

**Tentative First Observations of
Quasi-Stationary Zonal Flows in a
Spherical Tokamak**
&
Improvements in Image-Velocimetry Workflows

Yorick Walter Enters

Doctor of Philosophy

University of York,
Department of Physics, Engineering and Technology

February, 2024

Abstract

Turbulence drives anomalously high cross-field transport rates in tokamak plasmas. This drastically decreases the energy efficiency and presents a major issue in fusion research because the underlying plasma dynamics are not yet fully understood.

Phenomena known as ‘zonal flows’ directly suppress cross-field transport while siphoning energy from the turbulence, making them highly beneficial for enhancing plasma confinement. Despite the theoretically predicted ubiquity of zonal flows in tokamak plasmas, experimental measurements, particularly in spherical tokamaks, are remarkably scarce. This scarcity is largely attributed to a notorious difficulty in obtaining accurate velocity field measurements. To address this gap, this thesis presents improvements to velocity field inference techniques and preliminary measurements of zonal flows. To this end, the Beam Emission Spectroscopy (BES) diagnostic on the upgraded Mega-Ampere Spherical Tokamak (MAST-Upgrade) was utilised.

Image-velocimetry techniques, which are used to infer velocity fields from imaging diagnostics such as BES, were extensively performance tested for the first time. The two primary techniques, Cross-Correlation Time-Delay Estimation (CCTDE) and Dynamic Time-Warping (DTW), were investigated across the majority of reasonably achievable experimental conditions. Strongly nonlinear behaviour was identified in both techniques, indicating that testing was *required* for reliable velocimetry. Tests uncovered the effects and mitigation techniques of the barber pole illusion, the impact of sheared flows, optimal operational velocimetry parameters, the effect of varying the number of spatial channels, and more. Results were numerous, nuanced, and often unexpected; highlighting that thorough reading is recommended to anyone aiming to perform image-velocimetry analysis.

Velocimetry analysis of BES data from the second MAST-U campaign was performed. The aforementioned testing results were used to guide and optimise the velocimetry workflow while preemptively avoiding complications. However, precision of the inferred velocities fell short of expectations, a discrepancy later attributed to an incorrectly calibrated optical filter in the BES diagnostic. Consequently, only background carbon II emissions reached the BES detector in most shots, rendering velocimetry of the main species impossible.

Instead, attempts were made to perform velocimetry using the carbon background emissions, which may appear impossible at first glance due to localisation issues. After calibrations using a synthetic diagnostic, results demonstrated agreement with the Charge-Exchange Recombination Spectroscopy (CXRS) diagnostic, implying accurate velocimetry. Notably, this marked the first instance of successful CII velocimetry using BES, opening a promising new avenue for investigating impurity dynamics.

A single useful shot, #46459, emerged unaffected by the preceding optical filter issues. The data was predicted to be suitable for accurate velocimetry based on the velocimetry testing results. Upon comparison with CXRS, temporal velocity trends showed agreement with CCTDE velocimetry. CCTDE velocities were inferred with typical standard errors below 1 km/s at a temporal frequency of 4 kHz. The results unveiled radially localised velocity structures, intermittently drifting and re-forming on timescales of 1-10 ms. Further analysis identified this as a coherent localised shear mode featuring a characteristic frequency of 60 Hz and a radial wavelength of approximately 10 cm. The observation of this mode closely aligns with the expected attributes of quasi-stationary zonal flows, pointing to the first zonal flow detection in a high-performance spherical tokamak.

Turbulence drains force.
Zonal flow shields, stays on course,
confines the life source.

Contents

1	Introduction & Motivation	1
2	An Introduction to Fusion Energy and Plasma Confinement	3
2.1	Fundamentals of Fusion	3
2.2	Introducing: the Tokamak	4
2.3	Cross-Field Transport	7
2.3.1	Anomalous Transport	7
3	Plasma Turbulence and Zonal Flow Theory	10
3.1	Drift Wave Theory	11
3.1.1	The Drift Wave Eigenmode	11
3.1.2	Nonlinear Drift Wave Interactions	13
3.2	Zonal Flow Theory	14
3.2.1	The Zonal Flow Eigenmode	14
3.2.2	Zonal Flow Generation	18
3.2.3	Zonal Flow Damping	20
3.3	The Drift-Wave–Zonal-Flow System	22
3.3.1	DW-ZF Self-Consistent State	22
3.3.2	Turbulent Transport Suppression	27
3.4	Global Transport Organisation: $E \times B$ Staircases	28
3.4.1	Description of the Numerical Simulations	28
3.4.2	$E \times B$ Staircase Characteristics	29
4	A Brief Review of Zonal Flow Experimental Measurements	31
4.1	Prior Observations of Zonal Flows	32
4.1.1	Observation of Radially Localised, Sheared Flows with $n=m=0$ near Zero Frequency	32
4.1.2	Nonlinear Turbulence Drive of Zonal Flows	34
4.1.3	Sheared Flows and Confinement Transitions	36
4.1.4	Observations of the $E \times B$ Staircase	37
5	An Overview of the Tokamak, Diagnostics, and Analysis Techniques	41
5.1	The Upgraded Mega-Ampere Spherical Tokamak	41
5.2	Beam Emission Spectroscopy	44
5.3	Cross-Correlation Time-Delay Estimation	46
5.4	Dynamic Time-Warping	47

6	Testing Velocimetry Methods Using Synthetic Data	49
6.1	Generating Synthetic Turbulence Fluctuation Data	50
6.1.1	Isolated Density Features	50
6.1.2	Turbulent Density Fields	51
6.2	Testing Cross-Correlation Time-Delay Estimation	54
6.2.1	Quantification of Technique Uncertainties	54
6.2.2	Results with Isolated Density Features	55
6.2.3	Results with Turbulent Density Fields	57
6.2.4	Discussion and Summary of CCTDE Uncertainties	59
6.3	Testing Dynamic Time-Warping	61
6.3.1	Results - Isolated Density Features	61
6.3.2	Results - Turbulent Density Fields	63
6.3.3	Discussion and Summary of DTW Uncertainties	65
6.4	Comparison Between CCTDE and DTW Performance	67
6.5	Workflows for Predicting Velocimetry Uncertainty	68
6.6	Discussion and Conclusion	71
7	Delocalised Velocimetry Experiments on MAST-U	74
7.1	Preliminary Analysis of BES Density Fluctuations	74
7.1.1	Coherent Density Fluctuation Structures	74
7.1.2	Preparatory Analysis for Velocimetry	77
7.2	Initial Results and Challenges with DTW	80
7.3	Initial Results and Challenges with CCTDE	84
7.3.1	Preparations for CCTDE velocimetry	84
7.3.2	CCTDE velocimetry results	86
7.4	CCTDE Velocimetry of the CII Impurity Shell	87
7.4.1	Dominant CII emissions in BES signal	88
7.4.2	Calibrating BES Radii for CII Emissions	89
7.4.3	CII velocimetry cross-diagnostic comparison	92
7.5	Discussion and Conclusion	93
8	Velocimetry Experiments with Localised BES Signals	96
8.1	Slowly Varying (>ms) Velocity Comparison with Charge-Exchange Re- combination Spectroscopy	96
8.1.1	Initial Look at the Data	96
8.1.2	CCTDE Velocimetry Comparison with CXRS	98
8.1.3	Discussion and Conclusion	99
8.2	Tentative Zonal Flow Observation with Line-CCTDE Velocimetry	102
9	Summary and Discussion	106
9.1	The relevance and impact of velocimetry testing	106
9.2	The delocalisation issue	107
9.3	A preliminary detection of zonal flows	108
9.4	Future work	109

10 Appendices **118**

- 10.1 Equations and expressions 118
 - 10.1.1 $E \times B$ flow 118
 - 10.1.2 Flow measurement equations derivation 118
- 10.2 Source Code: Cross-Correlation Time-Delay Estimation 119

Acknowledgements

The work presented would not have been possible without the many people who have supported me throughout my PhD project. I'd like to thank my supervisor, István Cziegler, for his ever-watchful eye, kind words when needed, and his patience in letting me gather my thoughts after a barrage of incoherent half-sentences. Thanks to everyone in the offices at the Culham Centre for Fusion Energy for the many lunchtime discussions, and in particular Steven Thomas for indulging my frequent interruptions and questions.

The fusion CDTea peeps for the tea breaks, lunchtime yoga sessions, and nights in 'pop world' that make me look back on our times in York with fondness. My friends in Oxford for the many pub nights, galas, dinners and walks to clear the head. Arun Nutter for being a wonderful flatmate - always with a cheeky insult ready for me. My friends in Turkey for being some of the kindest and welcoming people I know, and for introducing me to such a wonderful culture. The OUMC for the beautiful rock climbing trips and lovely community that solidified my love for the sport.

Thanks to my parents and family for raising me with much love and to be confident in my own abilities. Mama, zonder jou had ik nooit natuurkunde gestudeerd. Je kan wel zeggen dat dit een hele goede keuze is geweest voor mij. Ik ben je eeuwig dankbaar voor alle liefde, steun en begeleiding door de jaren heen. Papa, jij bent altijd mijn rolmodel geweest. Dat je successvol kan zijn als een volwassene zonder het innerlijke kind uit het oog te verliezen. Ik kijk vol genot toe wanneer je s'winters speels gaat checken of het ijs al dik genoeg is of een sneeuwpop gaat maken. Esmee, mijn lieve zus en een goede vriendin. Ik weet dat ik altijd op je kan rekenen. Samen boulderen en af een toe wat eten is een lichtpunt geweest tijdens het schrijven van deze thesis.

To the HepburnHomies and pals, 'growing up' with you all has been a blast and I will always treasure our friendships. I hope that we continue our holidays and see each other as often as possible. The EggBeanBees for the many chaotic Sundays - a cherished tradition at this point. I hope we stick with it for many years to come and continue bee-in' a bunch o' beans.

Finally to Selin, I cannot imagine having done this PhD without you. Meeting you during this time has changed my life and I can't wait to explore the future together. Hadi geleceği birlikte keşfedelim hayatım benim!

Declaration

This thesis has not previously been accepted for any degree and is not being concurrently submitted in candidature for any degree other than Doctor of Philosophy of the University of York. This thesis is the result of my own investigations, except where otherwise stated. All other sources are acknowledged as references.

This work has been (part-) funded by the EPSRC Energy Programme [grant number EP/R034737/1].

Chapters 2 - 5 cover summaries of previous work and give context for the rest of this thesis. Chapters 6 - 8 are novel work performed as part of this PhD project.

Chapter 6 has been adapted from the following publication:

- Y. W. Enters, S. Thomas, M. Hill, I. Cziegler; *Testing image-velocimetry methods for turbulence diagnostics*. Rev. Sci. Instrum. 1 July 2023; 94 (7): 075101. <https://doi.org/10.1063/5.0133453>

Some of the results in chapter 7 have previously been presented in the following conference poster:

- Y. W. Enters, S. Thomas, I. Cziegler, and the MAST-Upgrade team; *Characterisation and velocimetry of the CII shell using beam emission spectroscopy diagnostic on MAST-U*. 49th EPS DPP Plasma Physics conference. 3-7 July 2023. Bordeaux, France.

Some of the results in chapter 8 have previously been presented in the following conference poster:

- Y. W. Enters, S. Thomas, I. Cziegler, and the MAST-Upgrade team; *Hints of quasi-zero-frequency zonal flows in MAST-U and improvements in image-velocimetry*. 27th Joint EU-US Transport Task Force Meeting. 11-15 September 2023. Nancy, France.

Chapter 1

Introduction & Motivation

Fusion energy. It has been powering our Sun and the stars for billions of years. Here on Earth it promises a clean and inherently safe energy source - ideal for a decarbonising world. The fuel is abundant, spread out throughout the world, and could theoretically support global energy demands for millions of years. Besides these benefits of fusion energy, there is one major drawback: it is really difficult to reach *net* energy production.

Achieving nuclear fusion reactions in itself has been possible since the 1930s [1], but these first experiments produced next to no energy compared to the amount of energy used. Since then, the invention of the 'tokamak' device has led to drastically improved energy efficiency, from practically zero to $Q = 0.67^1$, a record set in 1997 by the Joint European Torus [2]. Next-generation tokamaks are currently under construction, such as the internationally funded ITER project², expected to achieve $Q = 5-10$ by 2040. Additionally, private fusion companies such as Tokamak Energy and Commonwealth Fusion Systems have joined the race, both expecting net energy production within this decade.

The main way to improve the energy efficiency is by minimising the rate at which energy 'leaks' out of the machine. This 'leakage' rate, determined by the cross-field transport rate, is orders of magnitude faster than theoretical predictions based on collisional transport and the exact underlying mechanisms are still unknown. Improving our understanding of cross-field transport is currently slowed down significantly by a lack of high quality experimental measurements of the plasma dynamics. Measurements which are incidentally 'notoriously difficult in tokamaks' [3]. This leads to the two central aims of this thesis:

- To improve the reliability and performance of spatially resolved plasma dynamics measurements.
- To measure zonal flows in the upgraded Mega-Ampere Spherical Tokamak (MAST-U).

Zonal flows are a type of flow in fusion plasmas known to strongly suppress cross-field transport rates, which means that they are highly beneficial for the overall energy efficiency. On top of that, they are not some rare, highly situational phenomenon and are

¹'Q' measures the energy efficiency of the reaction. $Q = 0.5$ means that, compared to the amount of energy used, 50% as much is produced by the fusion reactions. $Q = 5$ means that 5 times more energy is produced than used.

²Not an acronym. Simply 'the way' in Latin

instead thought to be ubiquitous in tokamak plasmas. Despite their importance, experimental measurements are scarce, limiting our understanding and eventual exploitation of this beneficial phenomenon.

Chapters 2 - 5 cover summaries of previous work and give context for the rest of this thesis. Chapters 6 - 8 are novel work performed as part of this PhD project. In chapter 6, extensive testing is performed of image-velocimetry techniques. These techniques are key to measuring plasma flows but had previously not thoroughly been tested, putting previous velocimetry results into question. With an improved understanding of the velocimetry techniques, chapter 7 covers experimental analyses of MAST-U data. This chapter focuses on initial velocimetry results, cross-checking with other diagnostics, and optimising velocimetry performance. Additionally, the discovery of a major issue with the experimental data is addressed with some unexpected and novel outcomes. The final experimental chapter, chapter 8, analyses a shot which did not suffer from the issues addressed in the previous chapter. High precision and high temporal frequency velocimetry are achieved, leading to the identification of a mode that closely aligns with zonal flows.

Chapter 2

An Introduction to Fusion Energy and Plasma Confinement

This chapter serves as an introduction to fundamental concepts in plasma confinement. Its purpose is to provide context for the more specialised topics explored throughout this thesis and how the research can impact the energy efficiency of fusion plasmas. A basic understanding of plasma physics and fusion energy will be assumed and this chapter will not provide an exhaustive literature review or delve into detailed theoretical derivations.

We will outline the requirements for achieving a self-sustaining fusion reaction and explore the associated challenges. The importance of improving the energy-confinement-time will be highlighted, which will lead into a discussion on how the plasma particles and energy are confined in the magnetic geometry of two categories of fusion machine. At the forefront of these machines stands the ‘tokamak,’ currently the most advanced and elegant solution for plasma confinement, surpassing all other machine types in all-round performance. Additionally, there is the ‘stellarator’ approach, which may be considered a more promising option for plasma stability and long term performance. However, this type of machine is less mature and currently does not perform as well as the tokamak.

Moreover, we will delve into the concept of cross-field transport, a key process that constrains the confinement efficiency and correspondingly governs the overall energy efficiency in tokamaks. We will explore various mechanisms that drive cross-field transport, spanning from basic collisional theory to the elusive ‘anomalous transport’. This topic will be covered most thoroughly in this chapter as it forms the cornerstone of this thesis, emphasising the critical need to manage transport levels in tokamaks effectively.

2.1 Fundamentals of Fusion

In the 1950s, during a period of growing interest fusion energy production, J. Lawson conducted an analysis of the power balance in idealised fusion reactors [4]. This widely known work laid the foundation for defining when a deuterium-tritium (D-T) plasma enters a ‘burning’ state, where the fusion reaction becomes self-sustaining, eliminating the need for external heating. This criterion, now known as Lawson’s criterion, can be

expressed as:

$$nT\tau_E \geq 3 \times 10^{21} \text{ keV s m}^{-3} \quad (2.1)$$

In this equation, n represents number density, T signifies temperature, and τ_E denotes the energy confinement time. Equation 2.1 shows that the triple product of density, temperature, and energy confinement time must exceed the specified threshold for the plasma to reach a burning state. Increasing each of these three parameters leads to significant engineering and physics challenges. For example, if the density in the plasma is increased, the corresponding increase in density gradient drives plasma instabilities which can drastically reduce the energy confinement timescale or lead to complete disruption of the plasma. This is why it can often be counterproductive to increase the density beyond a certain threshold. The temperature, in addition to instigating gradient-driven instabilities, is also limited by the fact that the D-T reaction cross-section has an optimal temperature, shown in figure 2.1. Beyond this optimal value, increasing the temperature decreases the D-T reaction probability, rendering further increase in temperature less effective.

This brings us to the energy confinement timescale, which can hypothetically be increased without known drawbacks. This makes increasing the energy confinement time an attractive avenue for improving the triple product. However, a significant challenge lies in the current inability to directly tune the energy confinement timescale due to the complex interplay of numerous non-linearly interacting transport processes. While some of these processes are known and constitute an active field of research, it is likely that more are yet to be uncovered. To provide a basic understanding of transport processes, particularly in tokamak plasmas, we will delve into an overview in sections 2.2 and 2.3.

2.2 Introducing: the Tokamak

Before the widespread adoption of the tokamak approach in fusion research, the 1950s and 1960s saw a predominant focus on linear devices [5]. As depicted in figure 2.2a, these linear devices featured ‘open’ magnetic field lines, providing a rapid path for charged particles to escape from the central plasma region. Researchers explored various techniques, such as magnetic mirrors and field-reversed configurations, to reduce these transport rates and enhance energy confinement times [5].

While western countries concentrated their efforts on linear machines, an arguably more elegant approach was being developed in the former Soviet Union. Their innovation involved curving the magnetic field lines into a toroidal shape, as illustrated in figure 2.2b. This configuration eliminated ‘open’ field lines, and by the late 1960s, these ‘tokamak’ devices claimed to significantly outperform linear counterparts. The validation of these claims by a sceptical team of British scientists in 1969 [6] triggered a ‘veritable stampede’ of tokamak construction in the early 1970s [7].

We will now delve into a brief overview of the fundamental plasma dynamics and transport within tokamaks. Along the field lines, rapid transport of charged particles occurs on the order of thermal velocities. The helical rotation of the field lines generates a series of nested surfaces, as depicted in figure 2.2b. In terms of transport, this means that all points on these so-called flux-surfaces are connected through fast

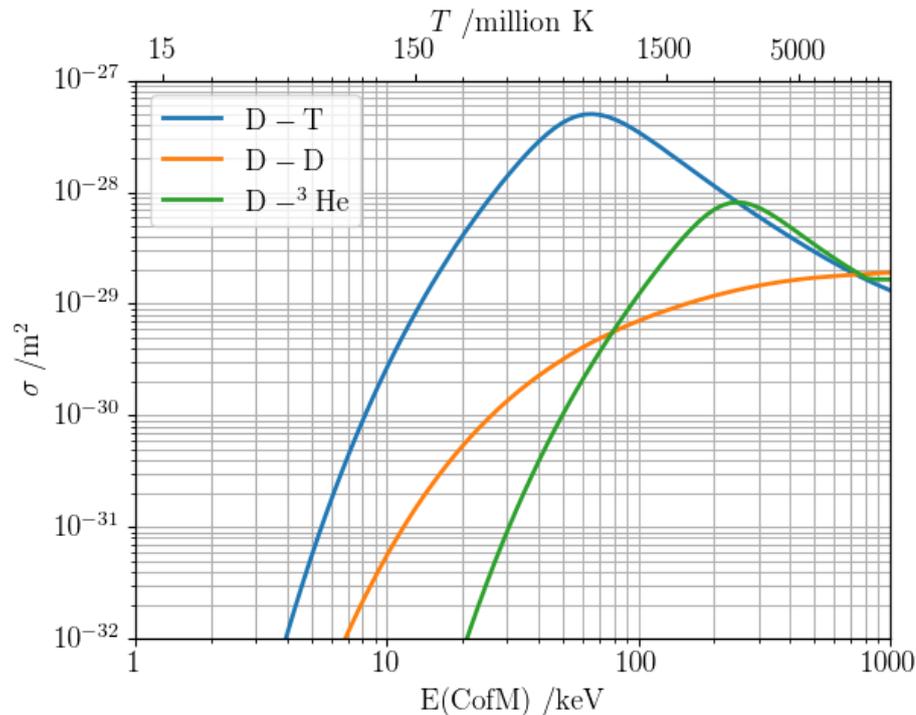


Figure 2.1: Nuclear cross-sections, σ , of several fusion candidate reactions. Data from IAEA's evaluated nuclear data services database. D-T peak cross-section observed at 60-70 keV.

parallel transport, rapidly equalising thermodynamic variables across these surfaces. Due to the toroidal symmetry in these devices, the three-dimensional magnetic geometry is often simplified to the two-dimensional radial-poloidal plane¹, as shown in figure 2.3a. In this two-dimensional view, flux-surfaces are represented by nested loops. The direction perpendicular to these flux-surfaces is referred to as the cross-field direction or the minor-radial direction².

Cross-field transport occurs much more slowly than parallel transport, enabling the development of significant thermodynamic gradients, as exemplified in the thermodynamic profiles displayed in figure 2.3b. High gradients can only be sustained when transport rates are low, and vice versa. Furthermore, different confinement 'modes' can be achieved in a plasma by adjusting tokamak operational parameters, with high-confinement (H-mode) and low-confinement modes (L-mode) illustrated in figure 2.3b. The control of confinement modes represents a central area of investigation in tokamak science, with a key focus on understanding the underlying transport mechanisms.

¹For visualisation purposes, I like to imagine this plane as the cross-section of a bagel that has been sliced in the *objectively* incorrect direction. Is this the best analogy one could make? Perhaps not, but it is certainly memorable...

²The cross-field and minor-radial directions are technically not quite the same if rotational symmetry around the magnetic axis is broken, but they are often conflated regardless for simplicity

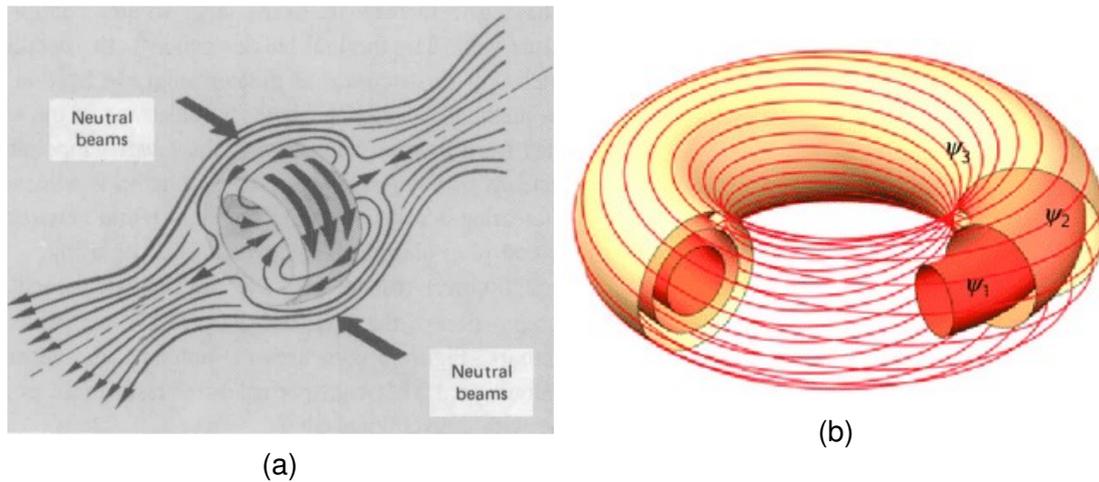


Figure 2.2: Comparison of magnetic field structure between linear devices (a) and tokamaks (b). Flux surfaces, ψ , shown in (b). Images (a) and (b) taken from publications [5] and [8] respectively.

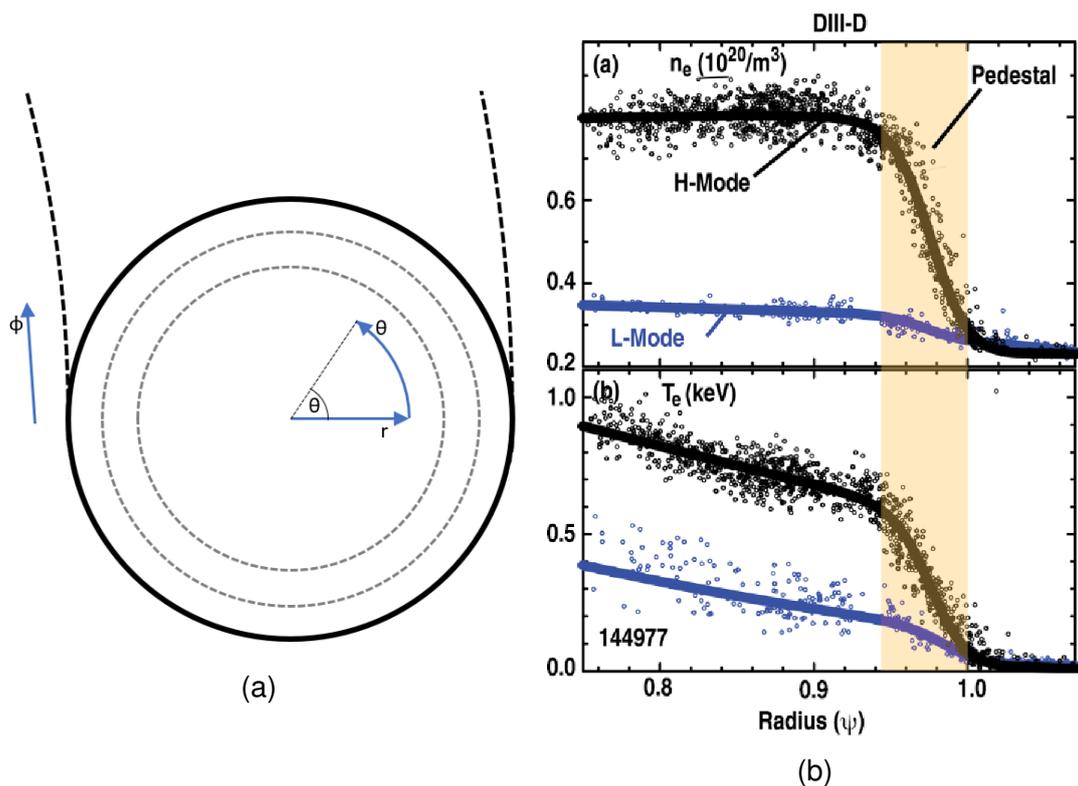


Figure 2.3: Radial-poloidal plane shown in a), with the outer solid black circle representing the separatrix. Grey dotted loops depict flux surfaces. Minor-radial and poloidal directions labelled θ and r respectively. Toroidal direction, ϕ , points into the page. Part b) shows example thermodynamic profiles from DIII-D. Steeper gradients in H-mode are due to suppressed cross-field transport. Image adapted from [9].

2.3 Cross-Field Transport

An essential advantage of the magnetic geometry within a tokamak lies in its ability to govern the energy confinement time through the relatively slow cross-field transport rates. However, it is crucial to emphasise that cross-field transport is far from negligible, and its suppression is key to enhancing plasma energy efficiency. This subsection will cover the primary cross-field transport processes, ranging from basic collisional theory to the realm of the initially unexpected ‘anomalous transport’ mechanisms.

A common starting point for comprehending cross-field transport is to delve into collisional theory. In a linear field geometry, cross-field single-particle motion and its associated diffusion lengthscale are confined to the gyro-radius, which is around 1 cm for protons in typical tokamak field strengths. If this so-called ‘classical’ transport were the sole form of transport in tokamaks, fusion reactors could theoretically be as compact as a couple of centimetres while maintaining effective energy confinement³. However, in a toroidal geometry, single-particle motion is no longer bound only to the gyro-radius and can also follow ‘banana orbits,’ governing cross-field diffusion at lengthscales an order of magnitude larger than the gyro-radius. This gives rise to ‘neoclassical’ transport, characterised by theoretically well-understood and predictable transport rates. Neoclassical transport rates, on their own, were thought to be sufficiently low for tokamaks like the Joint European Torus (JET) to achieve energy break-even. However, this prediction has evidently not materialised, as demonstrated by the absence of operational fusion plants in the world today. In reality, experimentally measured transport rates in tokamaks are orders of magnitude higher than neoclassical rates [11], underscoring the presence of other transport mechanisms, which we will explore further in the rest of this section.

MHD instabilities, if left unmitigated, can emerge as an example of strong drivers of cross-field transport levels. These instabilities have the potential to induce transport at rates many orders of magnitude above the neoclassical baseline and can even trigger full-scale plasma disruption. Fortunately, the underlying modes are typically well-described by MHD theory and evolve relatively slowly with the magnetic equilibrium. This characteristic enables analytical predictions to be made of when instabilities, such as the sawtooth or neoclassical tearing modes, might become unstable. Consequently, these instabilities can often be controlled by implementing various mitigation techniques [12]. The successful application of these mitigation techniques is critical in modern tokamaks and will be even more so in future fusion power plants. However, MHD mitigation falls outside the scope of this thesis. Instead, this thesis focuses on the transport processes that persist even when MHD-driven transport is maximally suppressed.

2.3.1 Anomalous Transport

In tokamak discharges characterised by minimal MHD activity, an unexpected phenomenon emerges: experimentally measured transport rates are consistently found to be up to two orders of magnitude higher than the predictions of neoclassical theory,

³Interestingly, this dimension matches the fictional Iron Man’s chest-piece ‘arc reactor’ [10].

as demonstrated by e.g. Connor *et al.* [11]. This revelation didn't stem from a single groundbreaking observation but rather from a multitude of experiments, each pointing increasingly toward the inadequacy of neoclassical theory in predicting transport rates. These unusually high transport rates were evident in both energy and particle transport across a diverse range of tokamak discharges, as reported by Artsimovich *et al.* [13]. Much like the enigmatic 'dark energy' and 'dark matter' in astrophysics, this so-called 'anomalous' transport was given a similarly vague name due to its unexpected discovery and the mystery surrounding its underlying physics.

However, amid these initial uncertainties, one key insight has consistently emerged from experimental observations: *anomalous transport can almost entirely be attributed to turbulent fluctuations* (Rowan *et al.* [14]; Wootton *et al.* [15]). While this observation provides a starting point for further investigations, it also introduces a host of complexities in pinpointing the exact mechanisms driving this anomalous transport.

First and foremost, turbulence presents a formidable challenge in theoretical modelling. Deriving analytical expressions for turbulent transport solely from first principles is even thought to be impossible. Although the transport levels due to turbulence cannot currently be modelled directly, there is a substantial theoretical body of knowledge surrounding the unstable modes which lead to turbulence [16]. A key mode in relation to this is the drift-wave mode, which will be covered in more detail in section 3.1.

Attempting to model turbulent transport empirically also poses challenges, primarily due to diagnostic limitations. Turbulent fluctuations manifest on spatial and temporal scales that are orders of magnitude smaller and faster than the plasma equilibrium. Turbulence diagnostics often require trade-offs between optimising temporal or spatial resolution to capture specific aspects of the turbulence.

Another approach involves employing nonlinear gyro-kinetic simulations to compute the driving modes and saturated states of turbulence. However, running simulations that encompass all relevant scales requires prohibitively high computational costs. Reduced models have been developed to mitigate computational expenses, yet these models often rely on precarious assumptions, and their predictive power beyond specific scenarios can be questionable. For example, an underlying assumption in a lot of computational work is that ion and electron scales evolve separately and independently, an assumption that recent state-of-the-art nonlinear simulations have challenged [17, 18]. This underscores the challenging task of producing experimentally relevant simulations of tokamak turbulence and transport, and furthermore highlights the importance of the verification of simulations with experimental observations. Nonetheless, computational endeavours have yielded numerous valuable insights which will be referenced throughout, where relevant, but a full review of computational work is outside the scope of this thesis.

Despite the aforementioned challenges, researchers have made noteworthy observations in both simulations and experiments. For instance, zonal flow modes, which are covered in detail in section 3.2, are known to act as internal transport barriers that exert a strong suppressive effect on cross-field transport. Streamers have been observed to enhance cross-field transport levels in localised poloidal regions, while avalanching processes intermittently project energy and matter radially outward. Furthermore, anomalous transport has been linked to impurity recycling further downstream in the scrape-off-layer [19, 20, 21, 22, 23]. These examples likely represent only a fraction

of the relevant mechanisms at play and a full review of the literature falls outside the scope of this thesis. Regardless, investigating individual mechanisms that are known to influence transport, and the interactions between them, will conceivably lead to a consistent understanding of cross-field transport in tokamak plasmas.

The focus of this thesis centres on one such system characterised by multiple interacting mechanisms: the drift-wave–zonal-flow (DW-ZF) system. At present, this system is a leading example of how multiple processes can organise into a globally consistent system, and we will delve into the theory in chapter [3.3](#).

Chapter 3

Plasma Turbulence and Zonal Flow Theory

In this chapter, the basic theories that provide context for the rest of this thesis will be covered. The target audience is individuals who may be new to the field of turbulent transport in tokamak plasmas or those who require a refresher on the importance of the drift-wave–zonal-flow system specifically. It is important to note that this chapter will not encompass a comprehensive review of the theories but will instead serve as an accessible introduction to the field, facilitating a heuristic understanding of the underlying physics.

In section 3.1, drift wave turbulence will be introduced, a phenomenon which is ubiquitous in tokamak plasmas. Although this mode is often subdominant to the ion-temperature-gradient (ITG), electron-temperature-gradient (ETG), or trapped-electron-mode (TEM) modes, the focus will be on the drift-wave because it is an essential prerequisite for zonal flow theory. Additionally, many key characteristics of drift-waves are analogous to ITG, ETG and TEM modes. The structure and evolution of a single drift-wave eigenmode will be described in part 3.1.1. Subsequently, the nonlinear interactions between drift waves will be covered in section 3.1.2, setting the stage for the development of a turbulent spectrum.

In section 3.2, zonal flows will be introduced, and after providing a bit of historical context, the structure and evolution of the zonal flow eigenmode will be examined. Emphasis will be placed on developing a mental image of zonal flows. It will be demonstrated how nonlinear coupling with drift waves is the sole driving mechanism for zonal flows, implying that zonal flows cannot directly tap into free energy reserves. Nevertheless, it will be shown that the zonal flow mode always grows in a tokamak plasma. Additionally, the only linear damping mechanism is through the notoriously slow collisional processes, which is why zonal flows are ubiquitously expected to grow to non-negligible amplitudes.

In section 3.3, the fact that drift-waves and zonal flows should not be seen as independent entities, but rather as two parts of a unified drift-wave–zonal-flow system, will be highlighted. As such, the DW-ZF system will be modelled using two simplified models, each of which are used to highlight specific aspects of the system’s characteristics. Additionally, numerical solutions to more generalised nonlinear diffusion equations will be provided to elucidate a more realistic picture of the system’s temporal evolution. Finally, the existence of a strongly-suppressed transport regime, also known as the

Dimits regime, will be pointed out during a general description of the radial transport due to the DW-ZF system.

In order to demonstrate how the DW-ZF system can be linked to a globally self-consistent transport system, the $E \times B$ staircase is introduced in section 3.4. This section introduces the results of gyrokinetic solutions that observed staircase-like structures consistently emerging in the thermodynamic profiles of simulated Tore Supra discharges. Incidentally, the regions of steep gradient in the staircases were found to be caused by zonal flows. The regions of near-flat gradients located between the zonal flows were caused by turbulent avalanching processes. The manner in which the zonal flows and avalanching regions self-organise, along with the temporal evolution and different regimes of the $E \times B$ staircase, will be described.

3.1 Drift Wave Theory

In this section, we will delve into the physics of drift-waves, which emerge as ubiquitous micro-instabilities within tokamak plasmas. These waves constitute one of the primary modes responsible for generating a turbulent spectrum, consequently driving anomalous transport rates. Our exploration will begin with an introduction to the spatial structure and evolution of individual drift-wave eigenmodes. To maintain clarity and brevity, we will adopt a heuristic derivation approach, which will necessarily sacrifice some mathematical rigour.

Following our exploration of individual drift-wave characteristics, we will briefly describe the nonlinear interactions between drift waves. This will shed light on the fact that drift waves should not be perceived as a collection of loosely interacting individual modes, but rather a unified spectrum of strongly interdependent drift-waves. These nonlinear interactions represent the key catalyst for the development of the drift-waves into a turbulent system.

For those seeking a more comprehensive and rigorous understanding of drift-waves, particularly within the context of cross-field transport rates, we highly recommend starting with the 1999 review paper authored by W. Horton [24].

3.1.1 The Drift Wave Eigenmode

We will begin with an exploration of the spatial structure and dynamics of the drift-wave eigenmode. Imagine a plasma in a magnetic field exclusively oriented in the toroidal direction. A radially-localised seed density perturbation is introduced with a non-zero poloidal wavenumber, which is shown in figure 3.1. From the perspective of the radial-poloidal plane (see figure 2.3a), one can think of this mode as aligning with the flux surfaces along the poloidal direction. Under the assumptions of massless electrons and quasi-neutrality, it can be demonstrated that the density perturbation gives rise to a corresponding electric potential perturbation:

$$\delta\phi = \frac{T_e}{e} \frac{\delta n}{n} \quad (3.1)$$

Here, T_e is the electron temperature. It should be highlighted that these electric field perturbations induce an $E \times B$ drift (see eq. 10.1). This is shown on the left-hand-side

of figure 3.2, where eddies encircle the potential perturbations in the radial-poloidal plane.

Another aspect to consider is the motion of the perturbations themselves. The $\mathbf{E} \times \mathbf{B}$ motions of the plasma particles impact the density perturbations, which sets up a feedback loop between the density and potential perturbations. The net result of the feedback loop is that particles undergo radial oscillations and the perturbations propagate in the poloidal direction. Consequently, the drift wave eigenmode can be thought of as a radially localised plane wave that travels in the poloidal direction:

$$\phi_k(\theta, t) = \phi_k'(t)e^{-ik\theta} \quad (3.2)$$

Here, θ is the poloidal coordinate, ϕ the electric potential and k the poloidal wavenumber. Although the perturbations have been shown to evolve as a plane-wave, the prolonged existence of these modes in the first place is unfounded until this point. Under the assumption of massless electron that led to equation 3.1, it can be shown that the seed perturbation does not grow [25, 26]. If instead we include a non-zero electron mass, a phase shift is introduced between the density and potential perturbations, which ubiquitously leads to the unstable growth of the drift-wave eigenmode [27, 28]. As a result, drift-waves will always grow from seed perturbations and are a pervasive presence in any tokamak plasma.

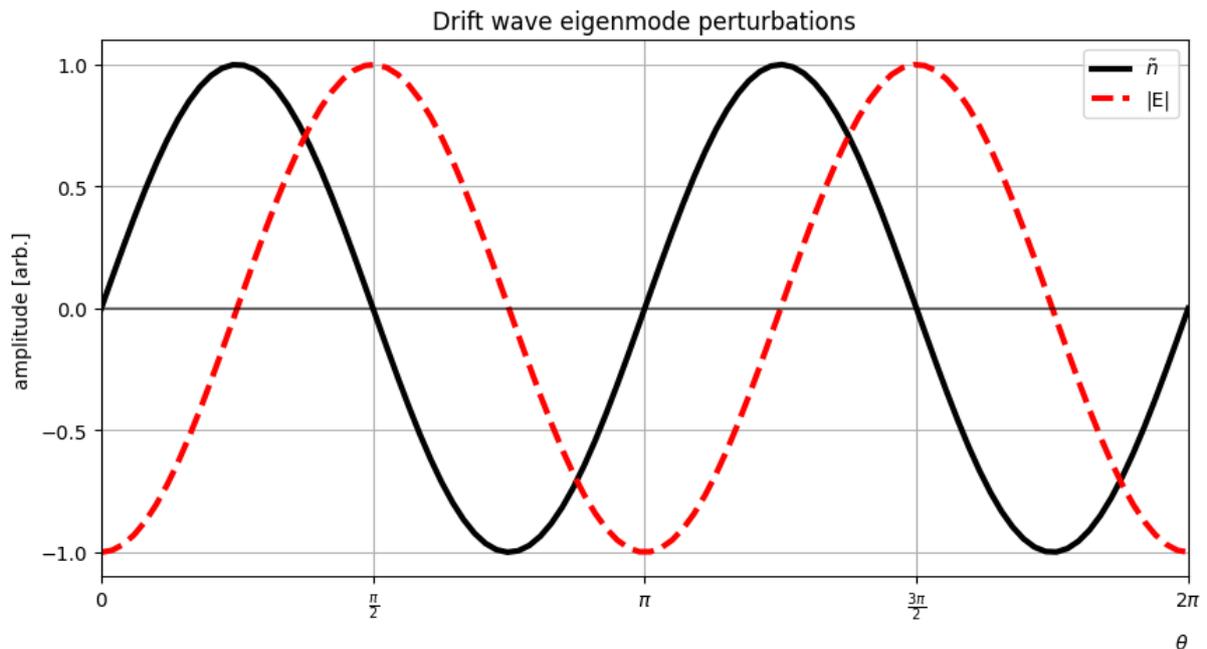


Figure 3.1: Example perturbations associated with the drift-wave eigenmode shown along the poloidal direction, θ . Density and electric field magnitude perturbations shown with a poloidal modenumber of 2.

Up until this point, we have mainly restricted our perspective to the radial-poloidal plane. Next, we will consider the structure of drift waves along the toroidal direction. Density perturbations couple strongly to the ion acoustic mode which travels at the speed of sound along the magnetic field lines. Consequently, the drift wave eigenmodes can extend up to the machine size along the field lines while remaining localised

within a few gyro-radii in the direction perpendicular to the field lines [24, 29]. An example of the elongated drift wave eigenmodes is presented in figure 3.2. Although these long, thin eddies can theoretically extend up to the machine size along the field lines, two main processes are known to limit the eigenmode length. Firstly, flow shearing can disrupt the modes at a rate faster than the Alfvén waves can traverse machine. Secondly, the eigenmodes are inherently unstable once they get to a certain length [24, 29].

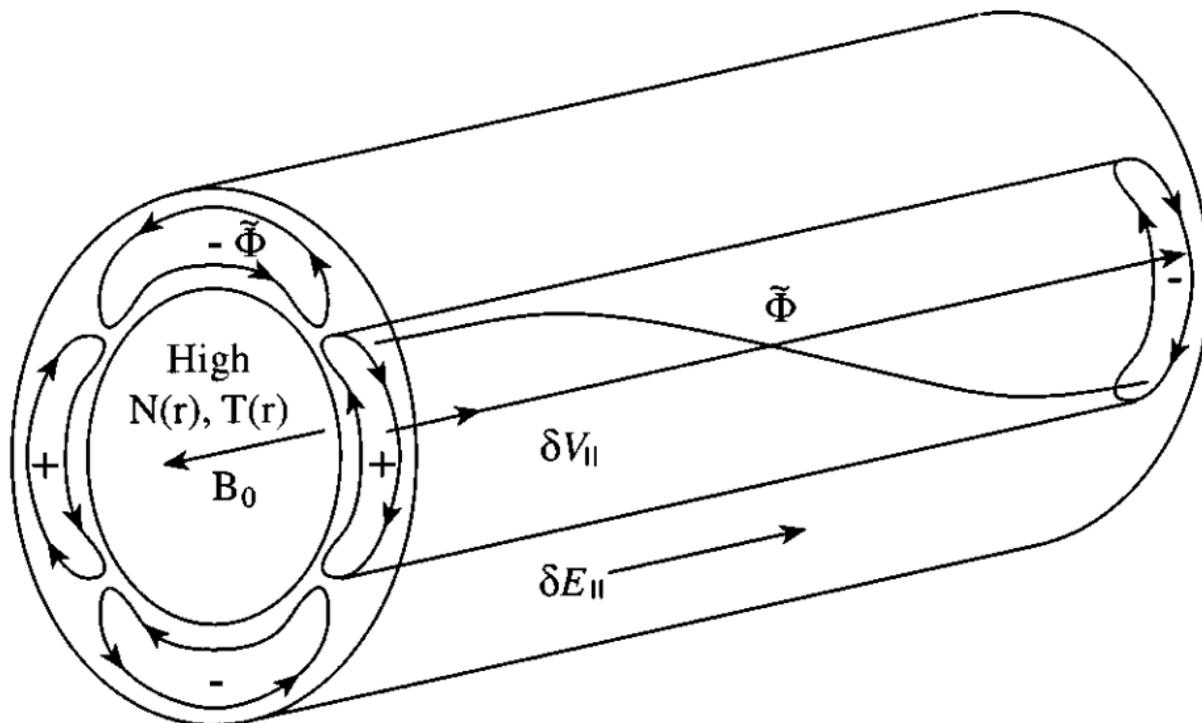


Figure 3.2: An example drift-wave eigenmode (with poloidal modenumber $m=2$) in a short toroidal section of a tokamak plasma. Positive and negative potential perturbations alternate along the poloidal direction. Perturbations are elongated along the magnetic field lines (toroidal direction). Adapted from Horton [24].

3.1.2 Nonlinear Drift Wave Interactions

We will start with a heuristic derivation of the time-evolution for an individual drift-wave eigenmode. It's essential to note that we'll skip some steps for the sake of brevity. Our starting point is the Hasegawa-Wakatani (H-W) equation [27], which is derived from the continuity equation in fluid dynamics and it crucially includes compressibility of the fluid. The H-W equation describes the evolution of the electric potential and density in a fluctuating magnetised plasma:

$$\frac{\partial}{\partial t}(n - \Delta_{\perp} \phi) + (\hat{b} \times \Delta_{\perp} \phi) \cdot \Delta_{\perp} (n - \Delta_{\perp} \phi) + \frac{\partial}{\partial \theta} \phi = 0 \quad (3.3)$$

Here, \hat{b} represents the direction of the magnetic field. Interestingly, the hypothesised plane-wave description of drift-waves, as shown in Equation 3.2, can indeed be shown

to be an eigenmode of the H-W equation by the assumption of an isothermal electron fluid and use of the electron equation of motion [27, 28]. The time evolution of such a plane-wave eigenmode can be expressed as:

$$\frac{d\phi_k}{dt} = -i\omega_k\phi_k + \sum_{k+k_1+k_2=0} \Lambda_{k_1,k_2}^k \phi_{k_1}^* \phi_{k_2}^* \quad (3.4)$$

Here, ω_k represents the frequency of mode k , and Λ_{k_1,k_2}^k denotes the nonlinear coupling coefficient between modes k , k_1 , and k_2 . This equation illustrates that the time evolution of mode k relies on its own linear evolution and on its nonlinear interaction with modes k_1 and k_2 .

Upon closer examination, it becomes evident that the drift-wave evolution is profoundly nonlinear [27], implying that the nonlinear term in Equation 3.4 exerts dominance over the linear term. These nonlinear interactions become increasingly important with increasing drift-wave amplitude and eventually lead to a fully turbulent spectrum [30]. Consequently, it's typically inappropriate to view drift-waves as an amalgamation of weakly-interacting individual drift-waves; instead, they often constitute a unified spectrum of strongly interdependent drift-waves. Numerous studies have corroborated this, demonstrating that the saturated drift-wave spectrum exhibits phenomena such as energy cascading, spectrum broadening, and notably, enhanced cross-field transport rates [27, 28, 26, 25, 29, 30].

3.2 Zonal Flow Theory

Zonal flows initially puzzled researchers when first observed in Jupiter's atmosphere [31]. These longitudinal bands, referred to as 'zonal' bands in atmospheric physics, consist of counter-propagating gas flows, as illustrated in figure 3.3. While these structures, characterised by high shear regions, may seem like an improbable phenomenon to emerge in the chaotic turbulence of the Jovian atmosphere, they have proven to be remarkably stable.

Curiously, what began as an intriguing discovery in the study of gas giants has now become directly relevant to the dynamics of tokamak plasmas right here on Earth.

This section is dedicated to providing a theoretical overview of zonal flows. Subsection 3.2.1 will delve into the spatial and temporal structure of the zonal flow eigenmode, aiming to provide a clear mental image of zonal flows for the reader. Subsection 3.2.2 will focus on elucidating the driving mechanisms behind zonal flows, emphasising that zonal flows are inherently expected to grow in any tokamak plasma. Finally, in subsection 3.2.3, we will explore the damping processes that affect zonal flows, highlighting the fact that the lack of fast collisionless processes will consistently lead to the saturation of zonal flows at significant amplitudes.

3.2.1 The Zonal Flow Eigenmode

The Spatial Structure of Zonal Flows

Zonal flow modes in tokamaks are radially localised, typically within the range of 1-10 cm, and display axisymmetry in both the poloidal and toroidal directions. In other

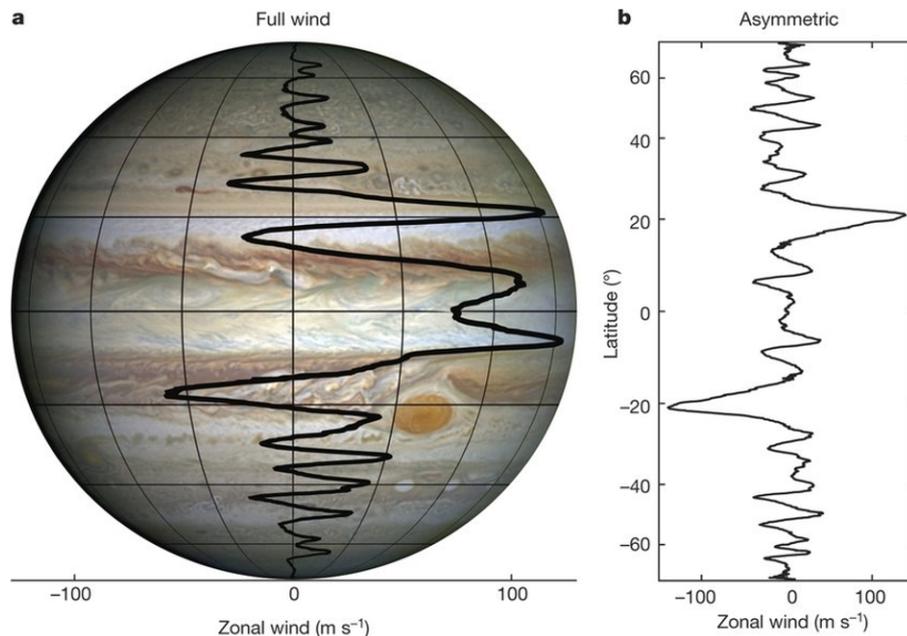


Figure 3.3: Gas flow velocities observed in the atmosphere of Jupiter. Counter-propagating zonal flows give Jupiter its banded appearance. Image adapted from [31].

words, the zonal flow mode takes on the appearance of a hollow toroid, with its primary flow oriented in the poloidal direction. For a visual representation of the spatial structure of zonal flows, please refer to figure 3.4a, where you can observe that zonal flows consist of a pair of poloidally counter-propagating flows, establishing a region with high flow shear between them.

Relation to Convective Cells

Given that turbulent flow fields typically include rotating eddy structures, it can be beneficial to conceptualise zonal flows as the strong asymmetric limit of such a convective cell; imagine taking a convective cell and wrapping it around the magnetic axis, as depicted in figure 3.4b. In this scenario, a pair of counter-propagating flows emerges. Whilst a rigorous proof has been omitted, this description was intended to provide an intuitive framework for visualising how counter-propagating poloidal flows can potentially emerge from a turbulent flow field.

A Heuristic Justification for the Emergence of Counter-Propagating Poloidal Flows

Let's dive deeper into understanding the zonal flow eigenmode by considering a specific scenario. Imagine an axisymmetric ($n=m=0$) electric potential perturbation within a tokamak plasma. For the sake of this discussion, the perturbation's origin is irrelevant. In section 3.2.2 we will justify the growth of such a mode but for now we will assume that it exists and remains constant with time.

The plasma response to this perturbation can be divided into two key components: first, a rapid, compressible element known as the oscillatory geodesic acoustic

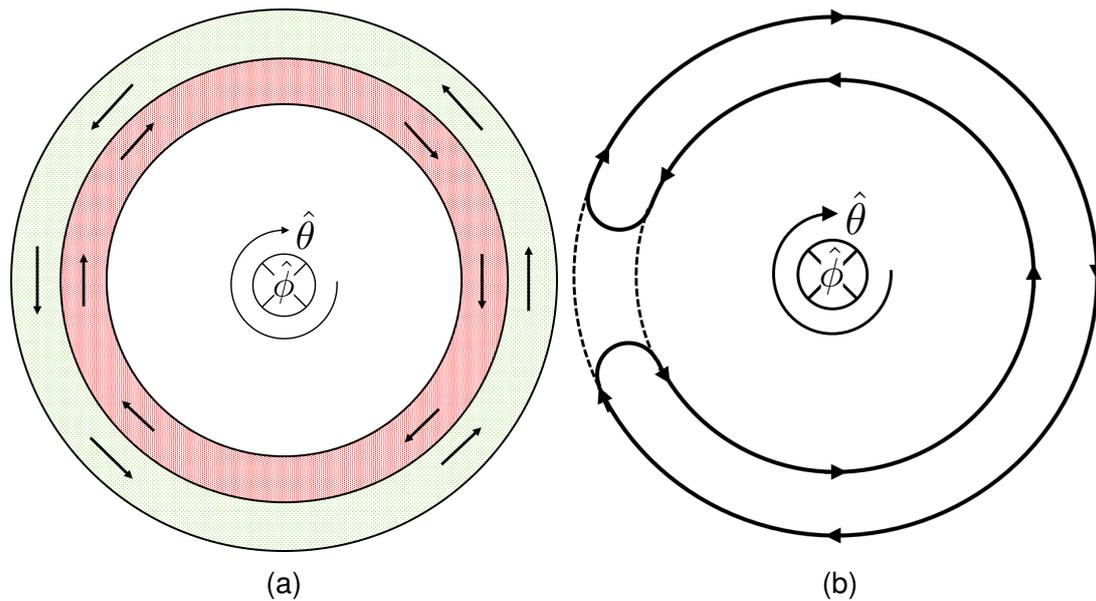


Figure 3.4: The spatial structure of zonal flows in the radial-poloidal cross-section. $\hat{\theta}$ represents the poloidal direction and $\hat{\phi}$ is the toroidal direction which points into the page. a) shows a zonal flow as a pair of counter-propagating flows which are radially localised. b) shows a convective flow cell wrapped around the magnetic axis, which is an alternative perspective of zonal flow structure.

mode (GAM), and second, a slower, in-compressible response referred to as the zero-frequency (or residual) zonal flow¹. In both cases, the plasma response takes the form of an $E \times B$ flow, as described by Equation 10.1. Considering the magnetic geometry of a large aspect ratio tokamak, the $E \times B$ direction primarily aligns with the poloidal direction, with only a minor component pointing in the toroidal direction. While this explains how a predominantly poloidal flow can be established within a tokamak, it doesn't yet clarify the emergence of a pair of counter-propagating flows.

Returning to our initial perturbation, we can deduce the existence of an equivalent perturbation but with the opposite sign, invoking the principle of quasi-neutrality of the plasma. This opposing perturbation must reside at an adjacent radius due to the axisymmetry observed in both the poloidal and toroidal directions.

In summary, any initial axisymmetric potential perturbation must come in pairs with opposite signs at adjacent radii. These perturbations induce $E \times B$ flows that run in parallel - but opposite - directions, precisely justifying the spatial structure of zonal flows.

Temporal Behaviour and Predator-Prey Oscillations

The zonal flow maintains a steady-state, and the mode is typically characterised as having a zero-frequency. This is evidenced by Equation 10.1, which implies that there

¹GAMs and residual zonal flows both technically fall under the umbrella of 'zonal flow'. However, the term 'zonal flow' is often colloquially used to refer to the residual component only. This convention will be used throughout this thesis for the sake of brevity, and GAMs will be referred to as such to distinguish the oscillatory branch.

is no acceleration of the flow if the fields are static.

In practice, the zonal flow mode is often referred to as a quasi-zero frequency phenomenon. This is due to a predator-prey relationship between the zonal flow and the drift-waves: Initially, the zonal flow mode grows through non-linear interactions with the drift-wave spectrum. Subsequently, the shear generated by the zonal flow suppresses the drift waves. The driving force behind the zonal flow diminishes, leading to the decay of the zonal flow mode. Eventually, as the drift-wave spectrum regains strength due to a lack of suppression, it kick-starts the next cycle of this oscillation. This cycle repeats at a characteristic frequency, typically on the order of kHz [32]. A further description of the predator-prey model is covered in section 3.3.1.

In conclusion, while zonal flows are theoretically zero-frequency, from an experimental measurement perspective, they are more accurately described as ‘quasi-zero frequency’ due to the predator-prey-like oscillations between zonal flow and drift-waves.

A Non-Axisymmetric Field Caveat

The magnetic field within a tokamak is inherently poloidally non-axisymmetric, with a stronger field strength closer to the central column. This non-axisymmetry should theoretically lead to a non-axisymmetric $\mathbf{E} \times \mathbf{B}$ flow. However, this presents a challenge since the zonal flow, by definition, is considered in-compressible[33]. Resolving this issue involves two plausible approaches:

- **radial flow component:** one possibility is the existence of a radial flow component with a return flow occurring at a different radius. This would eliminate the need for compression and allow the zonal flow to remain in-compressible.
- **toroidal Return Flow:** alternatively, a toroidal return flow could also mitigate the need for compression, preserving the in-compressibility of the zonal flow.

The exact nature of these return flows in actual tokamaks remains uncertain. Both scenarios are plausible, and it’s also possible that they coexist.

Exploring density perturbations as an experimental marker

Additionally, an intriguing avenue of investigation relates to the density perturbation associated with zonal flows. In cases featuring Boltzmann electrons, there is the usual expectation that density perturbations are directly proportional to potential perturbations. However, a more rigorous analysis reveals an expression for density perturbations, as shown in Equation 3.5, involving the zonal flow radial wavenumber (q_r) and the ion gyro-radius (ρ_i).

$$\frac{\tilde{n}_i}{n} \approx q_r^2 \rho_i^2 \frac{e\tilde{\phi}}{T_e} \quad (3.5)$$

This analysis indicates that density perturbations (\tilde{n}) are indeed linked to electric potential perturbations ($\tilde{\phi}$), but the pre-factor ($q_r^2 \rho_i^2$) is typically small in modern tokamaks [33]. Consequently, the density perturbation associated with zonal flows tends to be negligible in most cases, making it unsuitable as an experimental marker.

3.2.2 Zonal Flow Generation

Zonal flows emerge through non-linear interactions with drift waves, a phenomenon extensively examined in prior literature using both theoretical and computational methods [33]. Consequently, the fundamental theory underlying this process has been well-established, and the behaviour of zonal flow growth rates is theoretically well-defined. A thorough understanding of zonal flow growth necessitates a step-by-step derivation of its growth rate. However, a complete derivation lies beyond the scope of this thesis and instead we will focus on a heuristic approach to understand the growth of zonal flows. Consequently, this section provides an overview of zonal flow generation theory, prioritising clarity and comprehension over the inclusion of mathematical details. For a more detailed discussion of this theory and additional references, I refer the reader to P. Diamond's review paper on zonal flows [33].

First and foremost, it's important to highlight that zonal flows and drift waves occupy different spatial and temporal scales. Zonal flows are characterised by quasi-zero frequencies (typically \leq kHz) and mesoscale lengthscales (between the gyro-radius and machine size), while drift waves operate at considerably higher frequencies (typically in the range of 10-100 kHz) and lengthscales on the order of the gyro-radius. In turbulent systems, energy transfer between different modes typically occurs through cascading processes, where energy is continuously transferred through intermediate frequencies via local interactions.

However, a remarkable feature of zonal flows and drift waves is their non-local interaction in frequency and wavenumber space. In this context, energy transfer doesn't involve intermediate frequencies. To provide some general insight into non-local energy transfer phenomenon, let's delve into a transfer process known as the Reynolds stress. This process is prevalent in turbulent systems and is believed to be the underlying mechanism behind the interaction between drift waves and zonal flows.

Let's begin by considering a generalised incompressible fluid. We can express the relevant Navier-Stokes momentum equation as:

$$\rho \frac{Du_i}{Dt} = -\frac{\partial p}{\partial x_i} + \mu \frac{\partial^2 u_i}{\partial x_j^2} \quad (3.6)$$

In this equation, u represents the fluid flow velocity, x_i is a component of the coordinate vector \mathbf{x} , ρ is the fluid density, p stands for pressure, and μ denotes dynamic viscosity. Unsurprisingly, Equation 3.6 illustrates that the evolution of the fluid flow depends on the pressure gradient and viscosity.

Next, it's common practice to decompose the total fluid flow in a turbulent system into two components: the mean flow, denoted as \bar{u} , and the fluctuating flow, represented as u' . By employing the linear combination $u_i = \bar{u}_i + u'_i$, we can substitute this into Equation 3.6. Taking a time average of the expression and thus eliminating the fluctuating quantities, we arrive at the following equation:

$$\rho \frac{D\bar{u}_i}{Dt} = \rho \left(\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} \right) = -\frac{\partial \bar{p}}{\partial x_i} + \frac{\partial}{\partial x_j} \left(\mu \frac{\partial \bar{u}_i}{\partial x_j} - \rho \overline{u'_i u'_j} \right) \quad (3.7)$$

In Equation 3.7, we observe that the time evolution of the mean fluid flow still depends on the pressure gradient and a viscosity term. However, there is an additional term,

$\overline{\rho u_i' u_j'}$, which emerges. This new term serves to link the time evolution of the mean flow (zonal flow) to the nonlinear interaction with fluctuating flows (turbulence).

Having illustrated how energy can be transferred through nonlinear interactions between drift waves and zonal flows, let's delve into a more specific derivation. This derivation aims to establish an expression for the growth rate of zonal flows in the presence of a drift wave spectrum. The overall structure of the derivation will be as follows:

1. An expression for the evolution of a seed zonal flow in the presence of a single drift wave mode will be derived
2. The zonal flow evolution equation will be averaged over the interaction with an ensemble of drift waves
3. The response of the drift wave ensemble to the zonal flow shear will be derived to find the self-consistent evolution of the zonal flow
4. The growth rate of the seed zonal flow will be determined

We begin with the Charney-Hasegawa-Mima equation [25, 26], which serves as our framework for describing the evolution of electric potential in a turbulent plasma:

$$\frac{\partial}{\partial t}(n - \Delta_{\perp}\phi) + (\hat{b} \times \Delta_{\perp}\phi) \cdot \Delta_{\perp}(n - \Delta_{\perp}\phi) + \frac{\partial}{\partial y}\phi = 0 \quad (3.8)$$

Here, ϕ represents the electric potential, n denotes plasma density, \hat{b} signifies the magnetic field direction, and y is the direction perpendicular to both the magnetic field and the minor radial direction.

As our focus lies on drift waves and zonal flows, we define the following decomposition:

$$n = n_{zf} + n_d \quad \text{and} \quad \phi = \phi_{zf} + \phi_d \quad (3.9)$$

In these equations, the 'zf' subscripts refer to zonal flows, and 'd' signifies drift waves. Note that zonal flows exhibit negligible density perturbations ($n_{zf} = 0$), while the drift wave density is proportionally related to the drift wave potential, assuming adiabatic electrons ($n_d \propto \phi_d$). Consequently, we can eliminate density from Equation 3.9, yielding:

$$\frac{\partial}{\partial t}\nabla^2\phi_{zf} = -\frac{\partial}{\partial t}\langle \tilde{v}_{rd}\nabla^2\tilde{\phi}_d \rangle - \gamma_{damp}\nabla_r^2\phi_{zf} \quad (3.10)$$

Here, v_{rd} represents the drift wave velocity, and γ_{damp} is a generalised damping coefficient. **Equation 3.10 illustrates that the time evolution of the zonal flow potential is uniquely driven by the fluctuating drift wave potential.** Specifically, the zonal flow potential vorticity, $\nabla^2\phi_{zf}$, evolves based on the spatial flux of the drift wave potential vorticity, $\tilde{v}_{rd}\nabla^2\tilde{\phi}_d$. This demonstrates that zonal flow evolution is not a process of net vorticity generation, but rather a process of potential vorticity transfer.

Moving forward, we expand the scope of Equation 3.10 from the interaction with a single drift wave, $\hat{\phi}_d$, to the envelope of an ensemble of drift waves, $|\phi|^2$. This expansion is justified as a single zonal flow is likely to interact with multiple drift waves. Furthermore, we introduce the concept of drift wave density, denoted as N , which serves as a

conserved quantity representing the drift wave envelope. It can be related to the drift wave envelope by the expression $N \propto (1 + k_{\perp}^2 \rho_s^2) |\phi_k^d|^2$. This leads us to the following equation:

$$\frac{\partial}{\partial t} \tilde{V}_{zf'} = \frac{1}{B^2} \frac{\partial^2}{\partial r^2} \int d^2k \frac{k_r k_{\theta}}{(1 + k_{\perp}^2 \rho_s^2)^2} \frac{\delta N}{\delta \tilde{V}_{zf'}}(k, r, t) \tilde{V}_{zf'} - \gamma_{damp} [\tilde{V}_{zf'}] \quad (3.11)$$

In this equation, k_r , k_{θ} , and k_{\perp} represent the wavenumbers of the drift wave envelope $|\phi|^2$, ρ_s corresponds to the ion gyro-radius at T_e , and \tilde{V}'_{zf} denotes the zonal flow shear. Equation 3.11 illustrates the closed feedback loop and the self-consistent evolution of the zonal flow in the presence of a drift wave ensemble.

At this point, it becomes evident that zonal flow growth is expected if the drift wave modulational response to the zonal flow shear, $\partial N / \partial \tilde{V}'_{zf}$, is positive. The quantity $\partial N / \partial \tilde{V}'_{zf}$ is determined through linearisation of the relevant wave kinetic equation (see [33]). This leads to the final expression for the zonal flow growth rate:

$$\Gamma = \frac{-q^2}{B^2} \int dk^2 \frac{k_{\theta}^2}{(1 + k_{\perp}^2 \rho_s^2)^2} \frac{\gamma_k}{(\Omega - qv_g)^2 + \gamma_k} \left(\frac{k_r \partial \langle N \rangle}{\partial k_r} \right) - \gamma_{damp} \quad (3.12)$$

Here, Ω represents the instability eigenfrequency, γ_k is a linear damping term, Γ signifies the zonal flow growth rate, and q is the radial wavenumber of the zonal flow.

An essential observation from Equation 3.12 is that a positive growth rate necessitates $\partial \langle N \rangle / \partial k_r < 0$. This condition is met within any realistic drift wave spectrum, as the alternative would imply a population inversion of the drift waves. Additionally, it's worth noting that the damping rate, γ_d , is typically expected to be small in comparison to the drive term when the zonal flow starts with a small seed amplitude. **Consequently, we can conclude that the growth of zonal flow seed perturbations is positive in any tokamak.**

Furthermore, when investigating the corresponding transport coefficients, it can be found that the radial transport associated with zonal flows is orders of magnitude slower than the turbulent transport.

In conclusion, zonal flows are theoretically predicted to be widespread in tokamak plasmas and represent an improvement in confinement when compared to transport driven by drift wave turbulence. This is why zonal flows are considered a reliable and benign repository for free energy in tokamaks.

3.2.3 Zonal Flow Damping

The preceding subsection demonstrated that the zonal flow mode will invariably experience a positive initial growth rate due to its interaction with the drift wave spectrum. However, this alone doesn't guarantee that the zonal flow mode will saturate at a significant amplitude. If the damping rate scales more rapidly than the zonal flow growth rate, the zonal flow will remain negligible. Therefore, understanding the damping mechanisms is crucial to assessing the extent of the zonal flow's role in tokamak dynamics.

In any plasma, charged particles naturally counteract electric-potential-deviations from equilibrium through a process known as 'polarization shielding'. Typically, this phenomenon is highly efficient and has the capability to completely shield potential

perturbations, provided that there is an ample supply of charged particles. To assess the effectiveness of this shielding, it's crucial to consider both temporal and spatial scales of particle transport.

The zero-frequency nature of zonal flows means that the transport timescale isn't a limiting factor. Instead, the effectiveness of shielding hinges on the restricted spatial mobility of plasma particles. Notably, interactions with banana orbits can facilitate particle transport over a scale up to the banana width. Given that zonal flows are predicted to exhibit radial wavelengths comparable to the banana width, the extent to which polarization shielding suppresses the zonal flow potential remains uncertain.

This question was addressed theoretically by Rosenbluth and Hinton in a large-aspect-ratio tokamak with circular geometry [34]. They found that linear collisional damping had a negligible effect on zonal flows in the deep banana regime. Instead, they quantified the effect of neoclassical polarization shielding and discovered that the zonal flow potential is not entirely shielded, leaving a residual potential with the following form:

$$\frac{\phi_q(t)}{\phi_q(0)} = \frac{1}{1 + 1.6\varepsilon^{-1/2}q^2} \quad (3.13)$$

Here, $\varepsilon = r/R$ represents the inverse aspect ratio, and q is the radial wavenumber of the zonal flow. Equation 3.13 reveals that the aspect ratio of the tokamak plays a pivotal role in influencing the residual amplitude of zonal flows; Tokamaks with a tighter aspect ratio are expected to support zonal flows with a proportionally larger residual potential. Moreover, zonal flows with larger radial wavelengths exhibit a proportionally greater residual potential, presumably because spatially larger potentials are less effectively shielded. Additionally, other collisionless damping mechanisms, such as transit-time magnetic pumping and Landau damping, were shown to have no significant impact on the zonal flow mode [35, 36].

Instead, the primary damping of the residual zonal flow is driven by collisional mechanisms. However, this process is not simply collisional damping but is additionally associated with the asymmetry of the magnetic field. When a fluid element moves from the low-field side to the high-field side of the tokamak, it accelerates anisotropically due to the conservation of magnetic moment, leading to a departure from thermodynamic equilibrium. This anisotropy is subsequently thermalised through collisional processes, draining energy from the zonal flow mode and resulting in overall damping. Multiple more detailed collisional processes can thermalise the zonal flow mode, and their relative significance depends primarily on the ion-ion collision frequency. This dependence is shown in figure 3.5, which highlights that zonal flows are expected to have larger amplitudes at either very low collision rates or very high collision rates compared to the connection frequency. A detailed description of the underlying processes is omitted from this thesis but can be found in P. Diamond's review [33].

In summary, it has been demonstrated that polarization shielding is unable to entirely suppress zonal flows, leaving behind a residual zonal flow potential. Instead, the primary source of damping arises from collisional processes. In the absence of faster decay mechanisms, zonal flows have the capacity to grow to substantial amplitudes [33].

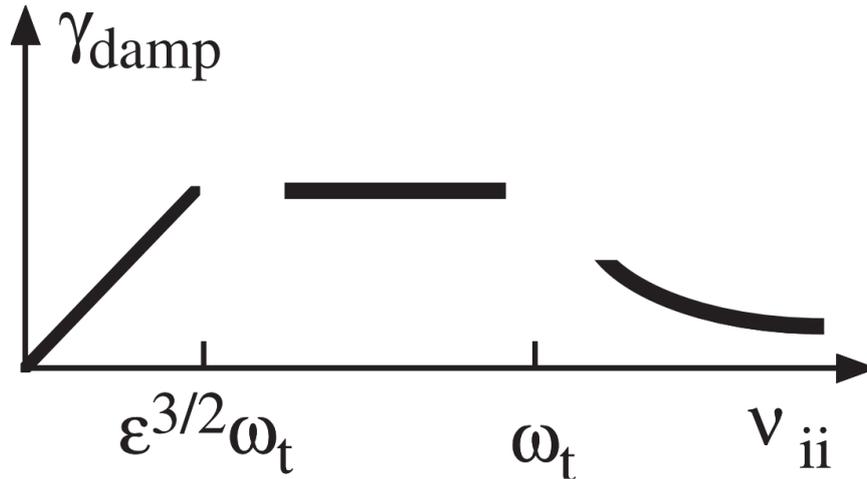


Figure 3.5: Theoretical zonal flow collisional damping rate dependence on ion-ion collision rate. ω_t is the connection frequency. Image adapted from [33].

3.3 The Drift-Wave–Zonal-Flow System

As previously discussed in sections 3.1 and 3.2, drift-waves, and subsequently zonal flows, are anticipated to be pervasive within tokamak plasmas. In this section, we will delve into the interplay between drift-waves and zonal flows, which forge a unified and self-consistent system. Our primary objective here is to furnish a basic overview of the system’s dynamics and equilibrium states.

To cover some specific aspects of the drift-wave–zonal-flow (DW-ZF) system, two toy models will be discussed. These models will highlight aspects like the oscillations between drift-wave and zonal flow amplitudes, and the existence of chaotic regimes. Furthermore, we will embark on a more comprehensive portrayal of the system by presenting numerical solutions to the relevant nonlinear diffusion equations. We will also address the levels of turbulent transport and its scaling in the DW-ZF system.

By considering these different perspectives, there are numerous themes that will consistently emerge. Specifically, we will observe the competition for energy between zonal flows and drift-waves, the inhibitory effect of zonal flows on turbulent transport, and the existence of a regime where the influence of zonal flows nearly extinguishes turbulent transport altogether.

For a more detailed discussion and in-depth information, I recommend referring to P. Diamond’s 2005 review paper [33].

3.3.1 DW-ZF Self-Consistent State

Predator-Prey Model

In this subsection we will cover a simple model capable of capturing some aspects of the DW-ZF system. This model supposes that the zonal flow is driven by the drift waves, while the drift waves, in turn, are suppressed by the presence zonal flows.

Analytically, the model is described in equation 3.14

$$\begin{aligned} \left(\frac{\partial}{\partial t} - \gamma_L + \gamma_{NL} \right) \langle N \rangle &= -\alpha \langle U^2 \rangle \langle N \rangle \\ \left(\frac{\partial}{\partial t} + \gamma_{damp} \right) \langle U^2 \rangle &= \alpha \langle U^2 \rangle \langle N \rangle \end{aligned} \quad (3.14)$$

Here, γ_L represents the linear drift-wave growth rate, and γ_{NL} represents the nonlinear self-saturation of drift-waves. γ_{damp} corresponds to the zonal flow damping term, α characterises the zonal flow's nonlinear drive coupling strength, $\langle U^2 \rangle$ stands for the zonal flow amplitude, and $\langle N \rangle$ represents the drift wave amplitude.

Now, let's have a look at some of the behaviours predicted by this model.

Predator-Prey Stable Points

There exist two steady-state solutions within the model, primarily governed by two key parameters: γ_L and γ_{damp} . Other parameters hold no substantial influence over the qualitative behaviour of these solutions. The transition between these two states becomes evident when varying γ_{damp} , as is shown in figure 3.6.

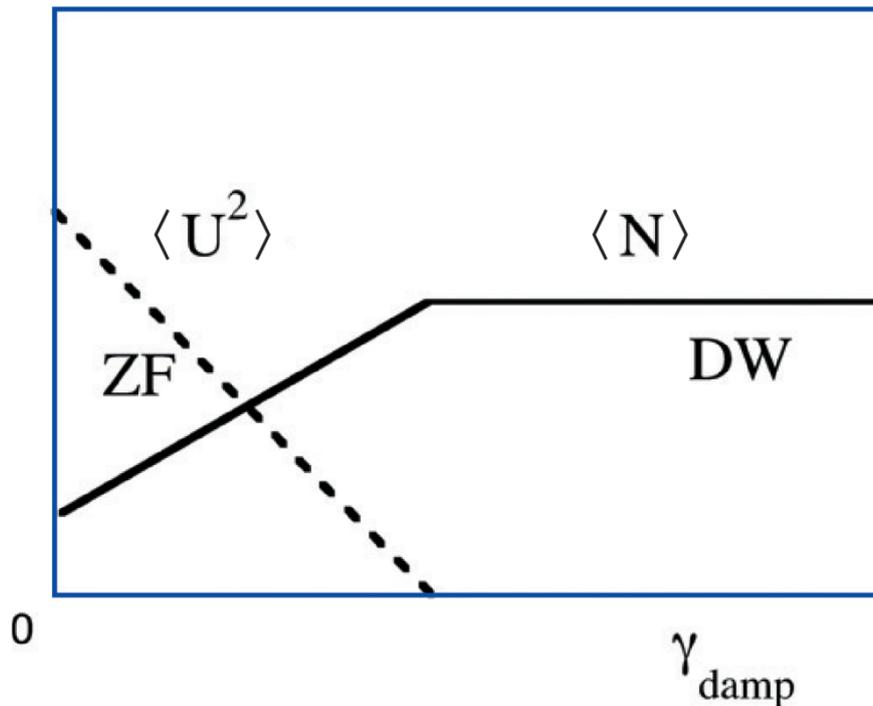


Figure 3.6: Steady state solutions to the drift wave-zonal flow predator-prey system. Y-axis represents the mode amplitude. Zonal flow mode shown by dotted line and drift wave mode shown by solid line. Image adapted from [33].

At high γ_{damp} , the model solution is given by the following:

$$\langle N \rangle = \gamma_L / \gamma_{NL} ; \langle U^2 \rangle = 0$$

It can be seen that only drift waves exist at nonzero amplitude in this regime. However, this regime is not expected to be observed experimentally because zonal flows are predicted to have low damping rates.

Conversely, when γ_{damp} is low ($0 < \gamma_{damp} < \frac{\alpha}{\gamma_{NL}}\gamma_L$), the solution is as follows:

$$\langle N \rangle = \gamma_{damp}/\alpha ; \langle U^2 \rangle = \alpha^{-1}(\gamma_L - \gamma_{NL}\gamma_{damp}\alpha^{-1})$$

In this regime, drift waves and zonal flows coexist with nonzero amplitudes, revealing several noteworthy features:

- As expected, the zonal flow population increases as the drift wave growth rate increases.
- The zonal flow and drift waves compete for energy. As γ_{damp} decreases, zonal flow amplitude grows while the drift wave amplitude shrinks.
- Interestingly, the drift wave population remained independent of γ_L and is instead entirely governed by γ_{damp} .

Predator-Prey Dynamics

In this system, any deviation from the stable states, $\{\langle N \rangle, \langle U^2 \rangle\}$, leads to a damped oscillation back to the stable solution. This decay may involve transient oscillations between drift-waves and zonal flows. Two example limiting cases of the dynamics are shown in figure 3.7.

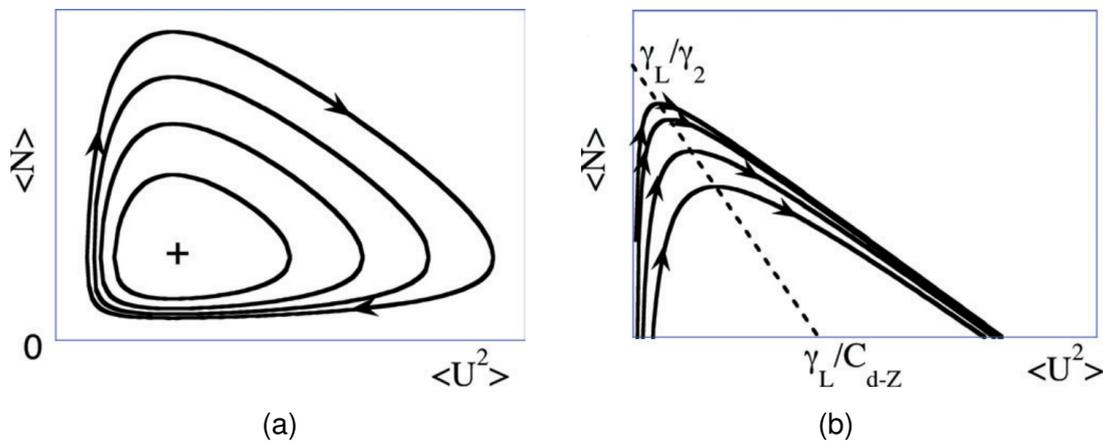


Figure 3.7: Phase space trajectories of the predator-prey model for the DW-ZF system. a) shows oscillatory behaviour when there are no nonlinear self-stabilisation mechanisms of the drift-waves, $\gamma_{NL} = 0$. b) shows the case when there is no zonal flow damping, $\gamma_{damp} = 0$. Single-burst quenching events are shown where the final state depends on initial conditions. Image adapted from [33].

Figure 3.7a illustrates one limiting case where the nonlinear self-stabilisation process of drift-waves is excluded. Oscillations occur between drift-waves and zonal flows around the stable point.

Figure 3.7b presents the other limiting case where zonal flows are not linearly damped. Here we observe single-burst quenching events of the drift-waves. Initially, the drift-waves grow, driving the zonal flow. As the zonal flow grows, the drift-waves decay and eventually reach zero amplitude. There is a complete quench, leaving only the zonal flow. This scenario is believed to be the underlying mechanism in the Dimits regime, where turbulent transport is very strongly suppressed by zonal flows [37].

Single-Instability Model

A system is considered where a primary drift-wave is unstable with two stable sidebands, in the presence of a zonal flow which is introduced artificially. Two observations are made which agree with the predator-prey model. Firstly, it is found that the drift-wave growth rate, γ_L , is entirely governed by the zonal flow damping rate, γ_{damp} . Secondly, the Dimits regime is reproduced in the limit of negligible zonal flow damping.

Interestingly, one novel observation is made in this model. As γ_{damp} is increased, bifurcations in the system occur and eventually lead to a chaotic system, as can be seen in figure 3.8.

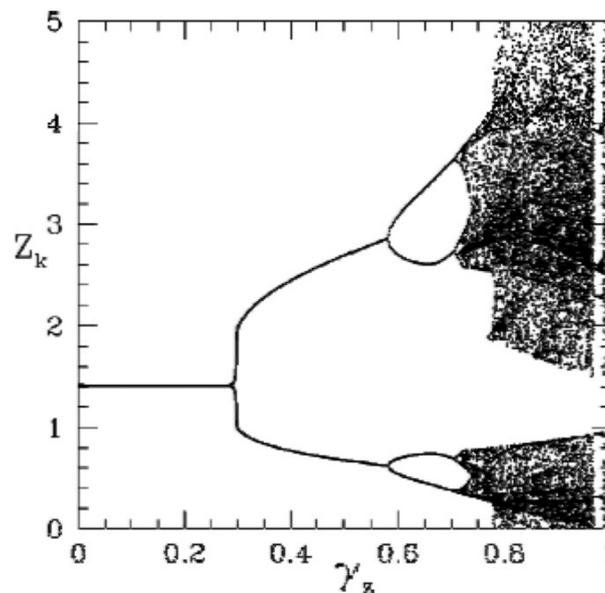


Figure 3.8: Bifurcations observed in the normalised zonal flow amplitude, Z_f , as the normalised damping rate, γ_z , increases. Bifurcations eventually lead to a fully chaotic system. Image adapted from [33].

Numerical Solutions

A set of nonlinear diffusion equations, which will not be shown here for brevity, can be derived to predict the self-consistent spectra of the drift-waves and zonal flow [38]. In this model, the zonal flow reaches steady state when it's nonlinear drive is balanced by collisional damping. Similarly, the drift waves reach saturation when its linear drive is balanced by linear damping, nonlinear damping and k-space diffusion induced

by the zonal flow shear. Upon solving these diffusion equations, the basic trends of the predator-prey and single instability models are qualitatively reproduced, although quantitative differences are observed.

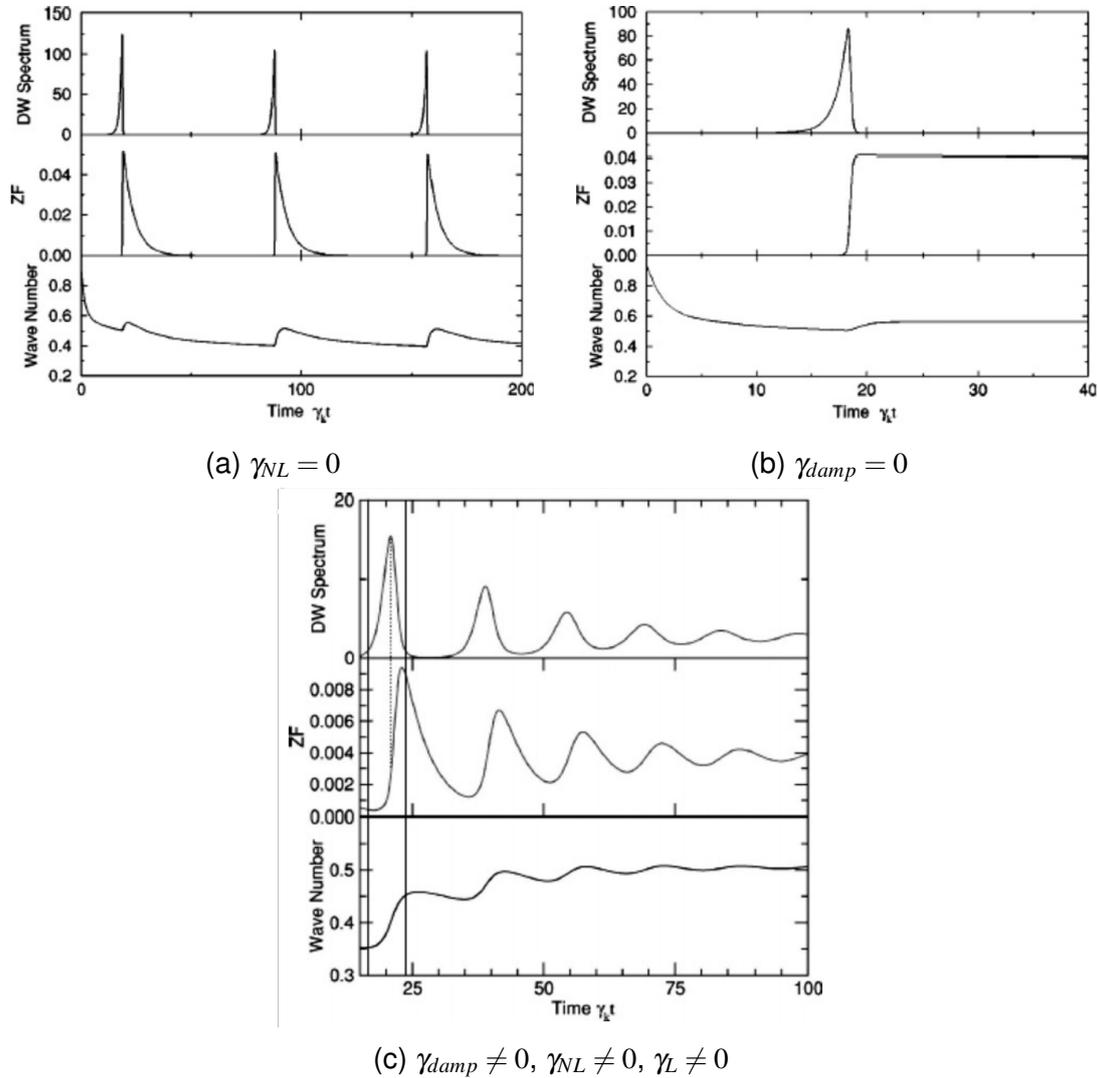


Figure 3.9: Time-evolution of the DW-ZF system according to nonlinear diffusion equations [38]. Three cases shown. Limit-cycle oscillation in a), single-burst quenching in b) and transient oscillations before saturation in c).

As shown in figure 3.9a, limit-cycle oscillations are observed, in agreement with the predator-prey model. Similarly, figure 3.9b exhibits single-burst quenching, which was also observed in the predator-prey model.

Finally, a more general case is shown in figure 3.9c. Here, it can be seen that a steady state is achieved after transient oscillations. Just like the toy models, the energy partition in the steady state between the zonal flows and drift-waves is governed by the zonal flow damping rate, γ_{damp} . Another interesting observation is that the average wavenumber of the system increases with time, a result of the drift-wave k-space diffusion induced by the zonal flow shearing.

3.3.2 Turbulent Transport Suppression

This subsection will briefly introduce some analytical expressions which highlight the suppressive effects of the zonal flow shear on turbulent transport. Two expressions will be shown, one with an artificially introduced shear flow and one that represents a self-consistently evolved DW-ZF system. Both expressions show a significant reduction in turbulent transport due to shear flow. Interestingly, the self-consistent model additionally captures the transition from the transport-suppressed Dimits regime to the strong turbulence regime.

The Effect of an Artificial Shear

Several analyses have been performed to investigate the impact that an artificially introduced shear flow can have on the transport levels [39, 40, 41]. These studies focused on ITG turbulence but the general findings are also applicable to drift-wave turbulence.

$$\chi_{turb} \simeq \frac{(\gamma_L - \omega_{E1} - \gamma_{*1})^{1/2} \gamma_d^{1/2}}{k_y^2} \quad (3.15)$$

Equation 3.15 presents an expression that illustrates the scaling of turbulent transport in the presence of shear flows. As expected, it can be seen that increasing the ITG growth rate, γ_L , and the zonal flow damping rate, γ_{damp} , causes an increase in the turbulent transport. Additionally, increasing the shearing rates leads to a reduction in turbulent transport.

Transport Scaling in the DW-ZF System

A less simplistic study, which did not rely on the introduction of an artificial shear, will now be covered [33]. Here, the DW-ZF system was evolved self-consistently, and analytical expressions for the transport scaling were determined. The exact scaling depended on the saturation mechanism of the zonal flow, but the general trend was consistent. One example scaling is shown in Equation 3.16.

$$\chi_{turb} = \frac{1}{1 + \tau_c^2 v_z^2 q_r^2} \chi_{turb,0} \quad (3.16)$$

Here, τ_c^2 is the turbulence decorrelation time, v_z^2 is the saturated zonal flow velocity, q_r^2 is the zonal flow radial wavenumber, and $\chi_{turb,0}$ is the turbulent transport level in the absence of zonal flows.

Once again, it is observed that zonal flows reduce transport levels. However, upon closer inspection, it can be seen that there are two regimes in Equation 3.16. One regime occurs where $\tau_c^2 v_z^2 q_r^2 \sim 1$, which is the case where strong turbulence dominates the system. In this regime, turbulent transport scales identically to a system without zonal flows. The other regime occurs near marginal stability of the turbulence and $\tau_c^2 v_z^2 q_r^2 \gg 1$. In this case, turbulent transport is strongly suppressed. This is another example of the Dimits regime, and the reduction in turbulent transport has been observed experimentally [37].

3.4 Global Transport Organisation: $E \times B$ Staircases

In this section, we will introduce a globally organised set of transport processes known as the $E \times B$ staircase. Much like the system discussed in section 3.3, this system undergoes self-consistent evolution, with zonal flows playing a key role in mediating transport levels.

The term ‘staircase’ accurately characterises this phenomenon, as it originates from initially smooth thermodynamic profiles observed in simulations, which subsequently transform into quasi-periodic, step-like profiles. These distinctive profiles, exemplified in figure 3.10, appear to be a prevalent feature in flux-driven simulations, provided they encompass the relevant physics [42, 3]. Additionally, experimental indications of these structures have been observed, and we will delve into this aspect in section 4.1.

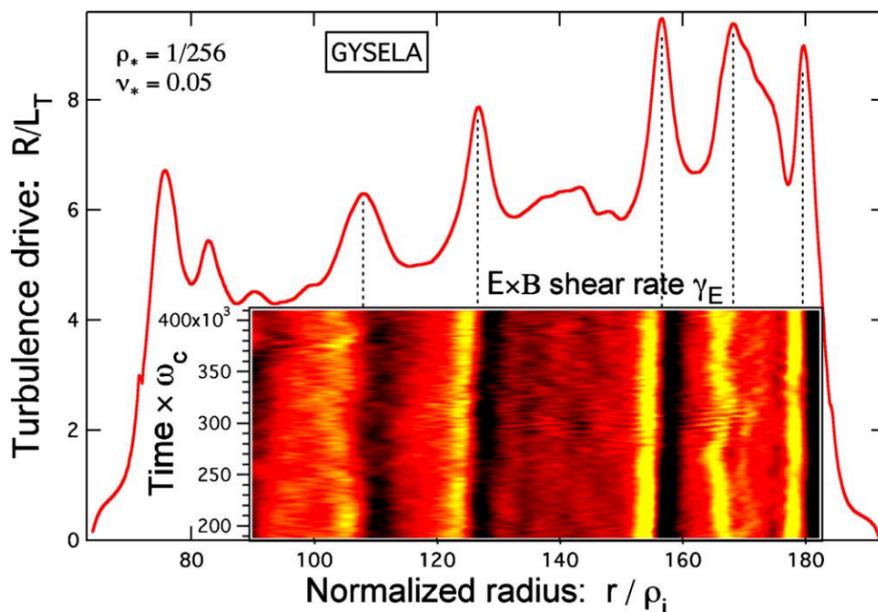


Figure 3.10: Temperature gradient profile corrugations observed in flux-driven GYSELA simulations of the Tore Supra tokamak. Localised steep profile gradients coincide with zonal flows (seen in the inset). Image adapted from [42].

3.4.1 Description of the Numerical Simulations

The simulations conducted aimed to replicate Tore Supra plasma geometry, enabling a direct comparison with experimental data. To achieve this, the GYSELA and XGCI codes were employed, running full-f, flux-driven gyrokinetic simulations [42, 3]. It’s worth emphasising the flux-driven nature of these simulations, as gradient-driven simulations are more commonly employed. In gradient-driven simulations, thermodynamic profile gradients are imposed, prohibiting the spontaneous emergence of staircase-like structures. In contrast, flux-driven simulations dictate the overall transport levels, allowing the profiles to evolve self-consistently and take any form.

A multitude of simulations were conducted, encompassing a wide range of parameters and scenarios:

- Various heating mechanisms were considered
- Collisionality was varied, spanning an order of magnitude within the banana regime
- Machine size was explored across the spectrum, from the smallest tokamaks in operation today to the scale of ITER
- The turbulence drive was also varied, ranging from scenarios of marginal stability to those characterised by strong turbulence

This comprehensive approach allowed for a thorough examination of the system's behaviour under diverse conditions and configurations.

3.4.2 $E \times B$ Staircase Characteristics

The aforementioned L-mode simulations consistently revealed the presence of the staircase pattern. In this section, we will delve into some noteworthy characteristics linked to this pattern. Crucially, this observation exhibited a remarkable level of robustness, showing no dependence on either machine size or the specific simulated experimental conditions. Consequently, it suggests that the $E \times B$ staircase phenomenon could be an inherent and nearly inevitable feature in tokamak plasma organisation.

One intriguing observation is the coincidence of strong profile gradients with poloidal shear flows, which were characterised by coherence times spanning from 1 to 100 milliseconds. These shear flows are typically spaced at intervals of 20 to 30 ion gyro-radii, with each flow having a width of approximately 10 gyro-radii. Moreover, these flows demonstrate a radial meandering behaviour, traversing the width of a shear flow within a timescale of 1 to 5 milliseconds. Additionally, the locations of these flows do not exhibit a direct connection to q -rationals, suggesting that they are not rooted in MHD processes.

In the intervals between zonal flows, nearly flat gradients are observed. These inter-shear regions exhibit periodic bursts of radial transport, which occur due to turbulent avalanche-like processes. A visual example of these avalanches is shown in figure 3.11, which was observed in fluid-model simulations [43]. These avalanching phenomena represent a rapid radial transport channel, justifying the locally nearly flat profile observations.

The zonal flows play a pivotal role as internal transport barriers, effectively blocking the avalanches that might otherwise propagate to the machine scale. This underscores the critical importance of zonal flows in suppressing turbulent transport. Without the presence of zonal flows, avalanches would likely lead to nearly flat profiles across the entire confined region of most tokamak plasmas. However, this blocking effect is semi-permeable, allowing larger avalanches to breach a zonal flow. Near states of marginal stability, the zonal flow may re-form but not necessarily at the same location. In cases where turbulence drive is strong, the zonal flow may not be given the opportunity to

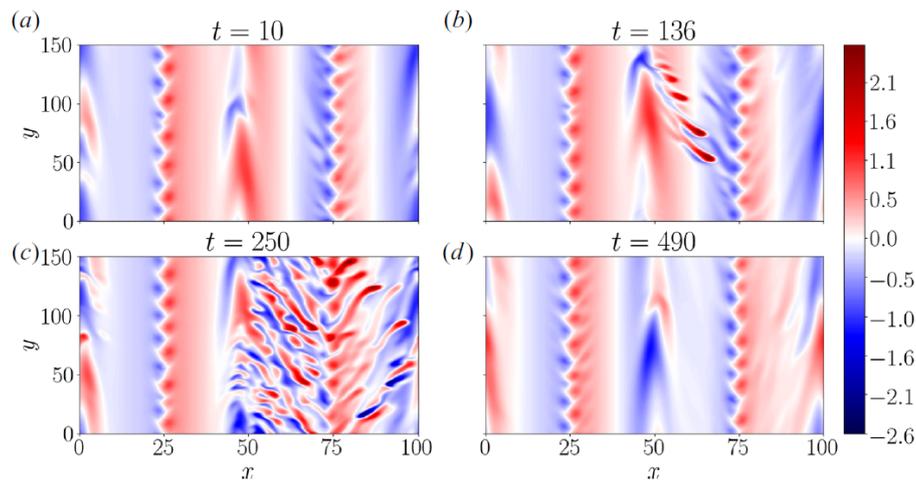


Figure 3.11: Two zonal flows shown at $x=25$ and $x=75$ in a fluid-model simulation, where the colour represents the flow velocity in the y -direction [all a.u.]. Intermittent turbulent structures known as ‘ferdinons’ form and propagate radially, temporarily breaching a zonal flow. Zonal flow reforms after breach. Image adapted from [43]

re-form, resulting in a regime dominated by turbulence. An illustrative example of intermittent turbulent structures breaching zonal flows is shown in figure 3.11.

The intricate interplay between zonal flows and avalanches gives rise to the concept of the mesoscale, represented by Δ . In numerical terms, Δ corresponds to the spacing between zonal flows. Below the scale of Δ , transport exhibits scale invariance and is primarily mediated by the presence of avalanches. Conversely, at scales surpassing Δ , avalanches become infrequent, especially near marginal stability, and transport characteristics are predominantly dictated by the properties of the zonal flows. This highlights importance of incorporating multi-scale effects into turbulence simulations, as the dominant transport mechanism fundamentally hinges on the scales encompassed within the simulation.

Finally, the $E \times B$ staircase system has been found to be generally advantageous for plasma confinement. Transport levels in a staircase-configured system are consistently measured to be 2 to 3 times lower than in plasmas that do not exhibit the staircase phenomenon.

Chapter 4

A Brief Review of Zonal Flow Experimental Measurements

The experimental detection of quasi-stationary zonal flows presents a formidable set of challenges. Part of this challenge originates from the intrinsic difficulty of measuring poloidal flows (or radial electric fields) in high-temperature plasmas. In addition to this, confirming a robust detection of the zonal flow mode requires the simultaneous execution of multiple measurements. Velocity measurements must be conducted simultaneously at several toroidal and poloidal positions to validate the flow mode's symmetry. This, in turn, necessitates the utilisation of (multiple) dedicated turbulence diagnostics, a task that is often hindered by cost limitations or the constraints of diagnostic space. The following ZF characteristics should be observable:

- Radially localised (i.e. sheared) poloidal flows
- Poloidal and toroidal mode symmetry ($n=m=0$)
- Near-zero frequency (sub-kHz)
- Nonlinearly driven by turbulence

The scientific literature is rife with hints of ZF characteristics, where papers claim to have 'observations consistent with quasi-stationary zonal flows.' Nevertheless, convincing direct identifications remain scarce.

Some early studies observed turbulent fluctuations with short radial correlation lengths and extended poloidal correlation lengths [44, 45]. Poloidally sheared flows have been documented in a range of fusion devices [46, 47, 48, 49]. GAMs, which are comparable to ZF in terms of spatial structure and the nonlinear turbulence drive, have repeatedly been detected [50, 51, 47, 52, 53, 54, 55]. However, GAMs have well-defined frequencies, typically on the order of ~ 10 kHz, which makes them significantly easier to isolate. Conversely, the longer timescales associated with the sub-kHz ZF requires longer steady-state plasmas to build up the statistics, which is often beyond experimental capabilities.

Compounding the situation is a matter of terminology: both quasi-stationary ZFs and GAMs technically exhibit flows with 'zonal' structures. Some research groups refer to both quasi-stationary zonal flows and GAMs as 'zonal flows', a practice which is not technically incorrect. Conversely, other groups reserve the term 'zonal flow' exclusively

for quasi-stationary zonal flows. It is emphasised that, in this thesis, the terms ‘zonal flow’ (ZF) and ‘geodesic-acoustic mode’ (GAM) refer specifically to the stationary and oscillatory branches of turbulence-driven shear flow, respectively.

4.1 Prior Observations of Zonal Flows

This section will cover a non-comprehensive overview of ZF related measurements. The intention was not to give a complete review of the field, but rather to point to some influential results of existing measurements and act as a starting point for further reading.

4.1.1 Observation of Radially Localised, Sheared Flows with $n=m=0$ near Zero Frequency

One early observation of zonal flows was reported in a 2004 Physical Review Letters publication by A. Fujisawa et al. [56]. They employed two Heavy-Ion-Beam Probe (HIBP) diagnostics on the Compact Helical System (CHS), which were toroidally separated by 90 degrees. Figure 4.1 provides a schematic of this setup. These diagnostics simultaneously provided local measurements of density and potential fluctuations at 500 kHz. The CHS was operated with a steady-state plasma to minimise MHD activity.

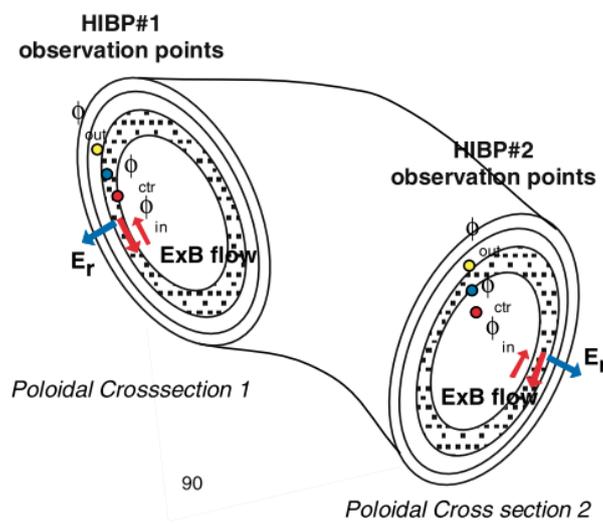


Figure 4.1: Geometry and observation points of the dual heavy ion beam probe set up in CHS. Each probe takes measurements from three different radial locations. Image adapted from [56].

Several observations supported the presence of zonal flows:

- The electric potential fluctuation power consistently peaked at frequencies less than 1 kHz, as depicted by the red line in figure 4.2a.

- Also in figure 4.2a, the blue line illustrated the coherence between potential fluctuations at different toroidal locations, which also peaked at low frequencies (<1 kHz).
- Although not shown here, the phase difference between potential fluctuations was found to be zero, not only when the probes were at the same radius but also when they were placed on the same field line. This confirmed toroidal symmetry across at least 90 degrees and poloidal symmetry from $\theta = 30$ degrees to $\theta = 50$ degrees.
- The peak coherence exhibited a strong variation when the radial position of one of the probes was adjusted, suggesting a significant shear flow. This dependency is depicted in figure 4.2b and implies a radial coherence length of $1.5 \text{ cm} \pm 0.5 \text{ cm}$ for the low frequency mode.
- The flow amplitude for this mode was approximately 0.1 km/s , in contrast to a background turbulent flow of around 1 km/s .

Summarising these observations, a radially sheared flow was identified with: toroidal and poloidal symmetry, a frequency in the range of $300\text{-}700 \text{ Hz}$, and radial localisation with a wavelength of $1.5 \text{ cm} \pm 0.5 \text{ cm}$. Taken together, these measurements strongly support the detection of the quasi-stationary zonal flow.

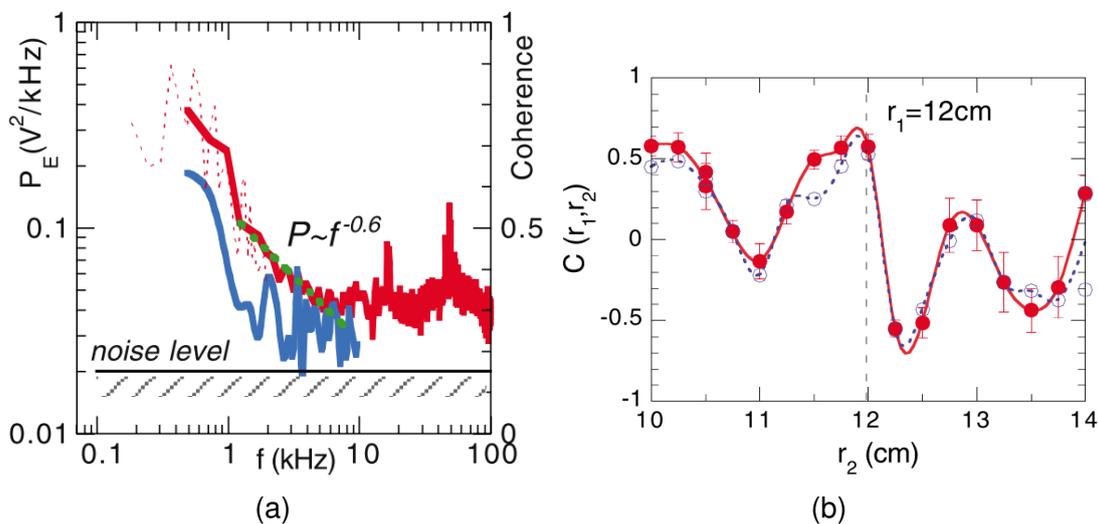


Figure 4.2: Results from CHS heavy ion beam probes detecting zonal flow. (a) Power spectra of potential difference (red), and coherence between potential differences (blue) at the two toroidal locations. Dashed red lines show the not-fully-converged power spectrum at longer segment lengths. (b) Radial correlation structure of zonal flow. Solid red circles represent the traditional correlation coefficient. Open blue circles represent coherent structure from FFT analysis. Both images adapted from [56].

4.1.2 Nonlinear Turbulence Drive of Zonal Flows

Various aspects of zonal flows were observed in the previously discussed paper by Fujisawa; however, the turbulence drive remained unaddressed. The experimental measurement of the nonlinear drive necessitates higher-order analyses, requiring longer integration times and higher signal-to-noise ratios to achieve statistical significance. This renders it a considerably more challenging measurement.

Early works involved the development of single-field nonlinear analyses capable of observing the nonlinear coupling strength between density fluctuations [57, 58]. For instance, these methods were employed to directly measure the nonlinear interactions between turbulent density fluctuations underlying the turbulent energy cascade in the Texas Experimental Tokamak (TEXT). However, nonlinear interactions between turbulent density fluctuations and flows could not be investigated in these experiments.

This section will discuss a paper that refined these early methods, enabling not only the determination of the *direction* of the energy transfer but also encompassing the multi-field nonlinear coupling between the density *and* velocity field [59]. Note that this thesis measured the nonlinear coupling between the GAM and the turbulence. However, these techniques have also been applied to ZF-turbulence coupling in subsequent studies [60, 61].

Measurements were conducted at the edge of a steady L-mode discharge of the DIII-D tokamak. The beam-emission spectroscopy diagnostic was utilised to infer 2D density fluctuation images at a 1 MHz sampling frequency. Velocity fields were inferred from the BES data using time-delay estimation techniques. In these particular shots, the GAM existed as an 18 kHz mode with significant coherence between the GAM and turbulence intensity.

The nonlinear interactions between density fluctuations and poloidal velocity fluctuations were quantified by measuring the transfer function, as represented in Equation 4.1. Here, $T_n^\theta(f', f)$ denotes the nonlinear transfer function between density fluctuations, $\tilde{n}(f)$, with frequency f , and density gradient fluctuations, $\frac{\partial \tilde{n}}{\partial \theta}(f')$, with frequency f' , mediated by the fluctuating poloidal flow V_θ . The angled brackets denote the short-time average. Positive $T_n^\theta(f', f)$ values indicate energy transfer from gradient fluctuations, $\frac{\partial \tilde{n}}{\partial \theta}(f')$, to density fluctuations, $\tilde{n}(f)$.

$$T_n^\theta(f', f) = -\text{Re} \left\langle \tilde{n}(f) \tilde{V}_\theta(f - f') \frac{\partial \tilde{n}}{\partial \theta}(f') \right\rangle \quad (4.1)$$

The quantities $\tilde{n}(f)$ and $\frac{\partial \tilde{n}}{\partial \theta}(f')$ could be directly extracted from the beam emission diagnostic, while the flow, V_θ , was inferred using time-delay estimation methods (see section 5.3). Thus, the transfer function could be calculated directly, and the result is shown in figure 4.3a. This figure illustrates that all density fluctuations with $f > 50$ kHz gain energy from the lower-frequency density gradient fluctuations. Conversely, the lower-frequency density gradient fluctuations lose energy to the higher-frequency density fluctuations. The frequency separation, $f - f'$, for these interactions is consistently equal to the GAM frequency of 18 kHz, suggesting that the GAM is mediating this energy transfer. Overall, the energy is progressively transferred to higher frequencies until it is dissipated at > 150 kHz.

For comparison with this experiment, nonlinear numerical simulations were con-

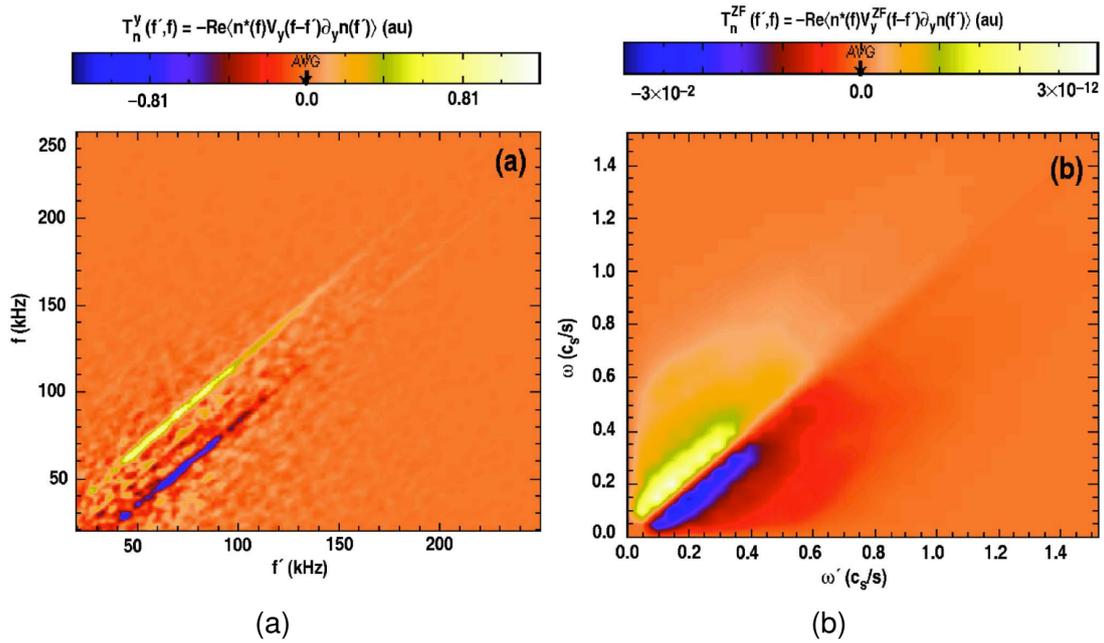


Figure 4.3: Nonlinear interactions between poloidal flow and turbulence. (a) Experimental bispectrum from BES diagnostic on DIII-D. Fluctuation energy is transferred to progressively higher frequencies through a mediating 18 kHz GAM. (b) Bispectrum of a simulated plasma from numerical simulations representative of the DIII-D experiment. No GAM was present in the simulations. Instead, a ZF mediates the nonlinear transfer. Images adopted from [59].

ducted using the gyrokinetic code ‘GYRO.’ A plasma similar to the experimental plasma was simulated, but instead of the GAM, quasi-stationary zonal flows were present. The corresponding transfer function is shown in figure 4.3b, which demonstrates a qualitative resemblance to the experimentally observed transfer function, with one difference: the two transfer lines are now adjacent to the $\omega = \omega'$ line, representing mediation of the quasi-zero-frequency zonal flow. Nonetheless, the overall pattern remains the same, with energy nonlinearly transferred to progressively higher frequencies until it is dissipated in the high-frequency range where the turbulence is linearly stable.

In summary, this thesis presents experimental measurements of the nonlinear interaction between GAM and turbulence. Fluctuation-energy transfer is mediated by the GAM and transferred to progressively higher frequencies until it is dissipated. Note the lack of direct evidence to identify the energy source of the GAM. Nevertheless, the experimental results were compared with numerical simulations, which showed qualitative agreement and suggested that ZF can play the same mediating role as the GAM. Since the ZF is not observed in these experiments, it is suggested that the GAM is instead the main regulator of turbulence.

4.1.3 Sheared Flows and Confinement Transitions

Sheared flows can play a pivotal role in confinement transitions and are believed to act as part of a triggering mechanism. Early observations indicate an increase in ZF or GAM power, coupled with a simultaneous decrease in turbulence power just before the L-H transition [48, 62, 63]. Subsequent studies on L-H transitions have solidified the critical role of sheared flows through phase delay observations and nonlinear transfer analyses [64, 65, 60, 66, 32, 67]. These findings have also been extended to I-mode confinement transitions, where GAMs were identified as playing a key role in the fluctuation dynamics [52, 66].

An illuminating narrative that underscores the role zonal flows can play in the L-H transition goes as follows. Let’s consider an L-mode plasma where turbulence and zonal flows exist in equilibrium. If the nonlinear drive of the zonal flow is temporarily enhanced, surpassing the turbulence drive, the zonal flow will rapidly absorb all turbulent energy. This large amplitude shear flow, combined with suppressed turbulence, results in strongly enhanced local gradients. If the local pressure gradient surpasses a critical threshold (30-100% larger than typical in L-mode [65, 32]), the gradient can self-stabilise. This is achieved through a gradient-driven neoclassical shear that suppresses transport and maintains the edge gradient. The zonal flow amplitude will decay due to a lack of turbulent drive, but the edge gradient will self-consistently remain, completing the transition to H-mode.

Note that large amplitude zonal flows are only one possible triggering mechanism. The transition into H-mode can be induced through various methods, provided a sufficiently prominent edge gradient, equivalent shear flow, or radial electric field well is initiated in the edge plasma. For example, confinement transitions have also been triggered by supersonic mass injection into the edge [68] and direct biasing [69, 70, 71], which initialised the edge gradient and E_r well, respectively.

4.1.4 Observations of the $E \times B$ Staircase

The emergence of the $E \times B$ staircase in numerical simulations has been discussed in section 3.4. Bolstering these numerical simulations, several experimental observations were made in the Tore Supra tokamak [3]. Radial profiles of turbulent fluctuations, along with the determination of radial correlation lengths of the turbulence, were facilitated through X-mode reflectometry. The $E \times B$ staircase is expected to be observable by imprinting its quasi-regular structure on the radial correlation length. As such, a corrugation of radial correlation lengths, consistent with the $E \times B$ staircase, was observed experimentally, as seen in figure 4.4. Additionally, the correlation length troughs did (mostly) not align with q -rationals, suggesting that MHD activity was not responsible for the observed structure. This experimental measurement, although indirect, provided a tantalising glimpse of the $E \times B$ staircase in an experimental plasma.

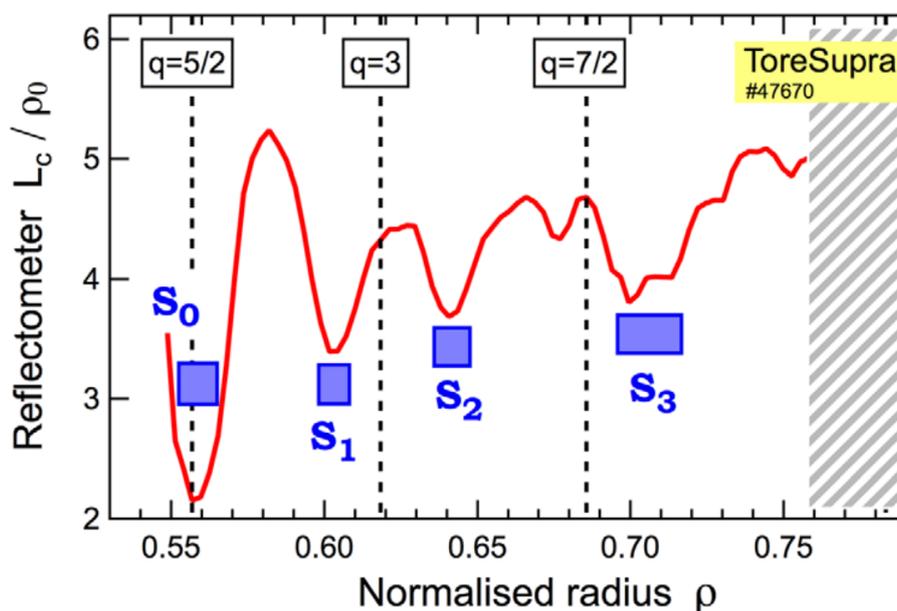


Figure 4.4: Reflectometer coherence length in Tore Supra showing corrugations as expected from the $E \times B$ staircase. S_1 , S_2 , and S_3 do not align with q -rationals, suggesting no MHD involvement. Image adapted from [3].

More recent experimental hints of the $E \times B$ staircase were presented in W. Liu's 2021 publication [72]. In these experiments, HL-2A, a large-aspect ratio tokamak, was utilised. The tokamak was operated in L-mode, with intermittent MHD activity observed in the relevant shots. However, significant steady-state periods were isolated between these MHD bursts for the purpose of their investigations.

To record data on density fluctuations, a BES system, featuring 2 (vertical) by 24 (radial) channels, was employed, encompassing minor radii ranging from 0.45 to 1.2 r/a . Some results from this diagnostic are presented in figure 4.5. Remarkably, interruptions in the radial structure of the fluctuations were intermittently observed at various locations, indicated by asterisks. The break in radial structure is most clear at the outermost asterisk, which corresponds to the last-closed flux surface (LCFS) and is therefore unrelated to the staircase. At the two innermost asterisks, one can

consistently observe radial structures extending beyond these positions, making these ‘breaks’ in structure less clear.

The structures in figure 4.5 *seem* to propagate radially but encounter barriers at these asterisk-marked locations, implying a potential connection with transport barriers. However, the barber pole illusion, known to cause misleading apparent motion, may be at play here due to the limited number of vertical channels [73]. Therefore, uncertainty exists regarding whether the observed propagation aligns with true motion or misleading illusory motion. Supplementary measurements, as portrayed in figure 4.6c), displayed a reduction in radial correlation lengths at the same asterisk-marked locations, although it is unclear if these corrugations are significant due to the lack of error margins. These observations may point to the existence of weak, semi-permeable transport barriers, but results remain inconclusive.

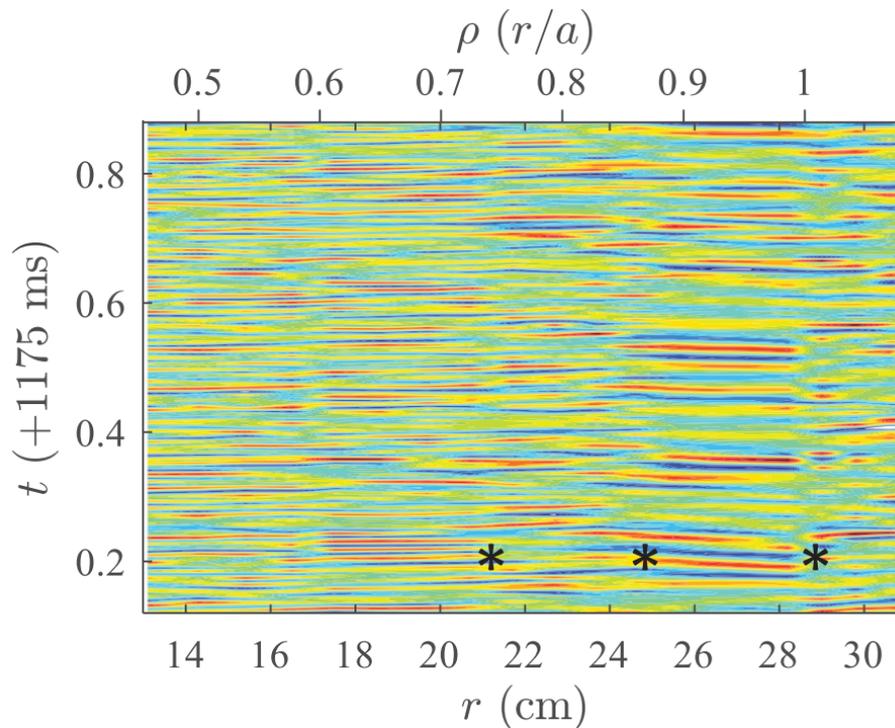


Figure 4.5: Radial-temporal evolution of normalised density fluctuations with a 25-120 kHz band-pass filter to isolate the turbulence. The asterisks indicate the pink box locations in figure 4.6. Image adapted from [72].

By examining the BES $k_r - k_\theta$ spectra, eddy tilting angles were investigated, although not shown here. The claim was made that increased eddy tilting was observed coinciding with the asterisk locations, hinting at the presence of localised shear flows. Although it is certainly true that structural disturbances are observed at these locations, it is unclear if these observations correspond to eddy tilting, especially considering the real possibility of aliasing with this BES set-up.

The $f - k_\theta$ spectra, which are also not shown here, show clear signs of aliasing at high wavenumbers. Despite this, reversals of the phase velocity were observed to occur at the asterisk locations. This is an interesting observation because flow reversals are not expected from the $E \times B$ staircase. Instead, localised shear flows

are expected to be observed on top of the background flow. The nature of these flow reversals remains unclear, although it is noteworthy that this structure could not be reproduced in measurements of the toroidal or neoclassical flows.

For the measurements of electron temperature and electron density profiles, the electron cyclotron emission and wave reflectometer diagnostics were relied upon, respectively. The results of these measurements can be found in figure 4.6a) and 4.6b). Local increases in the gradients are observed at the asterisk locations, which would be consistent with the staircase transport barriers. Contrary to this, the peaks are mostly within the given error margins, so it is unsure if these variations are statistically significant. Another point of contention is that the alleged gradient peaks coincide with q -rationals, making it challenging to disentangle the observations from MHD activity.

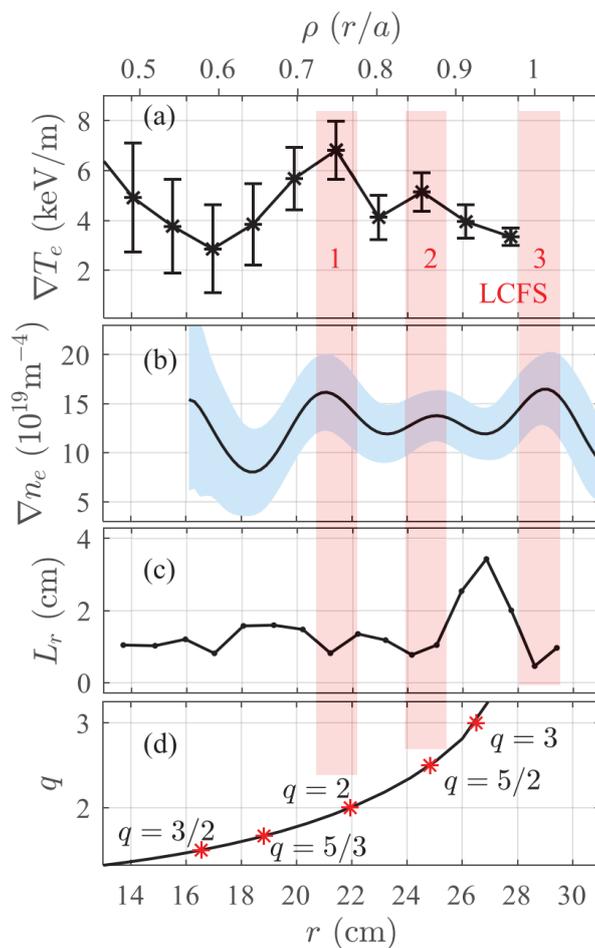


Figure 4.6: Profiles of (a) ∇T_e given by the ECE, (b) ∇n_e given by the FMCW reflectometer, (c) radial correlation length L_c of density fluctuations obtained from BES, and (d) safety factor, q , computed by the EFIT. Pink rectangle locations correspond to asterisk locations in figure 4.5. Image adapted from [72].

In conclusion, a tantalising glimpse of the $E \times B$ staircase was observed in Tore Supra, as shown in figure 4.4. Subsequent attempts (by Liu et. al. [72]) to replicate and build upon this observation have fallen short of achieving an unquestionably convincing result. Once again, tantalising indications of the $E \times B$ staircase were noted,

but a definitive observation remains elusive. Within these publications, the suggestion of employing methods such as velocimetry analysis for the identification and characterisation of sheared flows was repeatedly stated as a promising approach to detect the $E \times B$ staircase, which is a point that this thesis attempts to address.

Chapter 5

An Overview of the Tokamak, Diagnostics, and Analysis Techniques

This chapter introduces some experimental details and context necessary for the remainder of this thesis. The specific tokamak which was used to gather experimental data for this thesis is introduced in section 5.1. In addition to a basic overview of the machine, general pros and cons of spherical tokamaks are discussed.

The relevant diagnostic that was used to gather experimental data is also introduced in section 5.2, with a basic overview of the underlying physics and a description of the diagnostic capabilities.

Two velocimetry techniques commonly used in turbulence investigations are introduced in sections 5.3 and 5.4. Basic overviews of their underlying reasoning are given along with descriptions of step-by-step inner workings.

5.1 The Upgraded Mega-Ampere Spherical Tokamak

The Upgraded Mega-Ampere Spherical Tokamak (MAST-U) is currently the most powerful spherical tokamak in the world. Although referred to as an upgrade of the original MAST, over 90% of the components are new, resulting in large improvements to the machine capabilities. Higher performance and longer pulses can be achieved due to increased heating power and a stronger magnetic field. Additionally, a novel plasma exhaust system leads to decreased impurity contamination in the core whilst improving the divertor heat spreading. An overview of the machine parameters is given in Table 5.1 and a schematic is shown in figure 5.1. All in all, this machine is able to produce the most reactor-relevant plasmas of any spherical tokamak to date. This makes it useful for next-generation machines like ITER and especially relevant to the development of the UK-based Spherical Tokamak for Energy Production (STEP) project.

Spherical tokamaks can represent a significantly more compact design than conventional tokamaks, as seen in figure 5.2. This enables the use of smaller magnets, which represents a major cost saving in tokamak construction, and leads to more energy-efficient plasmas. Additionally, due to the magnetic geometry, the particles spend significantly more time on the ‘favourable’ inboard side of the plasma, which leads to enhanced plasma stability. This delayed onset of instabilities enables these machines to operate at higher plasma beta, which is a magnetic efficiency defined

as the ratio of the plasma pressure to the magnetic pressure. These benefits make spherical tokamaks a leading contender on the route to compact fusion power plants.

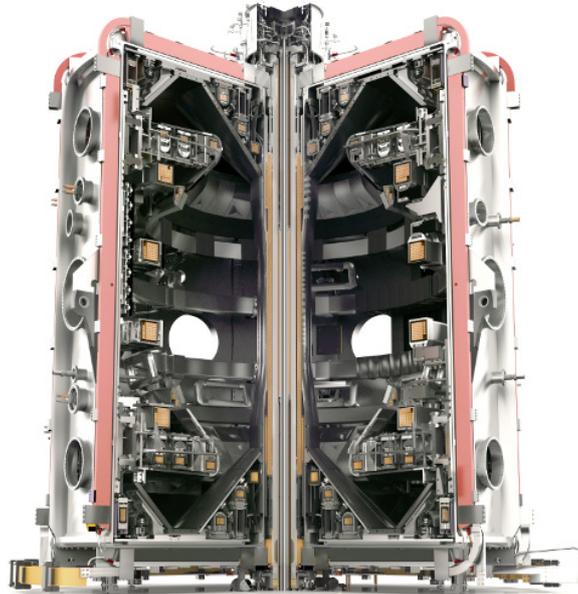


Figure 5.1: Rendering and cut-out of MAST-U.

MAST-U started operations with first plasma in October 2020. Campaign 1 ran from April 2021 until October 2021, mostly involving basic scenario development. Unfortunately, no BES data relevant to this thesis was produced during campaign 1.

Campaign 2 ran from May 2022 until January 2023. Throughout this campaign, the heating power, plasma current, magnetic field strength, and pulse length were *roughly* comparable to the plasmas observed in the old MAST. Additionally, the appearance of a locked mode resulted in limited rotation of the plasma, with typical toroidal velocities below 50 km/s, significantly below typical MAST velocities that exceeded 150 km/s. With the Beam Emission Spectroscopy diagnostic fully operational, this campaign was supposed to provide the bulk of raw data for analysis in this thesis. Unfortunately, all but 10 shots in campaign 2 had an optical filtering issue with the BES diagnostic, rendering the shots useless for ‘normal’ velocimetry analysis¹. Out of the remaining 10 shots, 9 disrupted early, leaving only shot #46459 as a suitable candidate for image-velocimetry analysis. This shot will be analysed in chapter 8. Campaign 3 was scheduled to start early summer 2023, which fell outside the timeline of this thesis.

¹I will self-indulge in stating that I was not involved in the operation of BES.

Parameter	Design value
Major and minor radii (m)	0.7/0.5
Pulse duration	typical 1-2s, max 5s
Electron density	$1.5 \times 10^{20} \text{ m}^{-3}$
Electron temperature	up to 3 keV
Toroidal rotation	typically 150 km/s
Normalised beta	3-6
First wall material	Carbon (graphite)
Maximum toroidal magnetic field (T)	0.8
Maximum plasma current (MA)	2.0
Maximum elongation	2.5
Triangularity range	up to 0.6
Diverter geometry	closed, pumped
Plasma heating	up to 6 MW ohmic + NBI

Table 5.1: MAST-U machine operational parameters and characteristics.

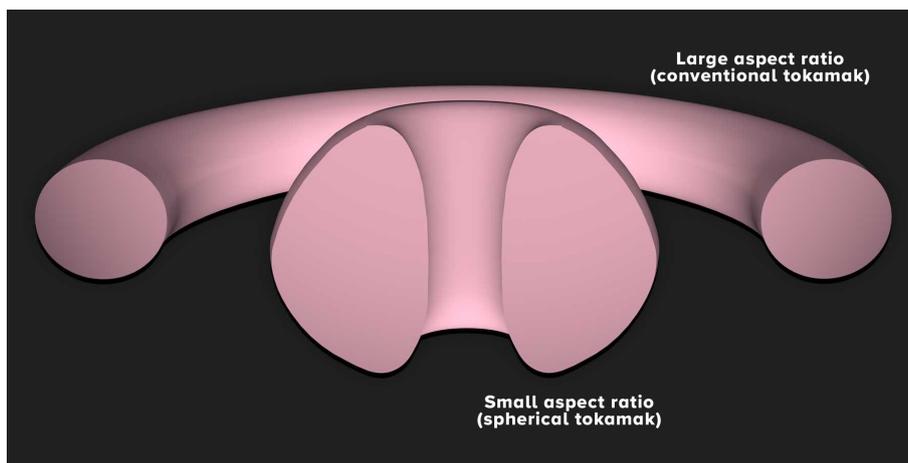
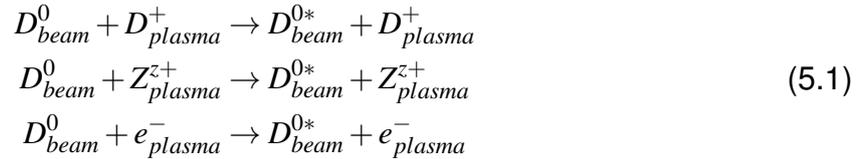


Figure 5.2: Plasma comparison between spherical tokamaks and conventional tokamaks.

5.2 Beam Emission Spectroscopy

The Beam Emission Spectroscopy (BES) diagnostic is typically employed to capture two-dimensional images of plasma density fluctuations [74, 75, 76]. It operates as an active diagnostic, requiring the simultaneous use of Neutral Beam Injection (NBI), which launches high-energy neutrals (usually atomic deuterium) into the plasma. These neutrals undergo collisional interactions with the main species ions, impurity ions, and, to a lesser extent, free electrons:



The resultant excited beam neutrals emit light which is doppler shifted according to the velocity of the incoming beam and the angle between the beamline and the viewing direction. The beam emission can therefore be localised and isolated from the background emissions. The emissions are rich in information and measurements thereof can be used to determine a number of metrics, such as the magnetic field strength and direction, and neutral particle density.

For example, the $\mathbf{v} \times \mathbf{B}$ Lorentz electric field experienced by the neutrals causes Stark splitting of the energy levels. The spacing between the energy levels can be measured to determine the magnitude of the magnetic field, given that the beam velocities are known. Similarly, the polarisation patterns of the emission lines can be compared to determine the local direction of the magnetic field.

The beam emission intensity fluctuations are generally dependent on local plasma temperature and density fluctuations. However, the plasma temperature is typically very small compared to the beam temperature and can be assumed negligible [77]. This means that the intensity fluctuations primarily depend on local density fluctuations, which is the measurement that will be exploited throughout this thesis.

In MAST-U, the BES diagnostic utilises an 8x8 array of Avalanche Photo-Diodes (APDs) to record emissions [74, 75, 76]. An optical filter is employed to isolate the Doppler-shifted emission, enhancing the contrast with background emissions and allowing for signal localisation along the line of sight of the detectors, as depicted in figure 5.3a. The 8x8 array of detectors generates two-dimensional images of density fluctuations in the radial-poloidal plane, as illustrated in figure 5.3b. The BES field of view can be adjusted radially through a motor in the internal optics, facilitating measurements from the core to the edge, as shown in figures 5.3a and 5.3b. The APD channels provide a spatial resolution of approximately 1.6 cm by 1.6 cm, enabling the resolution of the ion-scale turbulence. Additionally, the system boasts a refresh rate of 4 MHz, sufficiently fast to freeze the turbulent motions of the plasma.

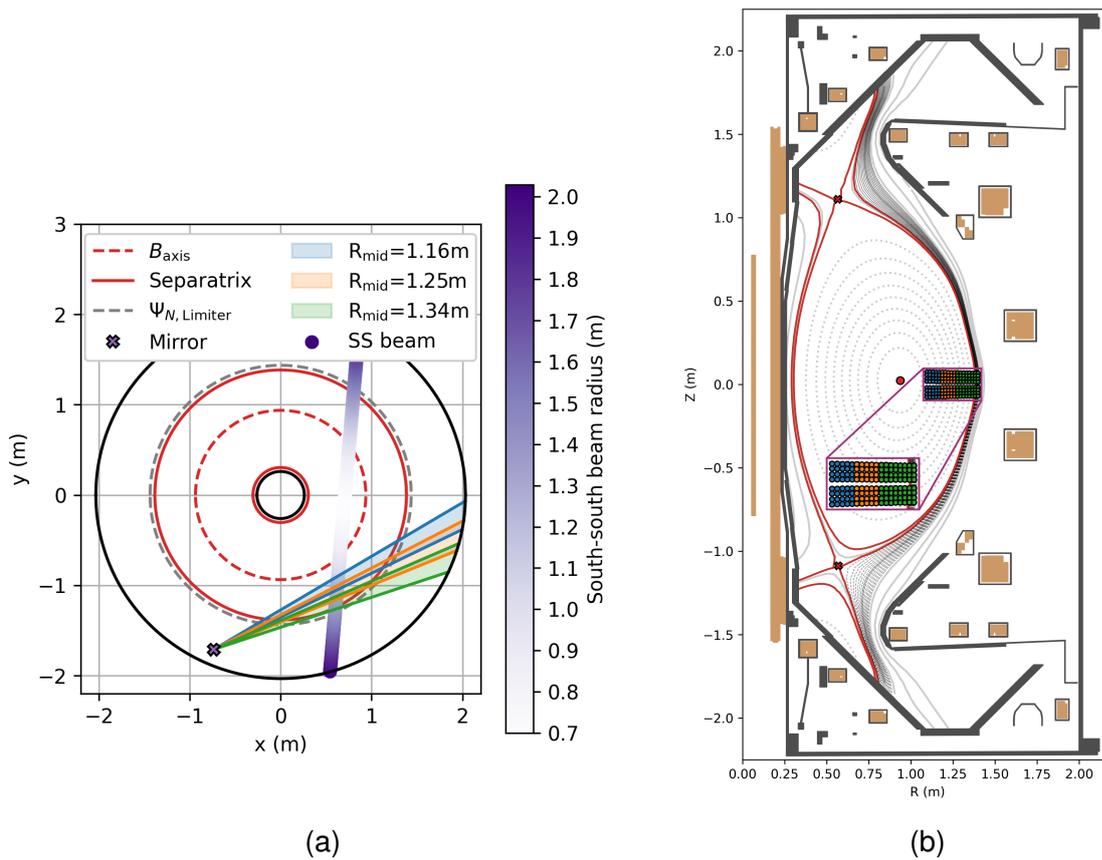


Figure 5.3: Geometry of the BES diagnostic on MAST-U. a) shows the top-down view of MAST-U, with the SS neutral beam, and BES sight-lines. b) shows the radial-poloidal view of MAST-U with the detector channel measurement locations. Three different radial viewing locations are shown in blue, orange and green. Image credit: S. Thomas, RO of MAST-U's BES.

5.3 Cross-Correlation Time-Delay Estimation

Cross-correlation time-delay estimation (CCTDE) is a technique used to estimate the velocity between two spatially separated time-signals. It is a specific implementation of the general cross-correlation-based Particle Image Velocimetry (PIV) approach [78]. This technique relies on finding the time-delay, denoted as τ_m , at which the cross-correlation between the two signals is maximised. When the maximum amplitude of the cross-correlation function (CCF) is close to one, it suggests that identical fluctuations exist in both signals and have travelled between the two spatial locations in time τ_m . Since the spatial separation between the two signals is typically known, one can infer velocity by dividing this known distance by τ_m . By repeating this process for various spatial locations, a velocity field can be constructed.

During the velocimetry testing in this thesis, the two-point CCTDE method is specifically focused on, which can be considered fundamental to most other CCTDE variations. Consequently, the results presented in section 6.2 can be extrapolated to more modern and elaborate techniques such as the line method [79] and the hybrid method [80]. The procedure for the two-point technique is outlined below:

- A time-series of spatially resolved images is loaded.
- Two spatially separated pixels are selected, and their time-signals are cross-correlated using the following function:

$$CCF(\tau) = \begin{cases} \frac{N-1}{N+\tau-1} \frac{\sum_{n=1}^{N-\tau} [f(n+\tau) - \bar{f}][g(n) - \bar{g}]}{\sqrt{\sum_{n=1}^N [f(n) - \bar{f}]^2 [g(n) - \bar{g}]^2}}, & \tau < 0 \\ \frac{N-1}{N-\tau-1} \frac{\sum_{n=1}^{N-\tau} [f(n) - \bar{f}][g(n+\tau) - \bar{g}]}{\sqrt{\sum_{n=1}^N [f(n) - \bar{f}]^2 [g(n) - \bar{g}]^2}}, & \tau \geq 0 \end{cases} \quad (5.2)$$

where N is the length of the time-series, τ is the time-delay between the two signals f and g , and \bar{f} and \bar{g} represent their means. The time-delay typically ranges from $-N$ to N frames. The pre-factor in the expressions is included to normalise the CCF based on the length of their overlap, which varies with τ .

- The time-delay τ_m at which the CCF peak occurs is determined. If the correlation peak exceeds a user-imposed threshold between 0 and 1, the velocity is inferred as $v = \Delta\ell / \tau_m$, where $\Delta\ell$ is the spatial separation between the two signals.
- This procedure is repeated for all pixel pairs in both orthogonal directions to produce two spatially resolved two-dimensional velocity fields.

An essential user-defined parameter in this context is the characteristic separation distance, denoted as $\Delta\ell$. This parameter signifies the distance between the selected pixels used for analysis. The implications of varying both the separation distance $\Delta\ell$ and the length of the time-series, denoted as N , are elaborated upon in section 6.2.

One noteworthy source of spurious velocity measurements is the ‘barber pole illusion,’ named after the apparent up/down motion of a spinning barber’s pole [74, 81], as illustrated in figure 5.4. This effect can occur especially when large, tilted density features propagate through the frame. Similar to a barber pole, tilted density features may appear to move in a different direction than their true underlying velocity. In such

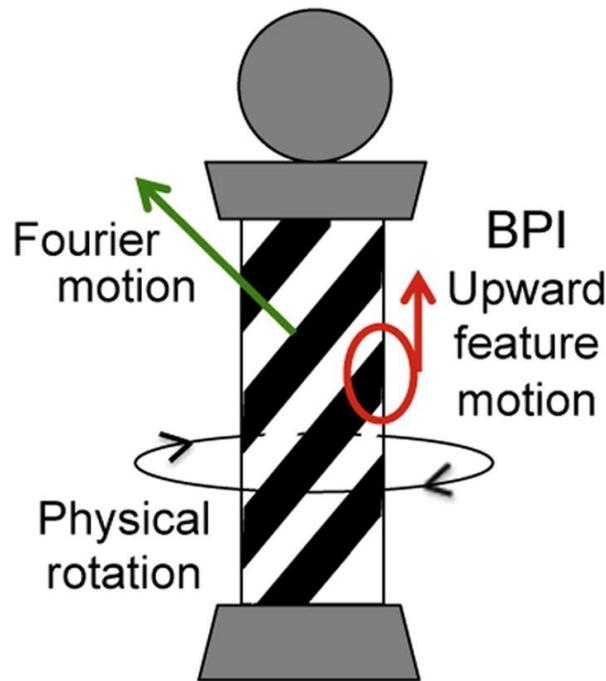


Figure 5.4: The apparent up/down motion of a barber pole as it spins. Adapted from [73].

cases, distinguishing between apparent motion and true motion can be challenging and sometimes impossible for *any* velocimetry method.

In tokamaks, the barber pole illusion can manifest in two primary ways. The first is where field-aligned density filaments are observed along the radial direction. The typical motion of the filaments is primarily in the toroidal direction but due to the tilting of the filaments, this can look like poloidal motion. For the purpose of this thesis, the previous manifestation will not be a problem because no measurements are made viewing along the radial direction. The second possibility is where density structures are tilted in the radial poloidal plane. This is frequently observed especially in high shear tokamak plasmas. When viewed along the toroidal direction, as is approximately the case with the BES diagnostic, the primarily poloidal motion can have the illusion of propagating in the radial direction. This version can result in spurious velocity measurements and the extent to which CCTDE is susceptible to the barber pole illusion is investigated in section 6.2.3.

5.4 Dynamic Time-Warping

Dynamic Time Warping (DTW) is a technique within the broader framework of the ‘optical flow’ approach in Particle Image Velocimetry (PIV), where the assumption is made that brightness is conserved along flow trajectories. An overview of optical flow is presented in the introductory book by Cai et al. [82]. The primary goal of DTW is to determine an optimised displacement field from one image to another [83, 84, 85, 86]. The spatial transformation is achieved through an iterative process that initially

calculates a displacement field based on the largest intensity features, followed by increasingly smaller corrections in subsequent iterations. The optimisation in each iteration is based on minimising the intensity difference between the two images.

DTW is a promising technique in image velocimetry due to its theoretical capability to determine accurate displacement fields, even in challenging scenarios such as turbulent flow fields. Furthermore, it offers the potential to deliver velocity fields at the diagnostic sampling rate, whereas other techniques, such as CCTDE, offer sampling rates typically at least an order of magnitude slower.

The operational details of the algorithm are thoroughly explained in G. Quenot's 1998 paper [83] and will be briefly summarised here. Two images, between which an optimised transformation is desired, are fed into the DTW algorithm. The two images are then divided into strips which overlap each other by half in the short direction, as can be seen in figure 5.5a. The number of pixels along the short direction is referred to as the strip width. Optimised displacements are determined along the long direction, or 'slicing direction,' of the strips, as depicted in figure 5.5b. The optimisation process is based on minimising the intensity difference between the two strips [83].

The division of the image into strips introduces a natural ordering of pixels, imposing a continuity constraint on the displacement search. By recombining the strips into the whole images, a full displacement field is constructed with pixel displacements *along the slicing direction* of the strips. This process is then repeated in the orthogonal direction, inferring a 2D displacement field.

In the first iteration, the algorithm infers a displacement field weighted towards intensity features with a spatial size comparable to the first strip width. Subsequent iterations progressively refine the displacement field by using smaller strip widths. In this thesis, the strip width was reduced by a factor of $\sqrt{2}$ in each iteration, following the convention used by Quenot [83]. The exact code utilised throughout this thesis is available in version 1.0.0 of S. Thomas' GitLab repository [87].

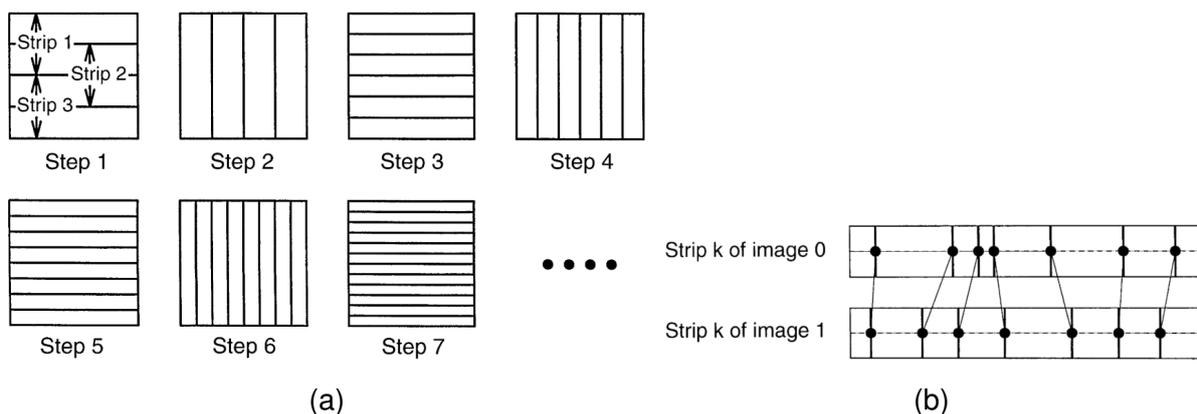


Figure 5.5: Strips used in the DTW algorithm. a) shows how the images are split into overlapping strips, with alternating slicing directions and increasingly small strip widths. b) shows displacements between strips of subsequent images. Images adapted from [83].

Chapter 6

Testing Velocimetry Methods Using Synthetic Data

This chapter focuses on the testing of velocimetry methods. Specifically, the two most commonly used methods: Cross-Correlation Time-Delay Estimation (CCTDE) and Dynamic Time-Warping (DTW) will be tested for accuracy and precision. These methods have typically been tested in previous literature for specific scenarios [88, 89, 80, 90, 91, 86]. However, as will become clear in this chapter, the results and observations under specific conditions should not be extrapolated outside the tested parameter range due to strongly nonlinear behaviours of both methods. The investigations in this chapter aimed to achieve the following:

- to quantify the accuracy and precision for both CCTDE and DTW under a broad range of conditions.
- to test both techniques to the limits of their operational parameter spaces.
- to provide examples of the reliable application of velocimetry techniques using the results in this chapter.
- to compare the techniques with each other under controlled conditions.

Synthetic data was generated to provide the controlled conditions under which both techniques could be tested, as is discussed in section 6.1.

CCTDE was tested in section 6.2, and DTW was tested in section 6.3. Note that the results sections for both CCTDE and DTW can be rather dense to read and may require several passes to fully grasp. Instead, it is recommended to start with the discussion and summary sections for an overview and return to the results sections if further details are desired. Note that digital copies of all results graphs, the vast majority of which are not shown in this thesis, are available from the authors upon reasonable request.

A cross-comparison of the techniques can be found in section 6.4, and an example workflow which applies the results in this chapter to the application of the techniques is covered in section 6.5. A general discussion and summary can be found in section 6.6.

6.1 Generating Synthetic Turbulence Fluctuation Data

Synthetic data were generated with the primary objective of accurately representing the common structures observed in turbulence diagnostics, such as gas-puff imaging [92] and beam-emission spectroscopy [75]. These diagnostic methods capture density fluctuation data in the form of a time-series of two-dimensional images. While the primary focus was on these specific diagnostic techniques, the synthesised data can effectively represent turbulent fluctuations found in any data-set.

Two distinct categories of fluctuation structures were generated. The first category consisted of ‘isolated density fluctuations,’ often observed as individual density fluctuations propagating through plasmas near marginal stability. We shall refer to this category as ‘isolated density feature’ (IDF) data. Note that, throughout this thesis, these isolated density features are occasionally denoted as ‘blobs.’ It is essential to clarify that, in the context of this chapter, the term ‘blob’ does not specifically relate to the scrape-off layer filaments observed in magnetically confined plasmas [20].

The second category comprised ‘turbulent density fields’ (TDF), representing the density structures commonly observed in fully-developed turbulent plasmas. It is noteworthy that these data types are distinguished not because they correspond to disparate turbulence regimes in real plasmas but rather due to their distinct data generation methods. In fact, one could conceptualise IDF data as the large feature-size limit of TDF data.

The synthetic data generation process was controlled by three primary user inputs:

- in-plane velocity field imposed on the density structures
- characteristic spatial scales associated with the density structures
- signal-to-noise ratio of the data

Additionally, the orientation of the fluctuations could be adjusted to investigate the impact of the barber pole illusion (as discussed in section 6.2.3). Furthermore, sheared flows could be introduced in the TDF data.

All variables in this thesis are expressed in terms of machine units. For example, velocities are given in pixels per frame, and lengthscales are represented in pixels. This approach ensures that the results can be applied to a wide range of diagnostic applications.

6.1.1 Isolated Density Features

Synthetic data were generated to simulate the behavior of isolated density features as they propagate through the field-of-view. Images of these isolated density features were created using a two-dimensional Gaussian function as described by Equation 6.1, where A and (x_0, y_0) represented the amplitude and centroid location of the blob, while (σ_x, σ_y) denoted the standard deviations of the Gaussian blob shape, offering control over the spatial size of the blobs.

$$z(x, y) = A \exp\left(-\frac{(x-x_0)^2}{2\sigma_x^2} - \frac{(y-y_0)^2}{2\sigma_y^2}\right) \quad (6.1)$$

However, true Gaussian features, as defined in Equation 6.1, extend infinitely in spatial extent, which is non-physical and unsuitable for testing purposes. To localise the density features, the following process was applied:

1. 25% of the maximum intensity was subtracted from all images
2. any resulting negative intensities were set to zero
3. the remaining intensities were scaled by a factor of 4/3 to restore the original maximum amplitude

This resulted in a modified two-dimensional Gaussian shape with a finite spatial extent that could be calculated in each direction using Equation 6.2, where $\Delta\lambda_{x,y}$ represented the full size of the blob in pixels in the x- or y-direction, often referred to as the blob y-size or x-size throughout this thesis. An example of a synthetic IDF image can be found in figure 6.1a.

$$\Delta\lambda_{x,y} = 2\sqrt{2\ln(1/0.25)}\sigma_{x,y} = 3.33\sigma_{x,y} \quad (6.2)$$

Synthetic time-series were generated by producing images of these blobs and shifting the blob's centroid according to the imposed velocity with each frame. All images had dimensions of 128 pixels by 128 pixels, and the blob's centroid location was initialised below the field-of-view in the negative y-direction. The starting location was constrained to prevent blobs from extending beyond the edges of the x-side of the field-of-view. Only one blob was present in the images at any given time, and the time-series generation ended once the blob exited the field-of-view entirely. Optionally, noise was introduced and added to each frame in the form of normally distributed, pixel-sized noise to simulate diagnostic electronic noise. The Signal-to-Noise Ratio (SNR) was defined as the ratio between the maximum blob amplitude and the root mean square (rms) of the added noise, as shown in Equation 6.3.

$$SNR_{blob} = \frac{A}{rms(noise)} \quad (6.3)$$

Synthetic data was generated with a range of blob y-sizes (1-100 pixels), blob velocities in the y-direction (0.1-60 pixels per frame), and signal-to-noise ratios (1-Inf). All data was generated with an arbitrary blob x-size of 25 pixels, chosen to be sufficiently large for registration by velocimetry techniques but small enough to minimise clipping. The x-size had negligible impact on velocimetry performance in initial tests, except when it approached the pixel size or frame size. Notably, in all tests conducted in this thesis, the velocity of the blobs was set to zero in the x-direction, and the implications of this choice are further discussed in section 6.2.4.

6.1.2 Turbulent Density Fields

Density fields observed in turbulence diagnostics often exhibit intricate structures that cannot be directly replicated using a simple function like the one demonstrated in section 6.1.1, Equation 6.1. However, when represented in wavenumber space, turbulent density fluctuations can often be approximated as broad, singly-peaked functions, such

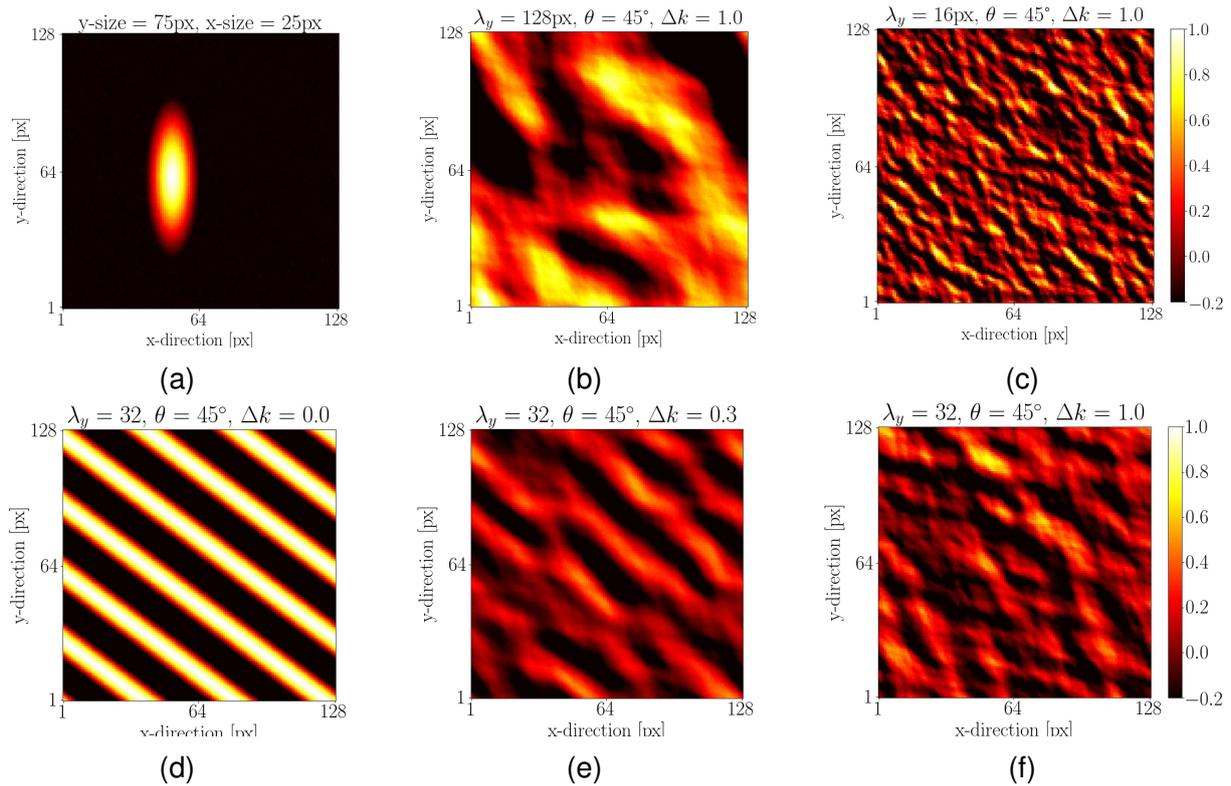


Figure 6.1: Example images of synthetic data showing: an isolated blob (a) and turbulent density fields with large (b) and small (c) spatial scales. (d-f) show the effect of varying Δk . All images were generated with $\text{SNR} = 100$ and colourbars were normalised to maximum intensity in time-series.

as Gaussians or Lorentzians [91]. To emulate this structure, distributions were initially generated in reciprocal space and then inverse Fourier transformed, resulting in real-space TDF images, as shown in figure 6.1b-6.1f.

In particular, arrays were generated following a Lorentzian distribution as expressed in Equation 6.4, where P represented the distribution amplitude, (k_{x0}, k_{y0}) indicated the centroid location of the Lorentzian, and Δk determined the function width. Randomised phase values were assigned to all elements, and the real component of the inverse Fourier transform yielded the final TDF image.

$$P(k_x, k_y) = \frac{\Delta k^2}{(k_x - k_{x0})^2 + (k_y - k_{y0})^2 + \Delta k^2} \quad (6.4)$$

The Lorentzian centroid location, (k_{x0}, k_{y0}) , could be adjusted to alter the spatial size and angle of the density features, as illustrated in figure 6.1b-6.1c. Throughout this thesis, all k_{y0} values were normalised to correspond to the number of full wavelengths in the y-direction per width of the image frame (wavelengths per 128px). For ease of interpretation, the real-space lengthscale of the fluctuations was defined as $\lambda_{y0} = 128 \text{ px}/k_{y0}$. k_{x0} is typically given as a fraction of k_{y0} , and the tilt angle of the density features, measured clockwise from horizontal, was defined as $\theta = \text{Arctan}(k_{x0}/k_{y0})$.

The Lorentzian width, Δk , could be increased from zero to enhance the ‘broken up’ appearance of the structures in the real-space image, as demonstrated in figure 6.1d-6.1f. Δk was also normalised to k_{y0} , ranging from 0 to 2 in increments of 0.3, with a typical value of 1.3 observed in density fluctuation diagnostics [91].

To generate a time-series with an imposed velocity field, the density fields were produced to be spatially larger than the final image dimensions. This allowed entire columns of pixels to be rigidly shifted up and down according to the imposed velocity field, defined by Equation 6.5, where $k_{v,y}$ represented the wavenumber of the velocity field sinusoid. All imposed velocity fields pointed purely in the y-direction, and $k_{v,y}$ varied from 1 to 8 wavelengths per 128px. Although especially computationally expensive at high velocities, this approach was chosen because it circumvented the use of periodic boundary conditions, which could introduce aliasing in the velocimetry analysis.

$$\mathbf{v}_{imposed} = v_0 \hat{\mathbf{y}} + v_1 \cos(k_{v,y} x) \hat{\mathbf{y}} \quad (6.5)$$

The time-series were generated with image dimensions of 128 px by 128 px and comprised 512 frames. Velocity fields with v_0 ranging from 0.1 to 60px/frame and v_1 from 1 to 15 px/frame were utilised. Additionally, normally distributed, pixel-sized noise was added to the images, resulting in SNR_{rms} values ranging from 1 to infinity and generated individually for each image frame, as defined in Equation 6.6. Please note that SNR_{blob} and SNR_{rms} represent significantly different definitions, often differing by more than an order of magnitude. Conversion factors are contingent solely on the blob size and can be found in figure 6.2.

$$SNR_{rms} = \frac{\text{rms}(\text{signal})}{\text{rms}(\text{noise})} \quad (6.6)$$

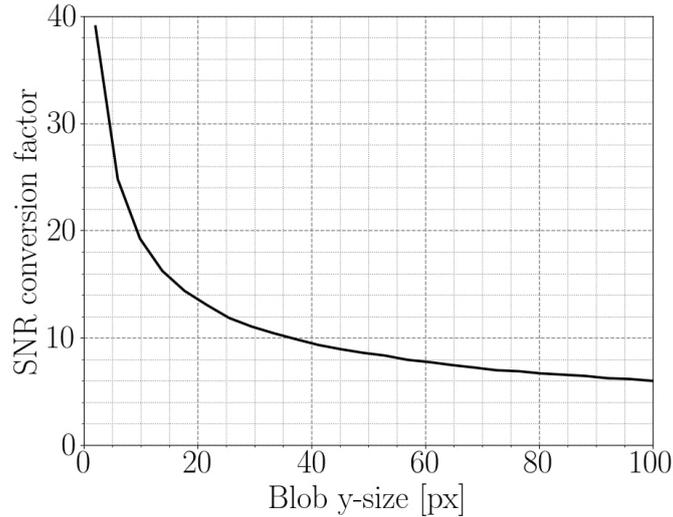


Figure 6.2: Conversion factor between SNR_{blob} and SNR_{rms} for IDF data. Factor defined as SNR_{blob}/SNR_{rms} . Blob x-size fixed at 25px.

6.2 Testing Cross-Correlation Time-Delay Estimation

This section is dedicated to testing the accuracy of the Cross-Correlation Time-Delay Estimation (CCTDE) technique. These tests encompass both IDF and TDF data, which represent turbulent data spanning from marginal stability to strongly developed turbulence. The metrics that were used to define the accuracy and precision were described in section 6.2.1.

For a smoother comprehension of this section, skipping the results sections and starting with the discussion and summary is recommended, can be found in section 6.2.4. The results sections can be rather dense and may require some time for proper interpretation. If one is interested in the more detailed results regardless, they can be found in section 6.2.2 and section 6.2.3.

6.2.1 Quantification of Technique Uncertainties

To facilitate a standardised comparison of results, the velocity fields derived through CCTDE were summarised using metrics designed to assess both overall accuracy and precision. To begin, the ‘percentage velocity-deviation field’, denoted as Δv_{meas} , was calculated. This field represents the percentage difference between the measured velocity field, v_{meas} , and the imposed velocity field, v_{imp} , and can be expressed as:

$$\Delta v_{meas} = 100\% \cdot \left(\frac{v_{meas} - v_{imp}}{v_{imp}} \right) \quad (6.7)$$

The metric for accuracy used in e.g. figure 6.3 was then taken to be the mean of Δv_{meas} . In the case of isolated density features, the average was performed over the area of the blob. In the turbulent density fluctuations case, the average was performed over the

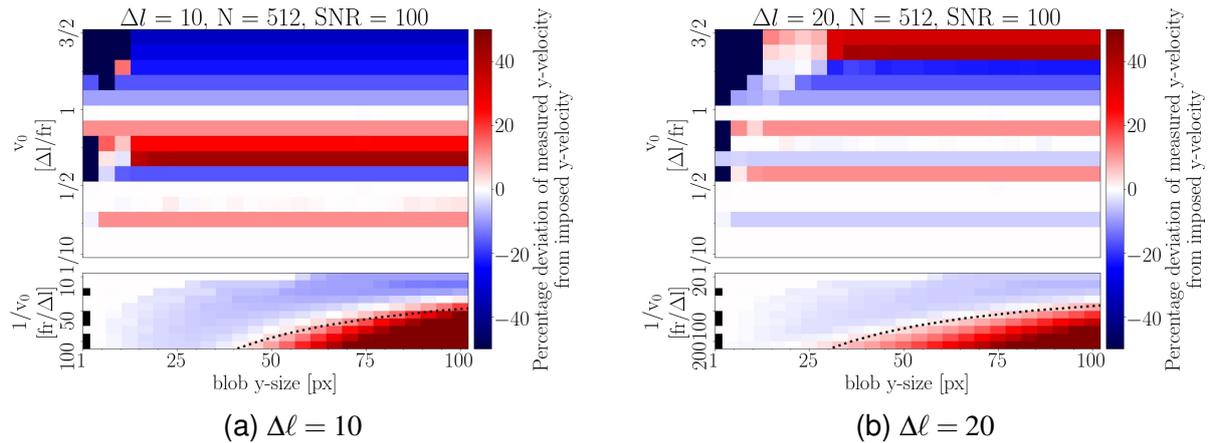


Figure 6.3: The accuracy of CCTDE with isolated density features at $\Delta\ell = 10$ (a) and $\Delta\ell = 20$ (b). Velocities are normalised to $\Delta\ell$. The bottom subplots show velocities below 1 px/frame. The dashed line denotes the predicted minimum measurable velocity that is imposed due to signal clipping by the finite length of the time-series, N . Accurate velocity fields (white) were typically measured with a standard deviation of $<1\%$. Precision was not a good proxy for accuracy.

entire field. Similarly, the precision was quantified by the root-mean-square of Δv_{meas} , with the same areas as above.

6.2.2 Results with Isolated Density Features

In this section, the performance of CCTDE in inferring velocities from data containing isolated density features is evaluated. Emphasis is placed on the assessment of velocimetry accuracy *parallel* to the imposed velocity. Investigations involving density features spanning spatial scales ranging from 1 pixel to 100 pixels were conducted. Additionally, the range of imposed velocities was varied from 1/10 pixel per frame to 30 pixels per frame.

As depicted in figure 6.3, the results reveal a negligible dependence of velocity measurement accuracy on the spatial size of the blobs. This outcome aligns with expectations for the CCTDE method, particularly in cases with minimal noise. However, an exception is noted in figure 6.3b, where, for blob y-sizes < 10 pixels and $0.5 < v_0 < 1.0 \Delta\ell/\text{frame}$, the small and fast-moving blobs can skip one of the measurement locations, rendering accurate velocity measurement impossible.

Furthermore, it was revealed in this investigation that varying the length of the time-series, denoted as N , had no discernible effect on velocimetry accuracy. Consequently, when dealing with isolated blobs, the choice of N need not be influenced by accuracy considerations.

Figure 6.3 also highlights a significant dependence of CCTDE's velocity measurement accuracy on the underlying velocity. This dependency aligns with the expectations of the two-point CCTDE method but has notably not been previously discussed in literature. Previous applications of CCTDE have implicitly assumed that this effect will be averaged out, the validity of which is further explored in section 6.2.4. No-

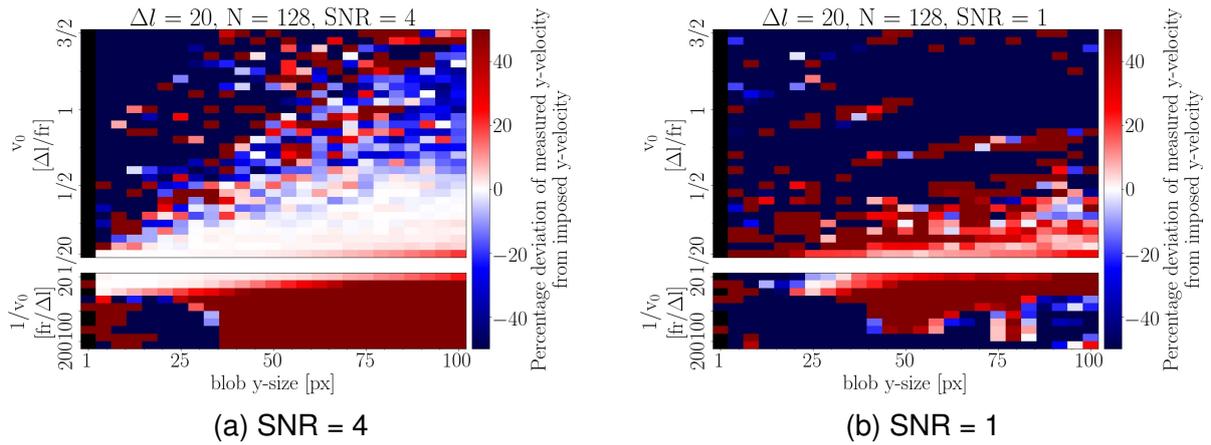


Figure 6.4: The accuracy of CCTDE with isolated density features. The SNR_{blob} was varied to assess the dependence of method accuracy on noise. Noise had little effect on accuracy for $SNR_{blob} > 10$ [not shown]. The standard deviation of the accurately measured velocity fields (in white) varied from $<20\%$ in a) to $>100\%$ in b).

tably, initial tests indicated that the strong accuracy dependence on imposed velocity magnitude could be mitigated by employing a ‘line method’, which involves performing the two-point method simultaneously for a range of $\Delta\ell$ values and selecting the measurement with the highest correlation. While this represents an advantage of the line method, spatial resolution is lost through this process. Extensive testing of the line method was beyond the scope of this thesis, but some considerations are discussed in section 6.2.4.

The investigation into the effect of varying $\Delta\ell$ revealed that increasing $\Delta\ell$ enables the accurate measurement of faster velocities, as one might expect. Additionally, the accurately measurable space expanded near slow velocities (around approximately 1 pixel/frame), as evident in figure 6.3. However, it’s important to note that maximising $\Delta\ell$ may not always be beneficial in practical applications, and this will be discussed further in section 6.2.4. Another noteworthy observation is the consistent pattern of accuracy, which remains constant with the normalised velocity, as shown in figure 6.3. Once again, this observation aligns with expectations and is elaborated on in section 6.2.4.

The results also reveal an often-overlooked fact: the finite length of the time-series imposes a limit on the minimum velocity that can be measured. This threshold is given analytically by:

$$v_{min} = \frac{\lambda_{y0} + \Delta\ell}{N} \quad (6.8)$$

where N represents the length of the time-series, λ_{y0} is the blob y-size, and v_{min} is the minimum velocity in the y-direction that can be measured. This expression was over-plotted in figure 6.3 and was found to accurately predict the minimum measurable velocity in all cases. Quantitatively, the threshold was observed to correspond approximately to the 25% velocity deviation mark. An additional observation was that the precision of velocity inference decreased with decreasing N , but this effect was not quantified.

The impact of increasing noise levels on the performance of CCTDE was also examined, as summarised in figure 6.4. Notably, the accuracy of CCTDE remained largely unaffected by noise at $SNR_{blob} > 10$, displaying accuracies similar to figure 6.3b. However, as SNR_{blob} was decreased from 10 to 4, significant signs of accuracy degradation due to noise emerged, as illustrated in figure 6.4a. This included increasingly unreliable accuracies and a shrinking parameter space of reliably accurate measurements. A further decrease in SNR_{blob} from 4 to 1 led to a steady reduction in the accurately measurable parameter space and increased standard deviation in the measured velocity fields. By $SNR_{blob} = 1$, no accurate velocities could be inferred, as shown in figure 6.4b. *Importantly, it was noted that the standard deviation of the inferred velocity was not necessarily a reliable predictor of accuracy.*

Subsequently, the ability to recover accurate CCTDE measurements from noisy data using frequency filtering was assessed. The synthetic data used for figure 6.4 underwent a two-way low-pass Butterworth filter prior to CCTDE analysis, with a cut-off frequency set at 0.3 times the Nyquist frequency to effectively filter out noise while preserving the signal. Initial tests indicated marginal improvements in accuracy; for instance, the $SNR_{blob} = 4$ case after filtering showed accuracies comparable to the unfiltered $SNR_{blob} = 100$ case seen in figure 6.3b. However, for data with $SNR_{blob} = 3$, only marginal recovery of accuracy was observed. Whether further optimisation of noise filtering techniques could fully recover accuracy for $SNR_{blob} = 3$ (and lower) appears unlikely based on the non-linearity of degradation observed in these tests.

Lastly, it was consistently observed in all tests within this subsection that the blob velocity orthogonal to the direction of blob propagation was accurately measured to be zero on average for imposed velocities greater than 1 pixel/frame. However, standard deviations of up to 20% relative to the measured parallel velocity were observed in this range. In contrast, orthogonal velocity inferences at $v_0 < 1$ pixel/frame were found to be susceptible to high statistical variation of the measurements. The standard deviation often exceeded that of the parallel velocity measurements, significantly reducing the reliability of the overall measurements in this region of velocity space.

6.2.3 Results with Turbulent Density Fields

This section assesses the performance of CCTDE when presented with turbulent density field data. The characteristic spatial scale, λ_{y0} , was varied from 1 pixels to 100 pixels, and the imposed velocity was varied from 0.1 pixels/frame to 30 pixels/frame. The angle of the density features was given a range from 0° to 75° clockwise from horizontal. Δk was given a range from 0 k_{y0} to 2 k_{y0} in 0.3 k_{y0} increments, while $\Delta \ell$ was varied from 1 pixels to 20 pixels.

As expected, variations in $\Delta \ell$, SNR_{rms} , and N had effects comparable to what was found in section 6.2.2. In summary, decreasing N was known to reduce precision and impose a minimum measurable velocity. The corresponding expression defined in equation 6.8 held for the turbulent density data. SNR_{rms} was found to have a negligible effect on accuracy for $SNR_{rms} > 1$, although reductions in precision were observed.

It was found that the barber pole illusion had a negligible impact on the measurement accuracy *parallel* to the flow direction. Conversely, spurious *perpendicular* velocity measurements were widely observed when $\Delta k \leq 0.3k_{y0}$. When Δk is this low, density

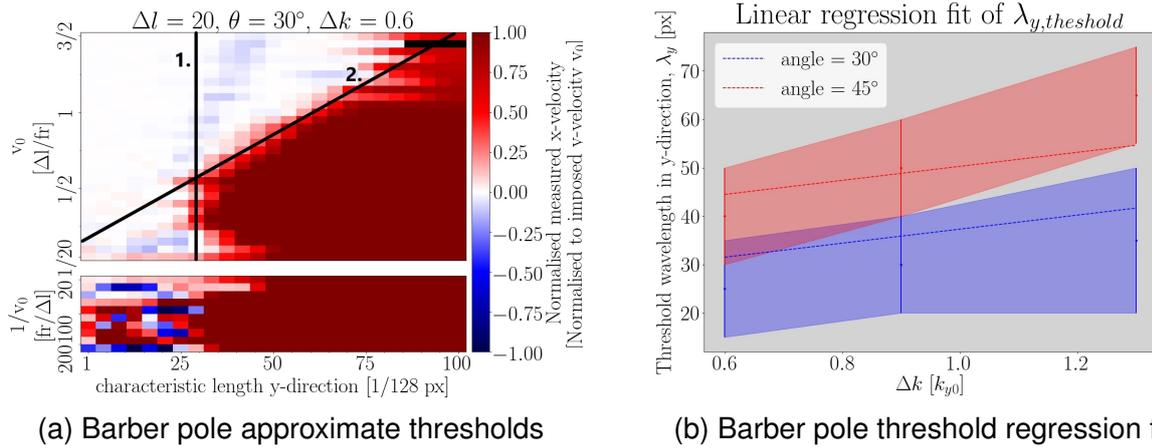


Figure 6.5: Defining a threshold at which barber poling becomes significant; a) thresholds drawn approximately at 25% deviation from v_0 ; b) examples from the multivariate linear regression aiming to predict threshold line 1. Data-points denote empirical measurements with estimated error bars, and the dashed lines denote the regression equation fits. R-squared was 0.9.

features often have a spatial extent comparable to the image field-of-view. Furthermore, negative velocities were observed, which were due to a combination of barber poling and aliasing. A subtle example of aliasing effects can be seen in the light-blue regions at $v_0 > 1$ in figure 6.5a. These spurious velocities could not consistently be predicted, and experimental CCTDE analysis in the region of $\Delta k \leq 0.3k_{y0}$ was considered to be unreliable without further in-depth testing.

In the region of $\Delta k \geq 0.6k_{y0}$, it was observed that a well-defined threshold emerges that predicts when the barber pole illusion becomes significant. An illustrative example is provided in figure 6.5a, where two distinct lines are approximately drawn at the 25% velocity deviation mark. Threshold 1 in figure 6.5a was found to depend solely on the spatial parameters: Δk , Δl , θ , and λ_{y0} . However, no analytical expression for the threshold could be determined from first principles. Instead, an empirical approach was employed, utilising a multivariate linear regression to define the threshold.

To estimate the characteristic length, λ_{y0} , at which threshold line 1 occurs, angles of 30° and 45° , Δk ranging from 0.6 to 1.3 k_{y0} , and Δl ranging from 5 pixels to 20 pixels were used. The threshold was determined as follows: the λ_{y0} values at which the percentage velocity deviation reached 25% were traced across v_0 . A five-point moving average was applied, resulting in a trace with a bi-linear form (approximately corresponding to lines one and two in figure 6.5a). The transition between the two linear regions did not always exhibit a sharp gradient change, and an approximate gradient transition region was defined visually. The λ_{y0} threshold was defined as the lowest λ_{y0} within the gradient transition region (roughly equivalent to the intersection of lines 1 and 2 in figure 6.5a). The error margin of the λ_{y0} threshold was defined to be equal to half the width of the gradient transition region.

Subsequently, a weighted multivariate linear regression was performed using the determined λ_{y0} thresholds as the dependent variable and Δk , Δl , and $Tan(\theta)$ as the independent variables, all of which exhibited approximate linear dependence. The fol-

lowing expression was constructed from the regression to approximate the λ_{y0} threshold:

$$\lambda_{y0,threshold} = 30 \cdot \tan(\theta) + 15 \cdot \Delta k + 2.2 \cdot \Delta \ell - 40 \quad (6.9)$$

For this regression involving 24 data points, all coefficients were found with a standard error of approximately 15%, and the R^2 was 0.9. Δk was normalised to k_{y0} , and $\Delta \ell$ was given in pixels. A good prediction (within errors) was also found when extrapolated to parameters $\theta = 15^\circ - 60^\circ$, $\Delta k = 0.6 - 2.0 k_{y0}$, and $\Delta \ell = 5-20$ pixels.

With increasing λ_{y0} , the perpendicular velocity increases and saturates at a well-defined value that can be calculated using the following expression:

$$v_{\perp} = v_{\parallel} / \tan(\theta) \quad (6.10)$$

The v_{\perp} calculated in Equation 6.10 strictly applies when there is no underlying velocity in the x-direction. It has not been explicitly tested whether the barber-pole-induced illusory-x-velocity is non-linearly affected by true underlying x-velocities.

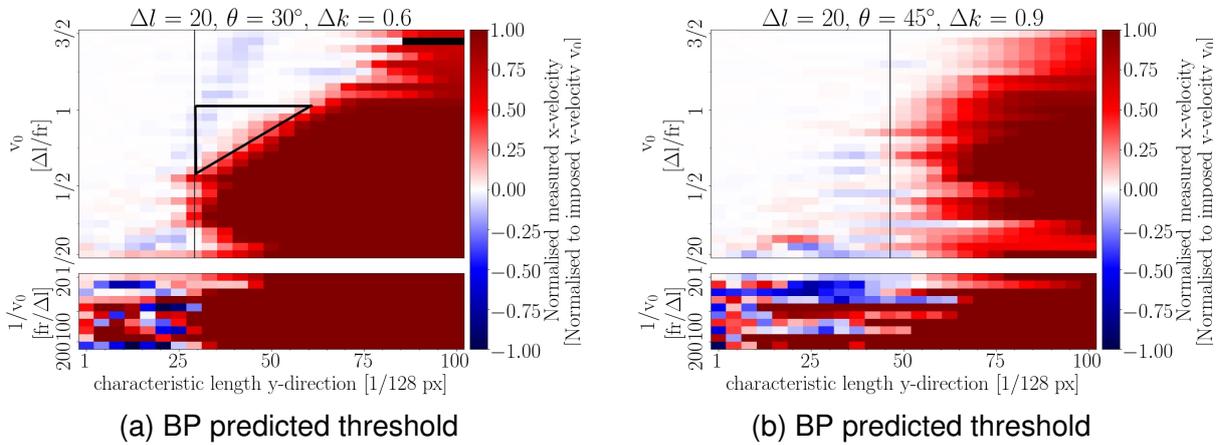


Figure 6.6: Two example plots showing the λ_{y0} threshold as predicted by the regression equation. Fit within error margins was found for most ($\sim 90\%$) cases tested. Triangular region in a) shows accurate area that can be gained by quantifying threshold 2. This area does not exist in b).

Threshold 2, as shown in figure 6.5a, was not investigated in detail because its quantification would have little impact. Measuring threshold 2 would not significantly expand the area in parameter space that can be measured accurately. The potential additional area is exemplified by the small triangular region in figure 6.6a. Conversely, no added benefit would be seen in figure 6.6b.

6.2.4 Discussion and Summary of CCTDE Uncertainties

Accuracy dependence on underlying velocity: The measurement accuracy of the two-point CCTDE method exhibits a strong dependence on the underlying velocity of the fluctuations, as illustrated in figure 6.3. This pattern aligns with the theoretical expectations of the method, although this has not been covered in previous literature. Only velocities that are exact factors of $\Delta \ell$ can be accurately measured. Velocities

falling between these factors are approximated to the nearest factor of $\Delta\ell$. However, this naturally poses a challenge as the underlying velocity is typically unknown before CCTDE analysis. Additionally, and crucially, it's worth highlighting that *the precision of the inferred velocity fields does not serve as a reliable predictor of measurement accuracy*. Therefore, it is important to use two-point CCTDE in conjunction with other velocity estimation techniques to ensure accurate results.

A lack of precision may be beneficial: In initial tests, it was observed that the line-CCTDE method's accuracy did not display the same velocity dependence as two-point CCTDE. It may not always be necessary, however, to employ this more computationally expensive method. This is due to a most welcome scenario where experimental noise introduces an unexpected but beneficial effect. Turbulent velocity fluctuations and optical jitter can cause fluctuations in the apparent velocity, leading to the inferred velocity 'flip-flopping' between factors of $\Delta\ell$ throughout the velocity field. This results in an average velocity field with improved accuracy but at the cost of reduced spatio-temporal resolution. This effect is particularly pronounced for underlying velocities that lie mid-way between two factors of $\Delta\ell$ and the effect can be enhanced by decreasing N to reduce measurement precision. While this beneficial averaging effect may be common, it should not currently be relied upon without further investigation.

Accuracy dependence on SNR: The impact of noise on CCTDE performance was found to be negligible for $\text{SNR} > 1$, in both IDF and TDF data. These findings appear to contradict previous literature, such as [89], which suggests a stronger dependence on SNR and defines a measurement limit around $\text{SNR} = 10$. However, due to the lack of detailed data generation descriptions and ambiguity of SNR definitions, meaningful comparisons with this prior work could not be made.

Choosing appropriate $\Delta\ell$: Increasing $\Delta\ell$ leads to an increase in the maximum accurately measurable velocity (up to $1 \Delta\ell/\text{frame}$) and expands the accurately measurable velocity space at low velocities near 1 pixel/frame (see figure 6.3). However, maximising $\Delta\ell$ may not be desirable when the decorrelation timescale is significant compared to the underlying velocity. In such cases, reducing $\Delta\ell$ can limit the impact of decorrelation. Importantly, $\Delta\ell$ can be independently defined in different orthogonal directions without significantly complicating the analysis because velocity inferences in different directions are independent of each other. It is re-iterated that, in this study, $\Delta\ell$ was kept symmetric in all tests.

The effect of reducing the length of the time-series, N : Reducing N can offer various advantages, including: i) lowering computational costs, ii) potentially beneficial averaging through reduced precision, and iii) increasing the temporal frequency of velocity inference. The reduction in precision with decreasing N in CCTDE has been quantified in existing literature [90]. One consideration that must be made is ensuring that the expected velocities do not fall below the minimum velocity limit defined in Equation 6.8 for the chosen N . This equation can be applied to both IDF and TDF data, with some extra care required to determine a representative λ_{y0} for TDF data. It was observed in this study that CCTDE can be effectively operated with N as low as 32 in all cases tested with little effect on measurement accuracy, whereas previous literature typically used $N \geq 256$ [89].

A simplified method to avoid barber pole illusions: The prevalence of spurious velocity measurements due to the barber pole illusion was quantified in section 6.2.3.

Due to this research, it can be predicted *before* CCTDE analysis if the barber pole illusion will significantly impact velocimetry. This prediction requires an estimation of the spatial parameters associated with fluctuations in the data, namely Δk and λ_{y0} . Equation 6.9 can be used to avoid spurious velocities by increasing $\Delta \ell$. An example of this simple check is provided in section 6.5. These findings differ from previous literature that focused on correcting spurious velocities post-velocimetry [88]. This previous correction process is typically laborious and requires assumptions to be made about the shape of underlying structures.

Precision in the low-velocity regime: For slow velocities, $v_0 < 1$ pixels/frame, it was observed that the statistical variation of the velocity measurements becomes significant. The standard deviation becomes comparable to the average inferred velocity. Therefore, extra care is recommended to ensure statistical convergence in this low-velocity regime.

6.3 Testing Dynamic Time-Warping

This section is dedicated to testing the accuracy of the Dynamic Time Warping (DTW) technique. These tests encompass both IDF and TDF data, which represent turbulent data spanning from marginal stability to strongly developed turbulence. The metrics used to define accuracy and precision were defined in section 6.2.1.

For a smoother comprehension of this section, skipping the results sections is recommended and instead one could start with the discussion and summary section 6.3.3. The results sections can be rather dense and may require some time for proper interpretation. If one is interested in the more detailed results regardless, they can be found in section 6.3.1 and section 6.3.2.

6.3.1 Results - Isolated Density Features

This section evaluates the performance of Dynamic Time Warping (DTW) when confronted with data featuring isolated density features. The accuracy and precision of the inferred velocity fields were assessed both parallel and perpendicular to the imposed velocity. The fluctuation spatial scales ranged from 1px to 100px, and the imposed velocities varied from 0.1px/frame to 60px/frame. The signal-to-noise ratio (SNR_{blob}) was varied from 1 to infinity. DTW was typically operated with seven iterations using the following strip widths: [32, 22, 16, 12, 8, 6, 4] px, unless specified otherwise. The initial slicing direction was set parallel to the known imposed velocity, a crucial step discussed in more detail in section 6.3.3.

As depicted in Fig. 6.7a, DTW exhibits highly accurate velocity measurements across most of the parameter space at $SNR_{blob} = 100$. The primary exception occurs at $v_0 \leq 2px$ and blob sizes $> 20px$, where there is consistent overestimation of the velocity. Aside from this small overestimation region, the accuracy of DTW is independent of the spatial size of the blobs or the underlying velocity. All perpendicular velocity measurements accurately yielded zero.

As SNR_{blob} decreases, the accuracy of DTW rapidly degrades. This is evident in Fig. 6.7b, where DTW fails to recover any velocities accurately at $SNR_{blob} = 2$. Further

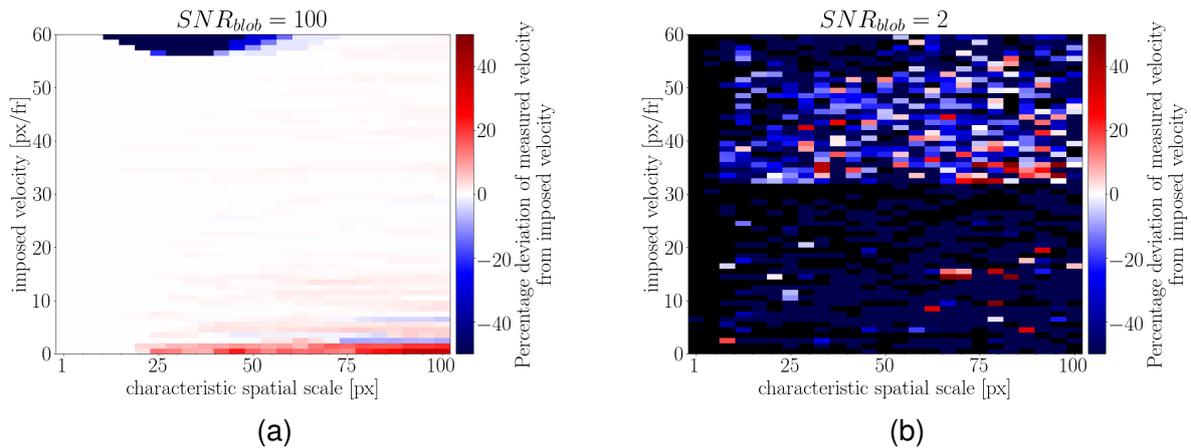


Figure 6.7: The effect of pixel-size noise on DTW velocity inference accuracy. $SNR = 100$ in (a) and $SNR = 2$ in (b). Input data was unsmoothed and 7 default DTW iterations were used.

investigation revealed that the displacement field was often accurate after the first iteration of DTW, and subsequent iterations would distort the originally accurate field. This effect stems from an operational quirk of DTW in situations where there are areas with no signal and only noise, as discussed in section 6.3.3. Follow-up tests with a single DTW iteration at strip width = 32px and varying SNR_{blob} from 100 to 1 were performed. Highly improved accuracies were found, typically with less than 10% velocity deviation, as shown in Fig. 6.8a, although full recovery of accuracy could not be achieved using this approach.

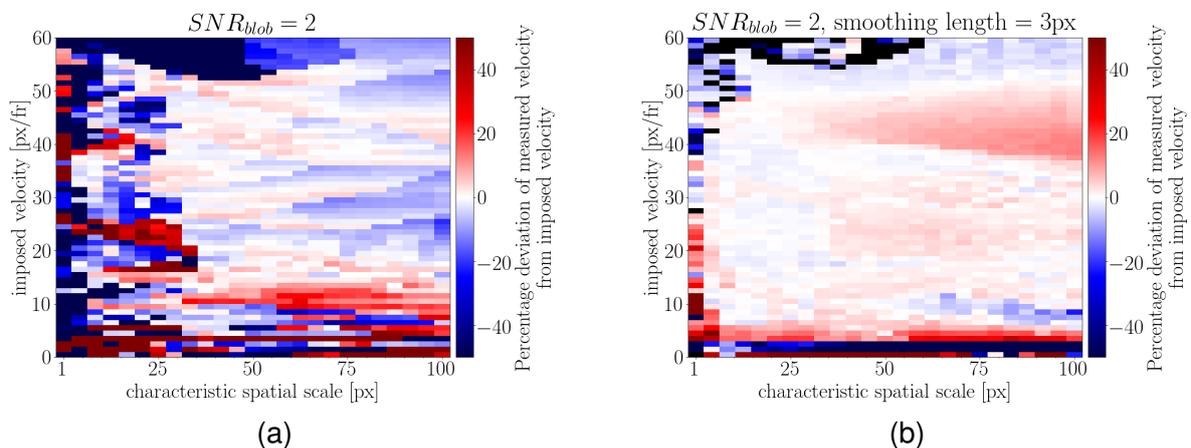


Figure 6.8: Two approaches for recovering from noise in IDF data. a) shows DTW performance with only one iteration. b) shows DTW performance when the data was pre-smoothed with a 3 px gaussian kernel.

To explore the effectiveness of spatially smoothing out the noise before DTW analysis, the input data images were passed through nested 1D convolution filters with a Gaussian kernel. The optimum Gaussian Full Width at Half Maximum (FWHM) was found to lie around 1-3px, where the noise was effectively smoothed out while leaving

the underlying signal relatively unaffected. The smoothed data were then analysed using DTW. Figure 6.8b illustrates that smoothing is highly effective across the majority of parameter space in recovering the accuracy of the DTW algorithm, allowing most velocities to be measured within 5% accuracy. However, at low velocities ($v_0 < 5px/fr$), accurate velocity inference could not be achieved using this approach.

6.3.2 Results - Turbulent Density Fields

This section evaluates the performance of DTW when confronted with data featuring turbulent density fields. The operational parameters for DTW remained identical to those detailed in section 6.3.1, unless otherwise specified. The characteristic spatial scale, λ_{y0} , ranged from 1 px to 100 px, and the imposed velocity in the y-direction varied from 1 px/frame to 60 px/frame. All tests in this section utilised TDF data with default parameters: $\Delta k = 1.3k_{y0}$, $\theta = 45^\circ$, unless stated otherwise.

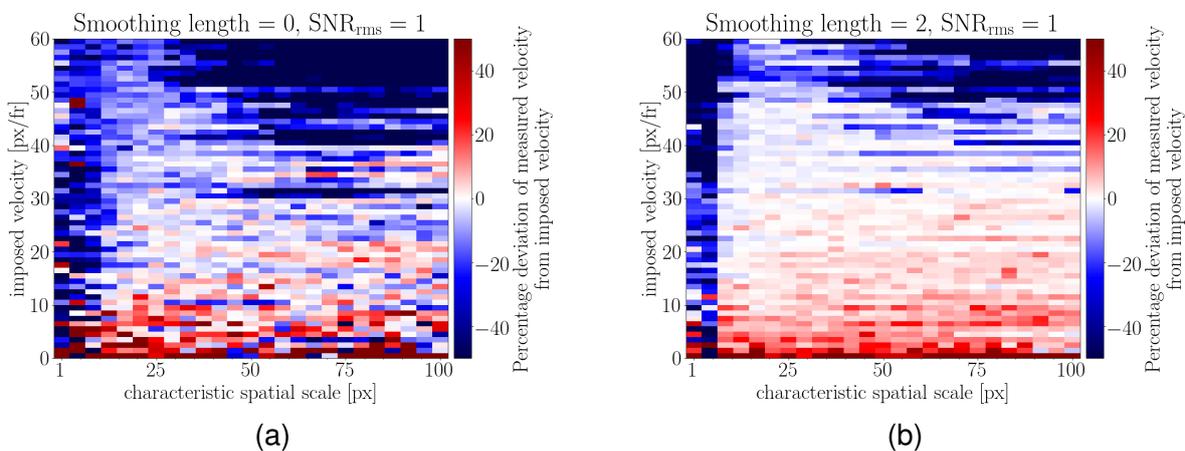


Figure 6.9: The effect of smoothing TDF input data on DTW performance with noise. No smoothing in (a) and with smoothing in (b).

Initial tests at $SNR_{rms} = 100$ demonstrated DTW's ability to accurately infer velocities across the entire spectrum of spatial sizes and imposed velocities, comparable to the IDF results shown in Fig. 6.7a. As SNR_{rms} was incrementally decreased to 1, the accuracy and precision generally diminished, although tests at $SNR_{rms} = 1$ still revealed accurately measurable regions in parameter space, as depicted in figure 6.9a. The standard deviation of the velocity fields in Fig. 6.9 was typically around 10%. Unlike the issues observed in section 6.3.1 when using multiple iterations, such problems were not reproduced with TDF data. Smoothing, applied using the same approach as described in section 6.3.1, substantially improved measurement precision. At $SNR_{rms} = 1$ and a 2px smoothing length, the standard deviation was reduced to within 5% of the average. However, a consistent overestimation of velocities below 10 px/frame was revealed, as shown in Fig. 6.9b.

The impact of reducing the number of spatial channels available in the input data was investigated. The original 128 by 128 channel images were downresolved onto new N_{ch} by N_{ch} channel images. This was done by splitting the original images into arrays of size Δd by Δd , where $\Delta d = 128/N_{ch}$. Thus, each channel in the new N_{ch} by N_{ch}

image corresponded to one Δd by Δd array in the original image. The intensities of the channels in the N_{ch} by N_{ch} images were calculated by taking weighted averages of the respective Δd by Δd arrays. The weighted average was calculated using a 2D Gaussian kernel with a FWHM equal to Δd , and were centred on the center of the Δd by Δd array. Exploratory tests showed negligible dependence on the shape of the averaging kernel. Nevertheless, the Gaussian shape was chosen to approximate the increased sensitivity in the center of the channels found in diagnostics such as BES [93]. Synthetic data with $N_{ch} = [4, 8, 16, 32, 64]$ was generated and pixel-size noise was re-introduced with $SNR_{rms} = [1, 2, 3, 4, 10, 100]$. Although DTW could technically run using data with down to 4 by 4 spatial channels, the accuracy of the velocity inferences (which was not shown here) was found to rapidly degrade with decreasing N_{ch} and SNR_{rms} . This degradation, although not quantified here, was sufficiently severe to recommend pre-processing of the low-channel data.

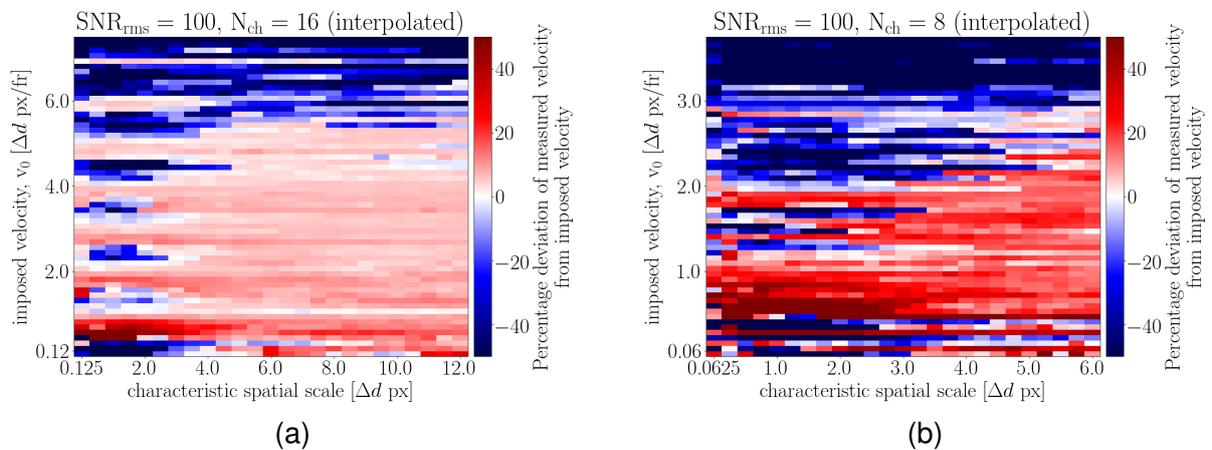


Figure 6.10: The effect of reducing the number of available channels on DTW performance. Images contained N_{ch} by N_{ch} channels. 16 by 16 channels in (a) and 8 by 8 channels in (b). All images were re-interpolated onto 128 by 128 channels before DTW analysis.

The downresolved data was re-interpolated onto the original 128 by 128 channel grid using bivariate cubic spline interpolation before DTW analysis [91]. At $SNR_{rms} = 100$, the accuracy was only marginally affected in the range $N_{ch} = 128 - 16$. Measured velocities typically deviated less than 10% from the imposed values. Figure 6.10 illustrates that accuracy deteriorated strongly going from $< 10\%$ deviation at $N_{ch} = 16$ to $\geq 20\%$ deviation at $N_{ch} = 8$. Additionally, there was a marked decrease in the reliable parameter space shown in figure 6.10. Decreasing N_{ch} also exhibited a strong decrease in the precision of DTW velocity fields, which could be exacerbated by simultaneously decreasing SNR_{rms} . The combined result of these effects is evident in Fig 6.11a, showing unreliable velocity measurements across all parameter space at $N_{ch} = 8$ and $SNR_{rms} = 2$, values which are not uncommon in BES diagnostic measurements. However, averaging 32 measurements from successive pairs of frames improved precision to a standard deviation of $< 10\%$, recovering considerable regions of accurate measurements, as shown in figure 6.11b. Averaging over multiple frames, while improving precision, reduced DTW temporal resolution, becoming comparable to typical

CCTDE temporal resolution.

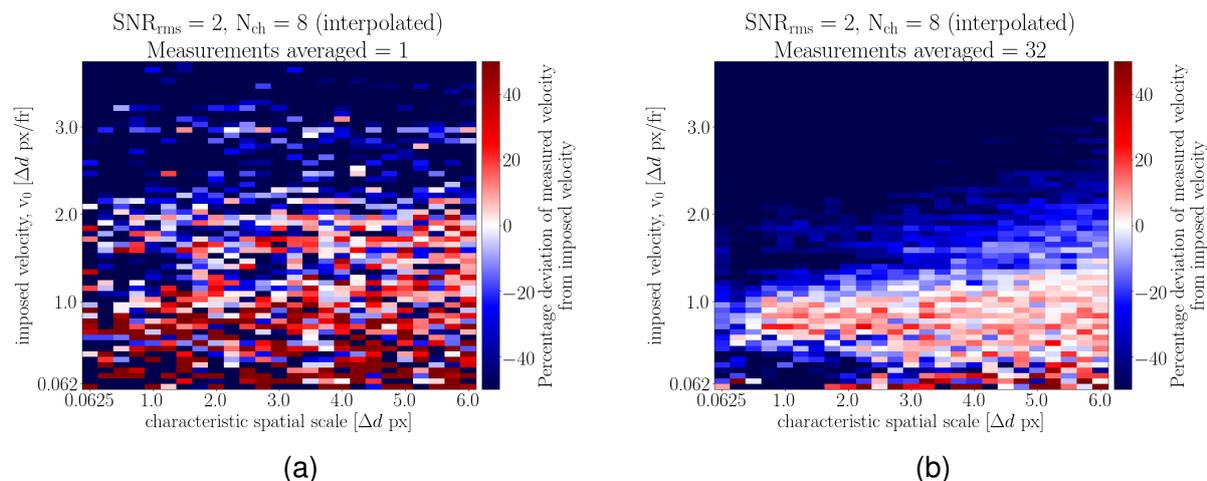


Figure 6.11: Assessing DTW performance at $N_{ch} = 8$ under noisy conditions. Reduced precision shown in fig a) is shown to converge upon averaging multiple measurements in b).

For 128 by 128 channel images with $\Delta k \geq 0.3k_{y0}$, the inferred velocity was zero in the direction perpendicular to the imposed velocity. At $\Delta k < 0.3k_{y0}$, spurious velocities were found due to the barber pole illusion.

Exploratory tests assessed the impact of sheared velocity fields on DTW performance, utilising TDF data with a characteristic fluctuation spatial scale of 7px. Velocity field parameters v_1 and $k_{v,y}$ (see equation 6.5) were varied from 1 to 15 px/frame and from 1 to 8 wavelengths per 128px, respectively. In literature, the same $k_{v,y}$ range had been investigated, but v_1 was not varied in those tests [91]. Independently increasing the amplitude and the wavenumber showed detrimental effects on accuracy and precision in both cases. Although thresholds were observed for the onset of these effects depending on v_1 and $k_{v,y}$, they were not quantified. It was found that the accuracy and precision remained constant when v_1 and $k_{v,y}$ were varied whilst keeping the maximum shear amplitude, $\partial_x(v_y)|_{max} = v_1 k_{v,y}$, constant. Conversely, increasing the maximum shear resulted in worse performance of the method, although this was not quantified. These results suggest that the maximum shear is a governing parameter for the performance of DTW.

6.3.3 Discussion and Summary of DTW Uncertainties

Accuracy Dependency on SNR. The accuracy of DTW decreases with increasing noise levels, as shown in figure 6.7. This was observed in 128px by 128px images for both IDF and TDF data. This loss of accuracy could largely be recovered through spatial smoothing of the images before velocimetry, as seen in figure 6.9. One caveat is observed in figure 6.9b, where a consistent overestimation of up to 20% is observed at low velocities, $v_0 < 10$ px/frame. Additionally, the accuracy at characteristic spatial scales below 5px could not be recovered using spatial smoothing. The precision of velocity fields also varied with signal-to-noise levels. Velocity fields with 1% standard

deviation were observed at $SNR_{rms} = 100$. The standard deviation increased to 10% at $SNR_{rms} = 1$. In conclusion, the reduced precision due to noise could be mostly, but not completely, counteracted by spatial smoothing of the images before DTW analysis.

An operational quirk with IDF data. This was observed in section 6.3.1 and attributed to regions in the images where no signal, only noise, was present. With a lack of signal outside of the blob area, DTW would transform the images here according to the noise. This distortion outside the blob area would affect the displacement field inside the blob area due to continuity constraints in the DTW method. Caution is advised for DTW velocimetry of IDF-like data. Two approaches to minimising detrimental effects were demonstrated in figure 6.8.

Performance Dependency on the Number of Spatial Channels. Without interpolation, it is important to note that DTW accuracy decreases rapidly below $N_{ch} = 32$, even when noise levels are negligible. The beneficial effect of interpolating the images onto a higher grid prior to DTW is ubiquitous, *but at best marginal*, for the accuracy of the velocimetry. Despite this improvement, DTW still increasingly and consistently overestimated velocities upon decreasing N_{ch} , which can be seen in figure 6.10. Spatial scales below approximately three times the channel size resulted in unreliable inferences, especially if v_0 was not an integer. The introduction of noise was investigated at $N_{ch} = 8$ and decreasing SNR_{rms} rapidly decreased the precision of DTW. At $SNR_{rms} = 2$, this effect could be considered fatal, as can be seen in figure 6.11a. Upon averaging multiple consecutive inferences, the precision improved and accurate velocity inferences were revealed in figure 6.11b. Specifically, a consistent overestimation of approximately 10% was found at $v_0 = 1$ px/frame. This shows that reliable velocity inferences can be obtained well below the measurement limit of $SNR = 10$ defined in previous literature [91].

Sheared Flow Fields. Exploratory tests investigated the ability of DTW to infer sheared velocity fields by varying the velocity field parameters v_1 and $k_{v,y}$ seen in equation 6.5. It was found that both the accuracy and precision were reduced by increasing *the maximum shear* past a threshold. Interestingly, the accuracy and precision did not vary locally with local shear amplitude. Instead the global maximum shear was found to be the parameter which governs the accuracy and precision. This is a result which generalises previous tests by Kriete et. al. [91], who found that shear flow wavenumber is *an* important parameter which affects DTW accuracy and precision. The reduction in accuracy presented itself as a reduction in amplitude of the measured velocity sinusoid, although shape and wavelength were conserved, which was consistent with previous findings [91].

The Limited Impact of the Barber Pole Illusion. All measurements in the direction perpendicular to the imposed velocity were accurately inferred to be zero at $\Delta k \geq 0.3$. Spurious perpendicular velocities due to the barber pole illusion occurred at $\Delta k < 0.3$. The onset of spurious velocities is hypothesised when the density features are tilted, extend past the diagnostic field of view, and if spatial variations within the density features are negligible compared to the noise levels.

The Main Complication of Using DTW. In all tests, the first slicing direction was set parallel to the direction of the imposed velocity. *Choosing the perpendicular direction instead was found to be catastrophic for DTW velocity inference.* The threshold at which misalignment between the initial slicing direction and flow direction becomes

an issue was not tested, although increasing the first strip width was hypothesised to accommodate larger misalignments. Nevertheless, it is critical that the direction of the velocity field be estimated *prior* to DTW analysis in a real experimental setting. This could be achieved by cross-comparison with other velocimetry methods like CCTDE. Additionally, sub-pixel velocities could not be accurately inferred by DTW. This shortcoming can be circumnavigated by increasing the temporal spacing between images that are analysed, although extra care in selecting the temporal spacing is advised.

6.4 Comparison Between CCTDE and DTW Performance

A fundamental distinction between DTW and CCTDE lies in their approach to velocity field inference. DTW relies on spatial information in images to identify velocity fields, while CCTDE leverages variations in temporal information within time-series data. This distinction implies that DTW theoretically has the capacity to infer velocity fields at a frequency equal to the frame rate of the diagnostic, whereas CCTDE typically operates at a measurement frequency at least an order of magnitude slower. This inherent drawback of CCTDE contrasts with its reduced reliance on spatial information, potentially resulting in higher accuracy and precision than DTW when the number of spatial channels is reduced.

The CCTDE inference frequency has historically been at least two orders of magnitude slower than the diagnostic frame-rate [89]. In these investigations, it was found that CCTDE can be operated with $N = 32$ in most cases, which represents at least an order of magnitude improvement compared to most previous literature. Additionally, DTW was found to require averaging of multiple subsequent inferences with noisy, low-spatial-channel data (e.g., $SNR_{rms} = 2$, $N_{ch} = 8$). This reduced the effective velocity-inference frequency of DTW, and in some cases, the inference frequency could be comparable between the two techniques.

Both techniques were susceptible to inferring spurious velocities due to the barber pole illusion. Neither technique reliably inferred accurate velocity fields at $\Delta k < 0.3$. At $\Delta k \geq 0.3$, DTW typically inferred accurate velocity fields, while CCTDE could also do so, assuming it passed a simple check using equation 6.9.

Shear in the velocity fields negatively impacted DTW accuracy once a threshold in the maximum local shear was surpassed. Shear cannot affect two-point CCTDE because the velocity inferences at each spatial location are independent of each other. For the same reason, CCTDE accuracy is unaffected by reducing the number of spatial channels, N_{ch} . On the other hand, reducing N_{ch} had detrimental effects on DTW accuracy and precision.

The accuracy of both techniques varied with the underlying velocity field. The *direction* of the velocity must be known before setting the initial slicing direction of DTW; if done incorrectly, DTW inferences were unreliable. In contrast, the *magnitude* of the velocity must be known to infer the accuracy of CCTDE. The salient issue is that the velocity is unknown prior to velocimetry, once again highlighting that a combination of velocimetry techniques must be used for accurate velocity inferences.

6.5 Workflows for Predicting Velocimetry Uncertainty

In this section, an illustrative workflow is presented to highlight how the preceding velocimetry testing results (section 6.2 and 6.3) could be utilised to guide the experimental application of velocimetry methods. Note that the intent of this workflow was not to be comprehensive but, instead, to serve as a simple example that highlights some main considerations that should be taken into account.

The velocimetry was conducted on two synthetic TDF time-series, denoted simply as A and B. Note that velocimetry applications employing *experimental* data are covered in chapters 7 and 8. The signal-to-noise ratio, denoted as $SNR_{rms} = 100$, was known prior to the analysis. Spatial parameters and the underlying velocity of the data were unknown before the analysis, but they could be subsequently retrieved from metadata for cross-verification. Despite the utilisation of synthetic data in this section, the example workflow remains directly applicable to experimental velocimetry.

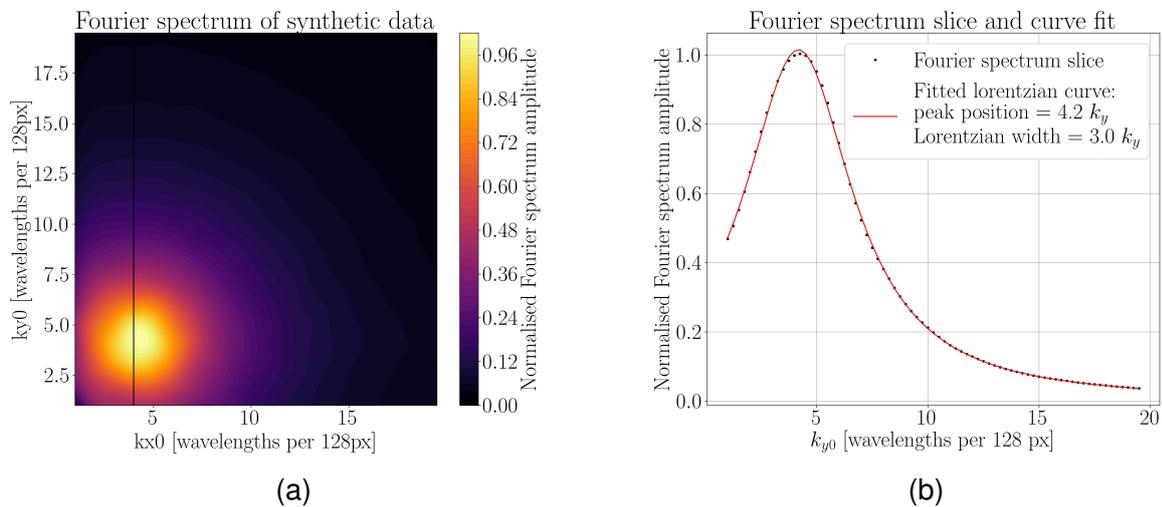


Figure 6.12: Example 2D spatial Fourier spectra of time-series A (a). 1D slice of the Fourier spectrum in the y-direction with Lorentzian fit (b). Slice location shown by black line in (a).

Before making any velocity inferences, preliminary analysis was conducted to evaluate the spatial characteristics of the fluctuations in the data. The magnitude of the 2D spatial Fourier transform of each image in the time-series was determined, with manual adjustment of the DC peaks to zero. The resulting Fourier images were summed together, and the outcome was normalised to the maximum amplitude. An illustration of such an aggregate Fourier spectrum is depicted in figure 6.12a, revealing a distinct, singular peak.

To quantify the spatial parameters, one-dimensional slices of the Fourier spectra were taken through the peak. These slices were then fitted through a least-squares routine to a Lorentzian, as presented in equation 6.4. A demonstration of the fit is shown in figure 6.12b. A summary of the final estimates for the spatial parameters of time-series A and B is provided in Table 6.1. Additionally, the minimum $\Delta\ell$ required to

avoid spurious velocity measurements induced by the barber pole illusion was calculated using equation. 6.9.

time-series	k_{x0}	k_{y0}	θ	Δk	$\Delta \ell_{\min}$
A	4.2	4.2	45°	0.7 k_{y0}	13 px
B	2.8	2.7	45°	0.7 k_{y0}	21 px

Table 6.1: Table summarising the estimated spatial parameters associated with the density fluctuations in the time-series. Wavenumbers given in wavelengths per image-width. Minimum $\Delta \ell$ was calculated from equation 6.9. All estimated values had uncertainty margins around 10%.

When applying CCTDE, multiple operational parameters must be considered. First of all, at $SNR_{rms} = 100$, which was well above the $SNR_{rms} \geq 1$ limit determined in sections 6.2 and 6.3, the influence of noise was thought to be negligible and spatial smoothing or filtering was presumed to be unnecessary. Furthermore, correlation parameters were computed using standard techniques (see [94] or section 7.1), revealing that decorrelation effects were negligible - which was of course expected from the synthetic data.

Subsequently, the selection of the time-series length, N , will be addressed. In general, it is preferable to minimise this parameter, as discussed in section 6.2.4. Nevertheless, it is essential to recognise that reducing N will decrease precision, potentially leading to extended averaging times and an overall adverse impact on the effective velocity inference frequency. Additionally, it noted that the minimum measurable velocity is influenced by N , as illustrated in Equation 6.8. $N = 128$ frames was arbitrarily chosen at this point, acknowledging that adjustments may be necessary.

The final parameter to be considered is the separation distance, $\Delta \ell$. Firstly, note that an increase in $\Delta \ell$ results in an elevation of the maximum measurable velocity and enhances the measured velocity resolution below $\Delta \ell$. While a larger $\Delta \ell$ is generally advantageous, caution is advised to avoid increasing $\Delta \ell$ to a degree where decorrelation effects become substantial. Given the insignificance of decorrelation effects in this data, $\Delta \ell$ could be adjusted freely.

As depicted in figure 6.3, the accuracy is anticipated to strongly vary with the underlying velocity. Consequently, $\Delta \ell$ had to be adjusted across multiple velocity inferences to hone in on an approximate quantification of velocity. After initial testing, a trial range of $\Delta \ell = 10\text{-}30$ px was selected. Now that both $\Delta \ell$ and N are established, the minimum measurable velocities were computed using Equation 6.8. Based on these calculations, velocities down to 1 px/frame can be measured accurately in both time-series.

Using the aforementioned parameters, velocities in both time-series A and B were estimated using CCTDE. The standard procedure outlined in section 6.2 was followed, and the results are summarised in Table 6.2. The inferred velocities exhibited significant variation with $\Delta \ell$, which was expected. No spatial variations were observed, so it was not considered worthwhile to present example velocity fields here.

For time-series A, only NaN velocities were recorded in the x-direction, and the velocity in the y-direction exhibited the highest correlation at 15px/frame. In time-series B, non-zero velocities in the x-direction were observed at $\Delta \ell \leq 15$. These were hypothesised to be spurious velocities induced by the barber pole illusion, as anticipated by

Table 6.1. This hypothesis could be cross-checked with DTW results at a later point. The velocity in the y-direction had multiple values with correlation close to unity, and the y-velocity was estimated to fall within the range of 10-20 px/frame. This range could be refined at a later point by increasing the $\Delta\ell$ resolution and through comparison with DTW.

time-series	$\Delta\ell$	v_x	$v_x \text{ corr.}$	v_y	$v_y \text{ corr.}$
A	10	nan	0.3	10	0.5
	15	nan	0.3	15	1.0
	20	nan	0.3	20	0.5
	25	nan	0.3	12.5	0.5
	30	nan	0.3	15	0.5
B	10	10	0.5	10	0.7
	15	15	0.4	15	1.0
	20	nan	0.4	20	0.9
	25	nan	0.3	25	0.6
	30	nan	0.3	15	0.9

Table 6.2: Summarising CCTDE inferred velocities for time-series A and B. $\Delta\ell$ was varied from 10 px to 30 px. Correlation amplitudes of the inferences are included. NaN velocities given when the time-delay was zero.

Subsequently, velocities will be inferred using DTW to enable cross-checking between the techniques. Given that the velocity fields were previously found to solely point in the y-direction, the initial slicing direction for DTW could be set accordingly. Moreover, no sheared flows were observed, and no spurious velocities due to the barber pole illusion should be anticipated from DTW at $\Delta k = 0.7k_{y0}$. DTW accuracy is also not significantly dependent on fluctuation spatial scale or underlying velocity at $SNR_{rms} = 100$. For these reasons, DTW could be straightforwardly applied without expected complications.

Strip widths of [32, 22, 12, 8, 6, 4, 2] px were employed for a total of 7 iterations, with the initial slicing direction set along the y-direction. As depicted in figure 6.13, velocity fields were successfully measured by DTW with standard deviations of less than 10% from the mean velocity. For time-series A, average velocities of 0 px/frame in the x-direction and 15 px/frame in the y-direction were determined. For time-series B, average velocities of 0 px/frame in the x-direction and 17 px/frame in the y-direction were inferred. These results generally align with the CCTDE estimates and support the hypothesis that the CCTDE-inferred non-zero velocities in the x-direction were indeed due to the barber pole illusion.

CCTDE was rerun for both time-series A and B while varying $\Delta\ell$ from 10px to 20px in 1px increments. In the case of time-series A, the velocity in the y-direction with the highest correlation remained unchanged at 15 px/frame. Conversely, for time-series B, the velocity in the y-direction with the highest correlation was determined to be 17 px/frame. Through the utilisation of both velocimetry methods and cross-checking results, dependable velocity field inferences were achieved for both time-series A and B.

In conclusion, the analysis revealed that fluctuations in both time-series exhib-

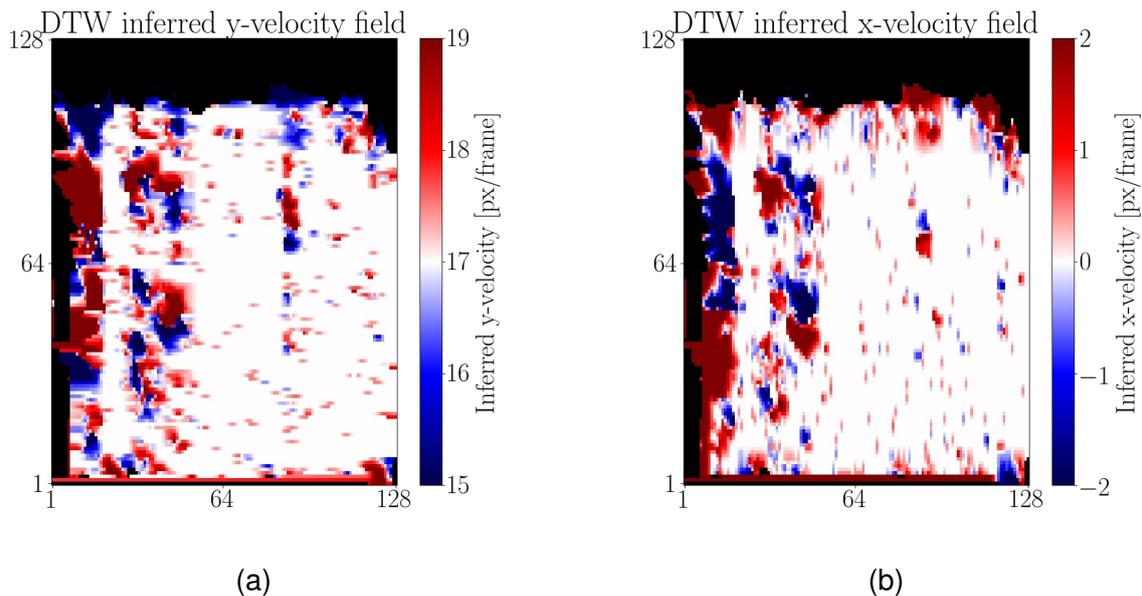


Figure 6.13: Example DTW velocity field inferences of time-series B. Velocity in the x-direction (a) averages to zero and velocity in the y-direction (b) averages to 17px/frame. Colourbars were centred on the averages. All velocity fields were measured with a standard deviation of $< 10\%$ from the average.

ited no significant velocity in the x-direction. The velocity in the y-direction for time-series A was determined to be $15 \pm 0.5 \text{ px/frame}$, while for time-series B, it was $17 \pm 0.5 \text{ px/frame}$. Uncertainties were defined as half the step size in $\Delta\ell$ for the final CCTDE inferences.

Imposed velocities were subsequently recovered from the metadata, confirming y-velocities of 15 px/frame and 17 px/frame for time-series A and B, respectively. The velocity in the x-direction remained at 0 px/frame in both cases.

This section demonstrated how fluctuation spatial parameters were instrumental in configuring the operational parameters of CCTDE, and accurately predicting spurious velocities due to the barber pole illusion. The direction of the CCTDE-inferred velocity fields was then utilised to set the slicing direction in the DTW analysis. By cross-checking results from both techniques and refining the CCTDE inferences, consistent outcomes between the two methods were achieved.

6.6 Discussion and Conclusion

The two main image velocimetry techniques, CCTDE and DTW, were tested extensively to quantify the dependencies of their accuracy and precision on key parameters in the underlying fields. Synthetic data was used to represent a range of fluctuation structures observed in turbulence diagnostics, namely, ranging from isolated density fluctuation structures to fully-developed turbulent density fields. Additionally, specific scenarios were investigated like the barber pole illusion, sheared velocity fields, and variation of the number of spatial channels. It was shown in sections 6.2 and 6.3 that

the accuracy of both techniques can exhibit strongly nonlinear behaviour. It is therefore ill-advised to extrapolate results from any such tests, including previous literature, beyond the investigated parameter range. Nevertheless, the scope of this thesis covers the typical data parameters for most plasma turbulence diagnostics.

Decorrelation effects were not imposed on the synthetic data. For DTW this is not thought to be an issue unless the decorrelation timescale is comparable to the diagnostic measurement frequency, which is the realm where any velocimetry method is doomed to fail anyway. CCTDE can be more strongly affected by decorrelation effects, especially if $\Delta\ell$ is increased. Nevertheless, decorrelation effects can often be mitigated by analysing the decorrelation timescale, as discussed in [94] and section 7.1.

This study only investigates velocities pointing in the orthogonal directions, thus preventing the investigation of, for example, rotational flows. However, the impact of such flows on two-point CCTDE and DTW is likely limited. If rotational flows are spatially larger than the spatial resolution, CCTDE will be unaffected. If rotations are smaller, the diagnostic will not record them anyway. It should also be noted that some more elaborate CCTDE methods can infer diagonal velocities [80]. It has been shown previously that DTW can perform accurately with rotational velocity fields [83].

It was found that CCTDE's accuracy strongly depends on the underlying velocity, which can unknowingly introduce inaccuracies if not carefully considered. A number of options which address this issue were discussed in section 6.2.4.

The signal-to-noise ratio and density fluctuation spatial scale did not have a significant impact on CCTDE accuracy in the majority of cases. Precision loss due to noise could mostly be counteracted via filtering.

The length of the time-series could be reduced down to 32 frames without significant impact on the CCTDE accuracy. This represents an inference frequency which is an order of magnitude faster compared to the typical ≥ 256 frames in previous literature.

Spurious CCTDE velocities due to the barber pole illusion were quantified and could largely be avoided through a simple analysis *before* velocimetry. This complication, which is highly prevalent in plasma turbulence studies, could previously not be circumvented without extensive additional analysis and bold assumptions [88].

The accuracy of DTW, with 128 by 128 channel images, was not strongly dependent on the spatial size of the fluctuations or the underlying velocity. Noise was detrimental to accuracy but these effects could largely be counteracted via spatial smoothing of the images. Sheared flows were detrimental to accuracy and a threshold in the maximum shear was observed. These tests confirmed the expected result that DTW is typically a reliable technique for such high-spatial-resolution images.

A major caveat with DTW is that *the direction of the underlying velocity field must be known prior to analysis*. The flow direction is used to set the initial slicing direction of DTW, which results in complete failure of the velocimetry if set incorrectly.

A more challenging test for DTW was to observe its performance when the number of spatial channels is reduced. In this regime, it was found that the accuracy and precision were strongly reduced in images with 16 by 16 channels or fewer. Additionally, noise had a strongly enhanced detrimental effect on the method precision and accuracy. Through averaging consecutive measurements, accurate velocity fields could be recovered in 8 by 8 channel images with $SNR_{rms} = 2$. This is novel evidence that DTW

can confidently be applied to high-time-resolution, but low number of spatial channels plasma diagnostics such as beam emission spectroscopy [75], which was a point of contention in previous literature. No spurious velocities due to the barber pole illusion were observed using DTW at $\Delta k \geq 0.3 k_{y0}$.

In conclusion, the accuracy and precision of both CCTDE and DTW were quantified under a broad range of conditions. Improvements on multiple fronts were made to the operational range of both techniques. It was found that two of CCTDE's main drawbacks were: the velocimetry accuracy can depend strongly on the **magnitude** of the underlying velocity, and CCTDE is prone to spurious inferences due to the barber pole illusion. Conversely, DTW accuracy was found to strongly depend on the **direction** of the underlying velocity and the number of spatial channels in the data. Additionally, DTW was much more robust against spurious inferences due to the barber pole illusion. In general, it is recommended that both CCTDE and DTW be used in conjunction with each other for accurate velocity field inferences. To this end, a basic example workflow was presented in section 6.5 which utilises both techniques and should significantly improve confidence in velocity estimates. In any case, the highly nonlinear and often unexpected behaviour of both techniques show that referral to tests like the ones in this chapter are **necessary** for confident and accurate image velocimetry.

Chapter 7

Delocalised Velocimetry Experiments on MAST-U

The first velocimetry applications in this thesis to *experimental* data are covered in this chapter. The focus was on applying the results shown in chapter 6 and investigating potential differences with the synthetic data tests. In section 7.1, the BES data is examined, and preparatory analysis for velocimetry is performed. The initial results from DTW and CCTDE velocimetry are covered in sections 7.2 and 7.3, respectively. In section 7.4, it is observed that de-localised carbon emissions dominate the BES signal in almost all MAST-U shots in the second campaign (MU02). At first glance, this discovery is devastating for *any* potential velocimetry, but, through further investigation, a method was developed to uncover the carbon II shell dynamics for the first time using BES.

7.1 Preliminary Analysis of BES Density Fluctuations

This section is dedicated to the analysis of density fluctuation structures captured by BES in shot 46883. Subsection 7.1.1 explores the existence of coherent density structures in both radial and poloidal directions of the BES field-of-view. Subsection 7.1.2 covers the preparatory analysis of the BES data required for confident velocimetry analysis. Specifically, the analysis investigates the signal-to-noise characteristics, spatial correlation lengths, eddy tilting angles, and decorrelation timescales of the density fluctuations.

7.1.1 Coherent Density Fluctuation Structures

In preparation for the application of BES density fluctuation data to velocimetry techniques, an analysis was undertaken to investigate the presence of coherent density structures in the data. The focus is on shot 46883, which was chosen due to a prolonged L-mode that features clean broadband turbulence signatures in the BES $k_\theta - f$ spectra. The shot encompasses an ELM-y H-mode between 200-300 ms, an ELM-free H-mode spanning 300-400 ms, a H-L backtransition around 400 ms, and a steady L-mode from 430-660 ms before disruption, which can be seen in figure 7.1. The L-mode is characterised by consistent broadband turbulence ranging from 10-30 kHz,

exhibiting a phase velocity of approximately ~ -7 km/s in the z-direction¹.

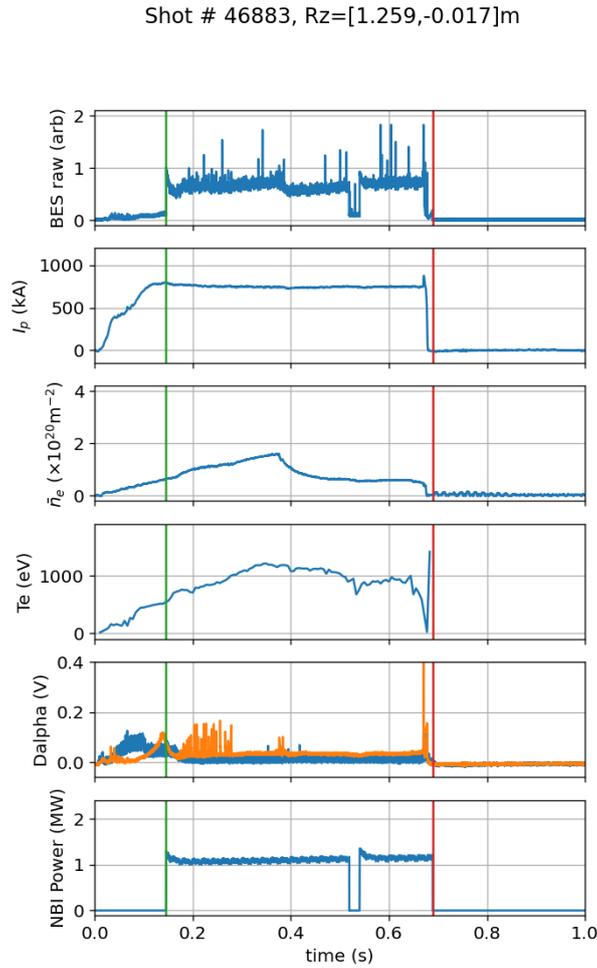


Figure 7.1: Overview of shot 46883. ELM-y H-mode between 200-300 ms, an ELM-free H-mode spanning 300-400 ms, a H-L backtransition around 400 ms, and a steady L-mode from 430-660 ms before disruption at ~ 700 ms. Top subplot shows a BES raw signal trace. Vertical green and red lines denote start and end of NBI power. Second, third and fourth subplots show plasma current, electron number density and electron temperature, respectively. Fifth subplot shows D_α emissions. Final subplot shows the SS NBI injection power.

To explore the presence of coherent structures, the magnitude-squared coherence between BES channels was calculated, as shown in Equation 7.1:

$$C_{xy}(f) = \frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)} \quad (7.1)$$

$C_{xy}(f)$ and $G_{xy}(f)$ represent the magnitude-squared coherence and spectral density, respectively. Both quantities are calculated with two time-signals x and y , and given at frequency, f . The time signals x and y are taken from BES channels (i,j) and $(i+\text{sep},j)$,

¹The z-direction of the BES array corresponds approximately to the -ve poloidal direction.

respectively, to investigate the poloidal coherence. Channels (i,j) and $(i,j+sep)$ are chosen to investigate the radial coherence. The separation value, 'sep', represents the poloidal (radial) distance between the BES channels. All possible pairs of channels are taken at constant separation value, and the subsequent coherence functions are averaged to maximise statistical convergence. Preliminary testing did not reveal significant differences between individual coherence spectra, justifying the use of averaging.

Examining figure 7.2, a poloidal coherence around unity is observed in the sub-kHz regime. At $sep = 2$ cm, increasing in frequency, the coherence gradually drops to ~ 0.8 at 30 kHz, after which a sharp drop in coherence is observed. Despite the sharp drop, significant coherent peaks are still observed above 100 kHz, although these peaks likely correspond to global noise that affect all BES channels simultaneously.

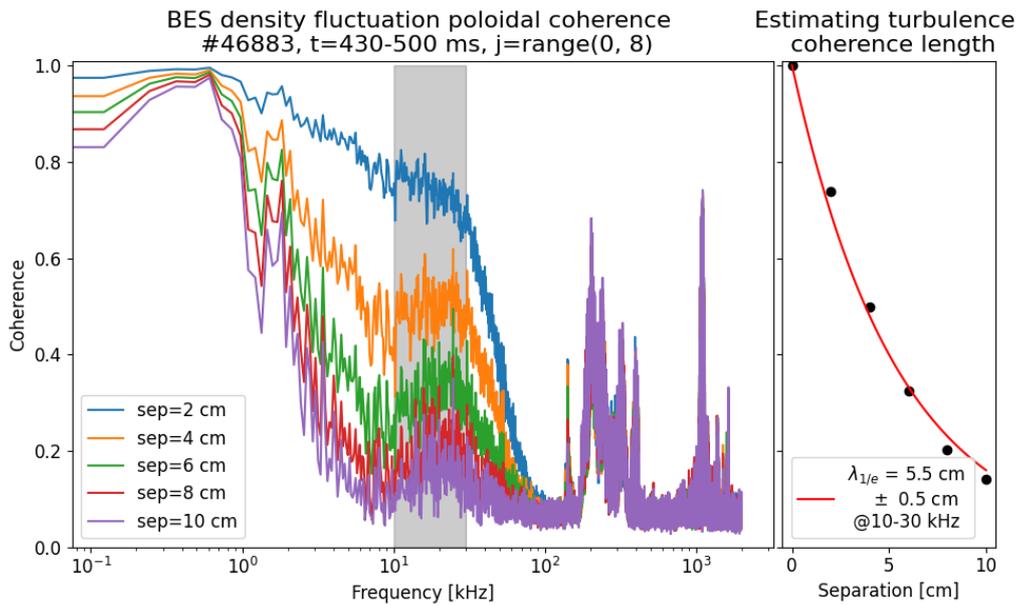


Figure 7.2: Poloidal coherence of the BES density fluctuations. The left plot shows the poloidal coherence spectra at different separation values. Separation values are given as average channel-pair separation in cm. Grey shaded area represents the broadband turbulence range. The right plot shows the median coherence of the broadband turbulence depending on poloidal separation distance. The e-folding length, $\lambda_{1/e}$, and uncertainty margin were estimated through a least-squares fit to an exponential decay function.

By varying the poloidal separation between channels, the poloidal coherence lengths of distinct features can be investigated. As seen in figure 7.2, a highly coherent peak is observed around 200-800 Hz which doesn't depend strongly on the poloidal separation distance. These frequencies correspond to the NBI frequency range and therefore likely represent coherent NBI modulation of the BES emissions. It can also be observed that the turbulence, with frequencies 10-30 kHz, has a significant poloidal coherence length, which was quantified to have an e-folding length of ~ 5.5 cm ± 0.5 cm. Finally, another noteworthy feature is a 1-2 kHz peak that displays a coherence length significantly longer than the turbulence. The nature of this peak was not clear, but further investigation fell outside the scope of this research.

A comparable analysis was performed investigating the radial coherence spectra, as shown in figure 7.3. Again, the NBI attenuation appears to have a long radial coherence, which was to be expected. Interestingly, the 1-2 kHz peak also displays a long radial coherence length. The turbulent fluctuations displayed a radial coherence e-folding length of $\sim 3.7 \text{ cm} \pm 0.4 \text{ cm}$, which was significantly shorter than the poloidal coherence length.

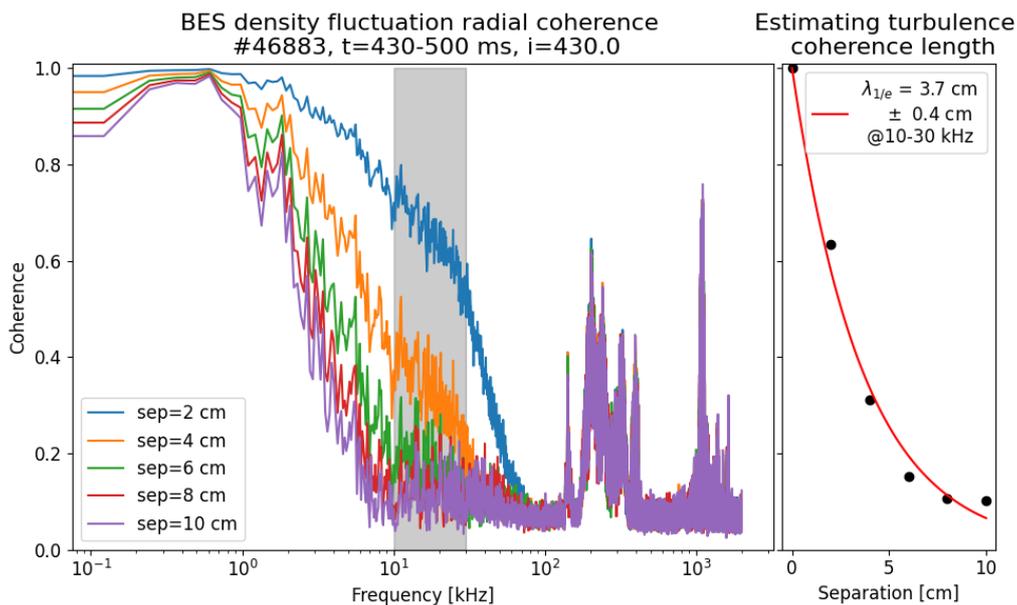


Figure 7.3: Radial coherence of the BES density fluctuations. The left plot shows the radial coherence spectra at different separation values. Separation values are given as the average channel-pair separation in cm. Grey shaded area represents the broadband turbulence range. The right plot shows the median coherence of the broadband turbulence depending on radial separation distance. The e-folding length, $\lambda_{1/e}$, and uncertainty margin were estimated through a least-squares fit to an exponential decay function.

In short, high coherence was observed both poloidally and radially in the NBI attenuation of the BES signal, as expected. An unidentified peak was identified around 1-2 kHz with consistently high poloidal and radial coherence. Finally, the background turbulence was found to be characterised by a poloidal coherence length of approximately 5.5 cm and a radial coherence length of approximately 3.7 cm. This result is not unsurprising because shear flows often decrease the radial extent of density structures whilst leaving the poloidal extent relatively unaffected. Although shear flows could explain the observation, concrete conclusions could not be drawn at this point.

7.1.2 Preparatory Analysis for Velocimetry

In the initial examination of density fluctuations within the BES data, conditional k_θ - f spectra were plotted, as depicted in figure 7.4. It can be seen that the density fluctuations exhibit a consistent phase velocity of approximately 6-7 km/s throughout the L-mode times. This observation remains consistent across the entire BES

radial range, with no significant temporal variations in velocity slower than a ~ 1 ms timescale. Faster timescales could not be investigated through this method due to convergence issues. Noise is present at frequencies above ~ 150 kHz and will be filtered out throughout further investigations. Furthermore, these plots reveal that the turbulent signal does not extend beyond ~ 50 kHz. Finally, coherence analyses between the BES fluctuations and the NBI fluctuations (which are not shown here) indicate a rapid drop-off in NBI modulation of the BES signal above 2 kHz, becoming negligible by 3 kHz. Consequently, a frequency range of 3-50 kHz was deemed appropriate to isolate turbulent density fluctuations.

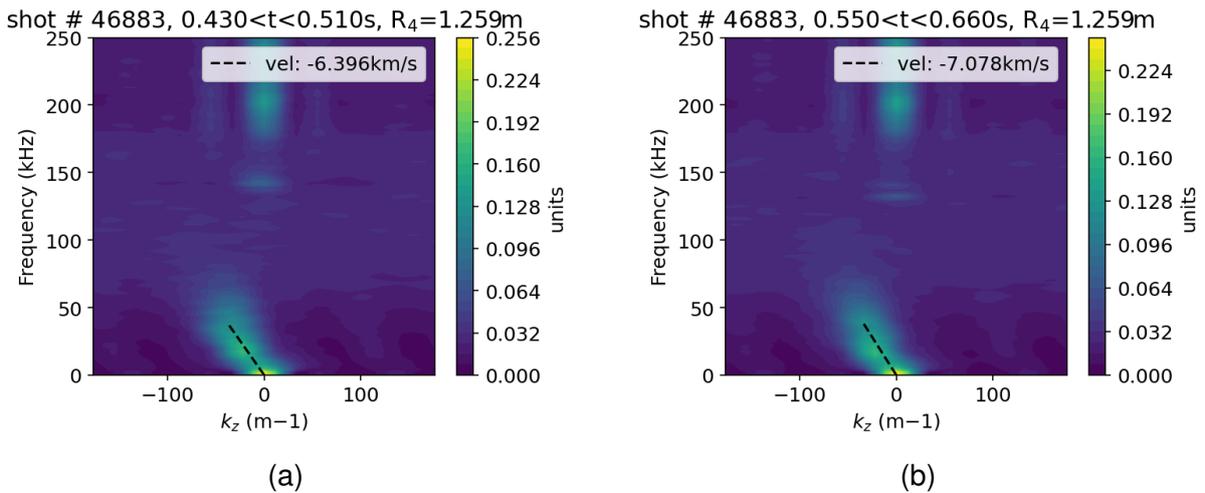


Figure 7.4: k_θ - f spectra of BES data. An NBI beam notch was present in the shot between graphs a) and b), at times 520-540 ms. Both graphs represent an L-mode plasma.

Two signal samples were acquired to estimate signal-to-noise characteristics: a raw signal sample taken during times 430-510 ms and a background sample during the NBI notch at 520-540 ms. The NBI notch, characterised by its brief duration to minimise changes in the plasma equilibrium, was thought to provide an appropriate background signal for BES measurements. The signal-to-background ratio (SBR) was computed by dividing the RMS of the raw signal by the RMS of the background. This calculation was conducted for all BES spatial channels, resulting in the graph illustrated in figure 7.5. Notably, the SBR exhibits significant variation across the field of view, likely corresponding to the variation of NBI amplitude. Additionally, the signal-to-noise ratio (SNR) was determined by filtering both the signal and background between 3-50 kHz and taking the ratio of fluctuation amplitudes from their respective means. The SNR, estimated to be $\text{SNR} = 1$, displayed no significant spatial variation. Comparing these estimates with the established noise limits in sections 6.2 and 6.3 suggests that noise may *marginally* affect velocimetry accuracy for both CCTDE and DTW, but velocimetry should generally be accurate.

A spatial Fourier analysis was attempted like the one performed in section 6.5, but no useful information could be resolved, likely due to the low number of spatial channels and low signal to noise ratio of the BES data. Instead, a spatial correlation approach was used to investigate the typical spatial structures of the density fluctuations, which

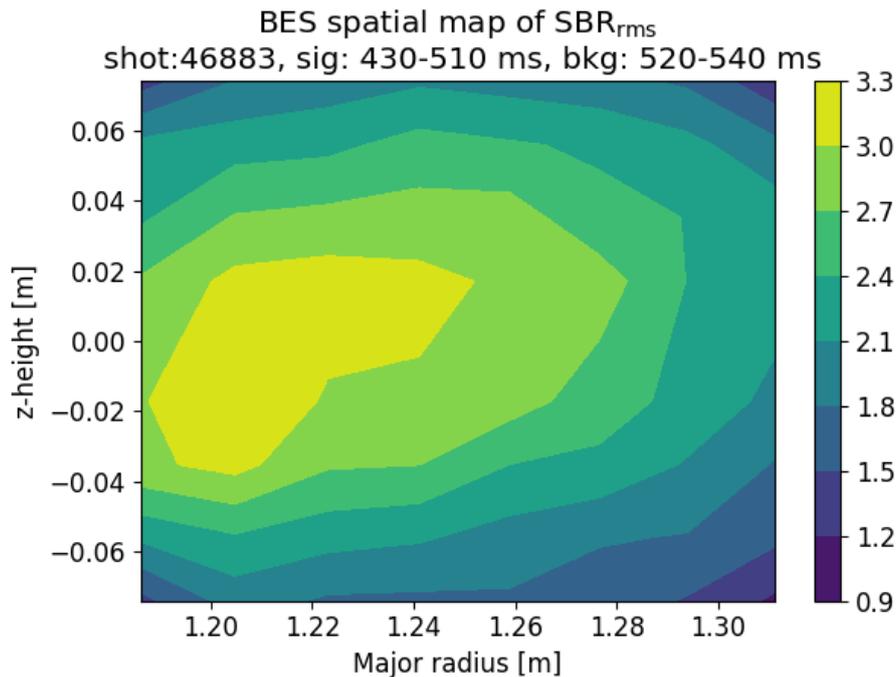


Figure 7.5: Spatial map of the signal-to-background ratio of the BES signal.

involved the following steps:

1. A reference BES channel was selected.
2. The reference time-signal was cross-correlated with signals from each of the other BES channels.
3. A spatial correlation map was built up by taking the zero-time-lag correlation values.
4. Gaussian least-square fits were taken through the reference in the radial and z-directions.
5. Average correlation lengths of the density structures were estimated based on the full width at half maximum (FWHM) of the Gaussian fits.
6. Steps 1-5 were repeated with different reference channels to build up a radial scan of the correlation lengths

The results in figure 7.6 suggest correlation lengths of approximately 8-10 cm in both radial and z-directions, with estimates at the edges disregarded due to edge clipping effects. Individual spatial correlation maps, not presented here, showed no evidence of density feature tilting. In short, the spatial correlation analysis estimated $\lambda_{z0} \sim \lambda_{R0} \sim 8 - 10 \text{ cm}$, $\theta = 0$, and no estimates could be made for Δk .

The typical decorrelation timescale of the density fluctuations was also investigated through a cross-correlation method, as outlined below:

- Two signals, denoted as f and g, were selected at BES indices (i, j) and (i+sep, j), respectively, where the parameter 'sep' was an integer initialised equal to 1.

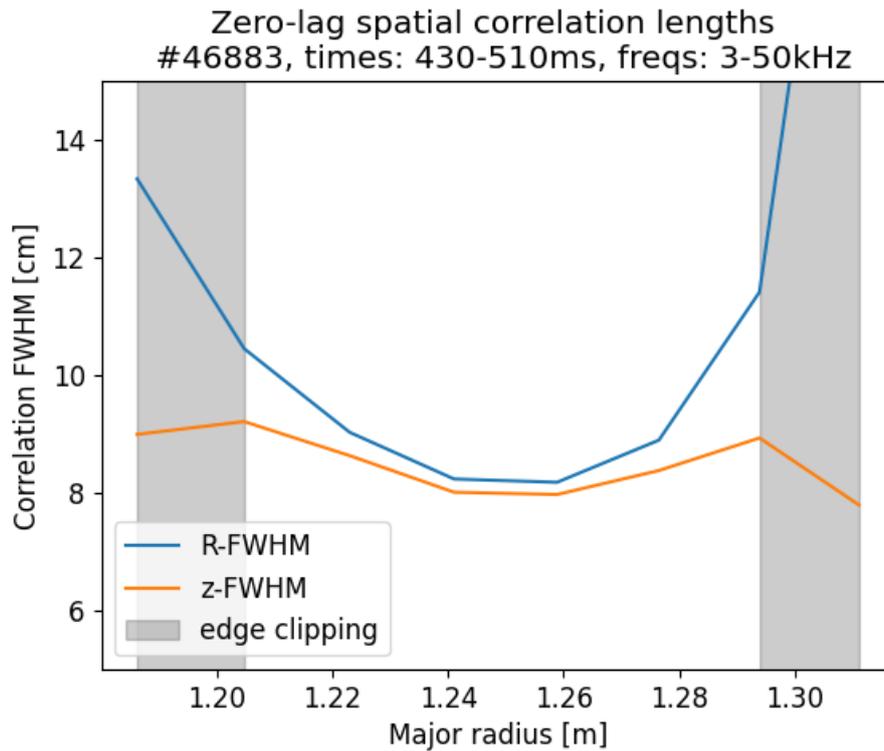


Figure 7.6: Spatial correlation lengths of BES density fluctuations. Estimates near the edges should be ignored due to clipping effects.

- The cross-correlation between signals f and g was computed.
- The peak of the cross-correlation function and its corresponding time-lag were located.
- This process was repeated with variations in the parameter 'sep'.
- An exponential decay function was then fitted through the peaks of the correlation function, providing an estimate of the decay time.

The outcomes, illustrated in figure 7.7, present the decorrelation of the BES fluctuations with time. The $1/e$ decay time was estimated at $15 \mu\text{s}$ or 59 frames. Correspondingly, at an estimated velocity of 6-7 km/s, this translates to a lengthscale of 9-10 cm. These findings suggest significant decorrelation effects across the BES field-of-view, yet they are unlikely to have a significant effect when analysing directly adjacent channels.

7.2 Initial Results and Challenges with DTW

This section encompasses initial tests which evaluate the performance of DTW when applied to BES data from MAST-U, shot 46883. The general methodology employed for inferring velocities using DTW is outlined as follows:

- The BES data is pre-processed to isolate turbulent fluctuations (as detailed in section 7.1).

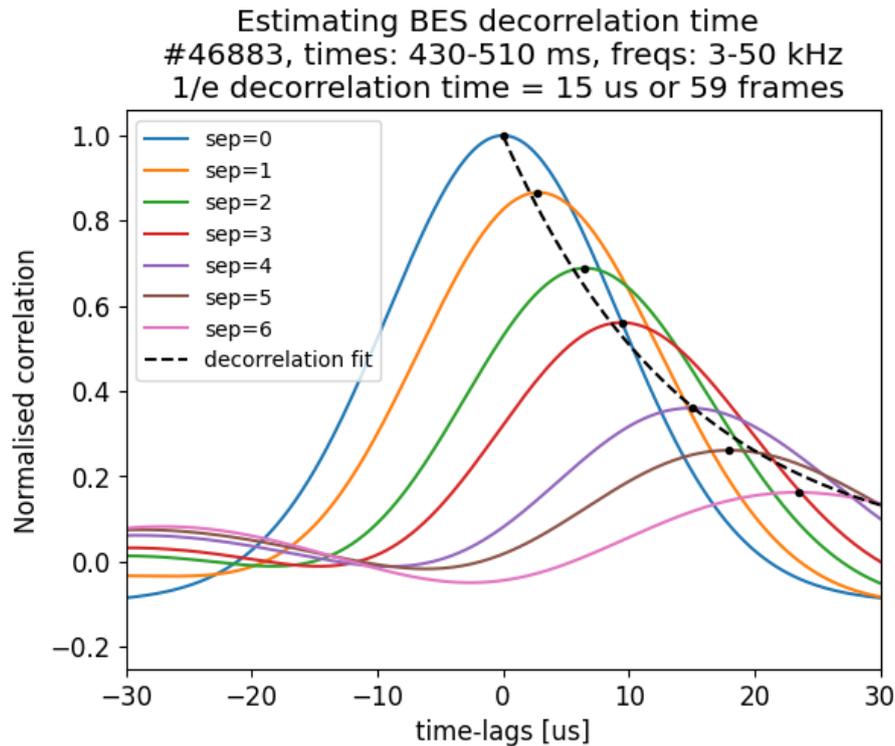


Figure 7.7: Estimating decorrelation timescale of the turbulence. Estimated timescales were given as the 1/e exponential decay time.

- The BES data is interpolated onto a regular 8 by 8 grid since irregular spatial maps are incompatible with the current implementation of DTW.
- Two images separated by 'shift' number of frames are selected and processed using DTW to estimate a displacement field (the rationale behind the shift parameter and its selection process will be discussed later in this section).
- The previous steps are repeated for the entire BES time series using subsequent pairs of frames.

In these investigations, BES data at times 430-510 ms and frequencies 3-50 kHz were utilised, and the parameter `sw_list` was set to two iterations with [4,2] channel strip-widths.

The estimated phase velocity in this time series was 6-7 km/s (see section 7.1.2), corresponding roughly to 0.1 channel separations per frame. However, this presents a challenge because, as established in section 6.3, DTW struggles to accurately infer velocities below 1 channel separation per frame (1 px/frame). To address this, instead of analysing adjacent BES images, the image at index t would be compared with image t +shift. This can artificially increase the apparent velocity inferred by DTW and bring the velocity into the accurately measurable regime. As depicted in figure 6.11, DTW can infer accurate velocities at ~ 1 channel separation per frame. Given the estimated velocity of 6-7 km/s, meaning that the density features travel one channel separation in ~ 12 frames, the shift parameter was set to 12 frames for initial tests. Since the

inclusion of the 'shift' parameter was not previously tested in section 6.3, the effect of varying the shift parameter was investigated later in this section.

Initial results yielded an average velocity of -2 km/s with a standard deviation of approximately 20 km/s. Apart from the low precision and the average velocity disagreement with the estimated phase velocity, several other complications were observed. Namely, edge effects were observed in the inferred velocity fields, with velocities at the edges of the field-of-view observed to be relatively large and point away from the edges. Further initial tests (not shown here) revealed that interpolating the BES data onto higher numbers of spatial channels exacerbated these edge effects, leading to the decision not to further explore this 'up-resolving' approach.

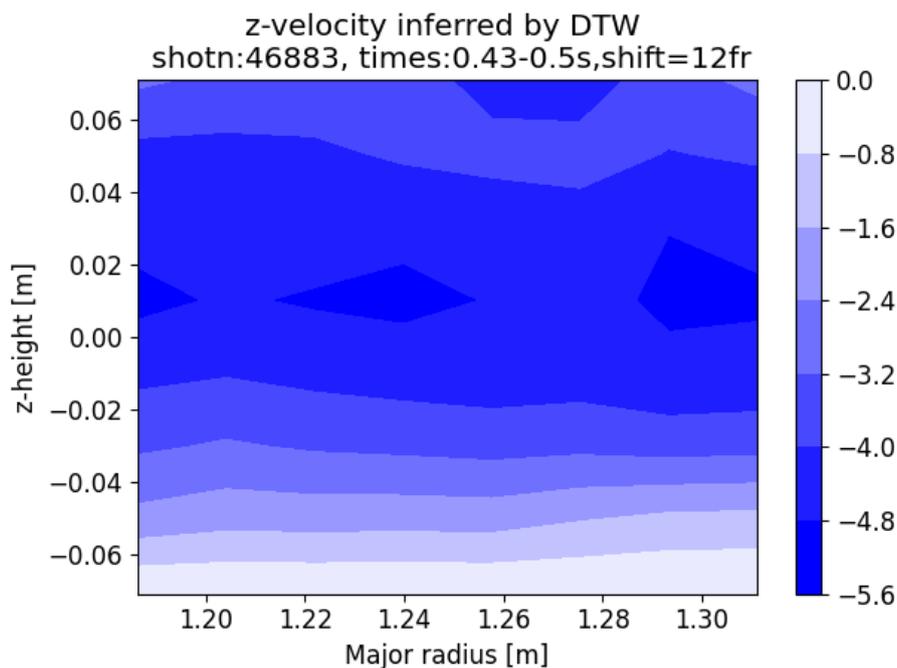


Figure 7.8: Time-averaged velocities in the z-direction inferred by DTW. Velocity gradient in the bottom half of the field-of-view was likely nonphysical and instead due to a currently unknown quirk of DTW.

An issue related to these edge effects is observable in figure 7.8. Velocities at the bottom of the frame are zero (and often even positive), while velocities in the top half of the frame were consistently negative. These DTW results imply that the density features are being compressed in the bottom half of the frame, although this is certainly not an accurate reflection of the true physical picture. Instead, this phenomenon is an expected quirk of DTW² that would not have been very significant in the testing of section 6.3; however, further testing of this effect was outside the scope of this investigation. Instead, the bottom four channels, as dictated by the largest strip width,

²This quirk is due to the fact that DTW cannot infer the velocities of features that leave the field of view. In other words, negative velocities cannot be measured at the very bottom of the frame. In theory this should only affect the bottom row of channels, but due to the way that DTW averages overlapping strips and due to the strip widths chosen, this effect bleeds into the bottom 4 channels. As such, re-running DTW with strip widths [2,1] channels could be considered for future work, but this falls outside of the scope of this investigation.

could be ignored to circumvent this edge effect. The remaining top half of the field-of-view appears to estimate velocities in rough agreement with the preliminary phase velocity estimations in figure 7.4.

The estimated optimal shift parameter for this data was previously evaluated to lie around 12 frames, but it remained unclear how sensitive DTW performance was to this shift. Consequently, the shift parameter was varied from 2 to 256, and the resultant average velocities were presented in figure 7.9. The inferred velocity at $z \sim -0.07$ m ranged between 0 and 1 km/s, consistent with the aforementioned ‘squishing’ problem. In the top half (+ve z) of the field-of-view, the inferred velocities varied strongly with varying shift. At shift ~ 2 frames, the inferred velocity was close to zero. Increasing the shift resulted in an increasingly negative velocity, until the average velocity (at +ve z) reached 5-7 km/s with shift = 12, which was in approximate agreement with the estimated phase velocity of 6-7 km/s. The average inferred velocities then remained approximately constant for shift 12-32 frames. Increasing the shift beyond this point brought the inferred velocity closer to zero again. Remarkably, this pattern of underestimation, roughly accurate velocities, and then underestimation again, is consistent with the testing shown in figure 6.11. To clarify this point, it is important to realise that increasing shift is akin to increasing v_0 in figure 6.11.

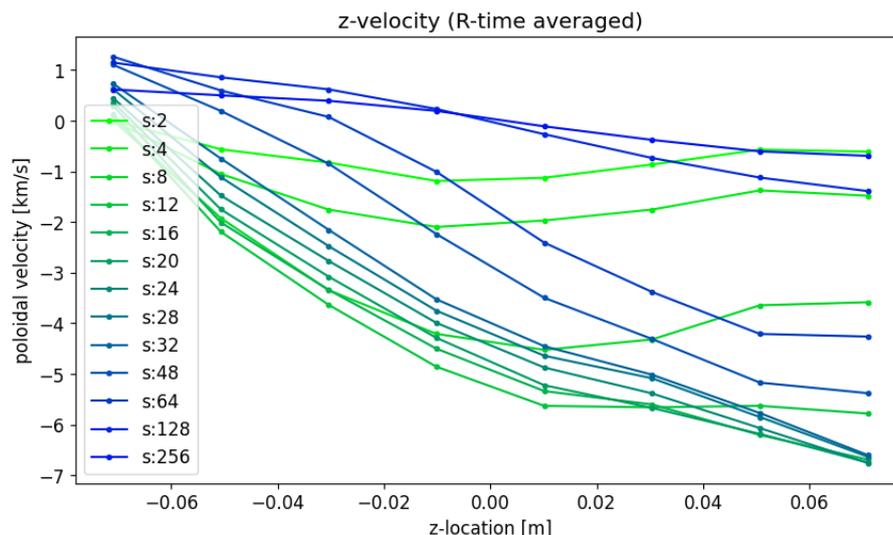


Figure 7.9: Investigating the effect of varying ‘shift’ on the velocities inferred by DTW. Velocities at $z \sim -0.07$ m were constrained to only positive velocities due to edge effects.

Finally, the effect of the DTW operational parameter ‘sw_list’ was investigated. In-depth results are not shown here because the effect of varying sw_list was rather limited, as expected from the testing in section 6.3. Regardless, it was observed that a strip-width of 8 was detrimental for this 8 by 8 array, and a strip-width of 1 did not yield any beneficial effects, despite requiring a significantly larger computational cost. Thus, a sw_list of [4,2] channels was chosen for these investigations.

In conclusion, an edge effect was found in this investigation that was not previously observed in the testing of section 6.3. This edge effect was an expected quirk of the DTW method, but further investigations aiming to minimise these edge effects were

outside the scope of this analysis. The edge effects do not impact the top half of the field-of-view for these investigations. In this area the DTW velocities were observed to be in agreement with the expected phase velocities. This agreement was observed for shift values ranging from 12-32 frames, and the introduction of the shift parameter enabled velocities to be inferred despite lying below the 1 channel separation per frame operational limit found in section 6.3. *Note that the introduction of the shift parameter also enables averaging of up to 'shift' number of velocity inferences without further impacting the effective inference frequency.* Despite these somewhat promising results which suggest that DTW *can* be successfully applied to BES data, more in-depth velocimetry analysis could not confidently be performed without significant further testing on the effect of varying the shift parameter and operation in this low velocity regime.

7.3 Initial Results and Challenges with CCTDE

This section presents the workflow and initial results of CCTDE velocimetry of shot 46883. Subsection 7.3.1 describes preparatory work performed to optimise the performance of CCTDE and identify possible challenges. Subsection 7.3.2 presents initial velocimetry results and investigates performance dependence on variations in $\Delta\ell$ and N.

7.3.1 Preparations for CCTDE velocimetry

The expected velocities of -6 to -7 km/s, as detailed in section 7.1.2, fall significantly below 1 channel separation per frame (approximately 80 km/s). In this regime, the accuracy dependence on the true underlying velocity was thought to be negligible, as can be seen in the bottom subplots of figure 6.3. Despite not facing complications related to the previous point, it was anticipated that precision may be decreased in this regime.

The decorrelation lengthscale was projected to be around 9 cm with expected velocities at 6-7 km/s. At a $\Delta\ell$ of 1.8 cm, the impact of decorrelation on the inferred velocities was deemed negligible. The impact of decorrelation on the inferred velocities can be quantified through analysis comparable to that in [94].

SNR and SBR estimates ranging from 1-3 approach the $\text{SNR} \geq 1$ limit defined in section 6.2. While a loss of accuracy was not expected due to the SNR based on the testing in chapter 6, a reduction in precision may be observed in the inferred velocities.

Operational CCTDE parameters $\Delta\ell$ and N were determined through consultation with figure 7.10. In this figure, the minimum measurable velocities, as predicted by equation 6.8, were plotted. It is evident that N = 32 and 64 frames cannot measure velocities down to the expected velocity of -6 km/s. N = 128 frames is the lowest N that may be able to detect expected velocities, provided $\Delta\ell$ is kept below ~ 8 cm. At N = 256 frames, minimum velocities range from 2-4 km/s, depending on the choice of $\Delta\ell$. With the selection of N = 256 frames, CCTDE should not only be able to capture the expected average velocity but also observe fluctuations of ~ 2 -4 km/s from a presumed average of -6 km/s.

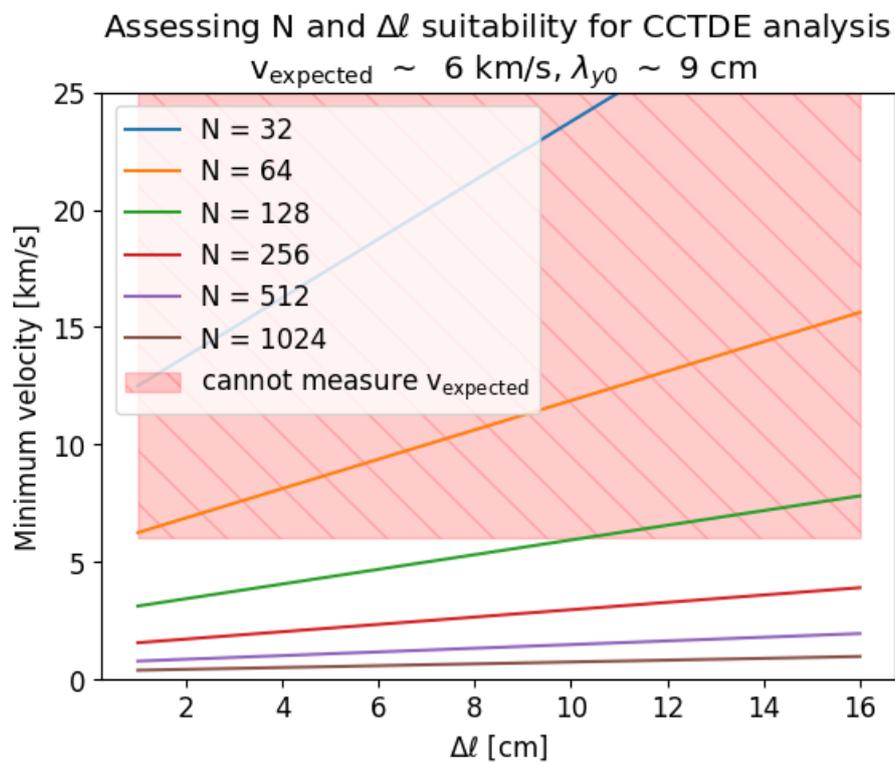


Figure 7.10: The minimum measurable velocities attainable by CCTDE, as predicted by equation 6.8. The red hatched area shows the parameter combinations which cannot measure down to the expected velocity of -6 km/s.

The impact of the barber pole illusion could not be estimated using the methods described in section 6.5 due to a lack of a Δk estimate. Nevertheless, given the absence of eddy tilting in the spatial correlation maps, it was assumed that the barber pole illusion would have a negligible impact.

7.3.2 CCTDE velocimetry results

Preliminary results were obtained from CCTDE using the following parameters: $N = 256$ frames and $\Delta \ell = 1.8$ cm. Figure 7.11 displays an example time-series of inferred velocities. The standard deviation of the velocities typically ranged from 15-20 km/s, significantly exceeding the magnitude of the estimated mean velocity of approximately -6 km/s. Notably, the substantial standard deviation is primarily attributed to large outliers in velocity, as reflected in the IQR estimate of around 6 km/s that excludes these outliers. Nevertheless, due to the low precision in the inferred velocities, long averaging times are necessary to ensure convergence of the standard error. For instance, averaging over the entire time series with 1250 realisations yields a standard error of 0.4 km/s.

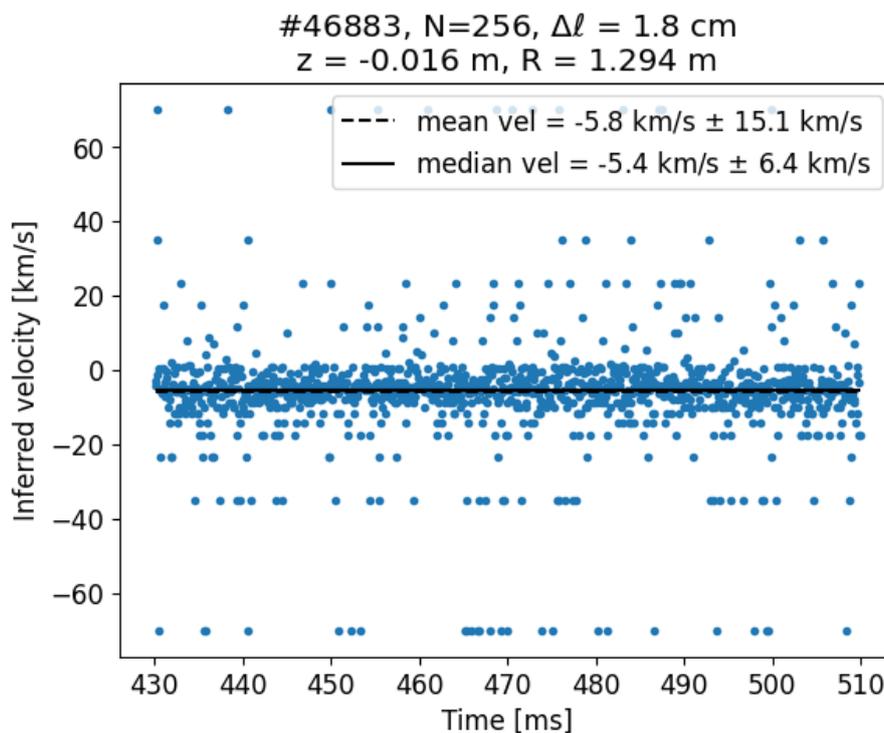


Figure 7.11: CCTDE velocimetry results at one example spatial location. Uncertainty margins shown for the mean and median were the standard deviation and IQR, respectively. Standard error for the mean was 0.4 km/s.

Figure 7.12 illustrates the dependence of velocimetry performance on $\Delta \ell$. Increasing $\Delta \ell$ substantially reduces precision, aligning with the drop in correlation due to longer $\Delta \ell$. The mean velocity tends towards zero with increasing $\Delta \ell$, which may be explained by the increased prevalence of velocity outliers, assuming outlier symmetry around

zero. The relatively unaffected median and IQR support this hypothesis, suggesting that inlier statistics are not strongly dependent on $\Delta\ell$. Therefore, using the median for velocity averaging could be considered a more robust metric than the mean. Additionally, applying an outlier removal scheme may prove beneficial for further investigations.

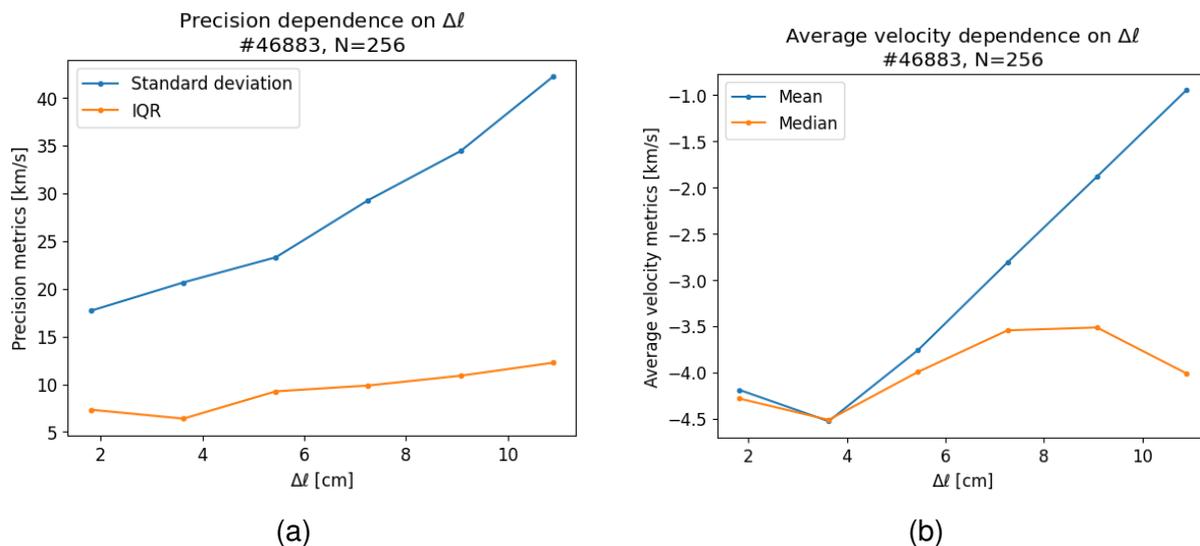


Figure 7.12: CCTDE performance depending on $\Delta\ell$. Results averaged spatially over the entire BES array and temporally from 430-510 ms.

The dependence of velocimetry performance on N is depicted in figure 7.13. Precision improves with increasing N , as anticipated from the testing in section 6.2. However, diminishing returns on IQR improvements are observed beyond $N = 512$ frames. Increasing N also converges the velocity mean and median, but strong diminishing returns are observed beyond $N = 1024$. Therefore, $N = 512$ - 1024 frames is considered the optimal range for maximising accuracy and precision while also maximising the number of independent realisations attainable in this specific time-series.

In summary, the presented results estimate a median velocity of -7 km/s with an IQR of 5 km/s at $N = 1024$ frames and $\Delta\ell = 1.8$ cm. This aligns with the phase velocity estimated in section 7.1.2. However, the precision was found to be too low for meaningful further investigations into temporal velocity variations or coherence analyses. The reduced precision was expected from the testing in section 6.2, but the discovery of large outliers was not previously observed. It was later found that these outliers likely resulted from a lack of spatial localisation of the BES signal, discussed in more detail in section 7.4.

7.4 CCTDE Velocimetry of the CII Impurity Shell

This chapter revolves around a significant challenge posed by the existing experimental data – specifically, *the absence of a substantial beam emission signal recorded in nearly all MAST-U shots by the BES diagnostic*. The examination and interpretation of the remaining background signal captured by BES are detailed in section 7.4.1. In section 7.4.2, this limitation is addressed by constructing a rudimentary synthetic BES

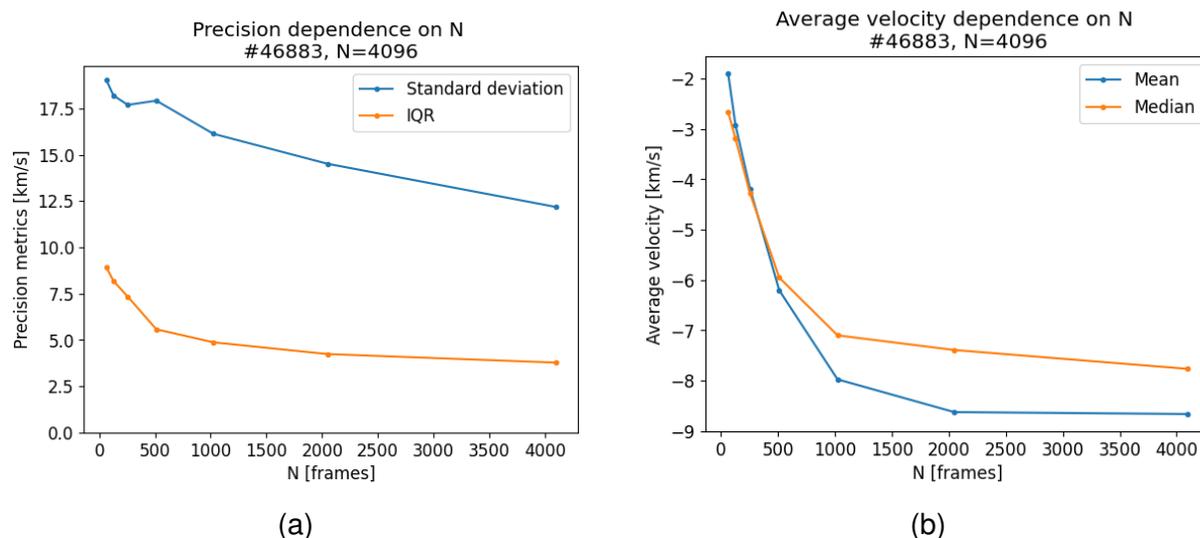


Figure 7.13: CCTDE performance depending on N. Results averaged spatially over the entire BES array and temporally from 430-510 ms.

diagnostic, aiming to calibrate the BES radii for the background emissions. Ultimately, in the cross-comparison of BES velocimetry with CXRS velocities discussed in section 7.4.3, an agreement between the two diagnostics is observed within acceptable error margins. This has never been achieved to this date, and could provide a promising new avenue through which to study carbon impurity dynamics in the future.

7.4.1 Dominant CII emissions in BES signal

Upon plotting the power spectra of the BES signal at different times, which are not shown here, it was observed that there was no discernible difference in BES turbulent signal power whether the Neutral Beam Injection (NBI) was switched on or off³. In other words, the diagnostic was not recording any signal other than the background emissions. This observation persisted across almost all shots in the second MAST-U campaign, including those shots most promising for L-mode velocimetry. Despite this disheartening discovery, this section explores the prospect of conducting velocimetry analysis on the background signal. To my knowledge, this has never been attempted before but could provide useful information, particularly in the study of impurity dynamics.

The prevailing background emissions that dominate the BES signal were thought to originate from Carbon-II (CII), as it was the only known emission line that would pass through the optical filter. The R-z coordinates assigned to each BES spatial channel were not calibrated for these CII emissions, and it can be assumed that the emissions are not accurately localised by the BES diagnostic. This lack of localisation

³The lack of beam signal was due to an unforeseen calibration issue with the BES optical filter. Instead of maximally filtering out background emissions and transmitting the beam emission, the beam emission was entirely blocked by the filter. This issue could easily be fixed through a heating element which shifts the bandpass frequencies of the optical filter, but this adjustment could unfortunately not be applied in time for the writing of this thesis.

poses a challenge for image-velocimetry analysis, which relies on spatially localised measurements.

CII emissivity typically forms a hollow shell in tokamak plasmas, where the temperature is sufficiently high for the first carbon ionisation but not high enough for further ionisation. In search of the shell location, a radial scan of the total BES emission intensity revealed a clear peak in CII emission, as illustrated in figure 7.14a. The observation of a single emission peak was deemed consistent with the location where the BES line-of-sight intersects the greatest volume of CII. This intersection would approximately occur where the BES views the CII shell tangentially or as close to tangential as possible. If this hypothesis is correct, at the location of $R = 1.31\text{-}1.35$ m, the emissions may be radially localised enough to yield meaningful velocimetry results.

CCTDE was applied to shot 47108 during notch times 620-635 ms, where the BES data was frequency filtered between 3-100 kHz. The lower frequency limit was selected for consistency and the upper limit was to remove high-frequency noise. $N = 512$ frames and $\Delta\ell = 1.8$ cm were used for the CCTDE operation. The results, depicted in figure 7.14b, exhibit a distinct velocity shear at approximately 1.34 m. Notably, this shear does not align with the separatrix location predicted by the equilibrium modelling tool ‘EFIT.’ This discrepancy is intriguing because a flow shear region typically coincides with the separatrix. Thus, considering that the BES radii were not calibrated for CII emissions, this misalignment is most likely attributed to inaccurate BES radii rather than an accurate reflection of the true physical picture.

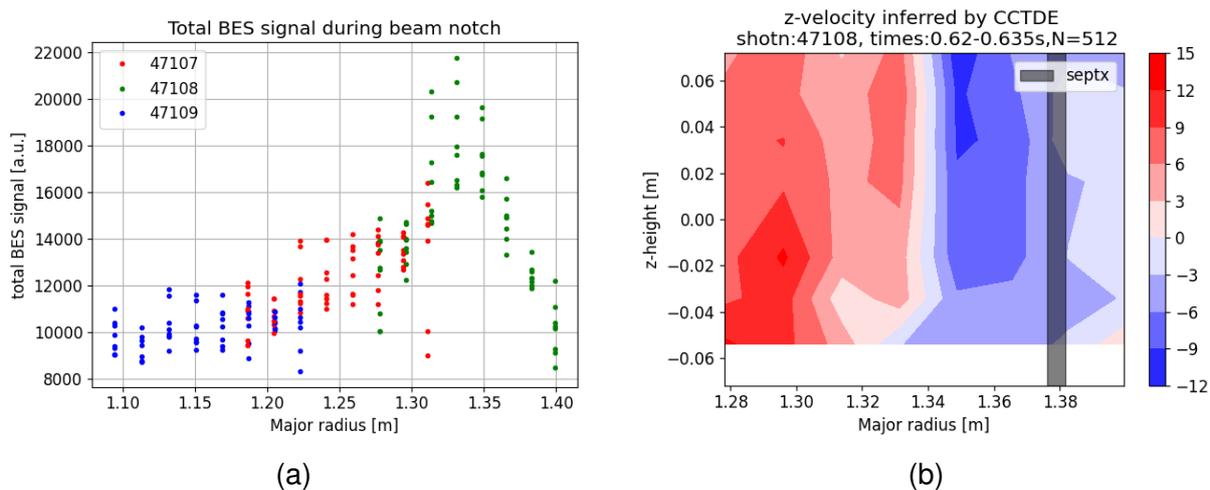


Figure 7.14: BES intensity profile of the CII background emissions shown in (a). The three shots that were used represent repeat shots with near-identical discharges. Velocimetry applied to shot 47108 shows a shear flow in (b). The mismatch of the shear flow location and the separatrix suggests that the BES radii are not accurate for CII emissions.

7.4.2 Calibrating BES Radii for CII Emissions

A rudimentary BES synthetic diagnostic was developed to provide improved estimates of the true radii from which the CII emissions originate, as observed by BES. Firstly,

the theoretically expected CII radial density distribution was estimated using Thomson scattering measurements and the Saha ionisation equation [95]. The result is shown in figure 7.15a, where it can be seen that the CII density may be expected to peak outside the separatrix.

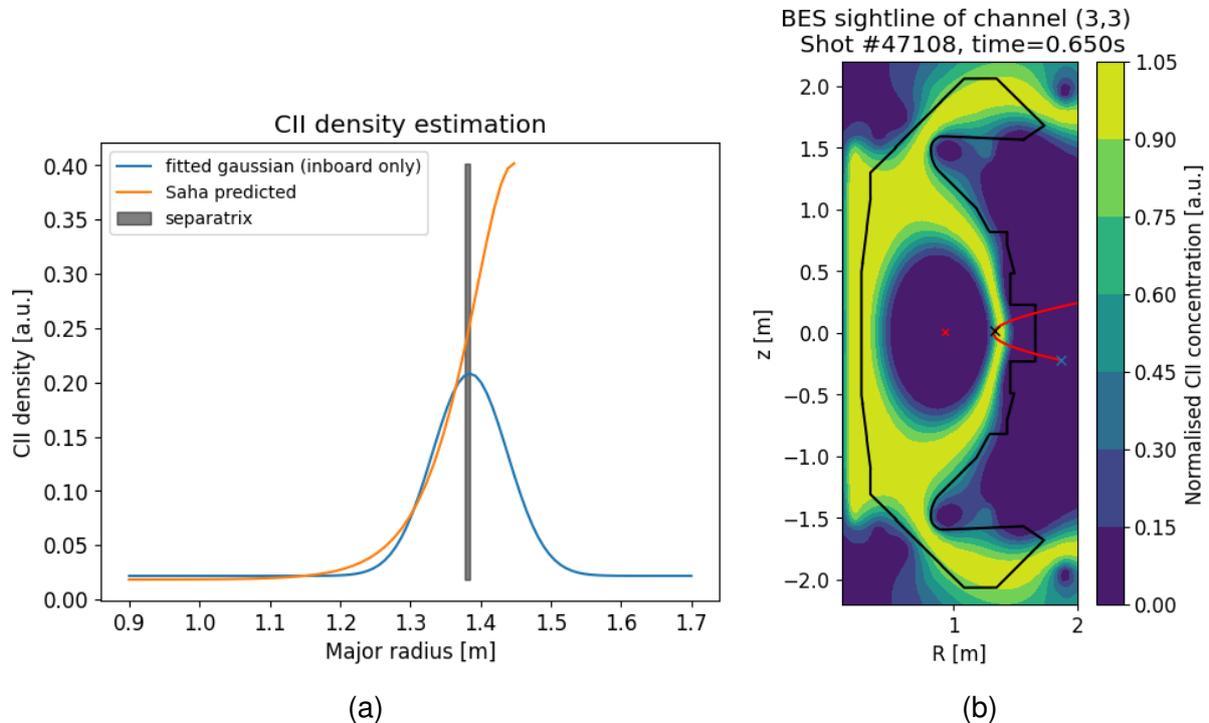


Figure 7.15: (a) shows an estimate of the CII radial distribution in blue. Estimate based on Saha ionisation curve shown in orange and SOL transport considerations that would likely deplete the CII concentration outside the separatrix. (b) shows the CII distribution extrapolated onto the R-z plane assuming equalised density along flux surfaces. BES sight-line of channel (3,3) shown as red line.

Considering the limitations of the Saha equation, which does not account for effects like transport or neutralisation, and recognising that SOL transport would likely rapidly deplete the CII concentration outside of the separatrix, a (slightly) more nuanced distribution was proposed. This distribution was estimated by fitting a Gaussian to the inboard part of the Saha equation fit, with the peak fixed at the separatrix location, as illustrated in figure 7.15a. Note that the precise CII distribution was not considered critical at this point, and the impact of varying the CII distribution was investigated later in this section.

To create a CII density map in the R-z plane, the radial distribution from figure 7.15a was extrapolated onto the MAST-U equilibrium of shot 47108, assuming CII density equalisation along flux surfaces. This assumption was thought to be reasonable near the outboard midplane, the region of interest, but less so in the divertor region. The resulting R-z density distribution is presented in figure 7.15b. BES lines of sight were then projected onto the R-z plane, with an example channel also illustrated in figure 7.15b. By integrating the CII density along the BES lines of sight, estimates were made for the total CII emission intensity observed by each BES channel. This assumes

that fluctuations in CII emission intensity are directly proportional to fluctuations in CII density only, a common assumption in BES data analysis. The result is an approximate image that would be observed by BES, given the emission from the estimated CII density profile, as is shown in the background of figure 7.16.

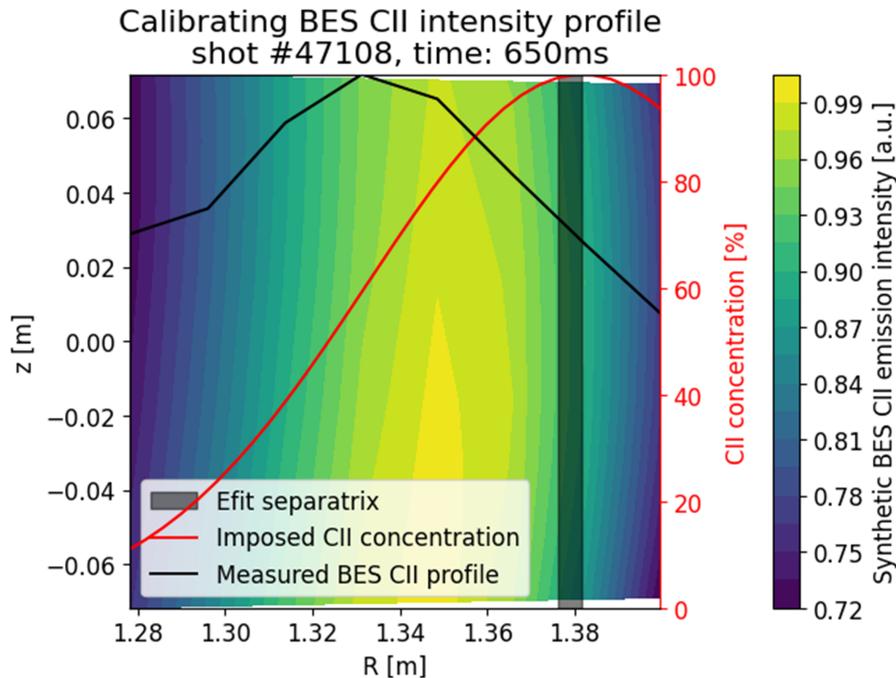


Figure 7.16: CII emission intensity distribution as observed by synthetic BES diagnostic shown in background. This result assumes a true CII distribution that is shown in red. The black curve shows the CII intensity profile as experimentally observed by BES.

The radial location where the intensity peaks in the synthetic diagnostic does align qualitatively, but does not align quantitatively with the experimental BES readings. It is unclear if the quantitative mismatch is due to the crudeness of the synthetic diagnostic or an inaccurate estimate of the true CII distribution. To gain further insight, tests were conducted by varying the peak location and Gaussian width of the theoretical CII distribution by ± 5 cm and $\pm 50\%$, respectively. In all cases, it was found that the synthetic-BES-predicted CII peak was radially *shifted inward* from the imposed CII density peak. The magnitude of this shift was found to increase with an increase of the Gaussian width of the imposed CII distribution, but no significant dependence was found on the imposed CII peak location. The magnitude of the shift was found to range between 1-5 cm.

In conclusion, these findings consistently support that the true CII density distribution is centred radially further outward than implied by experimental BES readings. The direction of this shift is an artefact of the viewing geometry of the BES optics, which was thought to be adequately captured by the synthetic diagnostic. The precise magnitude of this shift could not be confidently estimated due to the otherwise crude assumptions made for the synthetic diagnostic and CII distribution estimates, although values ranging from 1-5 cm were observed. Regardless, it is unlikely that the CII peak occurs outside the separatrix, noting the SOL transport considerations mentioned ear-

lier. Thus, the experimentally observed BES emission peak at $R = 1.33$ m can be conservatively estimated to originate from CII emissions peaking at $R = 1.33$ - 1.38 m.

7.4.3 CII velocimetry cross-diagnostic comparison

The objective of this section was to cross-compare the velocities estimated by BES at $R=1.33$ m with CXRS velocities at $R=1.34$ - 1.37 m, assuming a CII shift of 1-4 cm. CCTDE was applied to the entire L-mode portion of shot 47108 spanning 500-800 ms. Density fluctuations were frequency-filtered between 3-100 kHz, and CCTDE parameters $N=512$ frames and $\Delta\ell=1.8$ cm were chosen. Outlier velocities, differing from the long-term median velocity by more than twice the IQR, were removed. Next, the median of the velocities across all z -locations were taken at each time-point. Subsequently, the velocities - which pointed in the z -direction - were projected onto the diamagnetic drift direction (cross-field direction) using equation 7.2.

$$v_{\perp,BES} = -\frac{B_{\phi}}{|B|}v_{z,BES} \quad (7.2)$$

The resultant velocities are illustrated in figure 7.17. As anticipated from section 7.3, long-term averages were necessary for velocity convergence, with the standard error after averaging 10 ms segments being approximately 2 km/s.

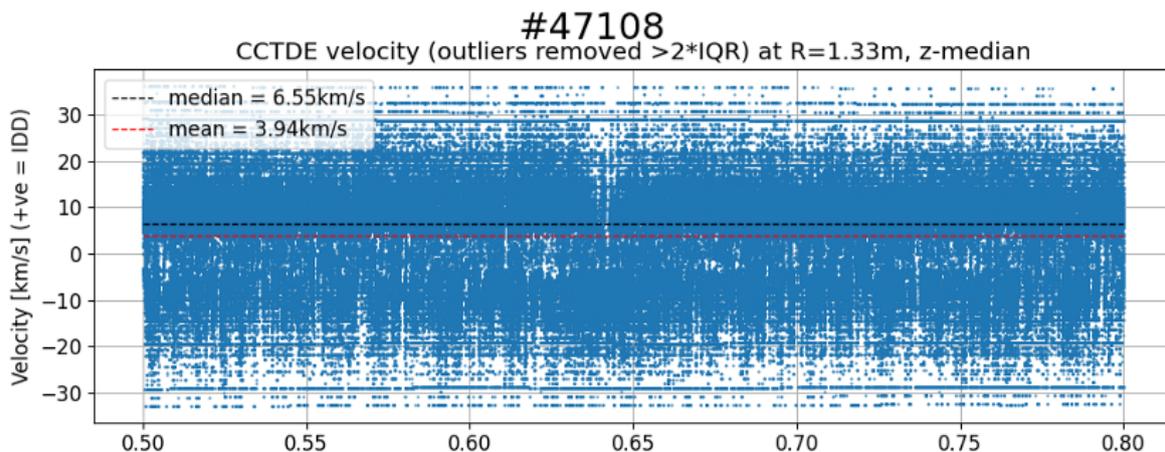


Figure 7.17: CCTDE-inferred velocities in the ion drift direction. Standard deviation of the velocities was 15 km/s. Averaging over 10 ms converges standard error to ~ 2 km/s. Time shown in seconds on the x-axis.

Cross-comparison between the diagnostics is not straightforward because the BES-inferred velocities point in the machine- z (approximately poloidal) direction, while the CXRS velocities point in the toroidal direction. As such, the ExB velocity will be estimated separately using both diagnostics since this quantity should be identical regardless of the diagnostic used. For both diagnostics, their inferred velocities were projected onto the diamagnetic drift direction and by subtracting the main species diamagnetic velocity, the ExB velocity could be estimated⁴. The expressions used to

⁴This assumes that the turbulent phase velocity observed by BES image-velocimetry techniques is

estimate the $E \times B$ velocities are detailed in equations 7.3 and 7.4, and a derivation can be found in section 10.1.

$$v_{E \times B, CXRS} = v_{*,e} - \frac{B_\theta}{|B|} v_{\phi, CXRS} \quad (7.3)$$

$$v_{E \times B, BES} = v_{*,e} - \frac{B_\phi}{|B|} v_{z, BES} \quad (7.4)$$

It is important to note that the derivation of equations 7.3 and 7.4 assumes that neo-classical transport theory holds. If this is not the case, a more detailed analysis can be followed as is shown, for example, in experimental measurements in C-Mod[96]. Regardless, the assumption was considered sufficient for the initial comparisons made in this thesis. On a separate note, the electron diamagnetic velocity was chosen over the ion velocity because electron profiles are more easily experimentally obtained, but the choice should not affect subsequent observations and conclusions.

The electron diamagnetic velocity, $v_{*,e}$, was calculated using the pressure profile obtained from combining the temperature and density measurements from Thomson scattering. Magnetic field strengths were obtained from the equilibrium reconstruction tool EFIT. The resultant $E \times B$ velocities for both BES and CXRS are presented in figure 7.18.

The two CXRS-predicted velocity bounds were based on the uncertainty of the BES radii with CII emissions. It is observed that the BES-predicted velocity largely falls between these bounds and both diagnostics predict the same trend of increasing $E \times B$ velocity with time. A notable exception is a large velocity excursion at 620-650 ms, which was likely due to transient flows generated by a NBI beam notch.

In conclusion, the fact that CII velocimetry agreement was observed between the diagnostics is a significant feat that, to the knowledge of the author, has never before been achieved. This approach to velocimetry of the CII shell may provide valuable information in the study of impurity transport and dynamics, which is an active and impactful field of research. Improvements are suggested in BES calibration for CII emissions or the development of a full BES synthetic diagnostic in order to further improve the reliability CII velocimetry with the BES system. In short, this proof-of-concept workflow has yielded results that suggest CII velocimetry is possible with BES, and a significant scope for improvements makes this a promising avenue for further exploration.

7.5 Discussion and Conclusion

The investigation of shot #46883 involved an examination of BES data, revealing density structures with typical poloidal coherence lengths of 5.5 ± 0.5 cm and radial coherence lengths of 3.7 ± 0.4 cm (see section 7.1). In both directions, correlation lengths of the density features were approximately 8-10 cm. The poloidal phase velocity, estimated by $k_\theta - f$ spectra, ranged from -6 to -7 km/s at turbulent frequencies between

equal to the bulk diamagnetic velocity. Similarly, for CXRS, the Carbon diamagnetic velocity is assumed to be equal to the main-species diamagnetic velocity. Although these assumptions do not necessarily hold, any resultant velocity mismatch is not expected to exceed the magnitude of the main-species diamagnetic velocity.

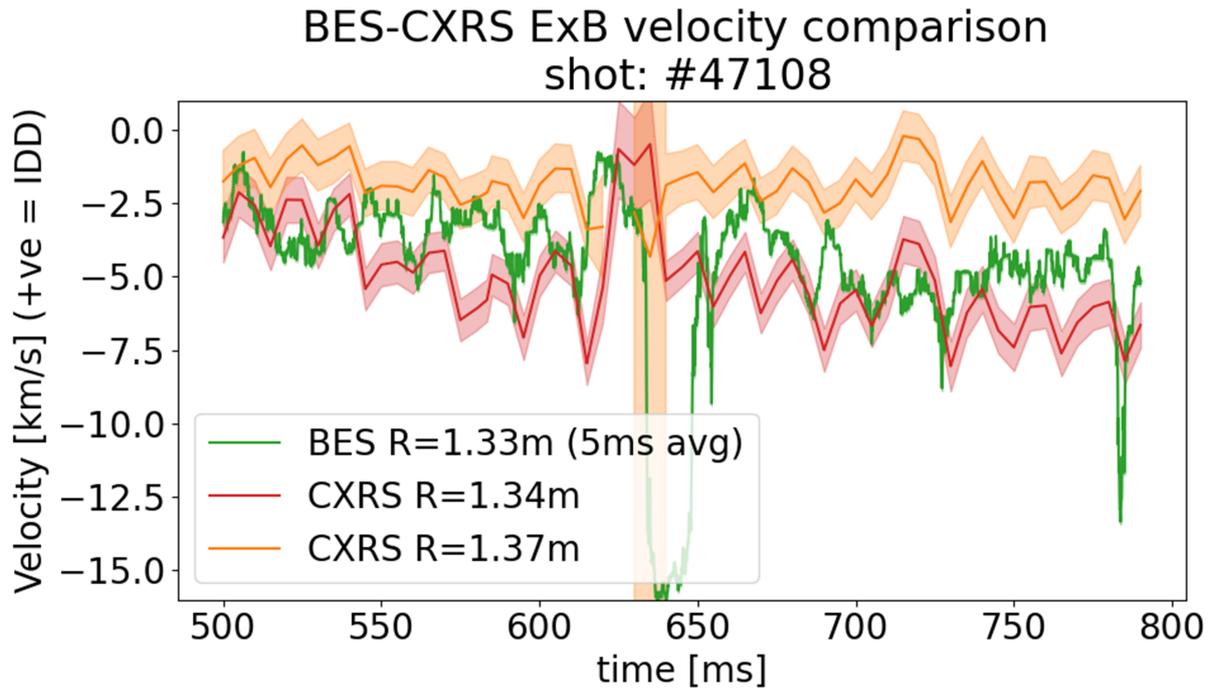


Figure 7.18: Estimated ExB velocities using the BES and CXRS diagnostics. Steady L-mode plasma present throughout times presented, except for an NBI beam notch at 620-640 ms that may cause transients. BES velocities with 5 ms averaging time resulted in standard errors of 2-3 km/s throughout the presented data.

3-50 kHz. The SNR was assessed to be around 1-3, and density features were observed to decorrelate within $15 \mu\text{s}$ or 59 frames. The estimation of these parameters, coupled with the testing in chapter 6, suggested that performing accurate velocimetry analysis on this data was feasible.

Velocimetry utilising the DTW algorithm was conducted in section 7.2. One challenge that was encountered was that the velocities estimated by the $k_\theta - f$ spectra were an order of magnitude too slow and DTW velocity inferences were predicted to be inaccurate. However, this issue was mitigated by introducing the ‘shift parameter’. Initial findings revealed additional challenges at the bottom edge of the BES field-of-view. These edge effects were expected, but strongly exacerbated by the limited number of spatial channels available to the DTW. Figure 7.9 depicts strong velocity underestimation in the bottom half of the BES field-of-view, while the top half exhibited rough but somewhat robust agreement with the anticipated velocities. This implies that DTW can yield reasonable velocity estimates, but further investigation is recommended for its performance in this low-velocity regime and near the edges of the field-of-view.

Initial CCTDE velocimetry analyses were conducted in section 7.3. After preparatory analysis and initial testing, optimal parameters were determined to be $N = 1024$ frames and $\Delta\ell = 1.8$ cm. With these parameters, the mean velocity was found to be -8.0 ± 0.4 km/s, roughly in line with expected velocities. However, this result required averaging over the entire 80 ms segment, which was due to the high standard deviation of 15.1 km/s. This high standard deviation was primarily attributed to large (± 100 km/s) outliers that were not observed in the testing in chapter 6.

The aforementioned outliers were attributed to localisation issues detailed in section 7.4. Specifically, the BES signal was dominated by background carbon emissions instead of the intended neutral beam emissions. Upon further inspection, it was determined that the background carbon emissions existed in a radially localised shell observable by the BES diagnostic. This raised the possibility of accurate velocimetry, granted that BES views the carbon shell tangentially. After ‘re-calibrating’ the BES radial coordinates for the carbon emissions, poloidal CCTDE velocimetry was performed on the radially localised carbon shell emissions.

To improve confidence in inferred velocities, a cross-comparison was undertaken with the charge-exchange-recombination spectroscopy diagnostic. Figure 7.18 illustrates that the magnitudes and temporal trends of the estimated ExB velocities agreed between the two diagnostics. This promising outcome suggests that the method is applicable to accurate velocity estimation of background carbon impurities, an achievement not previously realised and could prove valuable for impurity transport studies.

Chapter 8

Velocimetry Experiments with Localised BES Signals

This chapter investigates the only shot in the second MAST-U campaign that is suitable for main species velocimetry analysis - shot #46459. A preliminary analysis can be found in section 8.1, where the long-time (>10 ms) velocity trends are explored and compared with the CXRS diagnostic. In section 8.2, the focus shifts towards higher precision velocimetry at a faster inference frequency with the implementation of the line-CCTDE method.

8.1 Slowly Varying ($>$ ms) Velocity Comparison with Charge-Exchange Recombination Spectroscopy

In this section, velocimetry analysis was conducted on BES data presumed to originate from localised emissions. This became feasible following the identification of a sole shot (in MAST-U campaign 2) exhibiting substantial NBI beam emissions recorded by BES. CCTDE velocimetry analysis was then applied to this localised signal, presenting novel challenges that had not been encountered in chapter 6. Additionally, a cross-diagnostic comparison was undertaken between the $E \times B$ velocities predicted by the BES and CXRS diagnostics.

8.1.1 Initial Look at the Data

The BES system could use two different optical filters depending on the region of the plasma that was being imaged. The optical filtering issues outlined in section 7.4 were specific to the edge optical filter, which was used in the majority of experiments. In contrast, ten shots in MAST-U campaign 2 utilised a differently calibrated core filter, potentially recording significant levels of beam emission. Of those ten shots, nine experienced immediate or early disruptions, leaving only shot #46459 as a possible candidate for velocimetry.

In shot #46459, both SW and SS neutral beams were employed. Upon activating both beams, the plasma transitioned into an ELMy H-mode lasting from 210 to 270 ms.

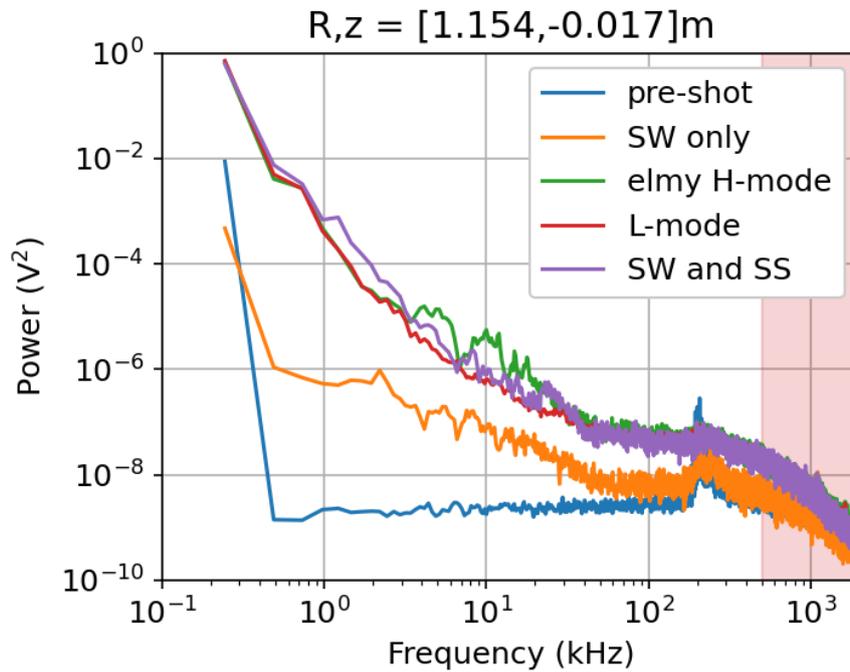


Figure 8.1: Power spectra of the BES signal at different times in shot #46459. Pre-shot: -100 to -1 ms, SW only: 100-130 ms, SW and SS: 140-180 ms, ELMy H-mode: 210-270 ms, L-mode: 350-550 ms.

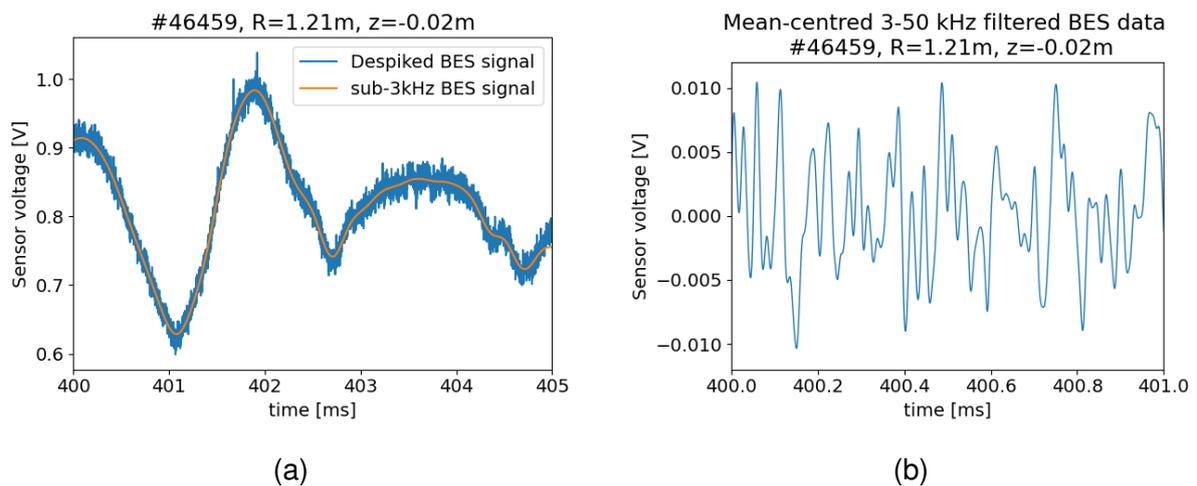


Figure 8.2: a) 'despiked' raw BES signal in a single channel. Signal spikes due to photon noise filtered out based on large spikes in signal gradient. sub-3kHz signal represents fluctuations primarily due to NBI fluctuations. b) 'cleaned' BES signal after despiking, frequency filtering and mean centring.

Subsequently deactivating the SW beam, thereby reducing plasma heating, led to a transition back to an L-mode from 350 to 550 ms until disruption.

Figure 8.1 displays the power spectra of the BES signal at different times in shot #46459. When the SS beam was activated, there was an almost order of magnitude increase in signal amplitude across frequencies of 5-100 kHz. This was in contrast to all other shots in the second campaign, where no such increase was observed. Additional investigations, using the same method covered in section 7.1.2, estimated a SNR ranging from roughly 1 to 5, depending on the time and channel. These findings indicate the presence of significant levels of beam emission signal in this particular shot¹.

Figure 8.2 shows example ‘despiked’ and ‘cleaned’ data from a single BES channel. Raw data contained large signal spikes due to neutrons striking the detectors. These were filtered out by removing datapoints that resulted in signal gradients that were orders of magnitude larger than the median gradient. Resultant ‘despiked’ but otherwise raw data is shown in figure 8.2a. Sub-3kHz signal fluctuations highlighted as resulting from fluctuations in NBI power. Figure 8.2b shows fully ‘cleaned’ data ready for velocimetry analysis. Raw data was despiked, frequency filtered between 3-50 kHz and mean-centred. Lower frequency limit determined by NBI fluctuation frequency. Upper frequency bound defined by electronic noise frequencies. Mean-centred to eliminate channel-to-channel signal variations.

The analysis of conditional k_{θ} - f spectra, though not presented here, unveiled two counter-propagating flows during the H-mode at 210-270 ms, with phase velocities of +15 km/s and -20 km/s, occurring at frequencies between 15 and 50 kHz. Throughout the L-mode at 350-550 ms, a singular flow component was observed at frequencies of 3-50 kHz, exhibiting an approximate phase velocity of -7 km/s.

Operational parameters for CCTDE were determined following the workflow described in section 7.3.1. It was estimated that the decorrelation timescale was around 16 μ s, or 62 frames, throughout the entire shot. The density fluctuation structures extended to 6-10 cm. These results, in conjunction with SNR estimates, guided the selection of $N = 512$ frames and $\Delta\ell = 1.8$ cm for accurate CCTDE velocimetry.

8.1.2 CCTDE Velocimetry Comparison with CXRS

Using the selected parameters of $N = 512$ frames and $\Delta\ell = 1.8$ cm, CCTDE was applied to times 150-550 ms. At times 150-350 ms, the data was frequency filtered to isolate the turbulent fluctuations with a 15-50 kHz bandpass filter. At times 350-550 ms, the frequencies 3-50 kHz were used. Velocities were determined in the z-direction across the full BES field-of-view.

Example velocity inferences, shown in figure 8.3, represent the typical results during H-mode (8.3a) and L-mode (8.3b). In both figures, it can be seen that the raw velocities have large variance due to ‘singularity-like’ phenomena, where velocities shoot off to infinity intermittently. These singularities could not easily be removed or circumvented, and their origin could not confidently be pinpointed². Regardless, the singularities were

¹This outcome was warmly received after the numerous experimental setbacks discussed in chapter 7

²Singularity frequency was found to increase with decreasing N and also with decreasing $\Delta\ell$. Tenuous

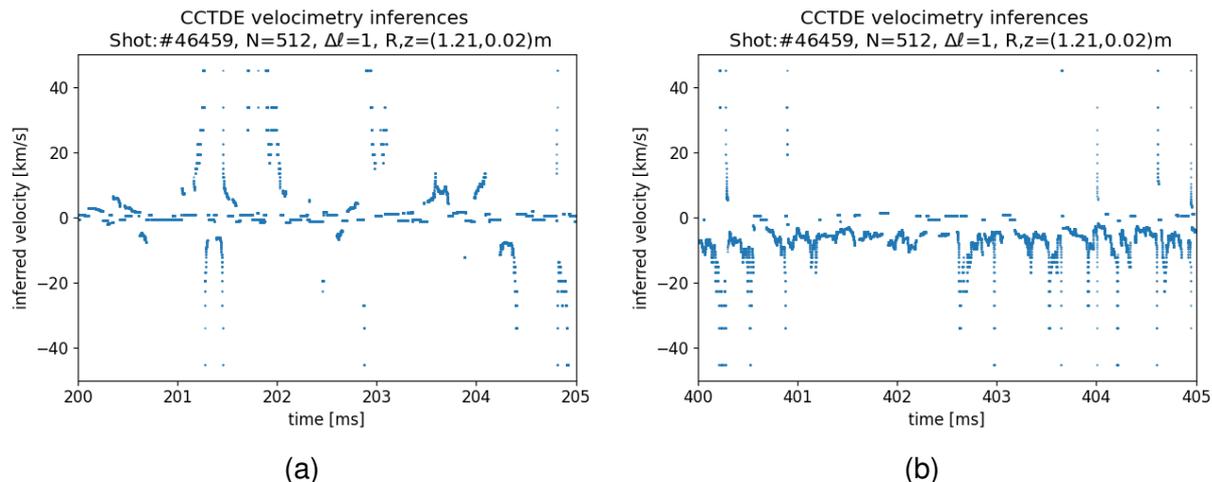


Figure 8.3: Example CCTDE velocimetry results in H-mode (a) and L-mode (b). Velocities point in the machine-z direction (approximately antiparallel to poloidal direction). CCTDE was expected to underestimate velocities by approximately 20% - 50%, as determined using (unshown) results from section 6.2.

statistical outliers and could be averaged out, which significantly reduced the effective measurement frequency, but an optimised measurement frequency was not required for this study anyway.

The following workflow was followed to process and clean the raw velocities:

1. In the H-mode segment, positive velocities were replaced with NaNs to isolate the mode propagating in the negative direction³. This was not necessary for the L-mode portion.
2. According to Equation 6.8, the minimum measurable velocity for these tests is 0.9 km/s. Thus, all velocities below 0.9 km/s in magnitude were replaced with NaNs as they are likely false positive measurements.
3. A rolling median was taken over the **reciprocal** of the velocities. Taking the reciprocal is equivalent to averaging over the time-lags determined by CCTDE and is necessary to avoid bias towards higher velocities.

Next, the ExB velocities were estimated using both the BES and the CXRS diagnostics, following the method described in section 7.4.3. Results can be seen in figure 8.4.

8.1.3 Discussion and Conclusion

The precision and effective temporal frequency of velocity inferences were drastically reduced by the velocity-singularities depicted in figure 8.3. Although the origin of the

links were found between singularities and the lack of signal peaks in the BES segments, but this could not fully explain singularity occurrence. The peak correlation amplitude of the cross-correlation functions did not show any links to the singularities. Large radially propagating features could also theoretically explain these singularities but these were not observed in the data. Notably, the singularity times varied significantly, although nonlinearly, with the z-location.

³Crude. Yes. Please refer to the section 8.1.3 for further discussion.

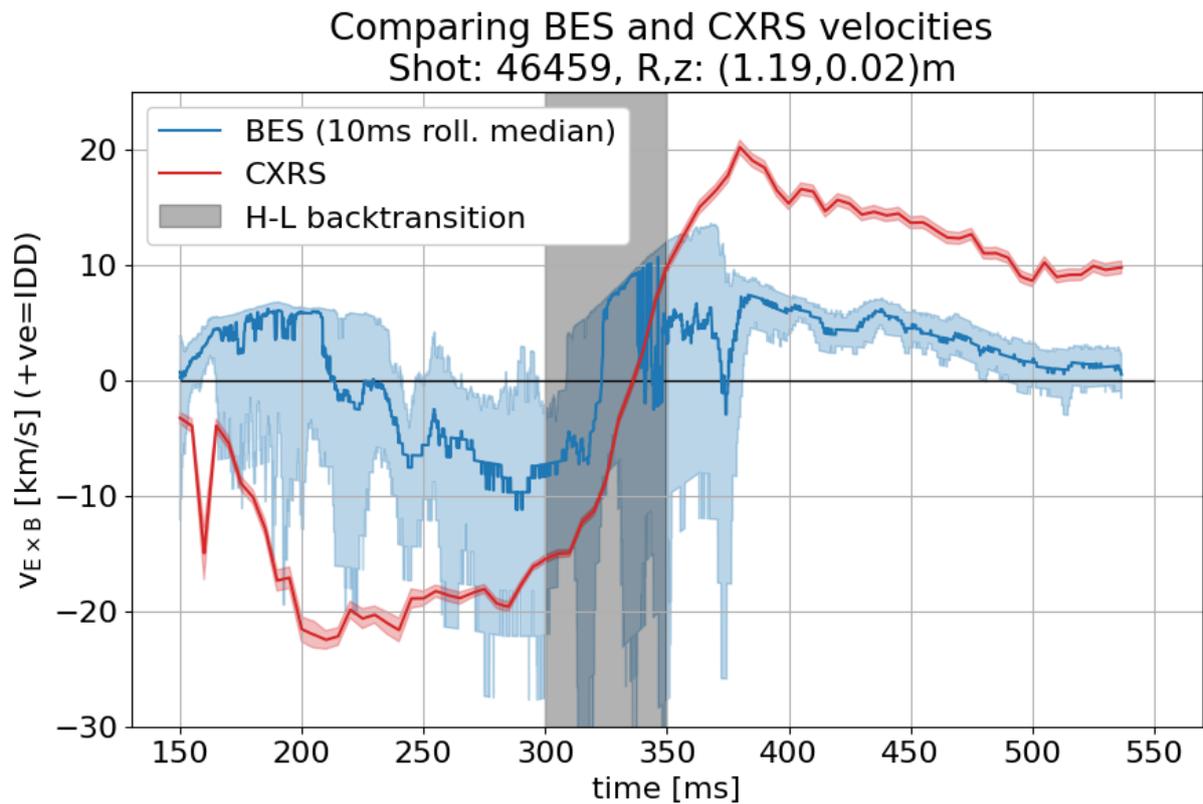


Figure 8.4: $E \times B$ velocities estimated using both BES and CXRS diagnostics. Error margin on the BES rolling median represents the inter-quartile-range of the cleaned CCTDE velocities. $E \times B$ velocities calculated using electron diamagnetic velocities and Equation 7.2.

singularities could not be pinpointed, a leading hypothesis suggests a link to the intermittency of large (3-6 channel) density features propagating through the BES field-of-view. This hypothesis was derived from visual inspections of the BES data, where singularities were exclusively observed during the absences of large density structures. During these times, velocities of the smaller (1-2 channel), rapidly decorrelating density structures are attempted to be determined by CCTDE. Regardless, the singularity behaviour was not observed in any of the testing in section 6.2, and further work is advised to explain this phenomenon.

The velocities presented in figure 8.4 demonstrated approximate agreement in the overall trends, although velocities estimated by the CXRS diagnostic typically exhibited larger magnitudes. This discrepancy could potentially be explained by predictions in chapter 6.2, where CCTDE was estimated to underestimate velocities by 20-50% in this regime.

Agreement in the velocity trend could be most clearly observed during the L-mode. During this time, relatively high precision was observed in the CCTDE velocity IQR of 2-3 km/s. However, this does not necessarily imply high accuracy of the CCTDE velocities, as cautioned in section 6.2.

The magnitudes of the velocities did not show agreement between the diagnostics within uncertainty margins. In addition to the CCTDE velocity underestimation of 20-50%, this mismatch may be due to assumptions regarding the non- $E \times B$ velocities picked up by the diagnostics. In this analysis, it was implicitly assumed that the turbulent phase velocity picked up by BES, v_{ph} , was equal to the Carbon diamagnetic velocity relevant to CXRS, v_C^* . Although a common approximation, this assumption does not generally hold, but the corresponding mismatch is not expected to exceed the main species diamagnetic velocity — in this case, 5-15 km/s. Notably, the $E \times B$ velocity estimates mostly agree within this additional error margin of ~ 15 km/s.

On top of the previous complication, it is not unlikely that the dominant turbulent mode changed during the H-L transition, resulting in a change of the turbulent phase velocity. This would further complicate the comparison of the non- $E \times B$ velocities between the diagnostics. Gyrokinetic simulations could be performed to identify the dominant turbulent modes throughout this shot, uncovering any possible mode changes and potentially explaining the discrepancy between the $E \times B$ velocity estimates. Unfortunately, this analysis fell outside the scope of this thesis.

It is worth noting that this shot, along with all other shots in the second MAST-U campaign, was not able to generate a high rotation plasma. Specifically, the toroidal rotation did not exceed 50 km/s in shot #46459, while values of 150 km/s were regularly achieved in pre-update MAST. This is relevant because the faster the rotation, the less significant the non- $E \times B$ flow components become in the velocity measurements by BES and CXRS. This means the assumptions made for the turbulent phase velocity and the carbon diamagnetic velocity would be much less critical to the cross-diagnostic comparison in the high rotation case, drastically simplifying the required analysis.

Finally, an improvement can be made to the velocity cleaning during the H-mode section. Instead of isolating the negative flow mode by manually removing the positive velocities post-CCTDE, the positive flow mode could be cut from the data in Fourier space prior to CCTDE analysis. This approach would likely yield higher precision and more statistics of the CCTDE inferences.

In conclusion, approximate agreement was found in the temporal trends of the velocities predicted by both the BES and CXRS diagnostics. The trends were especially converged during the L-mode, but the velocity magnitudes did not agree within uncertainty margins. Avenues for investigating the magnitude mismatch, along with general improvements, were suggested for future work.

8.2 Tentative Zonal Flow Observation with Line-CCTDE Velocimetry

In order to improve precision and circumvent the singularity issue seen in section 8.1, the more robust ‘Line-CCTDE’ method was employed. This method can achieve improved precision and exclude large outliers at the cost of sacrificing spatial resolution in the z-direction (or R-direction if desired). The line method builds upon the two-point approach, and the accuracy estimations found in chapter 6 are still directly relevant to the line method. The line method workings can be summarised as follows:

1. A reference channel (i, j) is chosen.
2. The reference channel is individually cross-correlated with all eligible other channels along one of the orthogonal directions. Only channels separated by less than $\Delta\ell_{max}$ are included.
3. The cross-correlation function peaks are calculated for each pair of channels, and time delay and distance between channels are stored.
4. A velocity and error are fitted through the time-delays and distance data points using a RANSAC estimator, as shown in figure 8.5. RANSAC was chosen over least-squares for improved outlier handling.
5. Steps 1-4 are repeated with a sliding time window to produce a temporally resolved velocity estimation.

Line-CCTDE can be applied to data using much of the same operational parameters as two-point CCTDE, with the only difference being that the $\Delta\ell$ parameter is replaced by a $\Delta\ell_{max}$ parameter, defining the maximum $\Delta\ell$ included in the analysis.

The L-mode portion of shot #46459 was analysed, covering times from 350 to 550 ms and frequencies from 3 to 50 kHz. A segment length of $N = 1024$ frames was predicted to be optimal based on a preliminary analysis similar to the one shown in section 7.3.1. The $\Delta\ell_{max}$ parameter was set to 10 cm, which was below the decorrelation length and the corresponding minimum measurable velocity was 3 km/s.

The Line-CCTDE method was applied with the aforementioned parameters, and the initial results, shown in figure 8.6, revealed typical uncertainty margins of <1 km/s. These raw inferred velocities were further processed by removing outliers in reciprocal space that were more than twice the inter-quartile range away from the long-term reciprocal median. Inspection of the velocity distributions showed that primarily positive- and low-velocity anomalies were removed through this scheme.

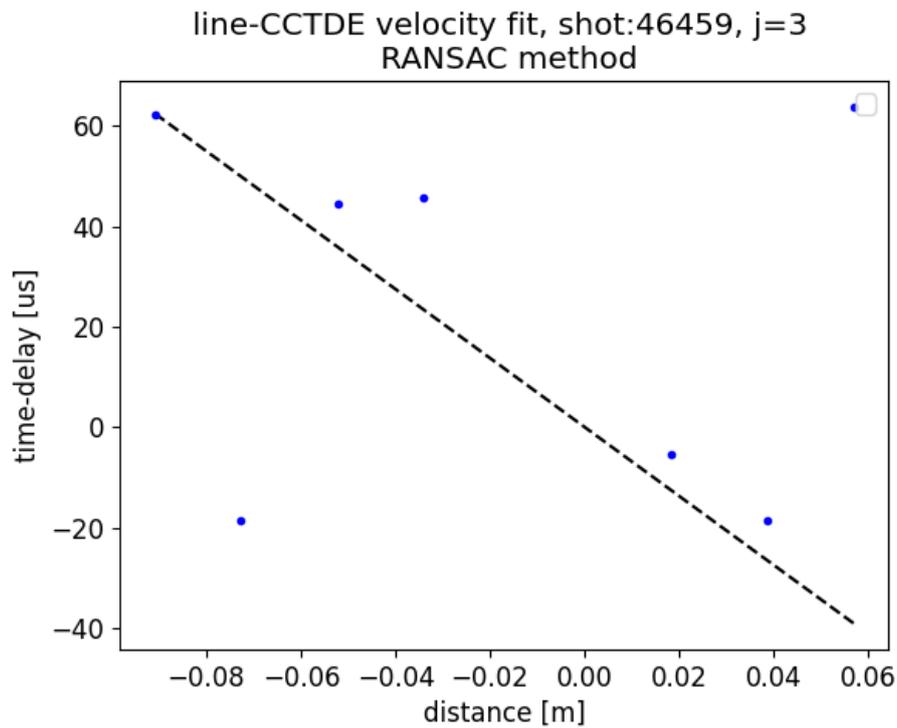


Figure 8.5: A Random Sample Consensus (RANSAC) estimator fit of the velocity. The RANSAC method estimates outliers in the data (in this case at -0.07 cm and $+0.06$ cm), and then estimates the best fit according to only the inliers. Gradient of the line of best fit corresponds to the velocity.

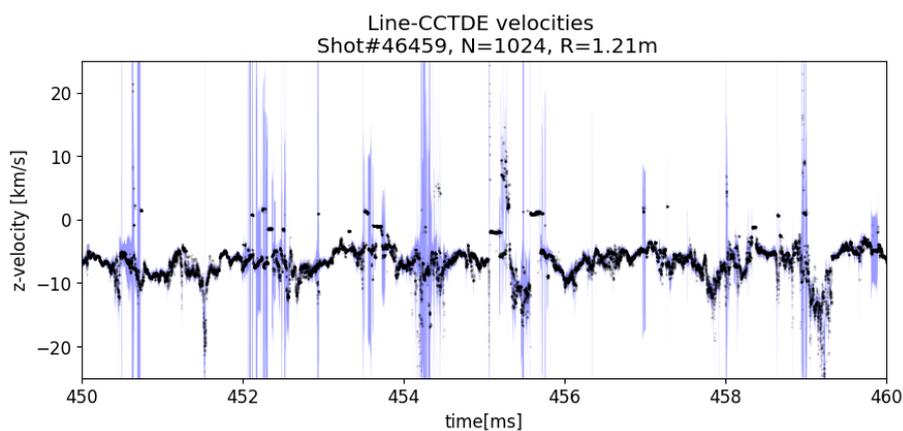


Figure 8.6: Example velocity estimates by line-CCTDE. Uncertainty margins - shown by blue shading - fall below 1 km/s on average.

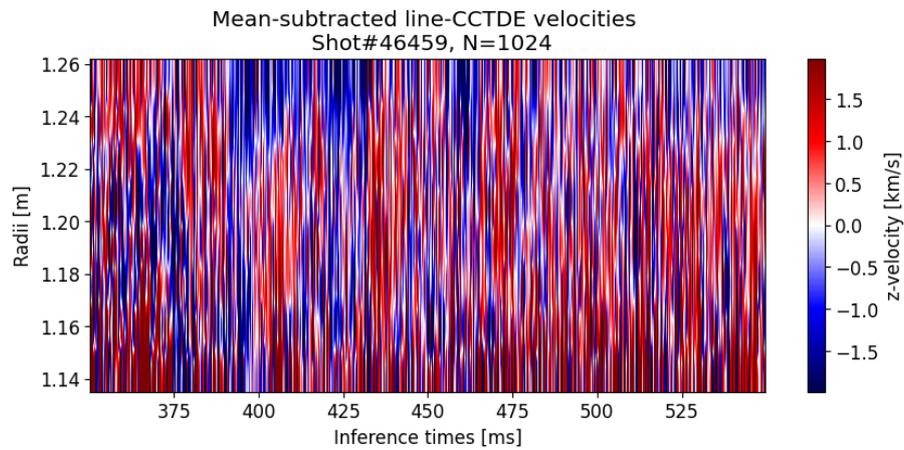


Figure 8.7: Radial-temporal plot of velocity inferences of L-mode plasma. Radii refers to the major radius. Mean velocity was -8.1 km/s in the z-direction.

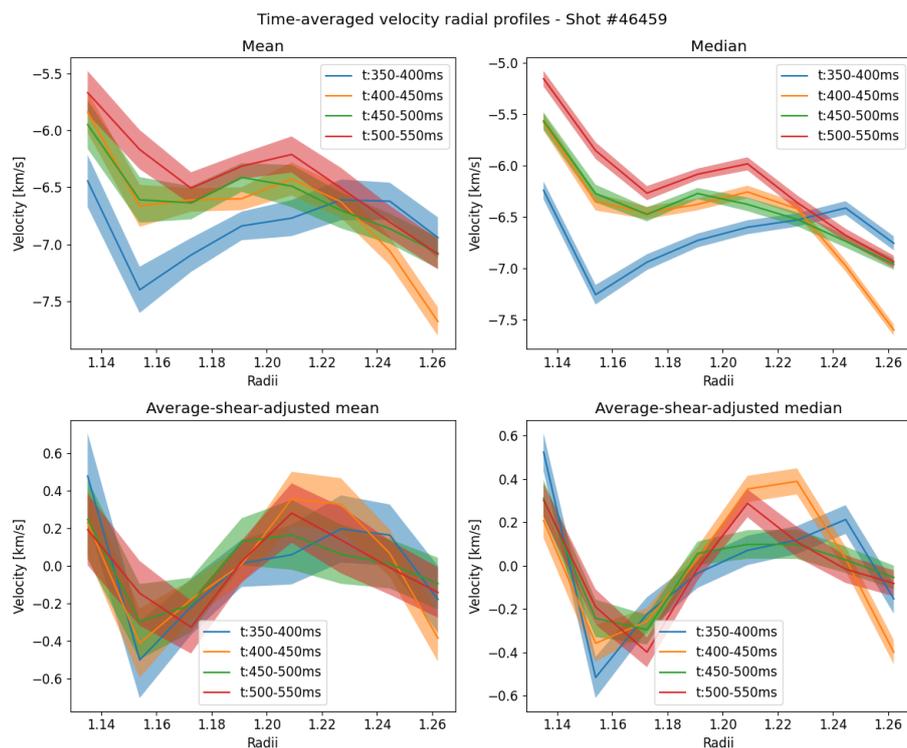


Figure 8.8: Major-radial profiles of the L-mode velocities in shot 46459. Uncertainty margins represent the standard deviation. Average shear estimated by least-squares routine.

The results, presented in figure 8.7, display the radio-temporally resolved velocities in the z-direction. These velocity inferences could be determined with an effective frequency of 4 kHz and an uncertainty margin of <1 km/s. Radially localised structures are clearly discernible here. In addition, radial drifting and intermittent re-formation of the structures are observed, both expected from zonal flow theory. Upon temporal averaging, as shown in figure 8.8, an average shear in the velocities becomes evident, with faster velocities on the radially outboard side. Furthermore, underlying the average shear is a velocity sinusoid with consistent wavenumber and phase. These observations were found to be robust with variations in segment length, N , $\Delta\ell_{max}$, time, and outlier removal scheme. Although not shown here, coherence analysis suggested that this sinusoid had a high radial coherence of 0.65, peaking at 60 Hz, with no other coherent modes observed in the velocity data.

The radial wavelength and coherence time of the quasi-stationary zonal flow are expected to lie around ~ 10 cm and $>ms$, respectively. As such, **the observed velocity sinusoid is consistent with quasi-stationary zonal flows described in section 3.2.** This is a novel result that has never before been observed in MAST(-U).

Future work is recommended to further investigate this flow mode:

- nonlinear flow-turbulence analysis should be performed to determine if the flow mode is nonlinearly driven by the turbulence.
- The short-time (0.1-10 ms) behaviour of the velocities could also be further investigated through the Line-CCTDE method.
- In-depth comparison with DTW velocimetry and also CXRS is recommended to test robustness.
- Repeat analyses of shots at varying radii, collisionality, confinement mode, etc.

These recommendations are thought to lie well within the performance capabilities of line-CCTDE. Additionally, at time of writing of this thesis, the third MAST-U campaign has concluded with many shots, including a 12 shot radial and toroidal field strength scan, ideal for performing velocimetry analysis. Unfortunately, the third MAST-U campaign shots fell outside the scope of this thesis due to time restrictions, but further velocimetry analysis of these shots is highly encouraged.

In conclusion, line-CCTDE successfully inferred velocities with high precision and high inference frequency from the BES data in shot #46459 (<1 km/s uncertainty at 4 kHz). Emerging from these velocity inferences were shear structures that were reminiscent of quasi-stationary zonal flows. The structures were observed to drift radially and re-form intermittently, which is expected behaviour of zonal flows. The long-time (>50 ms) average of the flows revealed a velocity sinusoid that was consistent with the wavenumber, frequency, and amplitude expected from zonal flows.

Chapter 9

Summary and Discussion

The first aim of this thesis: ‘to improve the reliability and performance of spatially resolved plasma dynamics measurements’, was successfully achieved. Through the testing in chapter 6, and applications in chapters 7 and 8, velocimetry techniques have been applied with improved confidence and performance. The work in this thesis is thought to provide a solid foundation for future image velocimetry endeavours, as is further summarised in section 9.1.

Unexpected challenges were encountered with the BES diagnostic in chapter 7, rendering velocimetry of the main-species impossible. Regardless, novel velocimetry was successfully performed of the carbon impurity ions, as is further described in section 9.2.

The second aim of this thesis: ‘to measure zonal flows in the upgraded Mega-Ampere Spherical Tokamak’, was also addressed successfully. This marks the first zonal flow measurement in MAST-U, which is expanded upon in section 9.3. Additionally, suggestions for future work are discussed in section 9.4.

9.1 The relevance and impact of velocimetry testing

Chapter 6 focuses on the performance testing of the two main image-velocimetry techniques employed in fusion research: Cross-Correlation Time-Delay Estimation and Dynamic Time-Warping. The accuracy and precision of both techniques were tested using synthetic data, spanning the majority of reasonably achievable experimental conditions. The intention was for this work to serve as a solid foundation to guide the implementation and interpretation of experimental image-velocimetry.

A broad range of novel, and sometimes unexpected, results were uncovered in the tests. For example, CCTDE’s paradoxical accuracy dependence on underlying velocity was uncovered, CCTDE’s minimum measurable velocities were quantified, the thresholds at which the barber pole illusion becomes significant were quantified. The critical DTW performance dependence on the direction of the underlying flow was uncovered, and DTW’s performance with a reduced number of spatial channels was investigated.

One limitation of this research was that decorrelation was not simulated in the synthetic data and therefore the corresponding impact on velocimetry performance could not be determined. Decorrelation was not thought to be a major issue for the experimental application of either technique, as long as the operational length- and

timescales of the methods were kept significantly below decorrelation length- and timescales. Although exact quantified thresholds are undefined, the impact of decorrelation on the experimental results in chapters 7 and 8 was thought to be small, if not negligible, as justified in section 6.6.

In-homogeneous and rotational velocity fields were generally not included in the testing either. The impact of such fields was, however, thought to be limited; CCTDE performance may only be expected to be impacted if the in-homogeneity or curvature lengthscales are comparable to $\Delta\ell$. In all other cases, CCTDE performance is thought to be accurately represented by the synthetic data tests. DTW has previously been shown to perform accurately with rotational velocity fields [83], as is expected based on the inner workings of the method.

One main conclusion is that **both CCTDE and DTW can display highly nonlinear and unexpected behaviour in their accuracy dependence**. This means that previous small-scale velocimetry test results should **not** be extrapolated to broader parameter regimes. The testing results in this thesis cover a broad parameter range that should be applicable to most conceivable experimental data not only from Beam Emission Spectroscopy and Gas-Puff Imaging diagnostics, but also from any diagnostic in which structures propagate across the field-of-view.

Without careful referral to these tests, image velocimetry cannot be relied upon to produce accurate results. High precision was explicitly found to be a bad proxy for accuracy, meaning that no intrinsic metrics are available to estimate uncertainty margins. This is why careful reading of chapter 6 is **necessary** for those wishing to reliably perform image velocimetry analysis.

9.2 The delocalisation issue

In the second campaign of MAST-U - the only campaign relevant to this thesis - all but 10 shots were conducted with an incorrectly calibrated optical filter in the BES diagnostic. This resulted in only carbon background emissions reaching the detector in most shots, rendering 'normal' velocimetry analysis of the main species impossible.

Despite this problem, the signals were further examined to determine if the data was suitable for velocimetry of carbon impurities. Initially, this may seem impossible because emission localisation is necessary for image-velocimetry analysis and the BES optics do not localise carbon emissions (because it was not designed to do so).

The carbon II ions, responsible for the background emissions, exist in a radially localised shell. Through radial scans of the total emission intensity, it was hypothesised that the carbon II shell is viewed approximately tangentially by BES, resulting in the approximate radial localisation of the emissions. This opened the door for velocimetry analysis, and after a rough radial 're-calibration' of the diagnostic, velocimetry analysis was carried out on the carbon II shell in section 7.4. Results indicated agreement with the velocities estimated by the Charge Exchange Recombination Spectroscopy diagnostic in terms of velocity magnitude and long-term temporal trends. This agreement suggests that **velocimetry can be accurately performed on the carbon II shell using BES**, a feat never before achieved. Although further work is recommended to optimise the method, this analysis has demonstrated that impurity dynamics studies

can be conducted with the BES diagnostic, offering a new avenue of exploration for the field.

In sections 7.2 and 7.3, velocimetry was performed on BES data *before* the discovery of the optical filter issue. Noteworthy is the fact that the lack of localisation was not immediately obvious in the BES data - most preliminary analyses of the data returned reasonable results. Unexpectedly low precision and velocity ‘singularities’ were observed during analysis, but the link to localisation issues was only made in hindsight.

Although the lack of localisation means that the velocimetry results in sections 7.2 and 7.3 are not representative of the true flows in the tokamak, they should not be entirely discounted either. It was found that the inferred velocities were consistent across three methods of inference: direct Fourier, CCTDE, and DTW. This suggests that the determined velocities were an accurate representation of the underlying ‘velocities’ in the BES data, despite the BES data not being representative of the bulk plasma dynamics. Additionally, these sections contain workflows to guide the application of velocimetry methods, making it a worthy read for individuals wishing to conduct image-velocimetry analysis.

9.3 A preliminary detection of zonal flows

Out of the ten shots with localised BES data, only one ran for the full time without disrupting. Initial investigations of this shot determined that the BES data was suitable for velocimetry, with CCTDE expected to perform accurately. Initial CCTDE velocity inferences were compared with the CXRS diagnostic, revealing agreement in the temporal trends, although CCTDE estimated consistently slower velocities than CXRS. These analyses, covered in section 8.1, suggested that reliable CCTDE performance could be achieved with this data.

The line-CCTDE method was introduced in section 8.2 for its significantly improved precision and enhanced inference frequency at the expense of spatial resolution in the z-direction. The line-method is based on the two-point method that was tested in chapter 6, making the testing results still directly applicable.

Line-CCTDE velocimetry was applied to the L-mode plasma of shot #46459, producing velocity inferences at an effective 4 kHz frequency and typical uncertainty margins below 1 km/s. These results revealed radially localised velocity structures observed to radially drift and re-form on 1-10 ms timescales, **behaviour that is expected from zonal flows**. Temporal averaging over 50 ms segments revealed a consistent average velocity shear, as expected in most tokamak plasmas. The novel observation was a sinusoidal velocity structure superposed on the average shear with a low ~ 0.5 km/s amplitude. Note that without the guide of chapter 6, convergence of this low-amplitude feature would likely not have been achieved. The mode was found to be coherent with a characteristic frequency of ~ 60 Hz and a radial wavelength of ~ 10 cm. Comparing with section 3.2, **the determined amplitude, frequency, and wavelength closely correspond to the expected characteristics of the quasi-stationary zonal flow mode**¹. This constitutes the first zonal flow observation in the MAST(-U) spherical

¹“If it looks like a duck, and quacks like a duck, we have at least to consider the possibility that we have a small aquatic bird of the family anatidae on our hands.” - Dirk Gently.

tokamak, and one of very few measurements in any fusion machine.

9.4 Future work

The velocimetry testing in chapter 6 is considered to be reasonably comprehensive and does not require further large-scale investigations. Regardless, further testing of DTW performance with sub-1 px/frame velocities could guide DTW velocimetry in this regime. DTW performance dependence on the shift parameter (see 7.2) and its effectiveness in recovering the sub-1 px/frame velocity regime would improve confidence in the method. Further in-depth investigations comparing the synthetic data with the experimental data may also prove insightful.

The CII velocimetry in section 7.4 was a proof of concept that showed promising results, but two main assumptions should be addressed further. Firstly, a more rigorous BES synthetic diagnostic should be developed, and secondly, more thorough estimates should be made of the CII shell radius. Addressing these points would likely minimise the main sources of uncertainty for the CII velocimetry analysis.

A common challenge faced in chapters 7 and 8 was the appearance of velocity ‘singularities’, where the CCTDE-inferred velocity periodically shot off to infinity. There was no obvious cause and this phenomenon was not previously observed in the testing. The singularities drastically reduced especially the precision of the velocity inferences, forcing the use of the line-method which sacrifices z-spatial resolution. Uncovering the cause and developing a method to circumvent the singularity issue would enable z-resolved velocimetry, a crucial requirement for poloidal coherence testing.

The highest impact future work recommendation is to further investigate the zonal-flow-like mode discovered in section 8.2. This ‘low-hanging fruit’ was not further analysed in this thesis due to time restrictions. No obvious technical challenges are expected. Shot #46459 could be further investigated to quantify the radial drifting and temporal evolution of the mode. H-mode velocimetry of the same shot could be performed as well as comparisons with DTW. Importantly, bispectral analysis should be performed to test if the mode is nonlinearly driven by turbulence. Measurement of the flow-turbulence nonlinear coupling would firmly confirm or reject the hypothesis that this mode is indeed a quasi-static zonal flow mode.

Hunting for zonal flow structures in other shots is also highly recommended. The third MAST-U campaign, carried out after the timeline of this PhD project, didn’t display localisation issues. Numerous shots dedicated to turbulent flow investigations are an ideal candidate for further zonal flow hunting. For example, 12 repeat shots were performed varying the BES radii and the toroidal field strength. If zonal flow modes are observed in these shots, zonal flow parametric dependencies could be experimentally determined. This would be a novel, high impact measurement that I would strongly encourage to be attempted.

Bibliography

- [1] Marcus Laurence Elwin Oliphant, Paul Harteck, and Ernest Rutherford. Transmutation effects observed with heavy hydrogen. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 144(853):692–703, 1934.
- [2] 60 years of progress. <https://www.iter.org/sci/BeyondITER>. Accessed: 01/02/2024.
- [3] G. Dif-Pradalier, G. Hornung, P. Ghendrih, Y. Sarazin, F. Clairet, L. Vermare, P. H. Diamond, J. Abiteboul, T. Cartier-Michaud, C. Ehrlacher, D. Esteve, X. Garbet, V. Grandgirard, O. D. Gurcan, P. Hennequin, Y. Kosuga, G. Latu, P. Maget, P. Morel, C. Norscini, R. Sabot, and A. Storelli. Finding the elusive $e \times b$ staircase in magnetized plasmas. *Physical Review Letters*, 114(8), 2015.
- [4] J. D. Lawson. Some criteria for a power producing thermonuclear reactor. *Proceedings of the Physical Society of London Section B*, 70(1):6–10, 1957.
- [5] R. F. Post. The magnetic-mirror approach to fusion. *Nuclear Fusion*, 27(10):1579–1731, 1987.
- [6] N. J. Peacock, D. C. Robinson, M. J. Forrest, P. D. Wilcock, and V. V. Sannikov. Measurement of electron temperature by thomson scattering in tokamak-t3. *Nature*, 224(5218):488–449, 1969.
- [7] M. Kenward. Fusion-research - temperature rises. *New Scientist*, 82(1156):626–630, 1979.
- [8] A. Fasoli, S. Brunner, W. A. Cooper, J. P. Graves, P. Ricci, O. Sauter, and L. Villard. Computational challenges in magnetic-confinement fusion physics. *Nature Physics*, 12(5):411–423, 2016.
- [9] A. W. Leonard. Edge-localized-modes in tokamaks. *Physics of Plasmas*, 21(9), 2014.
- [10] J. Favreau, M. Fergus, H. Ostby, and A. Marcum. Iron man. *Marvel Studios*, 2008.
- [11] J. W. Connor and H. R. Wilson. Survey of theories of anomalous transport. *Plasma Physics and Controlled Fusion*, 36(5):719–795, 1994.
- [12] I. T. Chapman, W. A. Cooper, J. P. Graves, M. P. Gryaznevich, R. J. Hastie, T. C. Hender, D. F. Howell, M. D. Hua, G. T. A. Huysmans, D. L. Keeling, Y. Q. Liu, H. F.

- Meyer, C. A. Michael, S. D. Pinches, S. Saarelma, S. A. Sabbagh, and MAST Team. Macroscopic stability of high beta mast plasmas. *Nuclear Fusion*, 51(7), 2011.
- [13] L. A. Artsimovich. Tokamak devices. *Nuclear Fusion*, 12(2):215–+, 1972.
- [14] W. L. Rowan, C. C. Klepper, C. P. Ritz, R. D. Bengtson, K. W. Gentle, P. E. Phillips, T. L. Rhodes, B. Richards, and A. J. Wootton. Global particle confinement in the texas experimental tokamak. *Nuclear Fusion*, 27(7):1105–1118, 1987.
- [15] A. J. Wootton, B. A. Carreras, H. Matsumoto, K. Mcguire, W. A. Peebles, C. P. Ritz, P. W. Terry, and S. J. Zweben. Fluctuations and anomalous transport in tokamaks. *Physics of Fluids B-Plasma Physics*, 2(12):2879–2903, 1990.
- [16] R. J. Bickerton. Magnetic turbulence and the transport of energy and particles in tokamaks. *Plasma Physics and Controlled Fusion*, 39(3):339–365, 1997.
- [17] S. Maeyama, Y. Idomura, T. H. Watanabe, M. Nakata, M. Yagi, N. Miyato, A. Ishizawa, and M. Nunami. Cross-scale interactions between electron and ion scale turbulence in a tokamak plasma. *Physical Review Letters*, 114(25), 2015.
- [18] N. T. Howard, C. Holland, A. E. White, M. Greenwald, and J. Candy. Multi-scale gyrokinetic simulation of tokamak plasmas: enhanced heat loss due to cross-scale coupling of plasma turbulence. *Nuclear Fusion*, 56(1), 2016.
- [19] P. Beyer, S. Benkadda, X. Garbet, and P. H. Diamond. Nondiffusive transport in tokamaks: Three-dimensional structure of bursts and the role of zonal flows. *Physical Review Letters*, 85(23):4892–4895, 2000.
- [20] D. A. D’Ippolito, J. R. Myra, and S. J. Zweben. Convective transport by intermittent blob-filaments: Comparison of theory and experiment. *Physics of Plasmas*, 18(6), 2011.
- [21] A. Ghizzo, D. Del Sarto, X. Garbet, and Y. Sarazin. Streamer-induced transport in the presence of trapped ion modes in tokamak plasmas. *Physics of Plasmas*, 17(9), 2010.
- [22] B. LaBombard, M. V. Umansky, R. L. Boivin, J. A. Goetz, J. Hughes, B. Lipschultz, D. Mossessian, C. S. Pitcher, J. L. Terry, and Alacator Grp. Cross-field plasma transport and main-chamber recycling in diverted plasmas on alcator c-mod. *Nuclear Fusion*, 40(12):2041–2060, 2000.
- [23] P. A. Politzer. Observation of avalanchelike phenomena in a magnetically confined plasma. *Physical Review Letters*, 84(6):1192–1195, 2000.
- [24] W. Horton. Drift waves and transport. *Reviews of Modern Physics*, 71(3):735–778, 1999.
- [25] A. Hasegawa and K. Mima. Stationary spectrum of strong turbulence in magnetized nonuniform plasma. *Physical Review Letters*, 39(4):205–208, 1977.

- [26] A. Hasegawa and K. Mima. Pseudo-3-dimensional turbulence in magnetized nonuniform plasma. *Physics of Fluids*, 21(1):87–92, 1978.
- [27] A. Hasegawa and M. Wakatani. Plasma edge turbulence. *Physical Review Letters*, 50(9):682–686, 1983.
- [28] M. Wakatani and A. Hasegawa. A collisional drift wave description of plasma edge turbulence. *Physics of Fluids*, 27(3):611–618, 1984.
- [29] S. C. Cowley, R. M. Kulsrud, and R. Sudan. Considerations of ion-temperature-gradient-driven turbulence. *Physics of Fluids B-Plasma Physics*, 3(10):2767–2782, 1991.
- [30] M. J. Burin, G. R. Tynan, G. Y. Antar, N. A. Crocker, and C. Holland. On the transition to drift turbulence in a magnetized plasma column. *Physics of Plasmas*, 12(5), 2005.
- [31] Y. Kaspi, E. Galanti, W. B. Hubbard, D. J. Stevenson, S. J. Bolton, L. Iess, T. Guillot, J. Bloxham, J. E. P. Connerney, H. Cao, D. Durante, W. M. Folkner, R. Helled, A. P. Ingersoll, S. M. Levin, J. I. Lunine, Y. Miguel, B. Militzer, M. Parisi, and S. M. Wahl. Jupiter’s atmospheric jet streams extend thousands of kilometres deep. *Nature*, 555(7695):223–+, 2018.
- [32] L. Schmitz, L. Zeng, T. L. Rhodes, J. C. Hillesheim, E. J. Doyle, R. J. Groebner, W. A. Peebles, K. H. Burrell, and G. Wang. Role of zonal flow predator-prey oscillations in triggering the transition to h-mode confinement. *Physical Review Letters*, 108(15), 2012.
- [33] P. H. Diamond, S. I. Itoh, K. Itoh, and T. S. Hahm. Zonal flows in plasma - a review. *Plasma Physics and Controlled Fusion*, 47(5):R35–R161, 2005.
- [34] M. N. Rosenbluth and F. L. Hinton. Poloidal flow driven by ion-temperature-gradient turbulence in tokamaks. *Physical Review Letters*, 80(4):724–727, 1998.
- [35] T. H. Stix. Decay of poloidal rotation in a tokamak plasma. *Physics of Fluids*, 16(8):1260–1267, 1973.
- [36] F. L. Hinton and M. N. Rosenbluth. Dynamics of axisymmetric ($e \times b$) and poloidal flows in tokamaks. *Plasma Physics and Controlled Fusion*, 41:A653–A662, 1999.
- [37] A. M. Dimits, T. J. Williams, J. A. Byers, and B. I. Cohen. Scalings of ion-temperature-gradient-driven anomalous transport in tokamaks. *Physical Review Letters*, 77(1):71–74, 1996.
- [38] M. A. Malkov and P. H. Diamond. Bifurcation and scaling of drift wave turbulence intensity with collisional zonal flow damping. *Physics of Plasmas*, 8(9):3996–4009, 2001.
- [39] S. Hamaguchi and W. Horton. Modeling of drift wave turbulence with a finite ion temperature-gradient. *Plasma Physics and Controlled Fusion*, 34(2):203–233, 1992.

- [40] J. Y. Kim, Y. Kishimoto, M. Wakatani, and T. Tajima. Poloidal shear flow effect on toroidal ion temperature gradient mode: A theory and simulation. *Physics of Plasmas*, 3(10):3689–3695, 1996.
- [41] R. E. Waltz, G. M. Staebler, W. Dorland, G. W. Hammett, M. Kotschenreuther, and J. A. Konings. A gyro-landau-fluid transport model. *Physics of Plasmas*, 4(7):2482–2496, 1997.
- [42] G. Dif-Pradalier, P. H. Diamond, V. Grandgirard, Y. Sarazin, J. Abiteboul, X. Garbet, P. Ghendrih, A. Strugarek, S. Ku, and C. S. Chang. On the validity of the local diffusive paradigm in turbulent plasma transport. *Physical Review E*, 82(2), 2010.
- [43] P. G. Ivanov, A. A. Schekochihin, W. Dorland, A. R. Field, and F. I. Parra. Zonally dominated dynamics and limits threshold in curvature-driven itg turbulence. *Journal of Plasma Physics*, 86(5), 2020.
- [44] S. Coda, M. Porkolab, and K. H. Burrell. Signature of turbulent zonal flows observed in the diii-d tokamak. *Physical Review Letters*, 86(21):4835–4838, 2001.
- [45] G. S. Xu, B. N. Wan, M. Song, and J. Li. Direct measurement of poloidal long-wavelength $e \times b$ flows in the ht-7 tokamak. *Physical Review Letters*, 91(12), 2003.
- [46] G. R. Tynan, C. Holland, J. H. Yu, A. James, D. Nishijima, M. Shimada, and N. Taheri. Observation of turbulent-driven shear flow in a cylindrical laboratory plasma device. *Plasma Physics and Controlled Fusion*, 48(4):S51–S73, 2006.
- [47] C. Holland, J. H. Yu, A. James, D. Nishijima, M. Shimada, N. Taheri, and G. R. Tynan. Observation of turbulent-driven shear flow in a cylindrical laboratory plasma device. *Physical Review Letters*, 96(19), 2006.
- [48] H. Xia, M. G. Shats, and H. Punzmann. Strong exb shear flows in the transport-barrier region in h-mode plasma. *Physical Review Letters*, 97(25), 2006.
- [49] J. C. Hillesheim, E. Delabie, H. Meyer, C. F. Maggi, L. Meneses, E. Poli, M. Abhangi, P. Abreu, M. Aftanas, M. Afzal, K. M. Aggarwal, L. Aho-Mantila, E. Ahoonen, M. Aints, M. Airila, R. Albanese, D. Alegre, E. Alessi, P. Aleynikov, A. Alfier, A. Alkseev, P. Allan, S. Almaguer, A. Alonso, B. Alper, I. Alsworth, D. Alves, G. Ambrosino, R. Ambrosino, V. Amosov, F. Andersson, E. Andersson Sunden, M. Angelone, A. Anghel, M. Anghel, C. Angioni, L. Appel, G. Apruzzese, P. Arena, M. Ariola, H. Arnichand, G. Arnoux, S. Arshad, A. Ash, E. Asp, O. Asunta, C. V. Atanasiu, Y. Austin, L. Avotina, M. D. Axton, C. Ayres, C. Bachmann, A. Baciero, D. Baiao, V. Bailescu, B. Baiocchi, A. Baker, R. A. Baker, I. Balboa, M. Balden, N. Balshaw, R. Bament, J. W. Banks, Y. F. Baranov, I. L. Barlow, M. A. Barnard, D. Barnes, R. Barnsley, A. Baron Wiechec, M. Baruzzo, V. Basiuk, M. Bassan, R. Bastow, A. Batista, P. Batistoni, R. Bauer, B. Bauvir, B. Bazylev, J. Beal, P. S. Beaumont, A. Becoulet, P. Bednarczyk, N. Bekris, M. Beldishevski, K. Bell, F. Belli, M. Bellinger, J. K. Belo, P. Belo, E. Belonohy, N. A. Benterman, H. Bergsaker,

- J. Bernardo, M. Bernert, M. Berry, L. Bertalot, M. N. A. Beurskens, B. Bieg, J. Bielecki, T. Biewer, et al. Stationary zonal flows during the formation of the edge transport barrier in the jet tokamak. *Physical Review Letters*, 116(6), 2016.
- [50] Y. Nagashima, K. Hoshino, A. Ejiri, K. Shinohara, Y. Takase, K. Tsuzuki, K. Uehara, H. Kawashima, H. Ogawa, T. Ido, Y. Kusama, and Y. Miura. Observation of nonlinear coupling between small-poloidal wave-number potential fluctuations and turbulent potential fluctuations in ohmically heated plasmas in the jft-2m tokamak. *Physical Review Letters*, 95(9), 2005.
- [51] D. K. Gupta, R. J. Fonck, G. R. McKee, D. J. Schlossberg, and M. W. Shafer. Detection of zero-mean-frequency zonal flows in the core of a high-temperature tokamak plasma. *Physical Review Letters*, 97(12), 2006.
- [52] I. Cziegler, P. H. Diamond, N. Fedorczak, P. Manz, G. R. Tynan, M. Xu, R. M. Churchill, A. E. Hubbard, B. Lipschultz, J. M. Sierchio, J. L. Terry, and C. Theiler. Fluctuating zonal flows in the i-mode regime in alcator c-mod. *Physics of Plasmas*, 20(5), 2013.
- [53] R. McKee, R. J. Fonck, M. Jakubowski, K. H. Burrell, K. Hallatschek, R. A. Moyer, W. Nevins, D. L. Rudakov, and X. Xu. Observation and characterization of radially sheared zonal flows in diii-d. *Plasma Physics and Controlled Fusion*, 45:A477–A485, 2003.
- [54] M. Jakubowski, R. J. Fonck, C. Fenzi, and G. R. McKee. Wavelet-based time-delay estimation for time-resolved turbulent flow analysis. *Review of Scientific Instruments*, 72(1):996–999, 2001.
- [55] G. D. Conway, B. Scott, J. Schirmer, M. Reich, A. Kendl, and ASDEX Upgrade Team. Direct measurement of zonal flows and geodesic acoustic mode oscillations in asdex upgrade using doppler reflectometry. *Plasma Physics and Controlled Fusion*, 47(8):1165–1185, 2005.
- [56] A. Fujisawa, K. Itoh, H. Iguchi, K. Matsuoka, S. Okamura, A. Shimizu, T. Minami, Y. Yoshimura, K. Nagaoka, C. Takahashi, M. Kojima, H. Nakano, S. Ohsima, S. Nishimura, M. Isobe, C. Suzuki, T. Akiyama, K. Ida, K. Toi, S. I. Itoh, and P. H. Diamond. Identification of zonal flows in a toroidal plasma. *Physical Review Letters*, 93(16), 2004.
- [57] C. P. Ritz, E. J. Powers, and R. D. Bengtson. Experimental-measurement of 3-wave coupling and energy cascading. *Physics of Fluids B-Plasma Physics*, 1(1):153–163, 1989.
- [58] J. S. Kim, R. D. Durst, R. J. Fonck, E. Fernandez, A. Ware, and P. W. Terry. Technique for the experimental estimation of nonlinear energy transfer in fully developed turbulence. *Physics of Plasmas*, 3(11):3998–4009, 1996.
- [59] C. Holland, G. R. Tynan, R. J. Fonck, G. R. McKee, J. Candy, and R. E. Waltz. Zonal-flow-driven nonlinear energy transfer in experiment and simulation. *Physics of Plasmas*, 14(5), 2007.

- [60] I. Cziegler, G. R. Tynan, P. H. Diamond, A. E. Hubbard, J. W. Hughes, J. Irby, and J. L. Terry. Nonlinear transfer in heated I-modes approaching the I-h transition threshold in Alcator C-Mod. *Nuclear Fusion*, 55(8), 2015.
- [61] M. Xu, G. R. Tynan, P. H. Diamond, P. Manz, C. Holland, N. Fedorczak, S. C. Thakur, J. H. Yu, K. J. Zhao, J. Q. Dong, J. Cheng, W. Y. Hong, L. W. Yan, Q. W. Yang, X. M. Song, Y. Huang, L. Z. Cai, W. L. Zhong, Z. B. Shi, X. T. Ding, X. R. Duan, Y. Liu, and HL-2A Team. Frequency-resolved nonlinear turbulent energy transfer into zonal flows in strongly heated -mode plasmas in the HL-2A tokamak. *Physical Review Letters*, 108(24), 2012.
- [62] K. Itoh, S. Toda, A. Fujisawa, S. I. Itoh, M. Yagi, A. Fukuyama, P. H. Diamond, and K. Ida. Physics of internal transport barrier of toroidal helical plasmas. *Physics of Plasmas*, 14(2), 2007.
- [63] G. D. Conway, C. Angioni, F. Ryter, P. Sauter, J. Vicente, and ASDEX Upgrade Team. Mean and oscillating plasma flows and turbulence interactions across the - confinement transition. *Physical Review Letters*, 106(6), 2011.
- [64] Z. Yan, G. R. McKee, R. Fonck, P. Gohil, R. J. Groebner, and T. H. Osborne. Observation of the I-h confinement bifurcation triggered by a turbulence-driven shear flow in a tokamak plasma. *Physical Review Letters*, 112(12), 2014.
- [65] I. Cziegler, G. R. Tynan, P. H. Diamond, A. E. Hubbard, J. W. Hughes, J. Irby, and J. L. Terry. Zonal flow production in the I-h transition in Alcator C-Mod. *Plasma Physics and Controlled Fusion*, 56(7), 2014.
- [66] I. Cziegler, A. E. Hubbard, J. W. Hughes, J. L. Terry, and G. R. Tynan. Turbulence nonlinearities shed light on geometric asymmetry in tokamak confinement transitions. *Physical Review Letters*, 118(10), 2017.
- [67] P. Manz, G. S. Xu, B. N. Wan, H. Q. Wang, H. Y. Guo, I. Cziegler, N. Fedorczak, C. Holland, S. H. Müller, S. C. Thakur, M. Xu, K. Miki, P. H. Diamond, and G. R. Tynan. Zonal flow triggers the I-h transition in the experimental advanced superconducting tokamak. *Physics of Plasmas*, 19(7), 2012.
- [68] W. L. Zhong, X. L. Zou, A. S. Liang, B. B. Feng, G. L. Xiao, C. Y. Chen, Z. B. Shi, W. W. Xiao, Z. C. Yang, P. W. Shi, Z. X. Wang, M. Jiang, J. Wen, K. R. Fang, J. Yin, X. M. Song, W. Chen, G. Z. Hao, X. Q. Ji, L. W. Yan, X. T. Ding, J. Q. Dong, D. L. Yu, Y. Liu, Q. W. Yang, M. Xu, and X. R. Duan. Stimulated effect of SMI on low-to-high confinement transition of tokamak plasmas. *Nuclear Fusion*, 60(8), 2020.
- [69] T. A. Carter and J. E. Maggs. Modifications of turbulence and turbulent transport associated with a bias-induced confinement transition in the large plasma device. *Physics of Plasmas*, 16(1), 2009.
- [70] G. R. Tynan, J. A. Boedo, D. S. Gray, R. Vannieuwenhove, G. Vanoost, and R. R. Weynants. Effects of radial electric-fields on the turbulence and transport in the TEXTOR edge and sol plasma. *Journal of Nuclear Materials*, 196:770–774, 1992.

- [71] G. R. Tynan, J. Liberati, P. Pribyl, R. J. Taylor, and B. Wells. Externally-driven h-mode studies in cct. *Plasma Physics and Controlled Fusion*, 38(8):1301–1305, 1996.
- [72] W. B. Liu, Y. H. Chen, R. Ke, G. McKee, Z. Yan, K. R. Fang, Z. C. Yang, Z. Gao, Y. Tan, and G. R. Tynan. Evidence of e b staircase in hl-2a l-mode tokamak discharges. *Physics of Plasmas*, 28(1), 2021.
- [73] P. Sun, C. Chubb, and G. Sperling. Two mechanisms that determine the barber-pole illusion. *Vision Research*, 111:43–54, 2015.
- [74] Y. C. Ghim, A. R. Field, D. Dunai, S. Zoletnik, L. Bardoczi, A. A. Schekochihin, and MAST Team. Measurement and physical interpretation of the mean motion of turbulent density patterns detected by the beam emission spectroscopy system on the mega amp spherical tokamak. *Plasma Physics and Controlled Fusion*, 54(9), 2012.
- [75] W. Mandl, R. C. Wolf, M. G. Vonhellermann, and H. P. Summers. Beam emission-spectroscopy as a comprehensive plasma diagnostic-tool. *Plasma Physics and Controlled Fusion*, 35(10):1373–1394, 1993.
- [76] D. Dunai and S. Thomas. The beam emission spectroscopy diagnostic on mast-u. *Review of Scientific Instruments*, in preparation.
- [77] I H Hutchinson. Excited-state populations in neutral beam emission. *Plasma Physics and Controlled Fusion*, 44(1):71, dec 2001.
- [78] M. Raffel, Christian Willert, Steve Wereley, and Juergen Kompenhans. *Particle Image Velocimetry: A Practical Guide*. Springer, 2007.
- [79] S. J. Zweben, J. L. Terry, M. Agostini, W. M. Davis, A. Diallo, R. A. Ellis, T. Golfinopoulos, O. Grulke, J. W. Hughes, B. LaBombard, M. Landreman, J. R. Myra, D. C. Pace, and D. P. Stotler. Comparison of edge turbulence imaging at two different poloidal locations in the scrape-off layer of alcator c-mod. *Physics of Plasmas*, 20(7), 2013.
- [80] J. M. Sierchio, I. Cziegler, J. L. Terry, A. E. White, and S. J. Zweben. Comparison of velocimetry techniques for turbulent structures in gas-puff imaging data. *Review of Scientific Instruments*, 87(2), 2016.
- [81] J. P. Guilford. Illusory movement from a rotating barber pole. *American Journal of Psychology*, 44:686–687, 1929.
- [82] S. Z. Cai, E. Memin, P. Derian, and C. Xu. Motion estimation under location uncertainty for turbulent fluid flows. *Experiments in Fluids*, 59(1), 2018.
- [83] G. M. Quenot, J. Pakleza, and T. A. Kowalewski. Particle image velocimetry with optical flow. *Experiments in Fluids*, 25(3):177–189, 1998.

- [84] H. Sakoe and S. Chiba. Dynamic-programming algorithm optimization for spoken word recognition. *Ieee Transactions on Acoustics Speech and Signal Processing*, 26(1):43–49, 1978.
- [85] G. M. Quenot. The orthogonal algorithm for optical-flow detection using dynamic-programming. *Icassp-92 - 1992 International Conference on Acoustics, Speech, and Signal Processing, Vols 1-5*, pages C249–C252, 1992.
- [86] G. R. McKee, R. J. Fonck, D. K. Gupta, D. J. Schlossberg, M. W. Shafer, C. Holland, and G. Tynan. Turbulence velocimetry of density fluctuation imaging data. *Review of Scientific Instruments*, 75(10):3490–3492, 2004.
- [87] Steven Thomas and István Cziegler. 2D Dynamic Time Warping (DTW) algorithm for python, 2 2023. v1.0.0 doi:10.5281/zenodo.7649022.
- [88] N. Fedorczak, P. Manz, S. C. Thakur, M. Xu, G. R. Tynan, G. S. Xu, and S. C. Liu. On physical interpretation of two dimensional time-correlations regarding time delay velocities and eddy shaping. *Physics of Plasmas*, 19(12), 2012.
- [89] C. Holland, G. R. Tynan, G. R. McKee, and R. J. Fonck. Investigation of the time-delay estimation method for turbulent velocity inference. *Review of Scientific Instruments*, 75(10):4278–4280, 2004.
- [90] B. Tal, A. Bencze, S. Zoletnik, G. Veres, and G. Por. Cross-correlation based time delay estimation for turbulent flow velocity measurements: Statistical considerations. *Physics of Plasmas*, 18(12), 2011.
- [91] D. M. Kriete, G. R. McKee, R. J. Fonck, D. R. Smith, G. G. Whelan, and Z. Yan. Extracting the turbulent flow-field from beam emission spectroscopy images using velocimetry. *Review of Scientific Instruments*, 89(10), 2018.
- [92] S. J. Zweben, J. L. Terry, D. P. Stotler, and R. J. Maqueda. Invited review article: Gas puff imaging diagnostics of edge plasma turbulence in magnetic fusion devices. *Review of Scientific Instruments*, 88(4), 2017.
- [93] Y. C. Ghim, A. R. Field, S. Zoletnik, and D. Dunai. Calculation of spatial response of 2d beam emission spectroscopy diagnostic on mast. *Review of Scientific Instruments*, 81(10), 2010.
- [94] M. F. J. Fox, A. R. Field, F. van Wyk, Y. C. Ghim, A. A. Schekochihin, and MAST Team. Experimental determination of the correlation properties of plasma turbulence using 2d bes systems. *Plasma Physics and Controlled Fusion*, 59(4), 2017.
- [95] M. N. Saha. On a physical theory of stellar spectra. *Proceedings of the Royal Society of London Series a-Containing Papers of a Mathematical and Physical Character*, 99(697):135–153, 1921.
- [96] C Theiler, RM Churchill, B Lipschultz, T Pütterich, A Diallo, A Hubbard, JW Hughes, E Marmar, ML Reinke, JL Terry, et al. Recent edge cxrs measurements in i-mode and elmy h-mode plasmas on c-mod. *54th Annual Meeting of the APS Division of Plasma Physics*, 2012.

Chapter 10

Appendices

10.1 Equations and expressions

10.1.1 $E \times B$ flow

$$u_{E \times B} = -\frac{\nabla \phi \times B}{B_0^2} \quad (10.1)$$

Equation 10.1 describes the $E \times B$ drift velocity, $u_{E \times B}$, with which a charged particle guiding center moves when placed in an electric and magnetic field.

10.1.2 Flow measurement equations derivation

Equations 7.3 and 7.4 are based on a consideration of radial force balance in tokamak plasmas, which can be described as follows:

$$-\nabla P_i = n_i q_i (E + u_i \times B) \quad (10.2)$$

where P_i , n_i , q_i , and u_i are the ion pressure, density, charge, and velocity, respectively. E and B are the electric and magnetic field strengths, respectively. Rearranging the equation and taking the minor-radial electric field projection gives:

$$E_r = \frac{\nabla P_i}{n_i q_i} - u_\phi B_\theta + u_\theta B_\phi \quad (10.3)$$

where θ , and ϕ are the poloidal and toroidal directions. The positive direction was defined along the positive minor-radial direction. Substituting the radial electric field into the $E \times B$ flow (eq. 10.1) gives us:

$$u_{E \times B} = \frac{-\nabla P_i}{n_i q_i B} - \frac{u_\phi B_\theta}{B} + \frac{u_\theta B_\phi}{B} = u_{*,i} - u_\perp \quad (10.4)$$

where $u_{*,i}$ is the ion diamagnetic velocity. u_\perp is the total flow in the ion drift direction and represents the directly experimentally measurable metric in this equation. Under neoclassical flow and transport theory, the poloidal flow can be assumed to go to zero at timescales longer than the neoclassical timescale. If this assumption does not hold, a more detailed analysis can be found in e.g. C-Mod experimental measurements[96].

10.2 Source Code: Cross-Correlation Time-Delay Estimation

```

1 import numpy as np
2 import warnings
3 import numpy.ma as ma
4 import matplotlib.cm as cm
5 import matplotlib
6 import multiprocessing as mp
7 from scipy import signal
8 from scipy.signal import butter, filtfilt
9 from scipy.ndimage import gaussian_filter
10 from sklearn.linear_model import RANSACRegressor, LinearRegression
11 from scipy.optimize import curve_fit
12 import itertools
13
14 #####
15 #####
16 # two-point method helper functions
17 #####
18 #####
19
20 def calc_norm_factor(f,g):
21     #calculates normalisation factor for ccf
22     #set precision to reduce overflow chance
23     f = f.astype(np.float32)
24     g = g.astype(np.float32)
25     #calculate sumsquares
26     sumsquare_f= (np.sum((f)**2))
27     sumsquare_g= (np.sum((g)**2))
28     # calculate normalisation factor
29     norm_factor = np.sqrt(sumsquare_f*sumsquare_g)
30     return norm_factor
31
32 def calc_distance(R1,R2,z1,z2):
33     #Calculates the distance between two locations (R1,z1) and (R2,z2).
34     distance = np.sqrt((R2-R1)**2 + (z2-z1)**2)
35     return distance
36
37 def reverse_direction_check(i1,i2,z1,z2):
38     # function which checks if the direction of the velocity needs to be
39     → reversed
40     # ensures that a positive velocity points in the positive z-direction
41     if i2>i1:
42         if z2>z1:
43             reflect_bool = False
44         elif z1>z2:
45             reflect_bool = True

```

```

45     else:
46         print('zero error. abort.')
47         print(1/0)
48 elif i1>i2:
49     if z2>z1:
50         reflect_bool = False
51     elif z1>z2:
52         reflect_bool = True
53     else:
54         print('zero error. abort.')
55         print(1/0)
56 elif i1==i2:
57     reflect_bool = False
58 return reflect_bool
59
60 def calc_corr_threshold(N,tolerance,print_bool = True):
61     '''
62     Determines the root mean square (rms) correlation amplitude expected
63     ↪ from
64     cross-correlating random noise. Correlation above this threshold is
65     ↪ considered
66     significant, while below is not.
67
68     Parameters
69     -----
70     N : int
71     Length of the random signals.
72     tolerance : float
73     Standard error tolerance used to calculate the rms correlation
74     ↪ amplitude precision.
75     print_bool : bool, optional
76     If True, prints the calculated correlation threshold. Default is
77     ↪ True.
78
79     Returns
80     -----
81     float
82     The calculated correlation threshold.
83
84     Notes
85     -----
86     This function generates random signals of length N, cross-correlates
87     ↪ them, and calculates
88     the rms of the cross-correlation function. It repeats this process
89     ↪ until a certain precision
90     level, defined by the standard error tolerance, is achieved.

```

```

86     The iteration limit is set to 1000000, and if the precision is not
87     ↪ achieved within this limit,
88     an error message is printed.
89
90     Examples
91     -----
92     To calculate rms correlation of 100 length random signals, to within
93     ↪ a 0.01 error margin:
94     >>> calc_corr_threshold(100, 0.01)
95     correlation threshold = 0.52
96     0.52
97     '''
98     #initialise
99     rms_ccfs = []
100    iteration_limit = 1000000
101    for i in range(iteration_limit):
102        # generate random signals of length N
103        sigA = np.random.normal(size=N)
104        sigB = np.random.normal(size=N)
105        # cross-correlate random signals
106        ccf,lags = calc_ccf(sigA,sigB)
107        # calculate rms of cross-correlation function
108        rms_ccf = np.sqrt(np.nanmean(ccf**2))
109        # add rms to storage list
110        rms_ccfs.append(rms_ccf)
111        # calculate the standard error of the rms_ccfs list
112        std_err = np.nanstd(rms_ccfs)/len(rms_ccfs)
113        if i>100:
114            # if standard error is below tolerance, precision level
115            # ↪ achieved. Break loop.
116            if std_err <tolerance:
117                break
118            if i == iteration_limit-1:
119                print('Err: iteration limit reached')
120    if print_bool:
121        print('correlation
122            ↪ threshold={0:.2f}'.format(np.nanmean(rms_ccfs)))
123    return np.nanmean(rms_ccfs)
124
125    #####
126    #####
127    # two-point method core functions (non-parallel)
128    #####
129    #####
130
131    def calc_ccf(f,g,norm_bool = True,plot_bool=False,
132                overlap_mode = 'same',method='fft'):
133        '''

```

```

130 Returns cross-correlation function and corresponding time-delays of
131     ↪ two one-dimensional arrays.
132 Arguments: (f,g,norm_bool = True,plot_bool=False,overlap_mode =
133     ↪ 'same',method='fft')
134 Returns: ccf,lags
135
136 Parameters
137 -----
138 f,g : 1D numpy array
139     two input array to be cross-correlated.
140
141 Keyword arguments
142 -----
143 norm_bool : boolean
144     turn cross-correlation function normalisation on or off, defaults
145     ↪ to True (on).
146 plot_bool : boolean
147     option to plot ccf, defaults to False.
148 overlap_mode : string
149     can change overlap mode of signal.correlate
150     ↪ ['same','full','valid']
151 method : string
152     can change signal.correlate method
153
154 Returns
155 -----
156 ccf : 1D numpy array
157     cross-correlation function of arrays f and g.
158 lags : 1D numpy array
159     lags corresponding to the ccf.
160
161 Notes
162 -----
163 :: ccf is calculated with scipy.signal.correlate.
164 :: Mode is set to 'same', meaning ccf output is same length as f
165 :: Method is set to 'fft', so correlation is calculated via fft method
166 '''
167 #zero center the signals
168 f = (f - np.nanmean(f))/np.nanstd(f)
169 g = (g - np.nanmean(g))/np.nanstd(g)
170 #calculate unnormalised ccf and lags
171 ccf = signal.correlate(g,f,mode=overlap_mode,method=method)
172 N = len(ccf)
173 lags = np.linspace(-N//2,N//2-1,N)
174 if norm_bool==True:
175     #normalise entire ccf based on ccf rms
176     norm_factor = calc_norm_factor(f,g)
177     ccf = np.divide(ccf,norm_factor)

```

```

174     if overlap_mode=='same':
175         #normalise each ccf value base on length array overlap (which
176         ↪ varies with the lag)
177         for i in range(len(ccf)):
178             lag = lags[i]
179             lag_norm = N/(N-np.abs(lag))
180             ccf[i] = ccf[i]/lag_norm
181     #plot ccf
182     if plot_bool==True:
183         index=np.where(np.max(ccf)==ccf)[0][0]
184         time_delay = lags[index]
185         print('time delay: {0} frames'.format(time_delay))
186         fig,ax=plt.subplots(3,figsize=(8,8))
187         ax[0].plot(lags,ccf,'.',ls='-')
188         ax[0].set_xlabel('time delay')
189         ax[0].set_ylabel('ccf')
190         ax[0].set_title('ccf')
191         ax[1].plot(f,'.',ls='-')
192         ax[1].set_title('f')
193         ax[2].plot(g,'.',ls='-')
194         ax[2].set_title('g')
195         fig.tight_layout()
196         plt.show()
197     return ccf,lags
198 def infer_velocity_two_point(sig1,sig2,times,R1,R2,z1,z2,
199                             correlation_threshold,mode='same',
200                             plot_bool=False,return_ccf=False):
201     """
202     Infers velocity from the cross-correlation time-lag between two
203     ↪ spatially separated signals.
204
205     Parameters
206     -----
207     sig1, sig2 : 1D numpy array
208     Two input arrays to be cross-correlated.
209     times : 1D numpy array
210     Array containing times at which samples were taken. Assumed to be
211     ↪ the same for sig1 and sig2.
212     Assumed to be in [seconds].
213     R1, R2, z1, z2 : floats
214     R- and z-locations of sig1 and sig2. Distances expected to be in
215     ↪ [meters].
216     correlation_threshold : float between 0 and 1
217     Defines the minimum correlation used for velocity inference.
218     If correlation below the threshold, then velocity defaults to
219     ↪ np.nan.

```

```

217 Keyword arguments:
218 -----
219 mode : str, optional
220     Overlap mode parameter for overlap between two signals. Default
221     → is 'same'.
222 plot_bool : bool, optional
223     Should the cross-correlation function (CCF) be plotted? Default
224     → is False.
225     WARNING! Make sure you're not inside several nested loops if set
226     → to True.
227 return_ccf : bool, optional
228     If True, returns velocity, correlation_max, CCF, and lags.
229     Default is False, returns velocity, correlation_max
230
231 Returns
232 -----
233 velocity : float
234     The inferred velocity [km/s] from signal 1 to signal 2.
235 correlation_max : float
236     The peak correlation value corresponding to the inferred velocity.
237
238 Notes
239 -----
240 If np.nan is returned, the correlation threshold was not surpassed OR
241 → time-lag was equal to zero.
242
243 Examples
244 -----
245 >>> infer_velocity_two_point(sig1, sig2, times, R1, R2, z1, z2,
246     → correlation_threshold=0.5)
247 (2.34, 0.75)
248 """
249 # calculate ccf
250 ccf,lags = calc_ccf(sig1,sig2,plot_bool=plot_bool,overlap_mode=mode)
251 if len(ccf) == 0:
252     # if ccf empty, set velocity to nan and correlation to 0
253     velocity = np.nan
254     correlation_max = 0.
255 #if ccf is not empty, proceed
256 else:
257     # find the peak of the cross-correlation function
258     correlation_max = np.max(ccf)
259     # correlation peak must exceed correlation threshold
260     if correlation_max>correlation_threshold:
261         #find time-delay at ccf peak
262         index=np.where(np.max(ccf)==ccf)[0][0]
263         time_delay = lags[index]
264         #account for zero-time delay scenario

```

```

260     if time_delay == 0.:
261         #manually set velocity to nan
262         velocity = np.nan
263     else:
264         #calculate unit conversion factors
265         distance = calc_distance(R1,R2,z1,z2)
266         t_sampling = np.mean(np.diff(times))
267         #calculate velocity
268         velocity = 1./time_delay * (distance/t_sampling)/1000.
269     else:
270         # set velocity to nan if below correlation threshold
271         velocity = np.nan
272 if plot_bool==True:
273     print('Velocity: {0}km/s \n Time: {1}s \n Correlation:
274         ↪ {2}'.format(velocity,np.mean(times),correlation_max))
275 if return_ccf==False:
276     return velocity,correlation_max
277 elif return_ccf==True:
278     return velocity,correlation_max,ccf,lags
279 else:
280     print('error')
281     print(1/0)
282 def infer_velocity_timeseries_two_point(sig1,sig2,times,R1,R2,z1,z2,N,
283                                         stepsize,correlation_threshold,
284                                         iterationlimit = 10000000,
285                                         mode='valid',plot_bool=False):
286     '''
287     Infer velocity timeseries from two signals. Shorter clips are taken
288     ↪ of the signals for velocity inference.
289
290     Parameters
291     -----
292     sig1, sig2 : 1D numpy array
293         Two input signals to be cross-correlated.
294     times : 1D numpy array
295         Array containing times at which samples were taken. Assumed to be
296         ↪ the same for sig1 and sig2. Assumed to be in [seconds].
297     R1, R2, z1, z2 : floats
298         R- and z-locations of sig1 and sig2. Distances expected to be in
299         ↪ [meters].
300     N : integer
301         The length of the individual time-series to be analysed [number
302         ↪ of frames].
303     stepsize : integer
304         The step size between clips moving through the time-series.
305         ↪ [number of frames]
306     correlation_threshold : float

```

```

302     Threshold of correlation below which the inferred velocity will
303     ↪ be ignored. [between 0 and 1].
304
305 Keyword Arguments
306 -----
307 iterationlimit : integer
308     Maximum number of velocity inferences to make.
309 mode : str
310     Overlap mode parameter for overlap between two signals. Default
311     ↪ is 'valid'.
312 plot_bool : bool
313     Should the cross-correlation function (CCF) be plotted? Default
314     ↪ is False.
315     WARNING! Make sure you're not inside several nested loops if set
316     ↪ to True.
317
318 Returns
319 -----
320 inferred_velocities : 1D numpy array
321     Array containing all the inferred velocities. [velocity output in
322     ↪ km/s].
323 inference_times : 1D numpy array
324     Contains the times at which the velocity inferences were taken.
325     ↪ Times taken as the middle of the time-series. Measured in [s].
326 inferred_correlations : 1D numpy array
327     Contains the peak correlation values corresponding to the
328     ↪ inferred velocities.
329
330 Notes
331 -----
332 :: The function takes two time series (sig1 and sig2), splits them
333     ↪ into consecutive shorter time-series of length N, and
334     ↪ cross-correlates each pair to infer velocities.
335 :: Velocity is inferred based on the peak correlation value exceeding
336     ↪ the specified correlation_threshold.
337 :: The output includes arrays of inferred velocities, corresponding
338     ↪ times, and peak correlation values.
339 :: If np.nan is returned in the velocities array, either the
340     ↪ correlation threshold was not surpassed or the time-lag was equal
341     ↪ to zero.
342
343 Examples
344 -----
345 >>> infer_velocity_timeseries_two_point(sig1, sig2, times, R1, R2,
346     ↪ z1, z2, N, stepsize, correlation_threshold=0.5)
347 (array([2.34, nan, ... ]), array([time1, time2, ... ]), array([0.75,
348     ↪ nan, ... ]))
349 '''

```

```

335     #initialise
336     more_data=True
337     i = 0
338     sig_length = len(sig1)
339     arr_length = int(sig_length/stepsize)+1
340     inferred_velocities = np.full(arr_length,np.nan)
341     inferred_correlations = np.full(arr_length,np.nan)
342     inference_times = np.full(arr_length,np.nan)
343     #loop until there is no more data
344     while more_data:
345         # use zero padding of input signals to calculate ccf
346         # output ccf will be of 'same' length as input signals
347         if mode == 'same':
348             #take slices of time-series
349             sliced_sig1 = sig1[i:i+N]
350             sliced_sig2 = sig2[i:i+N]
351             sliced_times = times[i:i+N]
352             #cross-correlate ts slices and infer velocity
353             velocity, maxcorr = infer_velocity_two_point(sliced_sig1,slic
↵ ed_sig2,sliced_times,R1,R2,z1,z2,correlation_threshold,mo
↵ de=mode,plot_bool=plot_bool)
354             #store velocity in array
355             inferred_velocities[int(i/stepsize)] = velocity
356             inferred_correlations[int(i/stepsize)] = maxcorr
357             inference_times[int(i/stepsize)] = np.mean(sliced_times)
358             # calculate ccf but don't allow zero padding
359             elif mode == 'valid':
360                 #exclude edge times
361                 if i < N//2:
362                     velocity, maxcorr = np.nan,np.nan
363                     inference_time = np.nan
364                 elif i > sig_length-N//2-N-1:
365                     velocity, maxcorr = np.nan,np.nan
366                     inference_time = np.nan
367                 else:
368                     #take slices of time-series
369                     sliced_sig1 = sig1[i:i+N]
370                     sliced_sig2 = sig2[i-N//2:i+N+N//2]
371                     sliced_times = times[i:i+N]
372                     inference_time = np.nanmean(sliced_times)
373                     #cross-correlate ts slices and infer velocity
374                     velocity, maxcorr = infer_velocity_two_point(sliced_sig1,
↵ sliced_sig2,sliced_times,R1,R2,z1,z2,correlation_thre
↵ shold,mode=mode,plot_bool=plot_bool)
375                     #store velocity in array
376                     inferred_velocities[int(i/stepsize)] = velocity
377                     inferred_correlations[int(i/stepsize)] = maxcorr
378                     inference_times[int(i/stepsize)] = inference_time

```

```

379     #move the current starting time index by 'stepsize' frames
380     i = i + stepsize
381     # abort loop if there is not enough data left in the time-series
382     if i+N >= len(sig1): more_data = False
383     # abort if iterationlimit is exceeded
384     if i/N > iterationlimit:
385         print('iteration limit exceeded!')
386         more_data = False
387     return inferred_velocities,inference_times,inferred_correlations
388
389     #####
390     #####
391     # core two-point cctde functions (parallel)
392     #####
393     #####
394
395 def z_velocity_scan_parallel_wrapper(i,j,signals,time,R,z,delta_ell,N,
396                                     stepsize,correlation_threshold):
397     # wrapper function to enable easy parallelisation of CCTDE analysis
398     # Initialise input parameters and signals for CCTDE
399     j1,j2 = (j,j)
400     i1,i2 = (i,i+delta_ell)
401     sig1 = signals[i1,j1]
402     R1,z1 = (R[i1,j1],z[i1,j1])
403     sig2 = signals[i2,j2]
404     R2,z2 = (R[i2,j2],z[i2,j2])
405     # Run CCTDE on one pair of signals
406     velocities_one_channel,inference_times,correlations_one_channel =
407     → infer_velocity_timeseries_two_point(sig1,sig2,time,R1,R2,z1,z2,N,
408     → stepsize,correlation_threshold,mode='valid')
409     # Reverse direction of inferred velocity if required
410     if reverse_direction_check(i1,i2,z1,z2): velocities_one_channel =
411     → np.multiply(velocities_one_channel,-1.)
412     return i,j,velocities_one_channel,inference_times,correlations_one_ch
413     → annel
414
415 def z_velocity_full_scan_parallelised(signals,time,j_range,i_range,R,z,N,
416                                     stepsize,correlation_threshold='auto',
417                                     o',
418                                     delta_ell = 1):
419     """
420     Scans field of view and performs velocimetry along the z (i)
421     → direction.
422     Scan channel numbers can be specified in both i and j
423
424     Arguments: (signals,time,j_range,i_range,R,z,N,correlation_threshold)
425     Returns: inferred_velocities,inference_times

```

```

421 Variables:
422 -----
423 signals: 3D numpy array [channel_i,channel_j,time]
424 The signals to be analysed. Time assumed to be in seconds
425 time: 1D numpy array [time]
426 The times at which the signal datapoints were sampled.
427 Time assumed to be in seconds.
428 j_range,i_range: list or numpy array of integers
429 The j/i channels to be scanned.
430 j_range can include the full range of j indices
431 i_range can include the full range, minus one, of i indices
432 R,z : 2D numpy array
433 the R,z coordinates corresponding to the j,i channel numbers.
434 N: integer
435 the length of the time-series clips to be analysed [number of
436 → frames]
437
438 Keyword arguments:
439 -----
440 correlation_threshold: string or float
441 threshold of correlation below which the inferred velocity will
442 → be ignored. [between 0 and 1, or 'auto']
443 if default 'auto', threshold determined by cross-correlation of
444 → random noise
445 delta_ell: integer
446 what should the distance be between analysed channels?
447
448 Returns:
449 -----
450 inferred_velocities: np array
451 an [i_range,j_range,time] array containing inferred velocities
452 inference_times: np array
453 contains the inference times of the velocities
454 inference_correlations: np array
455 an [i_range,j_range,time] array containing correlation values of
456 → inferred velocities
457
458 Notes:
459 -----
460 ::
461 '''
462 # Run CCTDE in parallel
463 nProcesses = len(i_range)*len(j_range)
464 # Calculate and set correlation threshold if it is 'auto'
465 if correlation_threshold=='auto':
466     correlation_threshold=calc_corr_threshold(N,0.01)
467 # Initialise a multiprocessing pool
468 with mp.Pool(processes= nProcesses) as pool:

```

```

465     # Execute z_velocity_scan_parallel_wrapper in parallel
466     i_indices, j_indices, inferred_velocities, inference_times, inferred_
    ↪ correlations =
    ↪ zip(*pool.starmap(z_velocity_scan_parallel_wrapper, [(i, j, sig_
    ↪ nals, time, R, z, delta_ell, N, stepsize, correlation_threshold) for
    ↪ i in i_range for j in j_range]))
467     # Convert outputs to numpy arrays
468     inferred_velocities = np.asarray(inferred_velocities)
469     inferred_correlations = np.asarray(inferred_correlations)
470     inference_times = np.asarray(inference_times)
471     # Reshape output arrays based on R, z coordinates and time stepsize
472     inferred_velocities_reshaped =
    ↪ np.full((R.shape[0], R.shape[1], int(len(time)/stepsize)+1), np.nan)
473     inferred_correlations_reshaped =
    ↪ np.full((R.shape[0], R.shape[1], int(len(time)/stepsize)+1), np.nan)
474     for k, i in enumerate(i_indices):
475         j = j_indices[k]
476         inferred_velocities_reshaped[i, j, :] = inferred_velocities[k, :]
477         inferred_correlations_reshaped[i, j, :] = inferred_correlations[k, :]
478     inference_times_reshaped = inference_times[0, :]
479     return inferred_velocities_reshaped, inference_times_reshaped, inferred_
    ↪ _correlations_reshaped
480
481 def R_velocity_scan_parallel_wrapper(i, j, signals, time, R, z, delta_ell, N,
482                                     stepsize, correlation_threshold):
483     # wrapper function to enable easy parallelisation of CCTDE analysis
484     # Initialise input parameters and signals for CCTDE
485     j1, j2 = (j, j+delta_ell)
486     i1, i2 = (i, i)
487     sig1 = signals[i1, j1]
488     R1, z1 = (R[i1, j1], z[i1, j1])
489     sig2 = signals[i2, j2]
490     R2, z2 = (R[i2, j2], z[i2, j2])
491     # Run CCTDE on one pair of signals
492     velocities_one_channel, inference_times, correlations_one_channel =
    ↪ infer_velocity_timeseries_two_point(sig1, sig2, time, R1, R2, z1, z2, N,
    ↪ stepsize, correlation_threshold, mode='valid')
493     # Reverse direction of inferred velocity if required
494     if reverse_direction_check(i1, i2, z1, z2): velocities_one_channel =
    ↪ np.multiply(velocities_one_channel, -1.)
495     return i, j, velocities_one_channel, inference_times, correlations_one_ch_
    ↪ annel
496
497 def R_velocity_full_scan_parallelised(signals, time, j_range, i_range, R, z, N,
498                                     stepsize, correlation_threshold='aut_
    ↪ o',
499                                     delta_ell = 1):
500     """

```

```

501 Scans field of view and performs velocimetry along the R (j)
    ↪ direction.
502 Scan channel numbers can be specified in both i and j
503
504 Arguments: (signals,time,j_range,i_range,R,z,N,correlation_threshold)
505 Returns: inferred_velocities,inference_times
506
507 Variables:
508 -----
509 signals: 3D numpy array [channel_i,channel_j,time]
    The signals to be analysed. Time assumed to be in seconds
510 time: 1D numpy array [time]
    The times at which the signal datapoints were sampled.
511 Time assumed to be in seconds.
512 j_range,i_range: list or numpy array of integers
    The j/i channels to be scanned.
513 j_range can include the full range, minus one, of j indices
514 i_range can include the full range of i indices
515 R,z : 2D numpy array
    the R,z coordinates corresponding to the j,i channel numbers.
516 N: integer
    the length of the time-series clips to be analysed [number of
517 ↪ frames]
518
519 Keyword arguments:
520 -----
521 correlation_threshold: string or float
    threshold of correlation below which the inferred velocity will
522 ↪ be ignored. [between 0 and 1, or 'auto']
523 if default 'auto', threshold determined by cross-correlation of
524 ↪ random noise
525 delta_ell: integer
    what should the distance be between analysed channels?
526
527 Returns:
528 -----
529 inferred_velocities: np array
    an [i_range,j_range,time] array containing inferred velocities
530 inference_times: np array
    contains the inference times of the velocities
531 inference_correlations: np array
    an [i_range,j_range,time] array containing correlation values of
532 ↪ inferred velocities
533
534 Notes:
535 -----
536 ::
537 '''
538

```

```

544 # Run CCTDE in parallel
545 nProcesses = len(i_range)*len(j_range)
546 # Calculate and set correlation threshold if it is 'auto'
547 if correlation_threshold=='auto':
548     correlation_threshold=calc_corr_threshold(N,0.01)
549 # Initialise a multiprocessing pool
550 with mp.Pool(processes= nProcesses) as pool:
551     # Execute R_velocity_scan_parallel_wrapper in parallel
552     i_indices,j_indices,inferred_velocities,inference_times,inferred_
    ↪ correlations =
    ↪ zip(*pool.starmap(R_velocity_scan_parallel_wrapper, [(i,j,sig_
    ↪ nals,time,R,z,delta_ell,N,stepsize,correlation_threshold) for
    ↪ i in i_range for j in j_range]))
553 # Convert outputs to numpy arrays
554 inferred_velocities = np.asarray(inferred_velocities)
555 inferred_correlations = np.asarray(inferred_correlations)
556 inference_times= np.asarray(inference_times)
557 # Reshape output arrays based on R, z coordinates and time stepsize
558 inferred_velocities_reshaped =
    ↪ np.full((R.shape[0],R.shape[1],int(len(time)/stepsize)+1),np.nan)
559 inferred_correlations_reshaped =
    ↪ np.full((R.shape[0],R.shape[1],int(len(time)/stepsize)+1),np.nan)
560 for k,i in enumerate(i_indices):
561     j = j_indices[k]
562     inferred_velocities_reshaped[i,j,:] = inferred_velocities[k,:]
563     inferred_correlations_reshaped[i,j,:] = inferred_correlations[k,:]
564 inference_times_reshaped = inference_times[0,:]
565 return inferred_velocities_reshaped,inference_times_reshaped,inferred_
    ↪ _correlations_reshaped
566
567 #####
568 #####
569 # line-CCTDE method helper functions
570 #####
571 #####
572
573 def myline(x,m):
574     #simple straight line function for fitting routines
575     #line is forced through 0,0
576     return m*x
577
578 def estimate_gradient_with_ransac_through_origin(distances, lag_peaks,
    ↪ max_trials=100):
579     """
580     Estimate best line of fit gradient through the origin using RANSAC
    ↪ (RANdom SAMple Consensus).
581
582     Parameters:

```

```

583 -----
584 distances : numpy.ndarray
585     Array of distances.
586 lag_peaks : numpy.ndarray
587     Array of corresponding lag peaks.
588 max_trials : int, optional
589     Maximum number of RANSAC trials. Default is 100.
590
591 Returns:
592 -----
593 Tuple (gradient, gradient_error):
594     gradient : float
595         Estimated gradient of the linear model.
596     gradient_error : float
597         Standard error of the estimated gradient.
598
599 Notes:
600 -----
601 - This function fits a RANSAC regressor to find the best linear model
602   → through the origin.
603 - It handles NaN values, reshapes input arrays, and performs RANSAC
604   → regression.
605 - Visualization code (optional) for inspecting the results is
606   → included but commented out.
607
608 Example:
609 -----
610 ```python
611 distances = np.array([1, 2, 3, 4, 5])
612 lag_peaks = np.array([2, 4, 6, 8, 10])
613 gradient, gradient_error =
614   → estimate_gradient_with_ransac_through_origin(distances, lag_peaks)
615 print(f"Estimated Gradient: {gradient}, Standard Error:
616   → {gradient_error}")
617 ```
618
619 References:
620 -----
621 - RANSACRegressor: https://scikit-learn.org/stable/modules/generated/
622   → sklearn.linear\_model.RANSACRegressor.html
623 - LinearRegression: https://scikit-learn.org/stable/modules/generated/
624   → sklearn.linear\_model.LinearRegression.html
625 """
626
627 # Check for NaN values and remove corresponding rows
628 nan_mask = np.logical_or(np.isnan(distances), np.isnan(lag_peaks))
629 distances = distances[~nan_mask]
630 lag_peaks = lag_peaks[~nan_mask]
631 # Reshape the arrays to 2D if they are 1D

```

```

624 if len(distances.shape) == 1:
625     distances = distances[:, np.newaxis]
626 if len(lag_peaks.shape) == 1:
627     lag_peaks = lag_peaks[:, np.newaxis]
628 # Initialize RANSAC regressor with a linear model without an intercept
629 ransac = RANSACRegressor(LinearRegression(fit_intercept=False),
    ↪ max_trials=max_trials)
630 # Fit the regressor
631 ransac.fit(distances, lag_peaks)
632 # Get the inlier mask
633 inliers = ransac.inlier_mask_
634 # Extract inlier points
635 inlier_distances = distances[inliers]
636 inlier_lag_peaks = lag_peaks[inliers]
637 # Fit the final model on the inliers
638 final_model = RANSACRegressor(LinearRegression(fit_intercept=False))
639 final_model.fit(inlier_distances, inlier_lag_peaks)
640 # Get the estimated gradient
641 gradient = final_model.estimator_.coef_[0]
642 # Calculate the residuals for the inliers
643 residuals = inlier_lag_peaks - final_model.predict(inlier_distances)
644 # Calculate the standard error of the gradient
645 n = len(inlier_distances)
646 mse = np.sum(residuals ** 2) / (n - 2) # Mean squared error
647 gradient_error = np.sqrt(mse / np.sum((inlier_distances -
    ↪ np.mean(inlier_distances)) ** 2))
648 # Visualization (optional)
649 # plt.scatter(distances, lag_peaks, color='b', label='Data points')
650 # plt.scatter(inlier_distances, inlier_lag_peaks, color='r',
    ↪ label='Inliers')
651 # plt.plot(distances, final_model.predict(distances), color='orange',
    ↪ label='RANSAC Model')
652 # plt.xlabel('Distances')
653 # plt.ylabel('Lag Peaks')
654 # plt.legend()
655 # plt.show()
656 return gradient, gradient_error
657
658 #####
659 #####
660 # line-CCTDE method core functions (non-parallel)
661 #####
662 #####
663
664 from scipy.stats import linregress
665
666 def infer_velocity_line_method(Linedata_segment, ref_j, zz,
    ↪ correlation_threshold, sampling_time,

```

```

667         max_distance=None, exclude_edges=True,
668         ↪ fit_method='ransac', plot_bool=False):
669     """
670     Infer velocity using the line method.
671
672     Parameters:
673     -----
674     Linedata_segment : numpy.ndarray
675         2D array of data segments.
676     ref_j : int
677         Reference channel index.
678     zz : numpy.ndarray
679         2D array of z coordinates.
680     correlation_threshold : float
681         Threshold for correlation to consider in the analysis.
682     sampling_time : float
683         Sampling time in seconds.
684     max_distance : float, optional
685         Maximum distance for analysis. Default is None.
686     exclude_edges : bool, optional
687         Flag to exclude edge indices. Default is True.
688     fit_method : str, optional
689         Method for fitting ('leastsq' or 'ransac'). Default is 'ransac'.
690     plot_bool : bool, optional
691         Flag for optional plotting. Default is False.
692
693     Returns:
694     -----
695     Tuple (velocity, velocity_err):
696         velocity : float
697             Inferred velocity in km/s.
698         velocity_err : float
699             Error in the inferred velocity.
700
701     Notes:
702     -----
703     - Uses cross-correlation function (CCF) to find time delays and
704       ↪ distances.
705     - Implements line fitting methods: 'leastsq' or 'ransac'. RANSAC is
706       ↪ default for its ability to ignore outliers.
707     - Optional plotting can be enabled with plot_bool.
708
709     Example:
710     -----
711     ```python
712     velocity, velocity_err = infer_velocity_line_method(Linedata_segment,
713     ↪ ref_j, zz, correlation_threshold,

```

```

710         sampling_time,
           ↪ max_distance ]
711         ↪ =None,
           ↪ exclude_edges=True ]
           ↪ ue,
           ↪ fit_method='
           ↪ ransac',
           ↪ plot_bool=False ]
           ↪ lse)

712     print(f"Inferred Velocity: {velocity} km/s, Velocity Error:
           ↪ {velocity_err} km/s")
713     """
714
715     # Find all combination pairs of indices
716     indices = range(Linedata_segment.shape[0])
717     combinations = np.array(list(itertools.combinations(indices, 2)))
718     index_exclusions = []
719     for index, (iref, icomp) in enumerate(combinations):
720         # Optionally filter out edge indices
721         if exclude_edges:
722             if iref == np.nanmin(indices) or icomp == np.nanmax(indices):
723                 index_exclusions.append(index)
724         # Filter out combinations with distances larger than max distance
725         if max_distance is not None:
726             distance = (zz[iref, ref_j] - zz[icomp, ref_j]) * 100 # cm
727             if distance > max_distance:
728                 index_exclusions.append(index)
729         # Remove combinations to be excluded
730         index_exclusions = list(set(index_exclusions.copy()))
731         combinations = np.delete(combinations, index_exclusions, axis=0)
732         # Initialise
733         N = Linedata_segment.shape[1] // 2
734         lag_peaks = np.full(combinations.shape[0], np.nan)
735         distances = np.full(combinations.shape[0], np.nan)
736         # Loop through all combinations
737         for index, (iref, icomp) in enumerate(combinations):
738             refsig = Linedata_segment[iref, N // 2:N + N // 2]
739             compsig = Linedata_segment[icomp, :]
740             # Calculate CCF
741             ccf, lags = calc_ccf(refsig, compsig, overlap_mode='valid',
           ↪ plot_bool=False)
742             # Check if CCF is empty
743             if len(ccf) == 0:
744                 # Leave lag_peak as nan and set correlation to 0
745                 correlation_max = 0.
746             else:
747                 # Find the peak of the cross-correlation function
748                 correlation_max = np.max(ccf)

```

```

749     # Correlation peak must exceed correlation threshold
750     if correlation_max > correlation_threshold:
751         # Find time-delay at CCF peak
752         peak_index = np.where(np.max(ccf) == ccf)[0][0]
753         time_delay = lags[peak_index]
754         # Account for zero-time delay scenario
755         if time_delay == 0.:
756             # Leave lag_peak as nan
757             pass
758         else:
759             # Store time_delay and distance at which the CCF peaks
760             lag_peaks[index] = time_delay * sampling_time * 1e6
761             distances[index] = -zz[ieref, ref_j] + zz[icomp, ref_j]
762     else:
763         # Leave lag_peak as nan if correlation threshold not
764         #   reached
765         pass
766     # Sort arrays for neatness
767     sorted_indices = np.argsort(distances)
768     distances = distances[sorted_indices].copy()
769     lag_peaks = lag_peaks[sorted_indices].copy()
770     # Line fitting (multiple methods)
771     if fit_method == 'leastsq':
772         mask = np.isfinite(lag_peaks)
773         distances = distances[mask]
774         lag_peaks = lag_peaks[mask]
775         grad, intercept, _, _, grad_err = linregress(distances, lag_peaks)
776     else:
777         grad, grad_err =
778         ↪ estimate_gradient_with_ransac_through_origin(distances,
779         ↪ lag_peaks)
780         grad = grad[0]
781     # Calculate velocity in km/s
782     velocity = 1. / (grad / 1000.)
783     velocity_err = velocity * (grad_err / grad)
784     # Optional plotting
785     if plot_bool:
786         # Fitted
787         plt.plot(distances, distances * grad, ls='--', c='k',
788                 label=f'v={velocity:.2f}(+/-){velocity_err:.2f} km/s')
789         # Raw
790         plt.plot(distances, lag_peaks, '.', c='b')
791         # Formatting
792         plt.title('line-CCTDE velocity fit, shot:46459, j=3 \n least
793                 ↪ squares method')
794         plt.xlabel('distance [m]')
795         plt.ylabel('time-delay [us]')
796         plt.legend()

```

```

793     plt.show()
794     return velocity, velocity_err
795
796 def infer_velocity_timeseries_line_method(Linedata, times, N, stepsize,
797     ↪ ref_j, zz, correlation_threshold,
798                                     max_distance=None,
799                                     ↪ exclude_edges=True,
800                                     ↪ plot_bool=False,
801                                     ↪ iterationlimit=1e9):
802
803     """
804     Analyze consecutive clips of density fluctuation data to infer
805     ↪ z-velocities.
806
807     Parameters:
808     Linedata (numpy.ndarray): 2D array containing density
809     ↪ fluctuations data.
810     times (numpy.ndarray): 1D array of timestamps corresponding to
811     ↪ the Linedata.
812     N (int): Size of the window for cross-correlation.
813     stepsize (int): Step size between consecutive clips of data.
814     ref_j (int): Index of the reference location (column).
815     zz (numpy.ndarray): 2D array of z-coordinates of the signals.
816     correlation_threshold (float): Threshold value for the
817     ↪ correlation coefficient to consider a valid peak.
818     max_distance (float, optional): Maximum distance for analysis.
819     ↪ Defaults to None.
820     exclude_edges (bool, optional): If True, exclude edge indices.
821     ↪ Defaults to True.
822     plot_bool (bool, optional): If True, plot the fitted line and raw
823     ↪ data points. Defaults to False.
824     iterationlimit (int, optional): Maximum number of iterations for
825     ↪ the analysis loop. Defaults to 1e9.
826
827     Returns:
828     tuple: A tuple containing the inference times, inferred
829     ↪ velocities, and their associated errors.
830
831     Note:
832     The function takes consecutive slices of the Linedata with a
833     ↪ specified window size (N) and step size
834     (stepsize). It performs cross-correlation on each slice to infer
835     ↪ the line velocities. The function
836     aborts if there is not enough data left in the time-series or if
837     ↪ the iteration limit is exceeded.
838
839     Example:

```

```

823     times, velocities, velocity_errors =
      ↪ infer_velocity_timeseries_line_method(Linedata, times, 100,
      ↪ 10, 4, zz, 0.7)
824     print("Inference Times:", times)
825     print("Inferred Velocities:", velocities)
826     print("Velocity Errors:", velocity_errors)
827     """
828     # Initialize
829     is_more_data = True
830     i = 0
831     sampling_time = np.nanmean(np.diff(times))
832     nframes = len(Linedata[i, :])
833     arr_length = int((nframes - N) / stepsize) + 1
834     inferred_velocities = np.full(arr_length, np.nan)
835     inferred_velocity_errors = np.full(arr_length, np.nan)
836     inference_times = np.full(arr_length, np.nan)
837     # Loop until there is no more data
838     while is_more_data:
839         # Take slices of time-series
840         Linedata_segment = Linedata[:, i:i + 2 * N]
841         sliced_times = times[i:i + 2 * N]
842         inference_time = np.nanmean(sliced_times)
843         try:
844             # Cross-correlate ts slices and infer velocity
845             velocity, velocity_err =
              ↪ infer_velocity_line_method(Linedata_segment, ref_j, zz,
              ↪ correlation_threshold, sampling_time,
              ↪ max_distance=max_distance, exclude_edges=exclude_edges,
              ↪ plot_bool=plot_bool)
846         except:
847             velocity = np.nan
848             velocity_err = np.nan
849         # Store velocity in array
850         inferred_velocities[int(i / stepsize)] = velocity
851         inferred_velocity_errors[int(i / stepsize)] = velocity_err
852         inference_times[int(i / stepsize)] = inference_time
853         # Move the current starting point
854         i = i + stepsize
855         # Abort loop if there is not enough data left in the time-series
856         if i + 2 * N >= nframes:
857             is_more_data = False
858         # Abort if iterationlimit is exceeded
859         if i / N > iterationlimit:
860             print('iteration limit exceeded!')
861             is_more_data = False
862     return inference_times, inferred_velocities, inferred_velocity_errors
863

```

```

864 def z_velocity_full_scan_line_method(BESdata, BEStimes, shotn, RR, zz, N,
    ↪ stepsize, j_indices, max_distance=None,
865         exclude_edges=True,
            ↪ correlation_threshold='auto',
            ↪ plot_bool=False):
866     """
867     Perform a full scan of z-velocities using the line method.
868
869     Parameters:
870         BESdata (numpy.ndarray): 3D array containing BES data.
871         BEStimes (numpy.ndarray): 1D array of timestamps corresponding to
            ↪ the BESdata.
872         shotn (int): Shot number.
873         RR (numpy.ndarray): 2D array of R coordinates.
874         zz (numpy.ndarray): 2D array of z coordinates.
875         N (int): Size of the window for cross-correlation.
876         stepsize (int): Step size between consecutive clips of data.
877         j_indices (list): List of indices for the j channels.
878         max_distance (float, optional): Maximum distance for analysis.
            ↪ Defaults to None.
879         exclude_edges (bool, optional): If True, exclude edge indices.
            ↪ Defaults to True.
880         correlation_threshold (float or 'auto', optional): Threshold
            ↪ value for the correlation coefficient
881             to consider a valid peak. If 'auto', the threshold is
            ↪ determined by calc_corr_threshold.
882         plot_bool (bool, optional): If True, plot the overview of
            ↪ inferred velocities. Defaults to False.
883
884     Returns:
885         tuple: A tuple containing the inferred velocities, their
            ↪ associated errors, and the inference times.
886
887     Example:
888         velocities, velocity_errors, inference_times =
            ↪ z_velocity_full_scan_line_method(BESdata, BEStimes, 123,RR,
            ↪ zz, 100, 10,[0, 1, 2],max_distance=500,exclude_edges=True,cor_
            ↪ relation_threshold='auto',plot_bool=True)
889         print("Inferred Velocities:", velocities)
890         print("Velocity Errors:", velocity_errors)
891         print("Inference Times:", inference_times)
892     """
893     # Initialize
894     all_inferred_velocities = np.full((BESdata.shape[1], (len(BEStimes) -
            ↪ N) // stepsize + 1), np.nan)
895     all_inferred_velocities_err = np.full((BESdata.shape[1],
            ↪ (len(BEStimes) - N) // stepsize + 1), np.nan)

```

```

896 all_inference_times = np.full((len(BEStimes) - N + 1) // stepsize,
    ↪ np.nan)
897 # Optionally calculate correlation threshold
898 if correlation_threshold == 'auto':
899     correlation_threshold = calc_corr_threshold(N, 0.01)
900 elif not np.issnumeric(correlation_threshold):
901     raise ValueError("correlation_threshold is not a number.")
902 # Loop over j indices
903 for j in j_indices:
904     BESdata_line = BESdata[:, j, :]
905     inference_times, inferred_velocities, inferred_velocity_errors =
    ↪ infer_velocity_timeseries_line_method(
906         BESdata_line, BEStimes, N, stepsize, j, zz,
    ↪ correlation_threshold, max_distance=max_distance,
907         exclude_edges=exclude_edges, plot_bool=False)
908     all_inferred_velocities[j, :] = inferred_velocities
909     all_inferred_velocities_err[j, :] = inferred_velocity_errors
910     all_inference_times = inference_times.copy()
911 # Plot overview if plot_bool is True
912 if plot_bool:
913     for j in j_indices:
914         plt.plot(all_inference_times * 1000.,
    ↪ all_inferred_velocities[j, :],
    ↪ label='R={:.2f}m'.format(RR[3, j]))
915     plt.legend()
916     plt.title('Overview of line-CCTDE inferred velocities \n
    ↪ shot:#{}, times {:.2f}-{:.2f}ms'.format(
917         shotn, np.nanmin(all_inference_times) * 1000.,
    ↪ np.nanmax(all_inference_times) * 1000.))
918     plt.ylabel('z-velocity [km/s]')
919     plt.xlabel('time [ms]')
920     plt.show()
921     return all_inferred_velocities, all_inferred_velocities_err,
    ↪ all_inference_times
922
923 #####
924 #####
925 # line-CCTDE method core functions (parallel)
926 #####
927 #####
928
929 def z_velocity_scan_line_method_parallel_wrapper(j, BESdata, BEStimes, N,
    ↪ stepsize, zz, correlation_threshold,
930
    ↪ max_distance=None,
    ↪ exclude_edges=True):
931     """
932     Wrapper function for parallelized z-velocity scanning using the line
    ↪ method.

```

```

933
934 Parameters:
935     j (int): Index for the j channel.
936     BESdata (numpy.ndarray): 3D array containing BES data.
937     BEStimes (numpy.ndarray): 1D array of timestamps corresponding to
938     ↪ the BESdata.
939     N (int): Size of the window for cross-correlation.
940     stepsize (int): Step size between consecutive clips of data.
941     zz (numpy.ndarray): 2D array of z coordinates.
942     correlation_threshold (float): Threshold value for the
943     ↪ correlation coefficient to consider a valid peak.
944     max_distance (float, optional): Maximum distance for analysis.
945     ↪ Defaults to None.
946     exclude_edges (bool, optional): If True, exclude edge indices.
947     ↪ Defaults to True.
948
949 Returns:
950     tuple: A tuple containing the index j, inferred velocities, their
951     ↪ associated errors, and the inference times.
952     """
953
954 BESdata_line = BESdata[:, j, :]
955 inference_times, inferred_velocities, inferred_velocity_errors =
956     ↪ infer_velocity_timeseries_line_method(
957         BESdata_line, BEStimes, N, stepsize, j, zz,
958         ↪ correlation_threshold, max_distance=max_distance,
959         exclude_edges=exclude_edges, plot_bool=False)
960 return j, inferred_velocities, inferred_velocity_errors,
961     ↪ inference_times
962
963 def z_velocity_full_scan_line_method_parallelised(BESdata, BEStimes,
964     ↪ shotn, RR, zz, N, stepsize, j_indices, max_distance=None,
965     ↪ exclude_edges=True, correlation_threshold='auto', plot_bool=False):
966     """
967     Perform a full scan of z-velocities using the line method in parallel.
968
969 Parameters:
970     BESdata (numpy.ndarray): 3D array containing BES data.
971     BEStimes (numpy.ndarray): 1D array of timestamps corresponding to
972     ↪ the BESdata.
973     shotn (int): Shot number.
974     RR (numpy.ndarray): 2D array of R coordinates.
975     zz (numpy.ndarray): 2D array of z coordinates.
976     N (int): Size of the window for cross-correlation.
977     stepsize (int): Step size between consecutive clips of data.
978     j_indices (list): List of indices for the j channels.
979     max_distance (float, optional): Maximum distance for analysis.
980     ↪ Defaults to None.

```

```

968     exclude_edges (bool, optional): If True, exclude edge indices.
969     ↪ Defaults to True.
970     correlation_threshold (float or 'auto', optional): Threshold
971     ↪ value for the correlation coefficient
972     ↪ to consider a valid peak. If 'auto', the threshold is
973     ↪ determined by calc_corr_threshold.
974     plot_bool (bool, optional): If True, plot the overview of
975     ↪ inferred velocities. Defaults to False.
976
977 Returns:
978     tuple: A tuple containing the inferred velocities, their
979     ↪ associated errors, and the inference times.
980
981 Example:
982     velocities, velocity_errors, inference_times =
983     ↪ z_velocity_full_scan_line_method_parallelised(BESdata,
984     ↪ BEStimes,123, RR, zz, 100, 10,[0, 1, 2],max_distance=500,excl
985     ↪ ude_edges=True,correlation_threshold='auto',plot_bool=True)
986     print("Inferred Velocities:", velocities)
987     print("Velocity Errors:", velocity_errors)
988     print("Inference Times:", inference_times)
989     """
990     all_inferred_velocities = np.full((BESdata.shape[1], (len(BEStimes) -
991     ↪ N) // stepsize + 1), np.nan)
992     all_inferred_velocities_err = np.full((BESdata.shape[1],
993     ↪ (len(BEStimes) - N) // stepsize + 1), np.nan)
994     all_inference_times = np.full((len(BEStimes) - N + 1) // stepsize,
995     ↪ np.nan)
996     # Optionally calculate correlation threshold
997     if correlation_threshold == 'auto':
998         correlation_threshold = calc_corr_threshold(N, 0.01)
999     elif not np.isscalar(correlation_threshold):
1000         raise ValueError("correlation_threshold is not a number.")
1001     num_processes = len(j_indices)
1002     with mp.Pool(processes=num_processes) as pool:
1003         results =
1004         ↪ pool.starmap(z_velocity_scan_line_method_parallel_wrapper,
1005         ↪ [(j, BESdata, BEStimes, N, stepsize, zz,
1006         ↪ correlation_threshold, max_distance,
1007         ↪ exclude_edges) for j in j_indices])
1008     for j, inferred_velocities, inferred_velocity_errors, inference_times
1009     ↪ in results:
1010         all_inferred_velocities[j, :] = inferred_velocities
1011         all_inferred_velocities_err[j, :] = inferred_velocity_errors
1012         all_inference_times = inference_times.copy()
1013     # Plot overview if plot_bool is True
1014     if plot_bool:
1015         for j in j_indices:

```

```
1001     plt.plot(all_inference_times * 1000.,
1002             ↪ all_inferred_velocities[j, :],
1003             ↪ label='R={:.2f}m'.format(RR[3, j]))
1004 plt.legend()
1005 plt.title('Overview of line-CCTDE inferred velocities \n
1006 ↪ shot:#{ }, times {:.2f}-{:.2f}ms'.format(
1007     shotn, np.nanmin(all_inference_times) * 1000.,
1008     ↪ np.nanmax(all_inference_times) * 1000.))
1009 plt.ylabel('z-velocity [km/s]')
1010 plt.xlabel('time [ms]')
1011 plt.show()
1012 return all_inferred_velocities, all_inferred_velocities_err,
1013 ↪ all_inference_times
```

