking inside m/n=2/1 magnetic islands in textor-94,” *Nuclear Fusion*, vol. 37, no. 11, p. 1641, 1997. [Online]. Available: <http://stacks.iop.org/0029-5515/37/i=11/a=I13>.
- [94] A. Y. Wong. (1977). Introduction to experimental plasma physics, Physics Department - University of California at Los Angeles, [Online]. Available: <http://www.physics.ucla.edu/plasmalab/images/> (visited on 09/20/2016).

- [95] N. Gorelenkov, E. Fredrickson, E. Belova, C. Cheng, D. Gates, S. Kaye, and R. White, “Theory and observations of high frequency alfvén eigenmodes in low aspect ratio plasmas,” *Nuclear Fusion*, vol. 43, no. 4, pp. 228–233, Mar. 2003. DOI: [10.1088/0029-5515/43/4/302](https://doi.org/10.1088/0029-5515/43/4/302). [Online]. Available: <http://stacks.iop.org/0029-5515/43/i=4/a=302>.
- [96] E. V. Belova, N. N. Gorelenkov, E. D. Fredrickson, K. Tritz, and N. A. Crocker, “Coupling of neutral-beam-driven compressional alfvén eigenmodes to kinetic alfvén waves in nstx tokamak and energy channeling,” *Phys. Rev. Lett.*, vol. 115, no. 1, p. 015 001, 1 Jun. 2015. DOI: [10.1103/physrevlett.115.015001](https://doi.org/10.1103/physrevlett.115.015001).
- [97] D. Stutman, L. Delgado-Aparicio, N. Gorelenkov, M. Finkenthal, E. Fredrickson, S. Kaye, E. Mazzucato, and K. Tritz, “Correlation between electron transport and shear alfvén activity in the national spherical torus experiment,” *Phys. Rev. Lett.*, vol. 102, no. 11, Mar. 2009. DOI: [10.1103/PhysRevLett.102.115002](https://doi.org/10.1103/PhysRevLett.102.115002).
- [98] J.-M. Noterdaeme, F. Wesner, M. Brambilla, R. Fritsch, H.-J. Kutsch, and M. Söll, “The ASDEX upgrade ICRH antenna,” *Fusion Engineering and Design*, vol. 24, no. 1-2, pp. 65–74, Feb. 1994. DOI: [10.1016/0920-3796\(94\)90037-x](https://doi.org/10.1016/0920-3796(94)90037-x). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/092037969490037X>.
- [99] F. Wagner, G. Becker, K. Behringer, D. Campbell, A. Eberhagen, W. Engelhardt, G. Fussmann, O. Gehre, J. Gernhardt, G. v. Gierke, G. Haas, M. Huang, F. Karger, M. Keilhacker, O. Klüber, M. Kornherr, K. Lackner, G. Lisitano, G. G. Lister, H. M. Mayer, D. Meisel, E. R. Müller, H. Murmann, H. Niedermeyer, W. Poschenrieder, H. Rapp, H. Röhr, F. Schneider, G. Siller, E. Speth, A. Stäbler, K. H. Steuer, G. Venus, O. Vollmer, and Z. Yü, “Regime of improved confinement and high beta in neutral-beam-heated divertor discharges of the ASDEX tokamak,” *Phys. Rev. Lett.*, vol. 49, no. 19, pp. 1408–1412, Nov. 1982. DOI: [10.1103/physrevlett.49.1408](https://doi.org/10.1103/physrevlett.49.1408). [Online]. Available: <http://dx.doi.org/10.1103/PhysRevLett.49.1408>.
- [100] F. Wagner, “A quarter-century of h-mode studies,” *Plasma Physics and Controlled Fusion*, vol. 49, no. 12B, B1, 2007. [Online]. Available: <http://stacks.iop.org/0741-3335/49/i=12B/a=S01>.
- [101] M. Wade, “Physics and engineering issues associated with edge localized mode control in ITER,” *Fusion Engineering and Design*, vol. 84, no. 2-6, pp. 178–185, Jun. 2009. DOI: [10.1016/j.fusengdes.2009.01.063](https://doi.org/10.1016/j.fusengdes.2009.01.063).
- [102] 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*, vol. 48, no. 5A, A71–A84, Apr. 2006. DOI: [10.1088/0741-3335/48/5a/s06](https://doi.org/10.1088/0741-3335/48/5a/s06). [Online]. Available: <http://dx.doi.org/10.1088/0741-3335/48/5a/s06>.

- [103] R. C. Isler, “An overview of charge-exchange spectroscopy as a plasma diagnostic,” *Plasma Physics and Controlled Fusion*, vol. 36, no. 2, pp. 171–208, Feb. 1994. DOI: [10.1088/0741-3335/36/2/001](https://doi.org/10.1088/0741-3335/36/2/001). [Online]. Available: <http://iopscience.iop.org/article/10.1088/0741-3335/36/2/001/meta;jsessionid=E6F1D66B55628FFD1F6146077A0BB0BE.c5.iopscience.cld.iop.org>.
- [104] J. L. Terry, E. S. Marmor, R. B. Howell, M. Bell, A. Cavallo, E. Fredrickson, A. Ramsey, G. L. Schmidt, B. Stratton, G. Taylor, and M. E. Mauel, “Measurement of internal magnetic field pitch using li pellet injection on TFTR (invited),” *Rev. Sci. Instrum.*, vol. 61, no. 10, p. 2908, 1990. DOI: [10.1063/1.1141775](https://doi.org/10.1063/1.1141775).
- [105] K. Kamiya, T. Fujita, A. Kojima, and H. Kubo, “Zeeman polarimetry measurement for edge current density determination using li-beam probe on JT-60u,” *Rev. Sci. Instrum.*, vol. 81, no. 3, p. 033 502, 2010. DOI: [10.1063/1.3309793](https://doi.org/10.1063/1.3309793).
- [106] F. M. Levinton, R. J. Fonck, G. M. Gammel, R. Kaita, H. W. Kugel, E. T. Powell, and D. W. Roberts, “Magnetic field pitch-angle measurements in the pbx-m tokamak using the motional stark effect,” *Phys. Rev. Lett.*, vol. 63, pp. 2060–2063, 19 Nov. 1989. DOI: [10.1103/PhysRevLett.63.2060](https://doi.org/10.1103/PhysRevLett.63.2060). [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevLett.63.2060>.
- [107] (Jul. 19, 2016). Very large array, National Radio Astronomy Observatory (NRAO), [Online]. Available: <http://www.vla.nrao.edu/>.
- [108] A. E. Costley, “Towards diagnostics for a fusion reactor,” *IEEE Trans. Plasma Sci.*, vol. 38, no. 10, pp. 2934–2943, Oct. 2010. DOI: [10.1109/tps.2010.2064182](https://doi.org/10.1109/tps.2010.2064182).
- [109] 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*, vol. 7, no. 10, P10016, 2012. [Online]. Available: <http://stacks.iop.org/1748-0221/7/i=10/a=P10016>.
- [110] S. Freethy, “Synthetic aperture imaging of b-x-o mode conversion,” PhD thesis, University of York, Mar. 2012.
- [111] D. Thomas, K. Brunner, S. Freethy, B. Huang, V. Shevchenko, and R. Vann, “2D Doppler backscattering using synthetic aperture microwave imaging of mast edge plasmas,” *Nuclear Fusion*, vol. 56, no. 2, p. 026 013, Jan. 15, 2016. [Online]. Available: <http://stacks.iop.org/0029-5515/56/i=2/a=026013>.
- [112] (Jul. 23, 2016). Ml605 hardware user guide, Ug534 (v1.8), Xilinx Inc., [Online]. Available: http://www.xilinx.com/support/documentation/boards_and_kits/ug534.pdf (visited on 07/23/2016).
- [113] B. Huang, R. Vann, S. Freethy, R. Myers, G. Naylor, R. Sharples, and V. Shevchenko, “FPGA-based embedded linux technology in fusion: The mast microwave imaging system,” *Fusion Engineering and Design*, vol. 87, no. 12, pp. 2106–2111, 2012, Proceedings of the 8th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research, ISSN: 0920-3796. DOI: [10.1016/j.fusengdes.2012.06.004](https://doi.org/10.1016/j.fusengdes.2012.06.004). [On-

- line]. Available: <http://www.sciencedirect.com/science/article/pii/S0920379612003079>.
- [114] J. C. Chorley, R. J. Akers, K. J. Brunner, N. A. Dipper, S. J. Freethy, R. M. Sharples, V. F. Shevchenko, D. A. Thomas, and R. G. L. Vann, “Gpu-based data processing for 2-d microwave imaging on mast,” *Fusion Science and Technology*, vol. 69, no. 3, pp. 643–654, 2016. DOI: [10.13182/FST15-188](https://doi.org/10.13182/FST15-188). [Online]. Available: <http://epubs.ans.org/?a=38361>.
- [115] 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*, vol. 87, no. 11, p. 11D902, Sep. 2016. DOI: [10.1063/1.4962253](https://doi.org/10.1063/1.4962253). [Online]. Available: <http://scitation.aip.org/content/aip/journal/rsi/87/11/10.1063/1.4962253>.
- [116] K. J. Brunner, J. C. Chorley, N. A. Dipper, G. Naylor, R. M. Sharples, G. Taylor, D. A. Thomas, and R. G. L. Vann, “Modifications to the synthetic aperture microwave imaging diagnostic,” *Review of Scientific Instruments*, vol. 87, no. 11, 11E129, 11E129, Sep. 2016. DOI: [10.1063/1.4961283](https://doi.org/10.1063/1.4961283). [Online]. Available: <http://scitation.aip.org/content/aip/journal/rsi/87/11/10.1063/1.4961283>.
- [117] T. Williams, “Full-wave simulation of high-frequency electromagnetic propagation through inhomogeneous plasma,” PhD thesis, University of York, Sep. 2014.
- [118] H. P. Laqua, “Electron bernstein wave heating and diagnostic,” *Plasma Physics and Controlled Fusion*, vol. 49, no. 4, R1, 2007. [Online]. Available: <http://stacks.iop.org/0741-3335/49/i=4/a=R01>.
- [119] G. D. Conway, J. Schirmer, S. Klenge, W. Suttrop, E. Holzhauser, and the ASDEX Upgrade Team, “Plasma rotation profile measurements using Doppler reflectometry,” *Plasma Phys. Control. Fusion*, vol. 46, no. 6, pp. 951–970, Apr. 2004. DOI: [10.1088/0741-3335/46/6/003](https://doi.org/10.1088/0741-3335/46/6/003).
- [120] M. Hirsch, E. Holzhauser, J. Baldzuhn, B. Kurzan, and B. Scott, “Doppler reflectometry for the investigation of propagating density perturbations,” *Plasma Phys. Control. Fusion*, vol. 43, no. 12, pp. 1641–1660, Oct. 2001. DOI: [10.1088/0741-3335/43/12/302](https://doi.org/10.1088/0741-3335/43/12/302).
- [121] J. Hillesheim, N. Crocker, W. Peebles, H. Meyer, A. Meakins, A. Field, D. Dunai, M. Carr, and N. Hawkes, “Doppler backscattering for spherical tokamaks and measurement of high-k-density fluctuation wavenumber spectrum in MAST,” *Nucl. Fusion*, vol. 55, no. 7, p. 073024, Jun. 2015. DOI: [10.1088/0029-5515/55/7/073024](https://doi.org/10.1088/0029-5515/55/7/073024).
- [122] D.-L. Liu and R. Waag, “About the application of the van cittert-zernike theorem in ultrasonic imaging,” *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, vol. 42, no. 4, pp. 590–601, Jul. 1995. DOI: [10.1109/58.393102](https://doi.org/10.1109/58.393102).

- [123] *First imagery generated by near-field real-time aperture synthesis passive millimetre wave imagers at 94 GHz and 183 GHz*, vol. 7837, 2010, pp. 78370I-78370I-10. DOI: [10.1117/12.860396](https://doi.org/10.1117/12.860396).
- [124] A. Luthra, S. Kassam, and R. Bernardi, “Body imaging using vectorial addition of acoustic reflection to achieve effect of scanning beam continuously focused in range,” US Patent 4,604,697, Aug. 1986. [Online]. Available: <https://www.google.com/patents/US4604697>.
- [125] M. Younis, C. Fischer, and W. Wiesbeck, “Digital beamforming in sar systems,” *IEEE Trans. Geosci. Remote Sensing*, vol. 41, no. 7, pp. 1735–1739, Jul. 2003. DOI: [10.1109/tgrs.2003.815662](https://doi.org/10.1109/tgrs.2003.815662).
- [126] Y. Zheng, R. Goubran, and M. El-Tanany, “Robust near-field adaptive beamforming with distance discrimination,” *IEEE Transactions on Speech and Audio Processing*, vol. 12, no. 5, pp. 478–488, Sep. 2004. DOI: [10.1109/tsa.2004.832982](https://doi.org/10.1109/tsa.2004.832982).
- [127] L. Lao, H. S. John, R. Stambaugh, A. Kellman, and W. Pfeiffer, “Reconstruction of current profile parameters and plasma shapes in tokamaks,” *Nuclear Fusion*, vol. 25, no. 11, pp. 1611–1622, Nov. 1985. DOI: [10.1088/0029-5515/25/11/007](https://doi.org/10.1088/0029-5515/25/11/007). [Online]. Available: <http://stacks.iop.org/0029-5515/25/i=11/a=007>.
- [128] G. Q. Li, Q. L. Ren, J. P. Qian, L. L. Lao, S. Y. Ding, Y. J. Chen, Z. X. Liu, B. Lu, and Q. Zang, “Kinetic equilibrium reconstruction on east tokamak,” *Plasma Physics and Controlled Fusion*, vol. 55, no. 12, p. 125008, Nov. 2013. DOI: [10.1088/0741-3335/55/12/125008](https://doi.org/10.1088/0741-3335/55/12/125008). [Online]. Available: <http://stacks.iop.org/0741-3335/55/i=12/a=125008>.
- [129] C. Zhang, J. Wu, H. Liu, and W. Y. Sun, “Near field imaging of synthetic aperture radiometer,” *Progress In Electromagnetics Research Symposium Online*, vol. 4, no. 8, pp. 886–890, 2008. DOI: [10.2529/PIERS080901081924](https://doi.org/10.2529/PIERS080901081924). [Online]. Available: <http://piers.org/piersonline/piers.php?volume=4&number=8&page=886>.
- [130] C. Zhang, J. Wu, H. Liu, and J. Yan, “Imaging algorithm for synthetic aperture interferometric radiometer in near field,” English, *Science China Technological Sciences*, vol. 54, no. 8, pp. 2224–2231, 2011, ISSN: 1674-7321. DOI: [10.1007/s11431-011-4403-3](https://doi.org/10.1007/s11431-011-4403-3).
- [131] R. C. Aster, B. Borchers, and C. H. Thurber, *Parameter estimation and inverse problems*, 1st ed. Academic Press, 2005, ISBN: 0120656043.
- [132] G. R. Tynan, A. Fujisawa, and G. McKee, “A review of experimental drift turbulence studies,” *Plasma Physics and Controlled Fusion*, vol. 51, no. 11, p. 113001, Oct. 2009. DOI: [10.1088/0741-3335/51/11/113001](https://doi.org/10.1088/0741-3335/51/11/113001). [Online]. Available: <http://stacks.iop.org/0741-3335/51/i=11/a=113001>.
- [133] E. Z. Gusakov and A. V. Surkov, “Spatial and wavenumber resolution of Doppler reflectometry,” *Plasma Phys. Control. Fusion*, vol. 46, no. 7, pp. 1143–1162, Jun. 2004. DOI: [10.1088/0741-3335/46/7/012](https://doi.org/10.1088/0741-3335/46/7/012).

- [134] E. Poli, A. Peeters, and G. Pereverzev, “Torbeam, a beam tracing code for electron-cyclotron waves in tokamak plasmas,” *Computer Physics Communications*, vol. 136, no. 1-2, pp. 90–104, May 2001, ISSN: 0010-4655. DOI: [10.1016/S0010-4655\(01\)00146-1](https://doi.org/10.1016/S0010-4655(01)00146-1). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0010465501001461>.
- [135] P. Gibson, “The vivaldi aerial,” in *9th European Microwave Conference, 1979*, Institute of Electrical & Electronics Engineers (IEEE), Oct. 1979. DOI: [10.1109/euma.1979.332681](https://doi.org/10.1109/euma.1979.332681).
- [136] (Aug. 26, 2016). Fmc108, 4DSP LLC, [Online]. Available: <http://www.4dsp.com/FMC108.php>.
- [137] B. Huang, “Development of FPGA controlled diagnostics on mast,” PhD thesis, Durham University, Mar. 2013.
- [138] *Microblaze processor reference guide - ug081 (v11.4)*, *Embedded development kit edk 12.4*, version 11.4, Xilinx Inc., Nov. 15, 2010. [Online]. Available: http://www.xilinx.com/support/documentation/sw_manuals/xilinx12_4/mb_ref_guide.pdf (visited on 09/16/2016).
- [139] J. G. Kwak, Y. K. Oh, K. P. Kim, S. W. Kim, S. H. Hong, Y. Chu, H. J. Lee, Y. O. Kim, J. Kim, S. L. Park, S. H. Hahn, M. K. Park, H. K. Kim, J. G. Bak, Y. S. Bae, W. H. Ko, S. G. Lee, J. H. Lee, J. I. Jung, K. D. Lee, Y. U. Nam, Y. M. Jeon, S. W. Yoon, A. England, W. C. Kim, K. L. Park, H. K. Na, H. L. Yang, M. Kwon, S. Takahiro, J. A. Leuer, N. W. Eidietis, A. W. Hyatt, M. Walker, D. Mueller, L. R. Grisham, J. K. Park, Y. S. Park, and S. A. Sabbagh, “Key features in the operation of KSTAR,” *IEEE Transactions on Plasma Science*, vol. 40, no. 3, pp. 697–704, Mar. 2012, ISSN: 0093-3813. DOI: [10.1109/tps.2011.2181958](https://doi.org/10.1109/tps.2011.2181958).
- [140] (Feb. 17, 2017). Virtex-6 fpga dsp48e1 user guide, Ug369 (v1.3), Xilinx Inc., [Online]. Available: https://www.xilinx.com/support/documentation/user_guides/ug369.pdf (visited on 07/23/2016).
- [141] *Axi reference guide (ug761)*, version 13.1, Xilinx Inc., Mar. 7, 2011. [Online]. Available: http://www.xilinx.com/support/documentation/ip_documentation/ug761_axi_reference_guide.pdf (visited on 09/16/2016).
- [142] C. Taylor, K. Luitjohan, B. Heim, L. Kollar, J. Allain, C. Skinner, H. Kugel, R. Kaita, A. Roquemore, and R. Maingi, “Surface chemistry analysis of lithium conditioned NSTX graphite tiles correlated to plasma performance,” *Fusion Engineering and Design*, vol. 88, no. 12, pp. 3157–3164, Dec. 2013. DOI: [10.1016/j.fusengdes.2013.09.007](https://doi.org/10.1016/j.fusengdes.2013.09.007).
- [143] S. J. Diem, G. Taylor, J. B. Caughman, P. C. Efthimion, H. Kugel, B. P. LeBlanc, C. K. Phillips, J. Preinhaelter, S. A. Sabbagh, and J. Urban, “Collisional damping of electron Bernstein waves and its mitigation by evaporated lithium conditioning in spherical-tokamak plasmas,” *Phys. Rev. Lett.*, vol. 103, no. 1, p. 015002, 1 Jun. 2009. DOI: [10.1103/physrevlett.103.015002](https://doi.org/10.1103/physrevlett.103.015002). [Online]. Available: <http://link.aps.org/doi/10.1103/PhysRevLett.103.015002>.

- [144] (Aug. 30, 2016). Experimental physics and industrial control system, Argonne National Laboratory, [Online]. Available: <http://www.aps.anl.gov/epics/>.
- [145] K. S. Saini and R. F. Bradley, "The sinuous antenna - a dual polarized element for wideband phased array feed application," National Radio Astronomy Observatory, Tech. Rep. 301, Feb. 13, 1996. [Online]. Available: <http://www.gb.nrao.edu/electronics/edir/edir301.pdf>.
- [146] J. Lovell, G. Naylor, A. Field, and R. M. Sharples, "An fpga-based bolometer for the mast-upgrade super-x divertor," in *Proceedings of the 1st EPS conference on Plasma Diagnostics (ECPD2015). 14-17 April 2015. Frascati, Italy*, Proceedings of Science, 2015, p. 137. [Online]. Available: http://pos.sissa.it/archive/conferences/240/137/ECPD2015_137.pdf.