A combined phase contrast imaging and heterodyne interferometer system for multiscale fluctuation measurements on DIII-D

by

Evan Michael Davis

B.S., Physics
B.S. Applied Mathematics
UNIVERSITY OF CALIFORNIA, RIVERSIDE, 2010

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2018

Author . . . .
Department of Physics
January 10, 2018

Certified by . . . .
Miklos Porkolab
Professor of Physics
Thesis Supervisor

Accepted by . . . .
Scott Hughes
Professor of Physics
Interim Associate Department Head

A combined phase contrast imaging and heterodyne interferometer system for multiscale fluctuation measurements on DIII-D

by

Evan Michael Davis

Abstract

A novel combined diagnostic capable of measuring multiscale density fluctuations that extend from magnetohydrodynamic (MHD) scales to the lower bound of the electron temperature gradient (ETG) mode has been designed, installed, and operated at the DIII-D tokamak. The combined diagnostic was constructed by adding a heterodyne interferometer to the pre-existing phase contrast imaging (PCI) system, both of which measure line-integrated electron-density fluctuations. The port-space footprint is minimized by using a single $10.6\,\mu m$ CO_2 laser and a single beampath. With temporal bandwidths in excess of $1\,MHz$, the PCI measures high-k ($1.5\,cm^{-1} < |k| \leqslant 25\,cm^{-1}$) fluctuations with sensitivity 3×10^{13} $m^{-2}/\sqrt{kHz}$, while the interferometer simultaneously measures low-k ($|k| < 5\,cm^{-1}$) fluctuations with sensitivity 3×10^{14} $m^{-2}/\sqrt{kHz}$. The intentional mid-k overlap has been empirically verified with sound-wave calibrations and should allow quantitative investigation of the cross-scale coupling predicted to be significant in the reactor-relevant $T_e \approx T_i$ regime.

The combined PCI-interferometer was operated during an experiment in which the ETG drive a/L_{T_e} and the ion temperature gradient (ITG) drive a/L_{T_i} were locally modified in an attempt to elicit a multiscale turbulent response. Numerous turbulent branches are observed. In particular, the interferometer measures a low-k electromagnetic mode driven unstable by collisionality, properties consistent with the micro-tearing mode (MTM), and the PCI measures a turbulent mode that exhibits distinct "spectral flattening" when increasing a/L_{T_e} relative to a/L_{T_i}, hypothesized to be a tell-tale signature of increased cross-scale coupling. Linear-stability analysis and quasilinear-transport modeling are performed with the trapped gyro-Landau fluid code TGLF, and qualitative agreement with the PCI-measured spectral flattening is obtained.

Further, via toroidal correlation with DIII-D's primary interferometer, the measurement of core-localized MHD toroidal mode numbers has been demonstrated. Where comparisons can be made with magnetic probes, the interferometer-measured toroidal mode numbers are typically in good agreement. Unfortunately, the $4\,\mathrm{cm}$ major-radial offset between the interferometer beam centers in DIII-D can bias the mode-number measurement, limiting widespread use of this capability until a robust compensation technique is developed.

Thesis Supervisor: Miklos Porkolab
Professor of Physics

between the cracks". Bob Pinsker fostered a collegiate atmosphere, regularly hosting DIII-D's student/postdoc seminar series, and he also lent us the high-power RF amplifier needed to drive our AOM. Lupe Cerda was always happy to help me navigate the DIII-D bureaucracy. The DIII-D "lunch crew", consisting of a hodge podge of graduate students, postdocs, and scientists from multiple institutions, always made for an entertaining lunch; of particular note are Matthias Knölker, Mitchell Clement, "Staff" Tyler Abrams, Ryan Sweeney, Brent Covele, and Wilkie Choi. The short-lived "Monday runday" group helped re-ignite my passion for running and introduced me to running along the picturesque bluffs of La Jolla.

The Stewardship Science Graduate Fellowship (SSGF) program generously supported the first 4 years of my graduate studies. The Krell Institute, which administered the SSGF program, and Lucille Kilmer in particular were always a pleasure to interact with.

From Riverside, to Boston, to San Diego, I have been blessed with some truly amazing friends. John Paul Issa, you were there at the beginning, and you were there at the end; I feel like we have both been to our own personal hell and back, and I am confident that the future is now bright. You have always told it to me straight, but you have never left me feeling bad. Riley and Julia Takano, you make being responsible adults look so easy; even a few hours with both of you has been enough to cure the lab blues. Steven Warneke, it has been a long time since our "Titans of Industry" talk and our paths have swerved a bit since then, but our friendship has only grown; Madeline, I am so happy you found him and are (mostly) keeping him out of trouble. Rachel Cheong, you know me better than most and still accept me for who I am; you are beautiful inside and out. Ian Faust, you made 881 Mass Ave home, and you truly helped me come into my own. From Coachella to Oktoberfest, you have never said no to adventure; "unfortunately", this has currently led you to Munich, but even this has not put a wedge between us. Alex Guo, you lead by example, and you have a heart of gold. Jason Ramapuram and Barbara Sandoval, the kindness and friendship you both showed me when I moved to San Diego far exceeded my expectations; your around-the-world adventures amaze and inspire me, but I selfishly still wish you were in San Diego. Matthias Knölker, from our off-hand introduction at lab to where we are now, who could have guessed? The only things that surprise me more than your intuition and your humor are your friendship and your loyalty. Theresa Wilks, your optimism is infections, and your laughter can move mountains. Darren Burton & co.

(i.e. Scott and Andrew), thanks for so many good times, good food, and laughter inducing situations over the last year.

Finally, my family has always been the cornerstone of my life. Cory, you've grown from being my little brother to one of my best friends. Your passions for astronomy and space exploration have helped me remember why I got into science in the first place. Thanks for never letting me forget where I came from. Mom and Dad, you have been my number one fans since day one. When I celebrated, you celebrated, and when I cried, you cried. At this point, you are probably almost as well-versed in this thesis as I am. Thank you for all of your love, support, and encouragement. This thesis is dedicated to both of you.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

INTRODUCTION

1.1 FUSION ENERGY

Fusion is the energy of the Sun and stars, and, for the better part of the last century, thousands of scientists and engineers have striven to replicate this process in a controlled manner here on Earth for a source of safe, clean, and virtually inexhaustible energy. The core principle is simple: combine (i.e. "fuse") the nuclei of two light elements to make the nucleus of a heavier element, with the total mass of the resulting products being slightly less than that of the reactants; this mass difference m is converted into an enormous amount of energy E via Einstein's famed $E = mc^2$, where c is the speed of light in vacuum [1, Ch. 14].

Practically speaking, however, fusion in the laboratory is extremely difficult. Even for deuterium (D) and tritium (T), the most reactive of nuclei, the probability of fusing (quantified by the reaction cross section) is orders-of-magnitude smaller than the probability of the positively charged nuclei repelling and scattering off of each other (quantified by the Coulomb cross section) [2, Sec. 9.3.4]. This relatively small fusion cross section explains why beam-target or beam-beam fusion are incapable of producing net energy. Instead of using a highly directed beam that is far from thermal equilibrium, an alternative approach is to employ the thermal energy of a material at or near thermal equilibrium to overcome the Coulomb barrier; because of its reliance on the thermal energy, such an approach is referred to as thermonuclear fusion [1, Ch. 14].

The temperatures required to initiate thermonuclear fusion dictate that the fuel exists as a hot, ionized gas referred to as plasma. Initially, the plasma must be externally heated to reach fusion-relevant conditions, but it is envisioned that at some threshold, much like the Sun, the fusion reactions will become self-sustaining; this threshold is referred to as *ignition*. Now, in addition to the thermal energy required to overcome the Coulomb barrier, the plasma must also be sufficiently dense and sufficiently well-confined to reach ignition. The threshold

for ignition in a D-T plasma is often quantified by the fusion triple product

$$n_i T_i \tau_E \geqslant 5 \times 10^{21} \, \text{m}^{-3} \cdot \text{keV} \cdot \text{s} = 8 \, \text{atm} \cdot \text{s}, \tag{1.1}$$

where n_i is the peak ion density, T_i is the peak ion temperature, and τ_E is the energy confinement time [3, Sec. 1.1]. This condition could be reached, for example, at $n_i = 10^{20} \, \text{m}^{-3}$, $T_i = 20 \, \text{keV}$, and $\tau_E = 2.5 \, \text{s}$.

It should be emphasized that $T_i = 20 \, \text{keV}$ corresponds to approximately 200 *million* degrees Celsius, approximately ten times hotter than the core of the Sun [4, Sec. 4.3]. No conventional container can confine such a plasma. While numerous confinement schemes exist, the only strategy considered in this thesis is that of tokamak magnetic confinement, briefly summarized in Section 1.1.1. While the tokamak's magnetic fields prevent the hot, core plasma from contacting any material walls, the outermost portions of the plasma (i.e. "the edge") do contact the reactor walls. To prevent damage to the reactor walls, the plasma temperature and density must decrease to tolerable levels at the edge; the resulting gradients in density and temperature provide free energy for numerous coherent and broadband instabilities, some of which are measured and characterized in this thesis.

1.1.1   *The tokamak approach to fusion*

The tokamak [3] is currently the leading configuration for a magnetic-confinement fusion reactor. The tokamak is an axisymmetric toroidal device characterized by a strong toroidal magnetic field B_ζ and a toroidal plasma current I_ζ. The toroidal plasma current produces a poloidal magnetic field B_θ, resulting in a total magnetic field $\boldsymbol{B} = B_\zeta \hat{\boldsymbol{\zeta}} + B_\theta \hat{\boldsymbol{\theta}}$ that wraps helically around the torus in a barber pole-like fashion. The magnetic force of this helical field balances the plasma pressure, allowing the establishment of an equilibrium; the large toroidal field B_ζ provides stability. Wesson provides a concise historical overview of tokamak research [3, Sec. 1.10]. Since the initial success of the USSR's T-3 tokamak in the 1960s, hundreds of tokamaks of numerous shapes, sizes, and field strengths have been built all over the world [3, Ch. 11,12] [5].

The work reported in this thesis was conducted at the DIII-D tokamak in San Diego, CA [3, Sec. 12.5]. DIII-D is a mid-size ($R = 1.67 \, \text{m}$, $a = 0.67 \, \text{m}$) diverted tokamak with maximum toroidal field $B_\zeta \leqslant 2.2 \, \text{T}$ and a maximum achieved plasma current of $I_\zeta \leqslant 3 \, \text{MA}$. Up to 20 MW of deuterium neutral beam injection (NBI) and 6 MW of electron cy-

clotron resonance heating (ECH) are available for auxiliary plasma heating. DIII-D is perhaps the best-diagnosed tokamak in the world, with an extensive suite of equilibrium and fluctuation diagnostics.

1.1.2 *Fluctuation-induced transport in tokamak plasmas*

As hinted at above, the gradients in a fusion plasma can drive instability. These instabilities may be broadband or coherent in nature, and they may be beneficial, benign, or detrimental to confinement. The brief discussion below is *not* a thorough or exhaustive survey of these instabilities but is instead intended to provide a fusion-energy context to the diagnostic development pursued in this thesis.

The radial transport of particles, heat, and momentum in a tokamak plasma is often larger than that predicted by collisional (i.e. neoclassical) theory. There is strong evidence that this "anomalous" transport results from drift-wave turbulence driven by the free energy in the plasma gradients [6, 7]. Due to the relatively large size of its eddies, ion-scale ($k_\theta \rho_i \lesssim 1$) turbulence is often considered to be the most detrimental to confinement, but electron-scale ($k_\theta \rho_e \lesssim 1$) turbulence may be capable of forming radially elongated "streamers" [8] capable of driving experimentally relevant electron heat transport [9]. Here, k_θ is a typical poloidal wavenumber of the turbulent mode, $\rho_j = v_{tj}/|\Omega_j|$ is the gyroradius of species j, $v_{tj} = \sqrt{T_j/m_j}$ is the thermal speed of species j, and $\Omega_j = q_j B/m_j$ is the angular cyclotron frequency of species j. For the reactor-relevant scenario $T_e \sim T_i$, the characteristic length scale of ion-scale turbulence is a factor $(m_D/m_e)^{1/2} \approx 60$ larger than the characteristic length scale of electron-scale turbulence, where m_D is the deuteron mass. Until the very recent work of Howard *et al.* [10, 11, 12, 13, 14], self-consistently and simultaneously simulating both ion- and electron-scale turbulence with realistic mass ratios was computationally intractable. By exploiting the latest advances in high-performance computing, Howard *et al.*'s multiscale simulations indicate that cross-scale coupling can drive experimentally relevant levels of electron heat flux [10, 14] and that this cross-scale coupling becomes stronger as the drive for the electron-scale turbulence increases relative to that for the ion-scale turbulence [11, 12].

In addition to broadband turbulence, coherent fluctuations can also drive transport in a tokamak plasma. Of note are Alfvén waves that are driven unstable by resonant interactions with superthermal energetic particles [3, 15]. Such energetic particles sit in the tail of the plasma's distribution function, and they are generated, for example, from the ionization of injected neutral-beam particles, the acceleration

of ions by intense radio-frequency fields, or the fusion of two fuel ions. Nonlinearly, the Alfvén instability results in the loss of energetic particles from the plasma, decreasing the fusion rate; further, these lost particles may damage components inside the reactor. (As an interesting aside, Alfvén spectroscopy is an effective diagnostic for measuring the minimum value of the plasma's safety factor, q_{min} [3, 16, 17]). Of much practical interest to stable tokamak operation is the neoclassical tearing mode [3, Sec. 7.3], which can "lock" to the vessel wall [3, Sec. 7.10] and result in "disruption", an abrupt, violent termination of the plasma discharge [3, Sec. 7.7-7.9].

1.2 OPTICAL INTERFEROMETRY

Following its invention by Michelson in the 1880s [18], optical interferometry has had a long and illustrious history in fundamental and applied sciences, ranging from the famed Michelson-Morley experiment disproving the existence of a luminiferous ether [19], to Zernike's phase-contrast method for imaging phase objects of wide importance in biology [20], to the recent observation of gravitational waves by the LIGO collaboration [21, 22]. The brief discussion of interferometry below is intended to provide an instrumentation context to the diagnostic development pursued in this thesis and to equip the reader with a foundation for the more detailed discussion in Chapter 2.

Interferometry exploits the interference of light to measure variations in optical path length, where optical path length is defined as the product of the geometric path length l and the index of refraction N of the intervening medium. The phase $\delta\phi$ acquired by light propagating through optical path length $\delta(Nl)$ is $\delta\phi = k\delta(NL)$, where k is the in-medium wavenumber of the light; integrating over the geometric path yields the total acquired phase, ϕ. Thus, an incident electric field E_0 will be phase shifted as $E_0 e^{i\phi}$ after transiting such an optical path. Note that this phase shift does *not* alter the field intensity $I \propto |E_0 e^{i\phi}|^2 = |E_0|^2$; as most detectors are square-law, intensity detectors, such naive detection of the phase-shifted radiation provides no information about the phase ϕ or the underlying optical path. However, by interfering the phase-shifted field with a reference field $E_R = E_0 e^{i\phi_R}$ of known amplitude E_0 and known phase ϕ_R, the total field becomes

$$E_{tot} = E_R + E_0 e^{i\phi} = E_0 \left(e^{i\phi_R} + e^{i\phi} \right), \qquad (1.2)$$

with corresponding intensity (averaged over an optical cycle)

$$I \propto |E_{tot}|^2 = 2E_0^2 \left[1 + \cos\left(\phi - \phi_R\right) \right].\tag{1.3}$$

Now, the measurable intensity I is a function of the phase ϕ. The specification of the reference field E_R determines the interferometric method. Below, two interferometric methods — heterodyne interferometry and phase contrast imaging (PCI) — are briefly discussed.

1.2.1 *Heterodyne interferometry*

Heterodyne interferometry [23] uses an external reference beam whose angular frequency has been shifted by $\Delta\omega_0$ (i.e. the reference field is $E_R = E_0 e^{-i\Delta\omega_0 t}$) to produce measurable intensity variations (averaged over an optical cycle)

$$I_{het} \propto 2E_0^2 \left[1 + \cos\left(\Delta\omega_0 t + \phi\right) \right].\tag{1.4}$$

Note that the desired baseband phase information ϕ is shifted to an intermediate frequency $\Delta\omega_0$; quadrature heterodyne detection can then be used to extract an absolute measurement of ϕ [24]. The technical complications of converting to and from the intermediate frequency are often justified by the resulting ability to overcome the shortcomings of homodyne interferometry [23, 25]. Heterodyne interferometry is an established technique for measuring both the bulk and the fluctuating components of plasma density in contemporary tokamaks [24, 26, 27, 28] and is expected to provide similar capabilities in ITER [29] and other next-step devices.

1.2.2 *Phase contrast imaging (PCI)*

To motivate the development of phase contrast imaging (PCI), consider a phase ϕ that consists of a uniform, bulk contribution $\overline{\phi}$ and a small, spatially fluctuating contribution $\widetilde{\phi} \ll 1$ such that $\phi = \overline{\phi} + \widetilde{\phi}$. To lowest order in $\widetilde{\phi}$, the heterodyne intensity (1.4) becomes

$$I_{het} \propto 2E_0^2 \left\{ 1 + \left[\cos\left(\Delta\omega_0 t + \overline{\phi}\right) - \widetilde{\phi} \sin\left(\Delta\omega_0 t + \overline{\phi}\right) \right] \right\}.\tag{1.5}$$

As $\widetilde{\phi} \ll 1$, the component of the intensity corresponding to the fluctuation $\widetilde{\phi}$ is only a *small fraction* of the total intensity. Because every physical detector has a maximum tolerable intensity I_{max}, this implies that the fluctuation-induced component of the signal will only occupy

a small fraction of the detector's dynamic range, establishing a fundamental limit on the fluctuation sensitivity of a heterodyne interferometer.

If measurement of the fluctuation $\widetilde{\phi}$ is of primary importance, the interference scheme can be reconfigured to provide better sensitivity to fluctuations. Examine the phase-shifted field

$$E_0 e^{i\phi} = E_0 e^{i(\overline{\phi}+\widetilde{\phi})} \approx E_0 e^{i\overline{\phi}} \left(1 + i\widetilde{\phi}\right), \tag{1.6}$$

where only the lowest-order term in $\widetilde{\phi}$ has been retained. In the rightmost expression, the unity term corresponds to the bulk-phase contribution, while the $i\widetilde{\phi}$ term corresponds to the fluctuating-phase contribution. Now, imagine prescribing a reference field

$$E_R = -E_0 e^{i\overline{\phi}} \left(1 - i\sqrt{\eta}\right) \tag{1.7}$$

for $0 < \eta \leqslant 1$ such that the total field becomes

$$E_{tot} = E_R + E_0 e^{i\phi} \approx i E_0 e^{i\overline{\phi}} \left(\sqrt{\eta} + \widetilde{\phi}\right), \tag{1.8}$$

with corresponding intensity (averaged over an optical cycle)

$$I_{PCI} \propto |E_{tot}|^2 \approx E_0^2 \left(\eta + 2\sqrt{\eta}\widetilde{\phi}\right), \tag{1.9}$$

where only the lowest-order terms in $\widetilde{\phi}$ have been retained in the expressions for both E_{tot} and I_{PCI}. For a given fluctuation $\widetilde{\phi}$ and a given detector, such an interference scheme has an amplitude sensitivity to fluctuations that is $2\pi\sqrt{2/\eta}$ *better* than that of a comparable heterodyne interferometer (here, the power loss in the heterodyne-interferometer signal that necessarily results from the intermediate frequency-to-baseband conversion has been accounted for). Note, however, that information about the bulk phase $\overline{\phi}$ is lost.

While prescription of the above reference field may seem rather academic, this is precisely the means by which PCI operates, as may be guessed from the suggestive notation I_{PCI}. PCI "prescribes" such a reference field by spatially filtering the radiation pattern $E_0 e^{i\phi}$ in the focal plane of a focusing optic. Typically, the spatial filtering is performed with a "phase plate", an optical element that selectively applies (due to its geometry and its location at the focal plane) an appropriate phase delay and attenuation to the component of the field corresponding to the bulk phase $\overline{\phi}$. Thus, in contrast to the heterodyne interferometry, PCI uses an *internal* reference beam. Because

diffraction limits the focal-plane spot size, this spatial filtering necessarily involves a low-k cutoff, which in the ideal, diffraction-limited case corresponds to $k_{min}^{PCI} = 2/w$, where w is the $1/e$ electric-field radius of the probe beam [30]. The principles of PCI will be discussed in more detail in Section 2.5.

PCI is an established technique for measuring plasma density fluctuations in contemporary tokamaks [31, 32, 33]. PCI measurements have characterized fluctuations in numerous high-performance regimes [34, 35, 36, 37, 38], constrained the q-profile evolution between sawteeth [16], and provided validation of gyrokinetic [39, 40, 41] and RF simulations [42]. While not pursued in this thesis, additional spatial filtering can provide some degree of spatial localization to the line-integrated PCI measurement [30, 43].

1.2.3 *Heterodyne interferometry vs. PCI*

Colloquially, heterodyne interferometry is considered a "low-k" technique, and PCI is considered a "high-k" technique. It is worth pausing here to note that, for the same probe beam and the same fluctuation $\widetilde{\phi}$, the laser-plasma interaction is *identical* for both techniques. Further, the high-k optical capabilities of both systems are governed by the size of the collection optics and finite sampling-volume effects [44]. Thus, there is nothing that intrinsically limits heterodyne interferometry to low-k measurements — a heterodyne interferometer's high-k limit can be just as high, if not higher, than that of a given PCI system. However, as outlined in Section 1.2.2, PCI is *more sensitive* to fluctuations than a comparable heterodyne interferometer, and, assuming a Kolmogorov-like fluctuation spectrum $S(k) \propto k^{-p}$ for some positive p, PCI's superior sensitivity may allow it to detect high-k fluctuations that are too weak to be seen by a heterodyne interferometer. In practice, measurement of the high-frequency heterodyne signal often requires the use of faster, noisier detectors than are required for measurement of the PCI signal, further increasing PCI's sensitivity relative to heterodyne interferometry. At the low-k side of the spectrum, PCI is diffraction-limited to $k \geqslant 2/w$, where w is the $1/e$ electric-field radius of the probe beam [30], whereas the heterodyne interferometer's external reference beam allows detection even at $k = 0$.

1.3 MOTIVATION FOR A COMBINED PCI−INTERFEROMETER

This thesis details the design, construction, and operation of a combined PCI–interferometer on the DIII-D tokamak. This first-of-a-kind

Parameter	PCI	Interferometer
probe beam	single CO_2 beam	single CO_2 beam
$\Delta \mathbf{f}$	$10\,\text{kHz} < f < 2\,\text{MHz}$	$10\,\text{kHz} < f < 1\,\text{MHz}$
$\Delta \mathbf{k}$	$1.5\,\text{cm}^{-1} < k \leqslant 25\,\text{cm}^{-1}$	$0\,\text{cm}^{-1} \leqslant k < 5\,\text{cm}^{-1}$
sensitivity	$3 \times 10^{13}\,\text{m}^{-2}/\sqrt{\text{kHz}}$	$3 \times 10^{14}\,\text{m}^{-2}/\sqrt{\text{kHz}}$

Table 1.1: PCI and interferometry have compatible probe beams, comparable temporal bandwidths Δf, and *complementary* spatial bandwidths Δk. All parameters are for DIII-D's currently implemented combined PCI–interferometer system.

system uses a single $10.6\,\mu\text{m}$ CO_2 probe beam, two interference schemes, and two detectors to measure electron density fluctuations at large spatiotemporal bandwidth [45]. The parameters of this combined system are shown in Table 1.1. The motivations for this work are threefold:

- DIII-D's pre-existing PCI system has a low-k cutoff $k_{min}^{PCI} = 1.5\,\text{cm}^{-1}$. For a $1\,\text{keV}$ temperature typical of DIII-D's pedestal, this corresponds to $k_\theta \rho_i \gtrsim 0.25$, with larger values in the higher-temperature core. The addition of a heterodyne interferometer extends the minimum detectable wavenumber to $k = 0$, allowing simultaneous measurement of ion- and electron-scale instabilities. Such a capability could provide direct experimental validation of the significant cross-scale coupling predicted by Howard *et al.* [10, 11].

- Correlating measurements from the interferometer channel with those from DIII-D's pre-existing, toroidally separated ($\Delta \zeta = 45°$) interferometer provides a direct toroidal-mode-number (n) measurement of low-n magnetohydrodynamic (MHD) modes. This technique is capable of measuring mode numbers of *core-localized* MHD that is invisible to the external magnetic probes.

- ITER and other next-step devices are expected to have harsh neutron environments and limited port space. While such devices will almost certainly rely on heterodyne interferometry for density control and measurement of Alfvén instabilities [29], the allocation of port space for a PCI system is less clear. It is expected that these next-step devices will operate in strongly α-heated regimes (i.e. at or near ignition), which may be subject to novel turbulent regimes. This first-of-a-kind combined PCI–

interferometer demonstrates that simultaneous interferometric and PCI measurements can be made with a single probe beam and a set of ports, potentially paving the way for sensitive PCI measurements with similar combined systems on ITER and other next-step devices.

1.4 THESIS OUTLINE

The remainder of this thesis is organized as follows:

- Chapter 2 discusses the optical foundations of interferometry and phase contrast imaging (PCI), with an emphasis on sensitivity and spatiotemporal bandwidths.

- Chapter 3 provides an exhaustive set of design considerations for a heterodyne interferometer.

- Chapter 4 details the addition of a heterodyne interferometer to the pre-existing PCI system on the DIII-D tokamak; both systems operate simultaneously, sharing a single $10.6\,\mu m$ probe beam through the plasma.

- Chapter 5 discusses the measurement of multiscale turbulent fluctuations with the combined PCI-interferometer; linear-stability analysis and quasilinear-transport modeling are performed to help interpret these measurements.

- Chapter 6 demonstrates toroidal mode-number measurement of core-localized MHD via correlation of the newly installed interferometer with DIII-D's toroidally separated, pre-existing V2 interferometer.

- Chapter 7 concludes the thesis with a summary of the primary results and suggestions for future work.

1.5 UNITS

Unless explicitly stated, all formulas in this thesis are written in SI units. The one notable exception, of course, is the suppression of the Boltzmann constant in favor of expressing temperatures in units of energy (i.e. eV).

BIBLIOGRAPHY

[1] K. S. Krane. *Introductory Nuclear Physics*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 1988.

[2] J. Freidberg. *Plasma Physics and Fusion Energy*. Cambridge University Press, Cambridge, UK, 2007.

[3] J. Wesson. *Tokamaks*. Oxford University Press, Oxford, UK, 4th edition, 2011.

[4] A. R. Choudhuri. *Astrophysics for Physicists*. Cambridge University Press, Cambridge, United Kingdom, 2010.

[5] All-the-world's tokamaks. http://www.tokamak.info/. Accessed: 2017-09-29.

[6] W. Horton. Drift waves and transport. *Rev. Mod. Phys.*, 71(3):735–778, 1999.

[7] G. R. Tynan, A. Fujisawa and G. McKee. A review of experimental drift turbulence studies. *Plasma Phys. Control. Fusion*, 51(11):113001, 2009.

[8] W. Dorland, F. Jenko, M. Kotschenreuther and B. N. Rogers. Electron temperature gradient turbulence. *Phys. Rev. Lett.*, 85(26):5579–5582, 2000.

[9] F. Jenko and W. Dorland. Prediction of significant tokamak turbulence at electron gyroradius scales. *Phys. Rev. Lett.*, 89(22):225001, 2002.

[10] N. T. Howard, C. Holland, A. E. White, M. Greenwald and J. Candy. Synergistic cross-scale coupling of turbulence in a tokamak plasma. *Phys. Plasmas*, 21(11):112510, 2014.

[11] 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. *Nucl. Fusion*, 56(1):014004, 2016.

[12] N. T. Howard, C. Holland, A. E. White, M. Greenwald, J. Candy et al. Multi-scale gyrokinetic simulations: comparison with experiment and implications for predicting turbulence and transport. *Phys. Plasmas*, 23(5):056109, 2016.

[13] C. Holland, N.T. Howard and B.A. Grierson. Gyrokinetic predictions of multiscale transport in a DIII-D ITER baseline discharge. *Nucl. Fusion*, 57(6):066043, 2017.

[14] N T Howard, C Holland, A E White, M Greenwald, P Rodriguez-Fernandez et al. Multi-scale gyrokinetic simulations of an Alcator C-Mod, ELM-y H-mode plasma. *Plasma Phys. Control. Fusion*, 60(1):014034, 2018.

[15] W. W. Heidbrink. Basic physics of Alfvén instabilities driven by energetic particles in toroidally confined plasmas. *Phys. Plasmas*, 15(5):055501, 2008.

[16] E. M. Edlund, M. Porkolab, G. J. Kramer, L. Lin, Y. Lin et al. Observation of reversed shear Alfvén eigenmodes between sawtooth crashes in the Alcator C-Mod tokamak. *Phys. Rev. Lett.*, 102(16):165003, 2009.

[17] B. N. Breizman, M. S. Pekker and S. E. Sharapov. Plasma pressure effect on Alfvén cascade eigenmodes. *Phys. Plasmas*, 12(11):112506, 2005.

[18] The Nobel Prize in Physics 1907. https://www.nobelprize.org/nobel_prizes/physics/laureates/1907/. Accessed: 2017-09-30.

[19] A. A. Michelson and E. W. Morley. On the relative motion of the Earth and the luminiferous ether. *Am. J. Sci.*, 34(203):333–345, 1887.

[20] The Nobel Prize in Physics 1953. https://www.nobelprize.org/nobel_prizes/physics/laureates/1953/. Accessed: 2017-09-30.

[21] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.

[22] The Nobel Prize in Physics 2017. https://www.nobelprize.org/nobel_prizes/physics/laureates/2017/. Accessed: 2017-10-03.

[23] I. H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, Cambridge, United Kingdom, 2nd edition, 2002.

[24] T. N. Carlstrom, D. R. Ahlgren and J. Crosbie. Real-time, vibration-compensated CO_2 interferometer operation on the DIII-D tokamak. *Rev. Sci. Instrum.*, 59(7):1063–1066, 1988.

[25] R. Nazikian and L. E. Sharp. CO2 laser scintillation interferometer for the measurement of density fluctuations in plasma confinement devices. *Rev. Sci. Instrum.*, 58(11):2086–2091, 1987.

[26] M. A. Van Zeeland, G. J. Kramer, R. Nazikian, H. L. Berk, T. N. Carlstrom et al. Alfvén eigenmode observations on DIII-D via two-colour CO_2 interferometry. *Plasma Phys. Control. Fusion*, 47(9):L31, 2005.

[27] A. Mlynek, G. Pautasso, M. Maraschek, H. Eixenberger and the ASDEX Upgrade Team. Infrared interferometry with submicrosecond time resolution in massive gas injection experiments on ASDEX Upgrade. *Fusion Sci. Tech.*, 61(4):290–300, 2012.

[28] C. P. Kasten, J. H. Irby, R. Murray, A. E. White and D. C. Pace. A new interferometry-based electron density fluctuation diagnostic on Alcator C-Mod. *Rev. Sci. Instrum.*, 83(10):10E301, 2012.

[29] M. A. Van Zeeland, R. L. Boivin, D. L. Brower, T. N. Carlstrom, J. A. Chavez et al. Conceptual design of the tangentially viewing combined interferometer-polarimeter for ITER density measurements. *Rev. Sci. Instrum.*, 84(4):043501, 2013.

[30] J. R. Dorris, J. C. Rost and M. Porkolab. Localized measurement of short wavelength plasma fluctuations with the DIII-D phase contrast imaging diagnostic. *Rev. Sci. Instrum.*, 80(2):023503, 2009.

[31] S. Coda. *An experimental study of turbulence by phase contrast imaging in the DIII-D tokamak*. PhD thesis, MIT, 1997.

[32] A. Mazurenko. *Phase contrast imaging on the Alcator C-Mod tokamak*. PhD thesis, MIT, 2001.

[33] A. Marinoni. *Plasma fluctuation studies in the TCV tokamak: modeling of shaping effects and advanced diagnostic development*. PhD thesis, EPFL, 2009.

[34] S Coda, M Porkolab and K.H Burrell. Decorrelation of edge plasma turbulence at the transition from low- to high-confinement mode in the DIII–D tokamak. *Phys. Lett. A*, 273(1):125 – 131, 2000.

[35] S. Coda, M. Porkolab and K.H. Burrell. Characterization of density fluctuations during elms in the diii-d tokamak. *Nucl. Fusion*, 41(12):1885, 2001.

[36] J. C. Rost, M. Porkolab, J. Dorris and K. H. Burrell. Short wavelength turbulence generated by shear in the quiescent H-mode edge on DIII–D. *Phys. Plasmas*, 21(6):062306, 2014.

[37] A. Marinoni, J.C. Rost, M. Porkolab, A.E. Hubbard, T.H. Osborne et al. Characterization of density fluctuations during the search for an I-mode regime on the DIII-D tokamak. *Nucl. Fusion*, 55(9):093019, 2015.

[38] A. Mazurenko, M. Porkolab, D. Mossessian, J. A. Snipes, X. Q. Xu et al. Experimental and theoretical study of quasicoherent fluctuations in enhanced D_α plasmas in the Alcator C-Mod tokamak. *Phys. Rev. Lett.*, 89(22):225004, 2002.

[39] L. Lin, M. Porkolab, E. M. Edlund, J. C. Rost, C. L. Fiore et al. Studies of turbulence and transport in Alcator C-Mod H-mode plasmas with phase contrast imaging and comparisons with GYRO. *Phys. Plasmas*, 16(1):012502, 2009.

[40] J. C. Rost, L. Lin and M. Porkolab. Development of a synthetic phase contrast imaging diagnostic. *Phys. Plasmas*, 17(6):062506, 2010.

[41] P. Ennever, M. Porkolab, J. Candy, G. Staebler, M. L. Reinke et al. The effects of dilution on turbulence and transport in C-Mod ohmic plasmas and comparisons with gyrokinetic simulations. *Phys. Plasmas*, 22(7):072507, 2015.

[42] N. Tsujii, M. Porkolab, P. T. Bonoli, E. M. Edlund, P. C. Ennever et al. Validation of full-wave simulations for mode conversion of waves in the ion cyclotron range of frequencies with phase contrast imaging in Alcator C-Mod. *Phys. Plasmas*, 22(8):082502, 2015.

[43] L. Lin, E. M. Edlund, M. Porkolab, Y. Lin and S. J. Wukitch. Vertical localization of phase contrast imaging diagnostic in Alcator C-Mod. *Rev. Sci. Instrum.*, 77(10):10E918, 2006.

[44] R. V. Bravenec and A. J. Wootton. Effects of limited spatial resolution on fluctuation measurements (invited). *Rev. Sci. Instrum.*, 66(1):802–805, 1995.

[45] E. M. Davis, J. C. Rost, M. Porkolab, A. Marinoni and M. A. Van Zeeland. A phase contrast imaging–interferometer system for detection of multiscale electron density fluctuations on DIII-D. *Rev. Sci. Instrum.*, 87(11):11E117, 2016.

2

INTERFEROMETRIC METHODS FOR TOKAMAK-PLASMA FLUCTUATIONS

Interferometric methods exploit the interaction of electromagnetic waves with a plasma to ascertain properties of the plasma's density. Because surveying every flavor of interferometric method exceeds the scope of this work, only the interferometric methods of direct relevance to this work — external reference-beam interferometry and phase contrast imaging (PCI) — are discussed in detail. Further, the emphasis is diagnosis of the plasma's density fluctuations rather than its equilibrium density. The overview of optical interferometry in Section 1.2 is mandatory background for this chapter. While Section 1.2 cavalierly assumes various forms for the interfering radiation fields, this chapter provides first-principles descriptions of the laser-plasma interaction and optical systems needed to produce these radiation fields.

Below, Section 2.1 examines the propagation of electromagnetic waves through a "cold" plasma and derives an expression for the plasma-induced phase delay. Section 2.2 reveals that plasma-density fluctuations weakly upscatter and downscatter an incident Gaussian probe beam, while Section 2.3 describes how an imaging system manipulates these scattered beams. Section 2.4 shows that interfering the imaged probe radiation with an external reference beam is an effective technique for diagnosing plasma-density fluctuations and discusses the relative merits of homodyne versus heterodyne detection. Section 2.5 details the principles of phase contrast imaging (PCI), which does away with an external reference beam and instead spatially filters the scattered and unscattered beams to produce an informative interference signal. Section 2.6 concludes the chapter by synthesizing the preceding results and discussing the strengths and limitations of each interferometric technique.

2.1 ELECTROMAGNETIC WAVES IN A PLASMA

A great deal can be learned about a plasma by probing it with electromagnetic waves. Sections 2.1.1— 2.1.3 derive the index of refraction N for a cold, homogeneous plasma. Section 2.1.4 extends these results to inhomogeneous plasmas via the WKB approximation and assesses the validity of this approach for a CO_2 probe beam in a typical toka-

mak plasma, and Section 2.1.5 computes the resulting plasma-induced phase delay.

2.1.1   *Derivation of the wave equation*

The electric field $\boldsymbol{E}$ and the magnetic field $\boldsymbol{B}$ of an electromagnetic wave are coupled via Faraday's law

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t} \tag{2.1}$$

and Ampere's law

$$\nabla \times \boldsymbol{B} = \mu_0 \boldsymbol{J} + \mu_0 \epsilon_0 \frac{\partial \boldsymbol{E}}{\partial t}, \tag{2.2}$$

where ϵ_0 is the permittivity of free space, μ_0 is the permeability of free space, and $\boldsymbol{J}$ is the current density [1, Sec. I.1]. Note that (2.1) and (2.2) are the vacuum formulations of Faraday's law and Ampere's law; however, they can still be used to describe an electromagnetic wave's propagation through a medium *if* that medium's electromagnetic properties are explicitly accounted for in the current density $\boldsymbol{J}$ [1, Sec. I.4][2, Sec. 4.1]. Now, taking the curl of Faraday's law and then using Ampere's law to eliminate $\boldsymbol{B}$ yields the electric field's wave equation

$$\nabla^2 \boldsymbol{E} - \frac{1}{c^2} \frac{\partial^2 \boldsymbol{E}}{\partial t^2} = \nabla(\nabla \cdot \boldsymbol{E}) + \mu_0 \frac{\partial \boldsymbol{J}}{\partial t}. \tag{2.3}$$

2.1.2   *Wave equation in a homogeneous medium*

Fourier decomposing the electric field and the current density

$$\boldsymbol{E}(\boldsymbol{r}, t) = \frac{1}{(2\pi)^4} \int \boldsymbol{E}(\boldsymbol{k}', \omega') e^{i(\boldsymbol{k}' \cdot \boldsymbol{r} - \omega' t)} d\boldsymbol{k}' \, d\omega', \tag{2.4}$$

$$\boldsymbol{J}(\boldsymbol{r}, t) = \frac{1}{(2\pi)^4} \int \boldsymbol{J}(\boldsymbol{k}', \omega') e^{i(\boldsymbol{k}' \cdot \boldsymbol{r} - \omega' t)} d\boldsymbol{k}' \, d\omega'. \tag{2.5}$$

reduces the wave equation (2.3) to an algebraic equation for each Fourier component. In particular, for the Fourier mode described by $\boldsymbol{E}_0 e^{i(\boldsymbol{k} \cdot \boldsymbol{r} - \omega t)}$ and $\boldsymbol{J}_0 e^{i(\boldsymbol{k} \cdot \boldsymbol{r} - \omega t)}$, the derivative operators become $\nabla \to i\boldsymbol{k}$ and $\partial/\partial t \to -i\omega$, and the wave equation reduces to

$$\left(\boldsymbol{N}\boldsymbol{N} - N^2 \boldsymbol{I} + \boldsymbol{I} \right) \cdot \boldsymbol{E}_0 = -\frac{i\boldsymbol{J}_0}{\epsilon_0 \omega}, \tag{2.6}$$

where

$$N \equiv \frac{c\boldsymbol{k}}{\omega} \qquad (2.7)$$

is the index of refraction seen by the given Fourier mode and $\boldsymbol{I}$ is the identity matrix. Now, assume that the medium is homogeneous in space and time [3, Sec. 3-2] such that the current density is easily related to the electric field via Ohm's law

$$\boldsymbol{J}(\boldsymbol{k},\omega) = \boldsymbol{\sigma}(\boldsymbol{k},\omega) \cdot \boldsymbol{E}(\boldsymbol{k},\omega), \qquad (2.8)$$

where $\boldsymbol{\sigma}$ is the conductivity of the surrounding medium [1, Sec. I.4]. Substituting the Ohm's law current density (2.8) into the Fourier decomposed wave equation (2.6) yields the eigenvalue equation

$$\left[\boldsymbol{N}\boldsymbol{N} - N^2 \boldsymbol{I} + \left(\boldsymbol{I} + \frac{i\boldsymbol{\sigma}}{\varepsilon_0 \omega} \right) \right] \cdot \boldsymbol{E}_0 = 0. \qquad (2.9)$$

To proceed further, a model is needed for the medium's conductivity.

2.1.3 *The cold-plasma index of refraction*

Although plasmas in contemporary fusion devices routinely approach temperatures $\lesssim 100$ million degrees Celsius ($\sim 10\times$ the temperature of the core of the sun!), the thermal velocities of the constituent particles are still far below the speed of light. The lightest, and consequently the fastest, of such a plasma's constituent particles are its electrons, which have thermal velocities $\lesssim 0.1c$. In contrast, as will be shown shortly, the electromagnetic waves used to make interferometric measurements in such plasmas have phase velocities very close to the speed of light. Thus, in the context of the wave-plasma interaction, the unperturbed plasma can be modeled as a collection of motionless (i.e. zero-temperature or "cold") charged particles. Consequently, the pressure of this model plasma is zero. For the present application, it is also appropriate to neglect the perturbed motion of the ions, whose inertia greatly exceeds that of the electrons, and to neglect collisions, which are relatively rare in fusion plasmas. The electron-fluid momentum equation for such a plasma is

$$m_e n_e \frac{d\boldsymbol{v}_e}{dt} = -e n_e \left(\boldsymbol{E} + \boldsymbol{v}_e \times \boldsymbol{B} \right), \qquad (2.10)$$

where n_e is the electron density, $\boldsymbol{v}_e$ is the perturbed electron-fluid velocity, and $d/dt = \partial/\partial t + (\boldsymbol{v}_e \cdot \nabla)$ is the advective time derivative.

Linearizing and Fourier analyzing the electron-fluid momentum equation (2.10) yields

$$i\omega m_e \boldsymbol{v}_e = e\left(\boldsymbol{E} + \boldsymbol{v}_e \times \boldsymbol{B}_0\right), \tag{2.11}$$

where $\boldsymbol{B}_0$ is the equilibrium magnetic field. Let $\boldsymbol{B}_0 = B_0 \hat{\boldsymbol{z}}$. Then, the perturbed electron-fluid velocity $\boldsymbol{v}_e$ is easily found to be [2, Sec. 4.1.2]

$$v_{e,x} = \frac{-ie}{\omega m_e} \frac{1}{1 - \Omega_e^2/\omega^2} \left(E_x - i\frac{\Omega_e}{\omega}E_y\right), \tag{2.12}$$

$$v_{e,y} = \frac{-ie}{\omega m_e} \frac{1}{1 - \Omega_e^2/\omega^2} \left(i\frac{\Omega_e}{\omega}E_x + E_y\right), \tag{2.13}$$

$$v_{e,z} = \frac{-ie}{\omega m_e}E_z, \tag{2.14}$$

where

$$\Omega_e \equiv \frac{eB_0}{m_e} \tag{2.15}$$

is the electron cyclotron frequency. Having neglected the ion's motion, the current density is simply

$$\boldsymbol{J} = -en_e \boldsymbol{v}_e. \tag{2.16}$$

Equating the current densities from (2.8) and (2.16) and substituting the solution for $\boldsymbol{v}_e$ yields the plasma's conductivity [2, Sec. 4.1.2]

$$\boldsymbol{\sigma} = \frac{in_e e^2}{m_e \omega} \frac{1}{1 - \Omega_e^2/\omega^2} \begin{pmatrix} 1 & -i\Omega_e/\omega & 0 \\ i\Omega_e/\omega & 1 & 0 \\ 0 & 0 & 1 - \Omega_e^2/\omega^2 \end{pmatrix}. \tag{2.17}$$

Choose axes such that

$$\boldsymbol{k} = k(\sin\theta\,\hat{\boldsymbol{y}} + \cos\theta\,\hat{\boldsymbol{z}}). \tag{2.18}$$

Then, substituting the cold-plasma conductivity (2.17) into the electric field's eigenvalue equation (2.9) and solving for the corresponding eigenvalues yields the well-known Appleton-Hartree formula for the plasma's index of refraction [2, Sec. 4.1.2]

$$N^2 = 1 - \frac{X(1-X)}{1 - X - \frac{1}{2}Y^2 \sin^2\theta \pm \Delta}, \tag{2.19}$$

where

$$X \equiv \frac{\omega_{pe}^2}{\omega^2}, \tag{2.20}$$

$$Y \equiv \frac{\Omega_e}{\omega}, \tag{2.21}$$

$$\Delta = \left[\left(\frac{1}{2} Y^2 \sin^2 \theta \right)^2 + (1 - X)^2 Y^2 \cos^2 \theta \right]^{1/2}, \tag{2.22}$$

$$\omega_{pe}^2 \equiv \frac{n_e e^2}{m_e \varepsilon_0}; \tag{2.23}$$

here, ω_{pe} is referred to as the angular electron plasma frequency.

The refractive index formula can be dramatically simplified for high-frequency electromagnetic waves propagating in typical tokamak plasmas. For example, a CO_2 beam ($\omega = 2\pi \cdot 28.3\,\text{THz}$) propagating in a typical DIII-D plasma sees

$$n_e \lesssim 10^{20}\,\text{m}^{-3} \quad \Rightarrow \quad X \lesssim 10^{-5},$$
$$B \lesssim 2\,\text{T} \quad \Rightarrow \quad Y \lesssim 2 \times 10^{-3}.$$

Now, the smallness of X and Y can be exploited to approximate the Appleton-Hartree index of refraction (2.19) by retaining only the terms that are linear in X or linear in Y; this yields $N^2 \approx 1 - X$ or, equivalently,

$$N \approx 1 - \frac{X}{2}. \tag{2.24}$$

Note that the corresponding phase velocity is $v_{ph} = c/N \approx c$. Thus, the cold-plasma assumption that the wave's phase velocity is much larger than the thermal velocities ($\lesssim 0.1c$) of the plasma's constituent particles is valid.

It is enlightening to also examine the corresponding electric-field eigenvectors. Substituting the index of refraction (2.24) into the corresponding eigenvalue equation (2.9) and solving to lowest order for the corresponding eigenvector yields

$$\boldsymbol{E}_0 \approx E_{0,x}\hat{\boldsymbol{x}} + E_{0,yz}(\hat{\boldsymbol{k}} \times \hat{\boldsymbol{x}}), \tag{2.25}$$

where $\hat{\boldsymbol{k}} = \boldsymbol{k}/|\boldsymbol{k}|$ is the unit vector corresponding to the wavevector defined in (2.18), and $E_{0,x}$ and $E_{0,yz}$ are arbitrary, uncoupled constants that are independent of the plasma's properties. Thus, to lowest order, a CO_2 beam in a tokamak plasma propagates as a transverse electromagnetic wave ($\boldsymbol{k} \cdot \boldsymbol{E}_0 \approx 0$) with near-constant polarization.

Finally, it should be noted for completeness that while cold-plasma theory is adequate for the current purposes, finite-temperature and relativistic effects will affect the interpretation of refractive index measurements in a fusion reactor. For example, the thermal correction to the cold-plasma interferometric phase in a $T_e \approx 10\,\text{keV}$ plasma is -3%, with slightly more substantial corrections for polarimetric measurements [4].

2.1.4   *Wave propagation in an inhomogeneous medium*

No physical plasma is truly homogeneous in space. While the above Fourier approach can still be employed, the inhomogeneities tend to couple the various modes together, greatly complicating the analysis. However, if the plasma varies sufficiently slowly, the plasma can be treated as *locally* uniform, and the wave field can be "stitched" together as the wave propagates from initial position $r^{(i)}$ to position r via the WKB approximation [3, Ch. 13][5, Ch. 8]

$$E(r, t) \approx E_0 \exp\left[i \left(\int_{r^{(i)}}^{r} k(r') \cdot dl' - \omega t \right) \right]. \tag{2.26}$$

Here, $k(r')$ is the local wavevector and the integration is performed along the wave's trajectory. Note that the amplitude variation and the reflected wave that are characteristic of the WKB approximation have been neglected because densities in a tokamak plasma are much less than a CO_2 beam's $\sim 10^{25}\,\text{m}^{-3}$ density cutoff.

What exactly is meant by "sufficiently slowly", though? To approximate the plasma as locally uniform, the change in the wavenumber δk over one wavelength λ should be small relative to the wavenumber k. Note that the change in wavenumber over one wavelength is $\delta k = \nabla k \cdot \lambda = 2\pi \nabla k / k$. Thus, the WKB validity criterion that $\delta k / k \ll 1$ becomes

$$\frac{|\nabla k|}{k^2} \ll \frac{1}{2\pi}. \tag{2.27}$$

Now, using the definition of N from (2.7), the cold-plasma index of refraction in (2.24) can be rewritten as $k = (\omega/c)(1 - X/2)$, and, to lowest order,

$$\nabla k \approx - \left(\frac{2\pi r_e}{k} \right) \nabla n_e, \tag{2.28}$$

where

$$r_e = \frac{e^2}{4\pi\varepsilon_0 m_e c^2} = 2.8 \times 10^{-15} \text{ m} \tag{2.29}$$

is the classical electron radius. The most extreme density gradients in a tokamak often occur in the so-called "pedestal", where the density changes by $\Delta n_e \lesssim 10^{20} \text{ m}^{-3}$ over a scale length $\Delta r \gtrsim 1\,\text{cm}$; for a CO_2 laser beam ($k = 5.9 \times 10^5 \text{ m}^{-1}$) propagating through such a pedestal,

$$\frac{|\nabla k|}{k^2} = \left(\frac{2\pi r_e}{k^3}\right) |\nabla n_e| \lesssim 10^{-9}$$

such that the criterion for WKB validity (2.27) is very well-satisfied. The WKB validity criterion can also be evaluated for a CO_2 beam propagating through plasma-density fluctuations; assuming $\widetilde{n}_e/\overline{n}_e \sim 10^{-3}$, $\overline{n}_e \lesssim 10^{20} \text{ m}^{-3}$, and density-fluctuation wavenumbers $\lesssim 30\,\text{cm}^{-1}$, $|\nabla k|/k^2 \lesssim 2 \times 10^{-11}$ such that the WKB validity criterion is also very well-satisfied for typical plasma-density fluctuations.

2.1.5 *Plasma-induced phase delay*

The WKB field solution (2.26) indicates that the wave's phase at a given point in space and time is determined by the properties of the medium that the wave has passed through. In particular, a CO_2 laser beam propagating through a tokamak plasma will acquire a phase shift ϕ relative to vacuum given by

$$\begin{aligned}
\phi &= \int \left[k(\boldsymbol{r}) - \frac{\omega}{c}\right] dl \\
&= \frac{\omega}{c} \int [N(\boldsymbol{r}) - 1]\, dl \\
&\approx \frac{\omega}{c} \int \left[\left(1 - \frac{X}{2}\right) - 1\right] dl \\
&= -r_e \lambda_0 \int n_e\, dl,
\end{aligned} \tag{2.30}$$

where the index of refraction has been approximated via (2.24), $\lambda_0 = 2\pi c/\omega = 10.6\,\mu\text{m}$ is the CO_2 beam's vacuum wavelength, and r_e is again the classical electron radius defined in (2.29). Now, if the plasma fluctuates about its equilibrium density $\overline{n}_e$ as $n_e = \overline{n}_e + \widetilde{n}_e$, the phase will similarly fluctuate about its equilibrium as $\phi = \overline{\phi} + \widetilde{\phi}$ where

$$\widetilde{\phi} = -r_e \lambda_0 \int \widetilde{n}_e\, dl. \tag{2.31}$$

It is precisely the intent of Section 2.2 to determine how such density fluctuations interact with an incident Gaussian probe beam. As a final note before leaving this section, accurate interpretation of very rapid phase-fluctuation measurements (such as those resulting from RF-wave perturbations) may require accounting for the beam's finite transit time through the plasma to the point of measurement [6, Sec. 3.1].

2.2   DIFFRACTION OF A GAUSSIAN PROBE BEAM

Lasers are almost always well-collimated enough that they are amenable to analysis in the paraxial limit. Gaussian beams are exact solutions to the paraxial wave equation in free space, and they are very good approximations to the eigenmodes observed in real lasers [7, Ch. 16]. For this reason, it is reasonable to investigate the interaction of a Gaussian probe beam with a plasma-density fluctuation.

2.2.1   *Definition of a Gaussian beam*

A Gaussian beam of angular frequency ω_0 propagating along the z-axis in a medium with index of refraction N has an electric field

$$E_G(\mathbf{r}, t) = E_G(\mathbf{r})e^{-i\omega_0 t} \tag{2.32}$$

with spatial dependence [7, Ch. 17]

$$
\begin{aligned}
E_G(\mathbf{r}) = E_0 \frac{w_0}{w(z)} &\exp\left[\frac{-\rho^2}{w(z)^2}\right] \\
&\times \exp\left\{i\left[Nk_0 z + \frac{Nk_0\rho^2}{2R(z)} - \psi(z)\right]\right\}.
\end{aligned}
\tag{2.33}
$$

Here, $\rho = (x^2 + y^2)^{1/2}$ is the transverse distance from the z-axis (i.e. the beam's "symmetry axis"), w_0 is the radius of the beam's waist, and $k_0 = \omega_0/c = 2\pi/\lambda_0$ is the beam's vacuum wavenumber. The beam's width $w(z)$, radius of curvature $R(z)$, and Gouy phase $\psi(z)$ are defined as

$$w(z) = w_0 \left[1 + \left(\frac{z}{z_R}\right)^2\right]^{1/2}, \tag{2.34}$$

$$R(z) = z\left[1 + \left(\frac{z_R}{z}\right)^2\right], \tag{2.35}$$

$$\psi(z) = \tan^{-1}\left(\frac{z}{z_R}\right), \tag{2.36}$$

where the Rayleigh range

$$z_R \equiv \left(\frac{\pi w_0^2}{\lambda_0} \right) N \qquad (2.37)$$

is the nominal division between the beam's near-field ($|z| \ll z_R$) and far-field ($|z| \gg z_R$) behaviors. Note that the beam's waist sits at $z = 0$.

From a Fourier perspective, a Gaussian beam can be decomposed into a set of infinite plane waves traveling in slightly different directions [7, Ch. 16.7]. The plane waves with the beam's nominal wavevector $\boldsymbol{k}_0 = k_0 \hat{\boldsymbol{z}}$ propagate along the beam's symmetry axis, while the plane waves with non-zero transverse wavevectors produce beam divergence (i.e. free-space diffraction of the beam). This behavior can be understood as an uncertainty principle, with the beam's finite transverse dimensions necessarily requiring a finite spread in transverse wavevectors.

2.2.2 *Diffraction of a Gaussian beam from plasma-density fluctuations*

As discussed in the text surrounding (2.25), a CO_2 probe beam in a tokamak plasma propagates as a transverse electromagnetic wave with near-constant polarization (any small changes to the beam polarization are of little practical interest to the present work). Thus, a *scalar* diffraction theory is sufficient to describe the relevant aspects of the beam-plasma interaction; the details of this theory are provided in Appendix A, and the relevant results are summarized below.

A Gaussian CO_2 probe beam propagating through a tokamak plasma acquires a plasma-induced phase delay $\phi(\boldsymbol{\rho}', t)$ given by (2.30), where $\boldsymbol{\rho}'$ corresponds to the beam's transverse dimensions. Explicitly dividing ϕ into bulk $\overline{\phi}(t)$ and spatially varying $\widetilde{\phi}(\boldsymbol{\rho}', t)$ components, the plasma-induced phase delay becomes

$$\phi(\boldsymbol{\rho}', t) = \overline{\phi}(t) + \widetilde{\phi}(\boldsymbol{\rho}', t). \qquad (2.38)$$

Typically, $\widetilde{\phi}$ varies on much faster time scales than $\overline{\phi}$, but this is not required. The spatial variation of the plasma-induced phase delay contributes to the diffraction of the incident Gaussian probe beam.

The response functions of the diagnostics investigated in Sections 2.4 and 2.5 will be shown to be linear in their regimes of relevance, so it is sufficient to examine diffraction from phase fluctuations $\widetilde{\phi}$ consisting of a single Fourier mode

$$\widetilde{\phi}(\boldsymbol{\rho}', t) = \widetilde{\phi}_0 \cos(k x' - \omega t). \qquad (2.39)$$

Here, without loss of generality, it is assumed that the CO_2 probe beam is propagating in the +z-direction, and the density fluctuation corresponding to $\widetilde{\phi}$ is propagating in the x-direction. Then, following the formalism pioneered by Raman and Nath [8, 9] and detailed in Appendix A, the diffracted electric field can be written (see (A.21)) as a discrete sum of scattered beams

$$E(\boldsymbol{r}, t) \approx e^{i\overline{\phi}} \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0) \right] E_G(\boldsymbol{r}_m) e^{-i(\omega_0 + m\omega)t}, \qquad (2.40)$$

where J_m is the m^{th} Bessel function of the first kind and $E_G(\boldsymbol{r}_m)$ is the spatial dependence of the m^{th} scattered beam. Note that $\boldsymbol{r}_m = [\boldsymbol{R}(\theta_m)]\boldsymbol{r}$ is the native coordinate system of the m^{th} scattered beam expressed in lab-frame coordinates; here, $\theta_m \approx mk/k_0$ is the angle at which the m^{th} scattered beam propagates relative to the lab-frame optical axis, and

$$\boldsymbol{R}(\theta) = \begin{pmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{pmatrix} \qquad (2.41)$$

is the rotation matrix that rotates the (x, z)-plane about the y-axis by angle θ.

It is worth pausing to discuss the physical significance of the diffracted electric field (2.40). The assumed sinusoidal phase modulation (2.39) diffracts an incident Gaussian beam predominantly into downscattered ($m = -1$), unscattered ($m = 0$), and upscattered ($m = 1$) Gaussian beams. The incident beam is coupled into the m^{th} scattered beam with strength $J_m(\widetilde{\phi}_0)$. The m^{th} scattered beam is Doppler shifted relative to the incident beam by $m\omega$ and propagates at an angle $\theta_m \approx mk/k_0$ relative to the lab-frame optical axis. The scattering is very nearly elastic (i.e. $\omega/\omega_0 \lesssim 1\,\text{GHz}/28.3\,\text{THz} \sim 10^{-5}$), so $|\boldsymbol{k}_{0,m}| = k_0$ is a very good approximation. This constraint of elasticity coupled with knowledge of the scattering angle θ_m allows determination of the scattered wavevector

$$\boldsymbol{k}_{0,m} = (mk)\hat{\boldsymbol{x}} + k_0 \left[1 - \left(\frac{mk}{k_0} \right)^2 \right]^{1/2} \hat{\boldsymbol{z}}. \qquad (2.42)$$

Finally, note that the simultaneous presence of both the upscattered and downscattered beams (a key prediction of the Raman-Nath for-

malism) under typical experimental conditions has been demonstrated empirically [10, Sec. 2.1].

2.2.3 *Wavenumber filtering of the diffracted field*

Some interferometric techniques, such as phase contrast imaging (PCI; discussed in Section 2.5), filter the Fourier wavenumber content of the diffracted electric field (2.40). The mathematical formalism for such filtering is detailed in Appendix A, but the relevant results are briefly summarized here. Assume that the filtering can be described by a transfer function $T(k_x)$, where k_x is the wavevector component in the lab-frame x-direction. Then, the filtered electric field can be written (see (A.36)) as

$$E(\boldsymbol{r}, t) \approx e^{i\overline{\phi}} \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0) \right] E_T(\boldsymbol{r}_m) e^{-i(\omega_0 + m\omega)t}, \qquad (2.43)$$

where

$$E_T(\boldsymbol{r}_m) \approx E_G(0, y_m, z_m) \cdot \mathcal{E}(\boldsymbol{r}_m, k), \qquad (2.44)$$

is the m^{th} scattered beam following the wavenumber-filtering transformation, and

$$\begin{aligned}
\mathcal{E}(\boldsymbol{r}_m, k) = &\frac{e^{-imkx_m}}{2\pi} \\
&\times \int dx' \exp\left[\frac{-x'^2}{w(z_m)^2} \right] \exp\left\{ i \left[mkx' + \frac{k_0 x'^2}{2R(z_m)} \right] \right\} \quad (2.45) \\
&\times \int dk_x \, T(k_x) e^{ik_x(x_m - x')}
\end{aligned}$$

is a complex-valued function that describes the amplitude and phase transformations that result from filtering the m^{th} scattered beam by $T(k_x)$. When there is no wavenumber filtering (i.e. $T(k_x) = 1$), the m^{th} transformed beam (2.44) readily reduces to the m^{th} scattered Gaussian beam (i.e. $E_T(\boldsymbol{r}_m) = E_G(\boldsymbol{r}_m)$), in agreement with expectations.

2.3 IMAGING OF THE DIFFRACTED FIELD

It is often desirable to *image* the above diffracted field in order to determine the spatiotemporal aspects of the responsible phase fluctuations. Below, the relevant aspects of imaging systems are briefly reviewed.

The imaged field is then computed and examined under typical experimental limits.

2.3.1 *Imaging systems*

Appendix B details the geometric optics and Gaussian-beam transformations of relevance to imaging systems; here, the directly applicable aspects are briefly summarized.

Let the optical axis of an arbitrary optical system lie along the z-axis, and let all optical rays lie in a plane with the optical axis. At a given position z_j, an optical ray is fully described by its transverse distance ρ to the optical axis and its slope $d\rho/dz$ [7, Ch. 15]. An imaging system $\mathcal{I}$, by definition, redirects all rays emanating from transverse position $\rho_\mathcal{O}$ in the object plane $S_\mathcal{O}$ to intersect at transverse position

$$\rho_\mathcal{I} = M\rho_\mathcal{O} \tag{2.46}$$

in the image plane $S_\mathcal{I}$. Here, M is the *magnification* of the imaging system, and $M < 0$ implies that the image is inverted relative to the object. Because the symmetry axis of a Gaussian beam behaves as a ray in the geometric-optics sense [11], the symmetry axes of the unscattered and scattered beams intersect in the image plane. Then, varying inversely with the spatial scale, the wavenumber k of the object-plane phase fluctuation (2.39) is imaged as

$$k_\mathcal{I} \equiv \frac{k}{M}. \tag{2.47}$$

The native coordinate systems of the m^th scattered beam and the unscattered beam do *not* align in the image plane (i.e. $r_{m,\mathcal{I}} \neq r_\mathcal{I}$). Considerations from geometric and Gaussian-beam optics show that the image-plane coordinate systems are related as

$$\begin{pmatrix} x_{m,\mathcal{I}} \\ y_{m,\mathcal{I}} \\ z_{m,\mathcal{I}} \end{pmatrix} = \begin{pmatrix} x_\mathcal{I} \cos\left(\frac{\theta_m}{M}\right) \\ y_\mathcal{I} \\ z_\mathcal{I} + x_\mathcal{I} \sin\left(\frac{\theta_m}{M}\right) \end{pmatrix} \approx \begin{pmatrix} x_\mathcal{I} \\ y_\mathcal{I} \\ z_\mathcal{I} + x_\mathcal{I}\left(\frac{\theta_m}{M}\right) \end{pmatrix}, \tag{2.48}$$

where the approximation is valid to first order in θ_m/M.

2.3.2 *The imaged field*

Let $\mathcal{I}$ image the object plane $S_\mathcal{O}$ such that the diffracted field (2.43) is imaged as

$$
\begin{aligned}
E(\boldsymbol{r}_\mathcal{I}, t) &= \mathcal{I}[E(\boldsymbol{r}_\mathcal{O}, t)] \\
&= e^{i\overline{\Phi}} \sum_{m=-1}^{1} \left\{ i^m \left[J_m(\widetilde{\phi}_0) \right] \right. \\
&\qquad\qquad \left. \times \mathcal{I}[E_T(\boldsymbol{r}_{m,\mathcal{O}})] e^{-i(\omega_0 + m\omega)t} \right\},
\end{aligned}
\tag{2.49}
$$

where the $\mathcal{I}$ operator has been brought within the summation because the scattered beams obey the superposition principle when propagating through the imaging system (it is assumed that typical beam intensities are far from any nonlinear thresholds in the imaging-system materials). The effect of $\mathcal{I}$ is to propagate each scattered beam from its object-plane coordinates $\boldsymbol{r}_{m,\mathcal{O}}$ to its corresponding image-plane coordinates $\boldsymbol{r}_{m,\mathcal{I}}$. Thus, referencing the definition of $E_T(\boldsymbol{r}_m)$ in (2.44), it readily follows that

$$
\begin{aligned}
\mathcal{I}[E_T(\boldsymbol{r}_{m,\mathcal{O}})] &= \mathcal{I}[E_G(0, y_{m,\mathcal{O}}, z_{m,\mathcal{O}}) \cdot \mathcal{E}(\boldsymbol{r}_{m,\mathcal{O}}, k)] \\
&= E_G(0, y_{m,\mathcal{I}}, z_{m,\mathcal{I}}) \cdot \mathcal{E}(\boldsymbol{r}_{m,\mathcal{I}}, k_\mathcal{I}) \\
&\approx E_G(0, y_\mathcal{I}, z_\mathcal{I}) e^{imk_\mathcal{I}x_\mathcal{I}} \cdot \mathcal{E}(\boldsymbol{r}_{m,\mathcal{I}}, k_\mathcal{I}),
\end{aligned}
\tag{2.50}
$$

where the last step naturally follows from the image-plane coordinate transformation (2.48) and the following approximations: $w(z_{m,\mathcal{I}}) \approx w(z_\mathcal{I})$, $R(z_{m,\mathcal{I}}) \approx R(z_\mathcal{I})$, and $\psi(z_{m,\mathcal{I}}) \approx \psi(z_\mathcal{I})$. Thus, the imaged field (2.49) becomes

$$
\begin{aligned}
E(\boldsymbol{r}_\mathcal{I}, t) &= E_G(0, y_\mathcal{I}, z_\mathcal{I}, t) e^{i\overline{\Phi}} \\
&\qquad \times \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0) \right] \mathcal{E}(\boldsymbol{r}_{m,\mathcal{I}}, k_\mathcal{I}) e^{im\nu},
\end{aligned}
\tag{2.51}
$$

where

$$
\nu = k_\mathcal{I} x_\mathcal{I} - \omega t
\tag{2.52}
$$

and $E_G(0, y_\mathcal{I}, z_\mathcal{I}, t) = E_G(0, y_\mathcal{I}, z_\mathcal{I}) e^{-i\omega_0 t}$.

2.3.3   *The weak-coupling limit*

Typically, the phase-fluctuation amplitude is very small ($\widetilde{\phi}_0 \ll 1$), and the Bessel function's small-argument limiting form [12] can be used

$$
\lim_{z \to 0} J_m(z) \sim
\begin{cases}
\frac{1}{\Gamma(m+1)} \left(\frac{z}{2}\right)^m, & m = 0, 1, 2, 3, \cdots, \\[2mm]
\frac{1}{\Gamma(|m|+1)} \left(\frac{-z}{2}\right)^{|m|}, & m = -1, -2, -3, \cdots.
\end{cases}
\tag{2.53}
$$

Here, Γ is the gamma function, and $\Gamma(m+1) = m!$ for positive integer m. Introducing the notational shorthand $\mathcal{E}_m \equiv \mathcal{E}(r_{m,J}, k_J)$, the imaged field to first order in $\widetilde{\phi}_0$ (2.51) becomes

$$
E(r_J, t) \approx E_G(0, y_J, z_J, t) e^{i\overline{\phi}}
$$
$$
\times \left\{ \mathcal{E}_0 + i\frac{\widetilde{\phi}_0}{2} \left[\mathcal{E}_1 e^{i\nu} + \mathcal{E}_{-1} e^{-i\nu} \right] \right\}.
\tag{2.54}
$$

When there is no wavenumber filtering of the diffracted field (i.e. $T(k_x) = 1$), the imaged field (2.54) readily reduces to

$$
E(r_J, t) \approx E_G(r_J, t) e^{i\overline{\phi}} \left[1 + i\widetilde{\phi}_0 \cos\nu \right],
\tag{2.55}
$$

which follows naturally from the image-plane coordinate transformation (2.48) and the following approximations: $w(z_{m,J}) \approx w(z_J)$ and $R(z_{m,J}) \approx R(z_J)$. Note that (2.55) is *linear* in the imaged phase fluctuation $\widetilde{\phi}_0 \cos\nu$. By the principle of linear superposition, then, (2.55) can be readily generalized to

$$
E(r_J, t) \approx E_G(r_J, t) e^{i\overline{\phi}} \left[1 + i\widetilde{\phi}(x_J, t) \right],
\tag{2.56}
$$

where $\widetilde{\phi}(x, t)$ is a real-valued but otherwise arbitrary phase fluctuation.

2.3.4   *The need for a reference beam*

Assume that there is no wavenumber filtering of the diffracted field such that (2.56) yields the imaged field. Now, most detectors of interest are square-law detectors in that they produce a response proportional to the square (i.e. intensity) of the impinging field. If such

a detector is placed at the image plane, the local intensity (averaged over an optical cycle) is

$$
\begin{aligned}
I(\boldsymbol{r}_{\mathfrak{I}}, t) &\equiv \frac{c\varepsilon_0}{2} |E(\boldsymbol{r}_{\mathfrak{I}}, t)|^2 \\
&= I_G(\boldsymbol{r}_{\mathfrak{I}}) \left[1 + \mathcal{O}\left(\left[\widetilde{\phi}(x_{\mathfrak{I}}, t) \right]^2 \right) \right],
\end{aligned}
\tag{2.57}
$$

where

$$
I_G(\boldsymbol{r}_{\mathfrak{I}}) = \frac{c\varepsilon_0 |E_G(\boldsymbol{r}_{\mathfrak{I}})|^2}{2}
\tag{2.58}
$$

is the intensity profile (averaged over an optical cycle) of the unscattered Gaussian beam. Thus, for the typical $|\widetilde{\phi}(x_{\mathfrak{I}}, t)| \ll 1$ limit, the measured response will be very weak if the detector is only exposed to the imaged radiation. Physically, this is attributable to the $m = \pm 1$ beams being $\pi/2$ out of phase with the unscattered $m = 0$ beam. As will be shown in Sections 2.4 and 2.5, the system response can be substantially improved by interfering the scattered beams with a reference beam. The generation of such a reference beam, then, is the trait that differentiates one interferometric method from another.

2.4 EXTERNAL REFERENCE-BEAM INTERFEROMETRY

In external reference-beam interferometry, there is no wavenumber filtering of the diffracted field such that (2.56) yields the imaged field. Thus, the probe radiation in the image plane is given as

$$
E_P(\boldsymbol{r}_{\mathfrak{I}}, t) \approx E_G(\boldsymbol{r}_{\mathfrak{I}}, t) e^{i\overline{\phi}} \left[1 + i\widetilde{\phi}(x_{\mathfrak{I}}, t) \right].
\tag{2.59}
$$

Now, assume that the imaged probe radiation is interfered with a reference beam of known phase ϕ_R

$$
E_R(\boldsymbol{r}_{\mathfrak{I}}, t) = E_G(\boldsymbol{r}_{\mathfrak{I}}, t) e^{i\phi_R}.
\tag{2.60}
$$

The total field impinging on the detector is then

$$
E(\boldsymbol{r}_{\mathfrak{I}}, t) = E_G(\boldsymbol{r}_{\mathfrak{I}}, t) \left\{ e^{i\phi_R} + e^{i\overline{\phi}} \left[1 + i\widetilde{\phi}(x_{\mathfrak{I}}, t) \right] \right\},
\tag{2.61}
$$

and, to first order in $\widetilde{\phi}(x_{\mathfrak{I}}, t)$, the corresponding intensity (averaged over an optical cycle) is

$$
I(\boldsymbol{r}_{\mathfrak{I}}, t) = 2I_G(\boldsymbol{r}_{\mathfrak{I}}) \left[1 + \cos(\phi_R - \overline{\phi}) + \widetilde{\phi}(x_{\mathfrak{I}}, t) \sin(\phi_R - \overline{\phi}) \right], \tag{2.62}
$$

where $I_G(r_J)$ is the intensity profile of the unscattered Gaussian beam on the detector as defined in (2.58). Reference-beam generation prescribes ϕ_R and consequently dictates the method of interferometric detection.

2.4.1 *Homodyne detection*

Homodyne detection results from using a reference phase ϕ_R that is constant (or nearly constant) in time such that the resulting intensity is

$$
\begin{aligned}
I_{\text{hom}}(r_J, t) = 2I_G(r_J)\big[1 + \cos(\phi_R - \overline{\phi}) \\
+ \widetilde{\phi}(x_J, t)\sin(\phi_R - \overline{\phi})\big].
\end{aligned}
\tag{2.63}
$$

The homodyne optical intensity can be explicitly separated into equilibrium and fluctuating components as $I_{\text{hom}} = \overline{I}_{\text{hom}} + \widetilde{I}_{\text{hom}}$, where

$$
\overline{I}_{\text{hom}}(r_J, t) = 2I_G(r_J)\left[1 + \cos(\phi_R - \overline{\phi})\right],
\tag{2.64}
$$

$$
\widetilde{I}_{\text{hom}}(r_J, t) = 2I_G(r_J) \cdot \sin(\phi_R - \overline{\phi}) \cdot \widetilde{\phi}(x_J, t).
\tag{2.65}
$$

Note that both $\overline{I}_{\text{hom}}$ and $\widetilde{I}_{\text{hom}}$ are functions of the phase difference $\phi_R - \overline{\phi}$. If measurement of the fluctuation $\widetilde{\phi}$ is of primary importance, $|\widetilde{I}_{\text{hom}}|$ should be maximized, which occurs when $\phi_R - \overline{\phi} = \pi/2 + m\pi$ for integer m. For concreteness in the following discussion, take $\phi_R - \overline{\phi} = \pi/2$. Physically, $\phi_R - \overline{\phi} = \pi/2$ means that the reference beam is out-of-phase with the unscattered beam but in-phase with the scattered beams. Note that interfering the scattered beams with this in-phase reference beam produces an intensity linear in $\widetilde{\phi}(x_J, t)$; this should be contrasted with the weak, quadratic intensity variation of (2.57) that occurs in the absence of a reference beam.

Now, a typical engineering constraint of such an interferometer is the saturation intensity of the detector elements. Beyond the linear saturation intensity I_{sat}, the detector's response ceases to be a linear function of the incident optical power. To make an "apples-to-apples" comparison of different interference schemes, it is useful to examine the ratio of the fluctuating intensity $\widetilde{I}_{\text{hom}}$ to the saturation intensity I_{sat}. Because $|\widetilde{\phi}(x_J, t)| \ll 1$

$$
I_{\text{hom}}(r_J, t) \approx \overline{I}_{\text{hom}}(r_J, t) \leqslant 2I_G(0) \quad \text{when} \quad \phi_R - \overline{\phi} = \frac{\pi}{2},
$$

where $I_G(0) = I_G(\rho_{\mathfrak{J}} = 0, z_{\mathfrak{J}})$ is convenient shorthand for the peak intensity of the unscattered Gaussian probe beam at the detector. To obtain optimal performance, select $I_G(0)$ such that

$$I_{sat} = 2I_G(0).$$

Then, the fraction of the detector's dynamic range occupied by the fluctuating signal is

$$\frac{\widetilde{I}_{hom}(\boldsymbol{r}_{\mathfrak{J}}, t)}{I_{sat}} = \frac{I_G(\boldsymbol{r}_{\mathfrak{J}})}{I_G(0)} \cdot T_{hom} \cdot \widetilde{\phi}(x_{\mathfrak{J}}, t), \tag{2.66}$$

where

$$T_{hom} \equiv 1 \quad \text{when} \quad \phi_R - \overline{\phi} = \frac{\pi}{2} \tag{2.67}$$

is the homodyne interferometer's transfer function. As presented here, T_{hom} is independent of the fluctuation wavenumber; however, the finite sampling-volume effects that accompany any real-world measurement introduce a wavenumber dependence, as discussed in Section 3.1.4.

The above derivation of the homodyne interferometer's transfer function requires $\phi_R - \overline{\phi} = \pi/2$. Practically speaking, however, it can be difficult to keep $\phi_R - \overline{\phi}$ fixed at $\pi/2$. First, a CO_2 beam passing through $\sim 1\,m$ of plasma with a density $n_e \sim 10^{20}\,m^{-3}$ will experience a bulk phase delay $\overline{\phi} \sim \pi$; thus, for constant ϕ_R, it will be impossible to operate the interferometer at $\phi_R - \overline{\phi} = \pi/2$ as the density evolves across the discharge. Second, fusion experiments are often characterized by large, pulsed electromagnets whose vibrations can change the path lengths of the interferometer's arms. A path-length change δl produces $\delta(\phi_R - \overline{\phi}) = k_0 \delta l$, where k_0 is the wavenumber of the probe radiation. For a $10.6\,\mu m$ CO_2 probe beam, a path-length variation $\sim 2.5\,\mu m$ is sufficient to produce $\delta(\phi_R - \overline{\phi}) \sim \pi/2$, pushing the homodyne interferometer from its configuration of peak fluctuation sensitivity into one of its nulls. Even vacuum-pump vibrations can provide such a push! On small fusion devices, actively controlled mirrors have been used in an attempt to account for the evolution of the equilibrium phase and to cancel vibrational path-length changes, minimizing excursions from $\phi_R - \overline{\phi} \approx \pi/2$ [13], but such an approach has not found application on larger (and presumably more vibration-prone) fusion experiments.

In addition to its variable sensitivity, the homodyne interferometer does *not* make an absolute measurement of the phase fluctuation

amplitude $\widetilde{\phi}_0$ [2, Sec. 4.2.2]. For example, vibration-induced misalignment or power fluctuations at the beam source can alter the intensity $I_G(\boldsymbol{r}_J)$. As a result, there are three potentially dynamic quantities: $\{\phi_R - \overline{\phi}, \widetilde{\phi}(x_J, t), I_G(\boldsymbol{r}_J)\}$, but there are only two measured quantities: the equilibrium and fluctuating homodyne powers. Thus, it is generally impossible to distinguish whether changes in the amplitude of the fluctuating power are attributable to real changes in $\widetilde{\phi}(x_J, t)$ or are simply an artifact of the system alignment or radiation source.

2.4.2 *Heterodyne detection*

To avoid the above-mentioned challenges of homodyne interferometry, the reference phase can be linearly ramped in time as $\phi_R = -\Delta\omega_0 t$ such that the intensity (2.62) becomes

$$I_{het}(\boldsymbol{r}_J, t) = 2I_G(\boldsymbol{r}_J)\big[1 + \cos(\Delta\omega_0 t + \overline{\phi})$$
$$- \widetilde{\phi}(x_J, t)\sin(\Delta\omega_0 t + \overline{\phi})\big]. \tag{2.68}$$

This approach is known as heterodyne interferometry, as the desired baseband phase information is shifted to an intermediate frequency $\Delta\omega_0$ satisfying $\omega_{max} \ll \Delta\omega_0 \ll \omega_0$, where ω_{max} is the maximum angular frequency in $\widetilde{\phi}(x_J, t)$. Note that the dominant Fourier components of I_{het} sit at the intermediate frequency $\pm\Delta\omega_0$ but that the phase-fluctuation term $\widetilde{\phi}(x_J, t)\sin(\Delta\omega_0 t + \overline{\phi})$ produces sidebands about $\pm\Delta\omega_0$; sideband production for a sinusoidal phase fluctuation of angular frequency ω is sketched in Figure 2.1. Practically, the ϕ_R ramp is accomplished by modestly Doppler shifting the reference beam relative to the plasma beam. (Note that $\phi_R = -\Delta\omega_0 t$ corresponds to a Doppler *upshifted* reference beam; i.e. the reference beam's temporal evolution is given by $e^{-i\omega_0 t + i\phi_R} = e^{-i(\omega_0 + \Delta\omega_0)t}$).

The heterodyne interference signal must be demodulated in order to retrieve the baseband phase-fluctuation information. Practically speaking, dedicated analog or digital electronics are used to demodulate the heterodyne signal; however, for the pedagogical purposes of this section, it is sufficient to consider the "equivalent optical intensities" corresponding to the demodulated signals. The so-called in-phase (I) and quadrature (Q) signals are obtained by mixing I_{het} with $(2\sqrt{2}/\pi) \cdot \mathrm{Re}(e^{-i\Delta\omega_0 t})$ and $(2\sqrt{2}/\pi) \cdot \mathrm{Im}(e^{-i\Delta\omega_0 t})$, respectively, and low-pass filtering the resulting signals; here, the normalization of the complex exponential is motivated by and is consistent with the physical processes that occur in a typical ring-diode double-balanced mixer, as discussed in Section 3.5.2, and, for simplicity, low-pass fil-
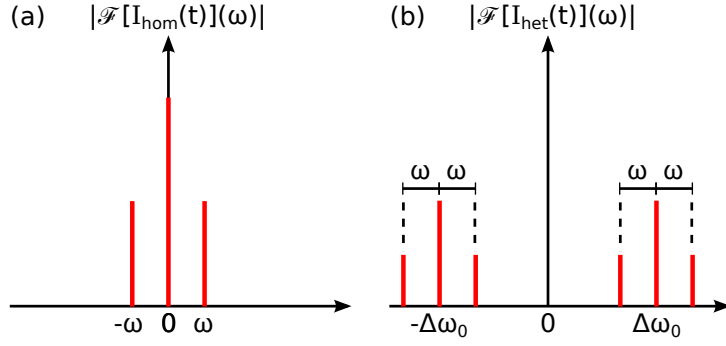
Figure 2.1: Schematic of (a) homodyne vs. (b) heterodyne signals in frequency space. A homodyne signal is predominantly DC, and a sinusoidal fluctuation with angular frequency ω produces sidebands at $\pm\omega$. In contrast, the power in a heterodyne signal is shifted to the intermediate frequencies $\pm\Delta\omega_0$, and a sinusoidal fluctuation with angular frequency ω produces sidebands at $\Delta\omega_0 \pm \omega$ and $-\Delta\omega_0 \pm \omega$.

tering is implemented by averaging over a cycle of the intermediate frequency. Thus, the equivalent I and Q optical intensities are defined as

$$I_I(\boldsymbol{r}_J, t) + i \cdot I_Q(\boldsymbol{r}_J, t) = \frac{2\sqrt{2}}{\pi} \langle e^{-i\Delta\omega_0 t} \cdot I_{het}(\boldsymbol{r}_J, t) \rangle_{\Delta\omega_0}$$

$$= \frac{2\sqrt{2}}{\pi} I_G(\boldsymbol{r}_J) e^{i\overline{\phi}} \left[1 + i\widetilde{\phi}(x_J, t) \right], \qquad (2.69)$$

where $\langle q \rangle_{\Delta\omega_0}$ denotes the average of quantity q over an intermediate-frequency cycle as

$$\langle q \rangle_{\Delta\omega_0} \equiv \frac{\Delta\omega_0}{2\pi} \int_0^{\Delta\omega_0/2\pi} q(t) dt. \qquad (2.70)$$

In contrast to the homodyne interferometer, the heterodyne interferometer makes an absolute measurement of the phase-fluctuation amplitude $\widetilde{\phi}_0$. To see this, note that I_I and I_Q (the real and imaginary components of (2.69), respectively) can be separated into equilibrium and fluctuating components as

$$\overline{I}_I(\boldsymbol{r}_J, t) = \frac{2\sqrt{2}}{\pi} I_G(\boldsymbol{r}_J) \cos\overline{\phi}, \qquad (2.71)$$

$$\overline{I}_Q(\boldsymbol{r}_J, t) = \frac{2\sqrt{2}}{\pi} I_G(\boldsymbol{r}_J) \sin\overline{\phi}, \qquad (2.72)$$

$$\widetilde{I}_I(\boldsymbol{r}_J, t) = -\frac{2\sqrt{2}}{\pi} I_G(\boldsymbol{r}_J) \sin\overline{\phi} \cdot \widetilde{\phi}(x_J, t), \qquad (2.73)$$

$$\widetilde{I}_Q(\boldsymbol{r}_J, t) = \frac{2\sqrt{2}}{\pi} I_G(\boldsymbol{r}_J) \cos\overline{\phi} \cdot \widetilde{\phi}(x_J, t). \qquad (2.74)$$

As is the case for the homodyne interferometer, there are three potentially dynamic quantities: $\{\overline{\phi},\ \widetilde{\phi}(x_J, t),\ I_G(\boldsymbol{r}_J)\}$; however, in contrast to the homodyne interferometer, there are now *four* measured quantities: $\{\overline{I}_I,\ \overline{I}_Q,\ \widetilde{I}_I,\ \widetilde{I}_Q\}$. Therefore, the number of measured quantities is sufficient to unambiguously determine $\{\overline{\phi},\ \widetilde{\phi}(x_J, t),\ I_G(\boldsymbol{r}_J)\}$ in absolute units.

Finally, it is useful to characterize the heterodyne interferometer's performance relative to the saturation limits of a given detector. Because $|\widetilde{\phi}(x_J, t)| \ll 1$, the heterodyne intensity (2.68) can be approximated as

$$I_{het}(\boldsymbol{r}_J, t) \approx 2I_G(\boldsymbol{r}_J)[1 + \cos(\Delta\omega_0 t + \overline{\phi})]$$
$$\leqslant 4I_G(0),$$

where $I_G(0) = I_G(\rho_J = 0, z_J)$ is convenient shorthand for the peak intensity of the unscattered Gaussian probe beam at the detector. To obtain optimal performance, select $I_G(0)$ such that

$$I_{sat} = 4I_G(0).$$

Further, define the total fluctuating intensity in the demodulated signals to be

$$\widetilde{I}_{IQ}(\boldsymbol{r}_J, t) \equiv \left\{[\widetilde{I}_I(\boldsymbol{r}_J, t)]^2 + [\widetilde{I}_Q(\boldsymbol{r}_J, t)]^2\right\}^{1/2}$$
$$= \frac{2\sqrt{2}}{\pi} I_G(\boldsymbol{r}_J) \cdot \widetilde{\phi}(x_J, t) \tag{2.75}$$

such that the fraction of the detector's dynamic range occupied by the fluctuating component in the demodulated signal is

$$\frac{\widetilde{I}_{IQ}(\boldsymbol{r}_J, t)}{I_{sat}} = \frac{I_G(\boldsymbol{r}_J)}{I_G(0)} \cdot T_{het} \cdot \widetilde{\phi}(x_J, t), \tag{2.76}$$

where

$$T_{het} \equiv \frac{1}{\sqrt{2} \cdot \pi} \tag{2.77}$$

is the heterodyne interferometer's transfer function. As presented here, T_{het} is independent of the fluctuation wavenumber; however, the finite sampling-volume effects that accompany any real-world measurement introduce a wavenumber dependence, as discussed in Section 3.1.4.

In contrast to the homodyne interferometer, the heterodyne interferometer's wavenumber transfer function is *not* a function of $\phi_R - \overline{\phi}$. Thus, the heterodyne interferometer always operates at its peak sensitivity, regardless of the bulk plasma phase or path-length vibrations. Robust sensitivity comes at a cost, however. Note that

$$\frac{T_{het}}{T_{hom}} = \frac{1}{\sqrt{2} \cdot \pi} \quad \text{for homodyne operation at } \phi_R - \overline{\phi} = \pi/2.$$

Thus, for a given detector, a heterodyne interferometer will be ($\sqrt{2} \cdot \pi$) times *less* sensitive than a homodyne interferometer operated in its optimal configuration ($\phi_R - \overline{\phi} = \pi/2$). There are two physical origins of the heterodyne interferometer's sensitivity deficit. First, the detector of a homodyne interferometer with $\phi_R - \overline{\phi} = \pi/2$ only sees small fluctuations about a DC offset, while the detector of a heterodyne interferometer sees the full sinusoidal waveform of the intermediate frequency; to ensure the detector is always within its saturation limits, the heterodyne interferometer must necessarily be operated with a mean intensity at the detector that is a factor of two lower than that for the homodyne interferometer. Second, a fraction of the power in the heterodyne interferometer's fluctuating signal (and its equilibrium signal) is lost during demodulation.

2.5 PHASE CONTRAST IMAGING (PCI)

As discussed in Section 2.3.4, imaging the probe radiation on a square-law detector produces a very weak response because the unscattered and scattered beams are $\pi/2$ out of phase with each other. To produce a measurable response, a traditional interferometer interferes the imaged radiation with an external reference beam. If the phase of the unscattered beam could be manipulated, though, the external reference beam would no longer be needed. This is the approach employed in phase contrast imaging (PCI). A typical PCI system is shown schematically in Figure 2.2.

2.5.1 *Reference-beam generation with a phase plate*

PCI uses an optical element known as a *phase plate* to delay the unscattered beam by $\pi/2$ relative to the scattered beams. The phase plate is typically a reflective optical element with a groove that is precisely fabricated to have a depth of $\lambda_0/8$; the unscattered beam reflects off of this groove, and the corresponding $\lambda_0/4$-increase in path length phase
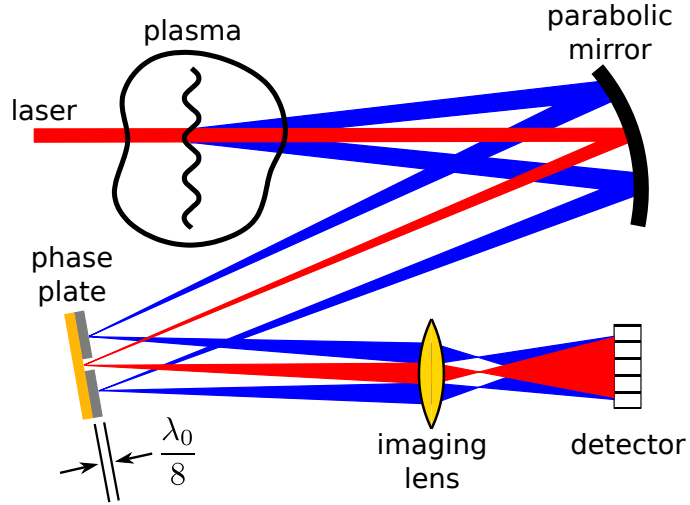
Figure 2.2: Schematic overview of a typical PCI system. A plasma-density fluctuation weakly scatters an incident probe beam, and an off-axis parabolic mirror focuses the scattered and unscattered beams. An optical element known as a phase plate is placed at the focal plane of the parabolic mirror. The phase plate has a narrow groove of depth $\lambda_0/8$ that imparts phase delay $\pi/2$ to the unscattered beam, effectively converting the unscattered beam into an "internal" reference beam against which the scattered beams can be interfered. A lens, whose object plane sits at the plasma midplane, is then used to image the resulting radiation onto a detector array.

delays the unscattered beam by $\pi/2$ relative to the scattered beams, which reflect off of the non-grooved portions (i.e. the "face") of the phase plate. To boost the relative size of the fluctuating signal, the phase groove typically reflects only a fraction $\eta < 1$ of the incident unscattered beam power, while the phase-plate face reflects all of the scattered beam power. Thus, by the action of the phase plate, $1 \to i\sqrt{\eta}$ in the imaged electric field (2.55) such that

$$E_{PCI}(\boldsymbol{r}_J, t) = iE_G(\boldsymbol{r}_J, t)e^{i\overline{\phi}} \left[\sqrt{\eta} + \widetilde{\phi}_0 \cos \nu \right], \tag{2.78}$$

and the corresponding intensity, averaged over an optical cycle and to first order in $\widetilde{\phi}_0$, is

$$I_{PCI}(\boldsymbol{r}_J, t) = I_G(\boldsymbol{r}_J) \left[\eta + 2\sqrt{\eta}\widetilde{\phi}_0 \cos \nu \right]. \tag{2.79}$$

Here, $E_G(\boldsymbol{r}_J)$ and $I_G(\boldsymbol{r}_J)$ are the field and intensity profiles of the unscattered Gaussian beam on the detector in the *absence* of the phase plate, with $I_G(\boldsymbol{r}_J)$ being explicitly defined in (2.58). Equation (2.79) should be contrasted with (2.57), which gives the image-plane intensity in the absence of the phase plate. Thus, the phase plate converts

the unscattered probe beam into an effective reference beam for the scattered beams.

2.5.2   *Focal-plane separation of scattered beams*

Implicit in the use of the phase plate is that the scattered and unscattered beams are well-separated in space such that the phase groove only affects the unscattered beam. The 1^{st}-order scattered beams are angularly separated from the unscattered beam by $\theta = k/k_0$, and, in the far field ($z \gg z_R$), the center of the scattered beam will fall outside of the unscattered beam's $1/e$ E radius if

$$|k| \geqslant \frac{2}{w_0}. \tag{2.80}$$

However, CO_2 laser beams used to probe tokamak plasmas often have $z_R \gg 10\,\text{m}$, so the beam's far field is not easily accessible in typical lab settings. Fortunately, the far-field diffraction pattern can be equivalently accessed in the focal plane of a focusing optic [14, Ch. 8].

The focal-plane location, beam size, and beam separation can be easily determined. Let the Gaussian probe beam have an in-vessel $1/e$ E waist radius of w_0, and place a focusing optic of focal length f a distance s downstream from the in-vessel beam waist. Then, the waist of the focused beam will be located a distance s' downstream of the focusing optic and will have $1/e$ E radius w_0' given as

$$s' = f\left(1 + \frac{s - f}{z_R}\right), \tag{2.81}$$

$$w_0' = \frac{w_0|f|}{\left[(s - f)^2 + z_R^2\right]^{1/2}}, \tag{2.82}$$

where z_R is the in-vessel Rayleigh length [15]. When $|s - f| \ll z_R$, as is typical for PCI, the expressions for s' and w_0' reduce to

$$s' \approx f, \tag{2.83}$$

$$w_0' \approx \frac{2|f|}{k_0 w_0}. \tag{2.84}$$

The spatial separation Δ of the scattered and unscattered beams in the focal plane is found by applying the appropriate ABCD ray matrices

from Table B.1 to a ray scattered in the plasma midplane by angle θ, i.e.

$$\begin{pmatrix} \Delta \\ \theta_{pp} \end{pmatrix} = \begin{pmatrix} 1 & s' \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ \theta \end{pmatrix},$$

which, upon substitution of of the focal plane location from (2.83) and the scattering angle $\theta = k/k_0$, simplifies to

$$\Delta \approx \frac{kf}{k_0}. \tag{2.85}$$

2.5.3 *Low-k cutoff of phase plate*

Now, let the phase-plate groove have a width d, as is shown in Figure 2.3. Finite PCI response requires that (most of) the scattered beams fall outside of the phase groove (i.e. $|\Delta| \geqslant d/2$). Application of the phase-plate beam-separation formula (2.85) then shows that there will be finite PCI response for $|k| \geqslant k_g$, where

$$k_g \equiv \frac{k_0 d}{2f}. \tag{2.86}$$

Here, the subscript g is in reference to the *groove* of the phase plate. Further, to provide the strongest phase contrast, the unscattered beam should fall wholly within the phase groove (i.e. $2w_0' \leqslant d$); substituting (2.84) for w_0' then yields a constraint on the phase groove width

$$d \geqslant \frac{4f}{k_0 w_0}, \tag{2.87}$$

and inserting (2.87) into (2.86) yields

$$k_g \geqslant \frac{2}{w_0}. \tag{2.88}$$

As finite response requires that $|k| \geqslant k_g$, it follows that (2.88) is equivalent to (2.80), which was derived by considering the far-field separation of the scattered and unscattered beams. Thus, PCI's low-k cutoff is ultimately constrained by the in-vessel beam size w_0, with diffraction being the constraining physical mechanism.
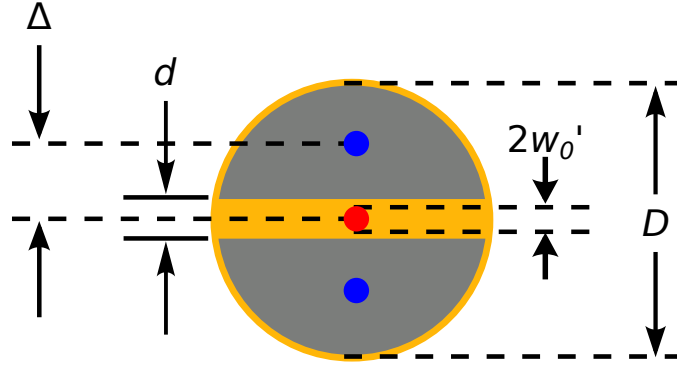
Figure 2.3: Transverse phase-plate dimensions. The unscattered beam is shown in red, while the scattered beams are shown in blue.

2.5.4 *High-k cutoff of phase plate*

Let the phase plate have a diameter D, as is shown in Figure 2.3. Detection of the scattered radiation requires that (most of) the scattered beam reflect from the face of the phase plate (e.g. $\Delta \leqslant D/2$). Application of the phase-plate beam-separation formula (2.85) then shows that there will be finite PCI response for $|k| \leqslant k_D$ where

$$k_D \equiv \frac{k_0 D}{2f}. \tag{2.89}$$

2.5.5 *Effect of phase plate on m^{th} scattered beam*

The effect of the PCI phase plate on the m^{th} scattered beam is given by the complex-valued function $\mathcal{E}(\boldsymbol{r}_m, k)$, which is derived and thoroughly discussed in Appendix C. The relevant results are briefly summarized here for completeness. Eq. (C.32) shows that PCI's image-plane $\mathcal{E}(\boldsymbol{r}_m, k)$ readily reduces to

$$\mathcal{E}(\boldsymbol{r}_{m,\mathfrak{I}}, k_\mathfrak{I}) = e^{-[x_{m,\mathfrak{I}}/w(z_{m,\mathfrak{I}})]^2} e^{imk_\mathfrak{I}x_\mathfrak{I}}$$
$$\times [F(\boldsymbol{r}_{m,\mathfrak{I}}, k_\mathfrak{I}) + G(\boldsymbol{r}_{m,\mathfrak{I}}, k_\mathfrak{I})], \tag{2.90}$$

where the phase-plate face acts on the m^{th} scattered beam via F, and the phase-plate groove acts on the m^{th} scattered beam via G. F and G are themselves defined in (C.33) and (C.34). Of particular note, F is Hermitian with respect to m

$$F(\boldsymbol{r}_{-m,\mathfrak{I}}, k_\mathfrak{I}) = F^*(\boldsymbol{r}_{m,\mathfrak{I}}, k_\mathfrak{I}), \tag{2.91}$$

while G is anti-Hermitian with respect to m

$$G(r_{-m,\jmath}, k_\jmath) = -G^*(r_{m,\jmath}, k_\jmath). \tag{2.92}$$

These symmetries imply that $F(r_{0,\jmath}, k_\jmath)$ is purely *real* and that $G(r_{0,\jmath}, k_\jmath)$ is purely *imaginary*.

2.5.6   *The imaged field and its intensity*

Introducing the notational shorthand $F_m \equiv F(r_{m,\jmath}, k_\jmath)$ and $G_m \equiv G(r_{m,\jmath}, k_\jmath)$, the weak-coupling ($\widetilde{\phi}_0 \ll 1$), image-plane electric field from (2.54) readily reduces to

$$E(r_\jmath, t) \approx E_G(r_\jmath, t)e^{i\overline{\phi}}\bigg\{F_0 + G_0 + i\frac{\widetilde{\phi}_0}{2}\bigg[(F_1 + G_1)e^{i\nu} \\ + (F_{-1} + G_{-1})e^{-i\nu}\bigg]\bigg\}, \tag{2.93}$$

where the approximation $w(z_m) \approx w(z)$ has been used, ν is defined in (2.52), and $E_G(r_\jmath, t)$ would be the image-plane electric field of the unscattered beam in the *absence* of the phase plate. Now, recall that F is Hermitian such that $F_{-1} = F_1^*$ and $F_0 = \text{Re}(F_0)$ and that G is anti-Hermitian such that $G_{-1} = -G_1^*$ and $G_0 = i \cdot \text{Im}(G_0)$; using these substitutions, the field further reduces to

$$E(r_\jmath, t) = E_G(r_\jmath, t)e^{i\overline{\phi}}\bigg\{\text{Re}(F_0) - \widetilde{\phi}_0\text{Im}(G_1 e^{i\nu}) \\ + i\left[\text{Im}(G_0) + \widetilde{\phi}_0\text{Re}(F_1 e^{i\nu})\right]\bigg\}, \tag{2.94}$$

and the corresponding intensity, averaged over an optical cycle and to first order in $\widetilde{\phi}_0$, is

$$I_{\text{pci}}(r_\jmath, t) = I_G(r_\jmath)\bigg\{|F_0|^2 + |G_0|^2 \\ + 2\widetilde{\phi}_0\left[\text{Im}(G_0)\text{Re}(F_1 e^{i\nu}) - \text{Re}(F_0)\text{Im}(G_1 e^{i\nu})\right]\bigg\}, \tag{2.95}$$

where $I_G(r_\jmath)$ would be the intensity profile (averaged over an optical cycle) of the unscattered beam in the *absence* of the phase plate. Using

the fact that $e^{i\nu} = \cos\nu + i\sin\nu$, $F_1 = \text{Re}(F_1) + i \cdot \text{Im}(F_1)$, and $G_1 = \text{Re}(G_1) + i \cdot \text{Im}(G_1)$, the image-plane intensity further reduces to

$$
\begin{aligned}
I_{\text{pci}}(\boldsymbol{r}_{\text{J}}, t) = I_G(\boldsymbol{r}_{\text{J}}) \Big\{ & |F_0|^2 + |G_0|^2 \\
& + 2\widetilde{\phi}_0 \left[\text{Im}(G_0)\text{Re}(F_1) - \text{Re}(F_0)\text{Im}(G_1) \right] \cos\nu \\
& - 2\widetilde{\phi}_0 \left[\text{Im}(G_0)\text{Im}(F_1) + \text{Re}(F_0)\text{Re}(G_1) \right] \sin\nu \Big\}.
\end{aligned}
\tag{2.96}
$$

Note that the linear combination $A_I \cos\nu - A_Q \sin\nu$ for real A_I, A_Q, and ν can be rewritten as

$$
A_I \cos\nu - A_Q \sin\nu = A \cos(\nu + \theta),
\tag{2.97}
$$

where $A = (A_I^2 + A_Q^2)^{1/2}$, $\theta = \text{atan2}(A_Q, A_I)$, and $\text{atan2}(A_Q, A_I)$ is the arctangent function of two arguments, which uses the signs of A_Q and A_I to correctly determine the quadrant corresponding to a tangent of A_Q/A_I. Note that the notation is mnemonic: A_I is the amplitude of the "in-phase" (I) component (i.e. proportional to the image of the assumed cosine phase fluctuation $\widetilde{\phi}(x, t)$ from (2.39)), and A_Q is the amplitude of the corresponding "quadrature" (Q) component.

As was the case for external reference-beam interferometry, it is useful to characterize PCI's performance relative to the saturation limits of a given detector. Because $\widetilde{\phi}_0 \ll 1$, the PCI optical intensity (2.96) satisfies

$$
I_{\text{pci}}(\boldsymbol{r}_{\text{J}}, t) \lesssim I_G(0) \left(|F_0|^2 + |G_0|^2 \right),
$$

where $I_G(0) = I_G(\rho_{\text{J}} = 0, z_{\text{J}})$ is convenient shorthand for the peak intensity of the unscattered Gaussian probe beam at the detector in the *absence* of the phase plate. To obtain optimal performance, select $I_G(0)$ such that

$$
I_{\text{sat}} = I_G(0) \left(|F_0|^2 + |G_0|^2 \right),
$$

where I_{sat} is the detector's linear saturation intensity. Then, taking inspiration from (2.97), the image-plane PCI intensity fluctuations from (2.96) can be rewritten as

$$
\frac{\widetilde{I}_{\text{pci}}(\boldsymbol{r}_{\text{J}}, t)}{I_{\text{sat}}} = \frac{I_G(\boldsymbol{r}_{\text{J}})}{I_G(0)} \cdot A_{\text{pci}}(k_{\text{J}}, x_{\text{J}}) \cdot \widetilde{\phi}_0 \cos \left[\nu + \theta_{\text{pci}}(k_{\text{J}}, x_{\text{J}}) \right],
\tag{2.98}
$$

where

$$A_{\text{pci}}(k_\jmath, x_\jmath) \equiv \frac{2A(k_\jmath, x_\jmath)}{|F_0|^2 + |G_0|^2}, \tag{2.99}$$

$$\theta_{\text{pci}}(k_\jmath, x_\jmath) \equiv \text{atan2}(A_Q, A_I), \tag{2.100}$$

and

$$A(k_\jmath, x_\jmath) \equiv (A_I^2 + A_Q^2)^{1/2}, \tag{2.101}$$

$$A_I(k_\jmath, x_\jmath) \equiv \text{Im}(G_0)\text{Re}(F_1) - \text{Re}(F_0)\text{Im}(G_1), \tag{2.102}$$

$$A_Q(k_\jmath, x_\jmath) \equiv \text{Im}(G_0)\text{Im}(F_1) + \text{Re}(F_0)\text{Re}(G_1). \tag{2.103}$$

Thus, the action of the phase plate results in both an amplitude response $A_{\text{pci}}(k_\jmath, x_\jmath)$ and a phase response $\theta_{\text{pci}}(k_\jmath, x_\jmath)$ in PCI's image-plane intensity. The on-axis responses $A_{\text{pci}}(k_\jmath, x_\jmath = 0)$ and $\theta_{\text{pci}}(k_\jmath, x_\jmath = 0) = 0$ are consistent with previous derivations at $x_\jmath = 0$, such as [16, Eq. 2.141] and [17, Eq. 20]. The present work, however, explicitly accounts for the spatial variation in the amplitude and phase responses across the face of the detector. Note that PCI operates as a nonlinear system if there are substantial spatial variations in either A_{pci} or θ_{pci}. The amplitude and phase responses can be easily evaluated numerically, and the results for typical system parameters are shown in Figure 2.4. As is colloquially understood, the amplitude response A_{pci} drops precipitously for $|k| \lesssim k_g$. (Additionally, the spatial variation in the amplitude response is minimal). However, perhaps less well known, is the fact that the phase response θ_{pci} exhibits dramatic spatial variation for $|k| \lesssim k_g$. This spatial variation biases the PCI-measured wavenumber k_{meas} away from the true wavenumber k, as shown in Figure 2.5. The physical explanation for this effect is relatively simple: a Gaussian beam can be decomposed into a set of infinite plane waves with a finite spread in transverse wavevectors (see Section 2.2.1), and only the components of the scattered beam with transverse lab-frame wavenumbers $|k| \geqslant k_g$ are reflected from the phase-plate face and produce measurable interference on the PCI detector.

Thus, for $|k| \lesssim k_g$, PCI operates as a nonlinear system that cannot be described via a transfer function. Practically speaking, PCI's nonlinear low-k operation prevents it from measuring the wavenumbers of low-k fluctuations, such as MHD and some ion-scale instabilities. (Of course, because the PCI amplitude response A_{pci} drops precipitously for $|k| \lesssim k_g$, such low-k fluctuations may be altogether invisible to PCI anyways). For $|k| \gtrsim k_g$, however, a transfer function can be defined. Specifically, for $|k| \gtrsim k_g$, the phase-plate face reflects the scattered beam ($F_1 \to 1$), while the phase-plate groove minimally affects
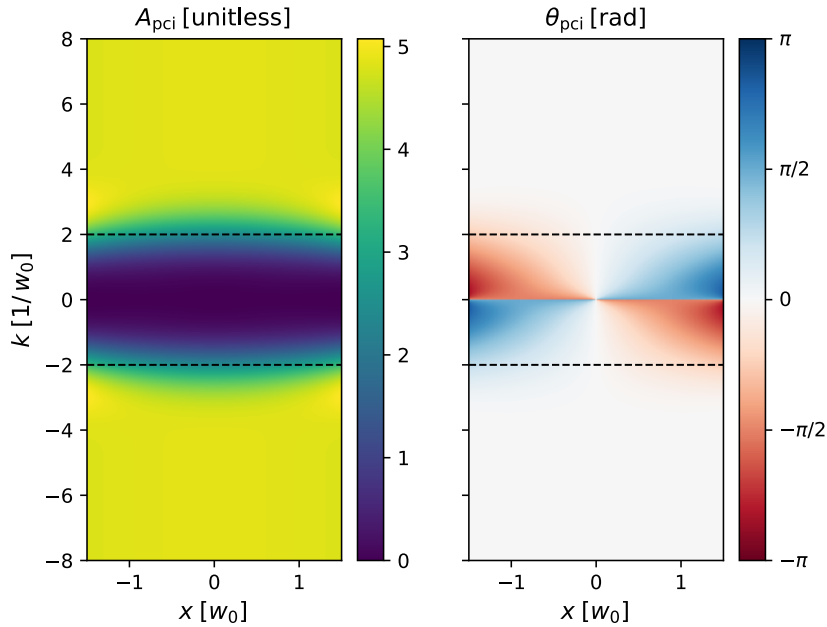
Figure 2.4: PCI amplitude response $A_{pci}(k, x)$ and phase response $\theta_{pci}(k, x)$ in object-plane coordinates. Spatial coordinates x and wavenumbers k are normalized to the $1/e$ E radius of the in-vessel probe beam, w_0. The system magnification is $M = 0.5$, a fairly typical value. The dashed horizontal lines indicate the low-k cutoff of the PCI phase-plate groove, k_g; here, $k_g = 2/w_0$, which is the minimum value allowed by diffraction, as discussed in (2.88). The phase-plate high-k cutoff, k_D, is taken to be infinite. The reflectivity of the phase groove is $\eta = 0.17$, which is characteristic of the ZnSe typically employed in 10.6 μm optics. Note that the low-k phase response θ_{pci} exhibits dramatic spatial variation, which results in nonlinear PCI operation for $|k| \lesssim k_g$.

the scattered beam ($G_1 \to 0$). Further, assuming the majority of the un-scattered beam falls within the phase groove (i.e. constraint (2.87)), the phase-plate groove attenuates and phase shifts the unscattered beam ($G_0 \to i\sqrt{\eta}$), while the phase-plate face minimally affects the unscattered beam ($F_0 \to 0$). Thus, $\theta_{pci} \to 0$ and $A_{pci} \to 2/\sqrt{\eta}$ such that the PCI transfer function can be defined as

$$
T_{pci}(k) = \begin{cases} \dfrac{2}{\sqrt{\eta}}, & |k| \gtrsim k_g \\ \text{undefined}, & \text{otherwise} \end{cases} .
\tag{2.104}
$$

As presented here, for $|k| \gtrsim k_g$, T_{pci} is independent of the fluctuation wavenumber; however, the finite sampling-volume effects that accompany any real-world measurement introduce a wavenumber dependence, as discussed in Section 3.1.4.
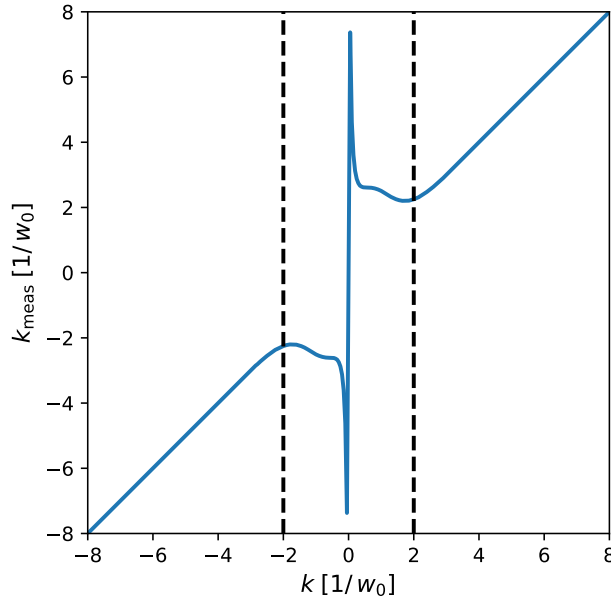
Figure 2.5: PCI-measured wavenumber k_{meas} vs. true wavenumber k. Wavenumbers are normalized to the $1/e$ E radius of the in-vessel probe beam, w_0. The system magnification is $M = 0.5$, a fairly typical value. The dashed vertical lines indicate the low-k cutoff of the PCI phase-plate groove, k_g; here, $k_g = 2/w_0$, which is the minimum value allowed by diffraction, as discussed in (2.88). For $|k| \gtrsim k_g$, k_{meas} has a $1:1$ linear relationship with k; however, for $|k| \lesssim k_g$, k_{meas} is *not* linearly related to k. Thus, a transfer-function description of the PCI operation is only appropriate for wavenumbers above the low-k cutoff ($|k| \gtrsim k_g$).

As is the case with the homodyne interferometer, the PCI technique does *not* make an absolute measurement of the phase-fluctuation amplitude $\widetilde{\phi}_0$. To see this, note that the PCI amplitude response A_{pci} depends very sensitively on the system alignment, with slight excursions of the unscattered beam from the partially reflective phase-plate groove onto the fully reflective phase-plate face resulting in macroscopic changes to the power reaching the detector. Further, power fluctuations at the beam source can alter $I_G(r_J)$. Thus, there are three potentially dynamic quantities: $\{\widetilde{\phi}_0, I_G(r_J), A_{pci}\}$, but there are only two potentially measurable quantities: the equilibrium and fluctuating powers. PCI systems on large, vibration-prone fusion devices typically employ feedback stabilization (see e.g. [16, Ch. 3.5]) in order to dynamically maintain the unscattered beam's alignment on the phase-plate groove, minimizing vibrational contamination of the PCI signal. It is then possible, after measuring a calibration constant, to *estimate* the phase-fluctuation amplitude $\widetilde{\phi}_0$ with PCI.

2.6 SELECTING AN INTERFEROMETRIC TECHNIQUE

Sections 2.4 and 2.5 detail external reference-beam interferometry and phase contrast imaging (PCI), respectively. The intent of this section is to synthesize these results and to discuss the strengths and limitations of these interferometric techniques so that a suitable method can be selected for a given application.

2.6.1 *Sensitivity*

The transfer functions in Sections 2.4 and 2.5 specify the fraction of a given detector's dynamic range that is occupied by the fluctuating signal. If identical detectors are used for each interferometric method, then an "apples-to-apples" comparison of fluctuation sensitivities can be made by examining the amplitudes of the corresponding transfer functions. (See Section 2.6.3 for practical considerations regarding detector selection). Just such a comparison is shown in Figure 2.6. Clearly, for a given fluctuation ($|k| \gtrsim k_g$) and a given detector, PCI has the best sensitivity. Specifically, PCI is $T_{pci}/T_{hom} = 2/\sqrt{\eta}$ more sensitive than a comparable homodyne interferometer ($\phi_R - \overline{\phi} = \pi/2$) and $T_{pci}/T_{het} = 2\pi\sqrt{2/\eta}$ more sensitive than a comparable heterodyne interferometer. Relative to homodyne interferometry, PCI's enhanced sensitivity is wholly attributable to the partial reflectivity ($\eta < 1$) of the phase-plate groove: the decreased power in the unscattered beam increases the fraction of the detector's dynamic range occupied by the fluctuating signal. The sensitivity deficit of heterodyne interferometry relative to homodyne interferometry has two physical origins: first, the heterodyne interferometer must capture the full sinusoidal waveform of the heterodyne interference signal, which mandates reduction of the mean optical intensity (and, correspondingly, the fluctuating optical intensity) at the detector by a factor of two relative to that of the homodyne interferometer; second, demodulation of the heterodyne interference signal additionally attenuates the fluctuating signal.

The enhanced sensitivity of PCI and homodyne interferometry relative to that of heterodyne interferometry comes with several costs, however. First, as previously discussed, neither PCI nor homodyne interferometry measure the absolute scale of phase fluctuations, while heterodyne interferometry does. Second, PCI depends sensitively on the position of the unscattered beam relative to the phase-plate groove, and homodyne interferometry depends sensitively on $\phi_R - \overline{\phi}$; often, feedback stabilization is required to dynamically maintain the optimal configuration [16, Ch. 3.5][13]. While feedback stabilization is an
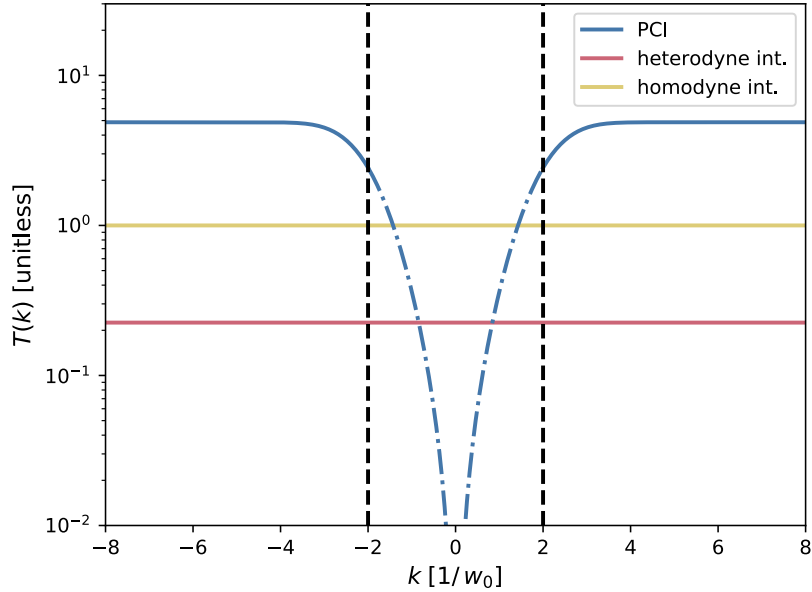
Figure 2.6: A comparison of the transfer functions for PCI (2.104), heterodyne interferometry (2.77), and homodyne interferometry (2.67, in the optimal fluctuation configuration with $\phi_R - \overline{\phi} = \pi/2$). Object-plane wavenumbers k are normalized to the $1/e$ E radius of the in-vessel probe beam, w_0. The magnification of each system is taken to be $|M| = 0.5$, which is a representative "typical" value. The vertical, dashed lines indicate the low-k cutoff of the PCI phase-plate groove, k_g; here, $k_g = 2/w_0$, which is the minimum value allowed by diffraction, as discussed in (2.88). The phase-plate high-k cutoff, k_D, is taken to be infinite. The reflectivity of the PCI phase groove is $\eta = 0.17$, which is characteristic of the ZnSe typically employed in $10.6\,\mu m$ optics. Because the PCI transfer function is not defined for $|k| \lesssim k_g$, the low-k PCI amplitude response (2.99) is indicated by the dash-dot curves instead.

added technical complication, it should be noted that such feedback is expected to become more commonplace for laser diagnostics on large fusion devices, such as ITER's heterodyne interferometer [18]. Finally, PCI cannot measure the equilibrium phase $\overline{\phi}$. This makes intuitive sense: $\overline{\phi}$ uniformly affects both the scattered and unscattered beams, the PCI phase delays the unscattered beam to generate an "internal" reference beam, and the $\overline{\phi}$ information cancels in the resulting interference. A homodyne interferometer operated with $\phi_R - \overline{\phi} = \pi/2$ is similarly unable to measure $\overline{\phi}$. In contrast, a heterodyne interferometer can always measure both the equilibrium phase $\overline{\phi}$ and the fluctuating phase $\widetilde{\phi}$.

2.6.2 *Spatial bandwidth*

PCI's sensitivity comes at the additional expense of spatial bandwidth. Specifically, the creation of an "internal" reference beam via spatial filtering produces a low-k cutoff in the PCI response, as shown in Figure 2.6. The minimum size of this cutoff is set by diffraction of the in-vessel probe beam (2.88), but it is not uncommon for the realized cutoff (2.86) to be $\sim 2 - 3\times$ larger than the diffraction limit. As shown in Figures 2.4 and 2.5, PCI operates as a nonlinear system below its low-k cutoff, preventing a transfer-function description of its low-k behavior. In contrast, by using an external reference beam, homodyne and heterodyne interferometers operate as linear systems over the entire wavenumber spectrum and are even capable of making measurements at $k = 0$.

Now, colloquially, interferometry is considered a "low-k" technique, and PCI is considered a "high-k" technique. However, as detailed in Section 2.2.2, for a given probe beam and a given fluctuation $\widetilde{\phi}$, the laser-plasma interaction is *identical* for both interferometry and PCI. Further, the high-k optical capabilities of interferometry and PCI are governed by the size of the collection optics and finite sampling-volume effects [19]. Thus, there is nothing that intrinsically limits interferometry to low-k measurements — an interferometer's high-k limit can be just as high, if not higher, than that of a given PCI system. However, as discussed in Section 2.6.1, PCI is *more sensitive* to fluctuations than a comparable interferometer, and, assuming a Kolmogorov-like fluctuation spectrum $S(k) \propto k^{-p}$ for some positive p, PCI's superior sensitivity may allow it to detect high-k fluctuations that are too weak to be seen by an interferometer.

2.6.3 *Temporal bandwidth*

Detector bandwidth is often the dominant constraint of an interferometric system's temporal bandwidth. Homodyne interferometry and PCI require

$$\omega_{\text{det}} > \omega, \qquad \text{homodyne interferometry, PCI,} \qquad (2.105)$$

where ω_{det} is the angular cutoff frequency of the detector and ω is the angular frequency of the fluctuation. Heterodyne interferometry, however, requires

$$\omega_{\text{det}} > \Delta\omega_0 + \omega, \qquad \text{heterodyne interferometry,} \qquad (2.106)$$

where $\Delta\omega_0$ is the (angular) intermediate frequency. Further, proper reconstruction of the baseband signal from the heterodyne interference signal requires $\omega < \Delta\omega_0$.

Thus, for a given fluctuation $\omega \ll \Delta\omega_0$, heterodyne interferometry requires a much faster detector than homodyne interferometry or PCI. For the HgCdTe detectors typically used at $10.6\,\mu m$, cooling the active element of the detector reduces detector noise and increases the detector response at the expense of reduced ω_{det}. Thus, for low–bandwidth applications ($\omega \ll \Delta\omega_0$), homodyne interferometry and PCI may be able to use slower, cooled detectors that are less noisy than the comparable faster, warmer detectors required for heterodyne interferometry. The use of less noisy detectors may produce sensitivity gains for PCI and homodyne interferometry relative to heterodyne interferometry. Although use of a "slow" detector prevents measurements of broadband fluctuations beyond the detector cutoff, it *is* possible to measure coherent, high-frequency fluctuations well beyond the detector cutoff ($\omega \gg \omega_{\text{det}}$) by rapidly modulating the intensity of the probe beam [6, Sec. 3.3.1].

BIBLIOGRAPHY

[1] J. D. Jackson. *Classical Electrodynamics*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 3rd edition, 1998.

[2] I. H. Hutchinson. *Principles of Plasma Diagnostics*. Cambridge University Press, Cambridge, United Kingdom, 2nd edition, 2002.

[3] T. H. Stix. *Waves in Plasmas*. Springer-Verlag, New York, NY, USA, 1992.

[4] V. V. Mirnov, W. X. Ding, D. L. Brower, M. A. Van Zeeland and T. N. Carlstrom. Finite electron temperature effects on interferometric and polarimetric measurements in fusion plasmas. *Phys. Plasmas*, 14(10):102105, 2007.

[5] D. J. Griffiths. *Introduction to Quantum Mechanics*. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, USA, 1st edition, 1995.

[6] N. Tsujii. *Studies of ICRF mode conversion with phase contrast imaging and full-wave simulations in Alcator C-Mod*. PhD thesis, MIT, 2012.

[7] A. E. Siegman. *Lasers*. University Science Books, Sausalito, CA, USA, 1986.

[8] C. V. Raman and N. S. N Nath. The diffraction of light by high frequency sound waves: Part I. *Proc. Indian Acad. Sci.*, A2(4):406–412, 1936.

[9] C. V. Raman and N. S. N Nath. The diffraction of light by high frequency sound waves: Part III. *Proc. Indian Acad. Sci.*, 3(1):75–84, 1936.

[10] J. R. Dorris. *Phase contrast imaging measurements and modeling of short wavelength turbulence in the DIII-D tokamak*. PhD thesis, MIT, 2010.

[11] A. A. Tovar and L. W. Casperson. Generalized beam matrices. IV. Optical system design. *J. Opt. Soc. Am. A*, 14(4):882–894, 1997.

[12] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*. U.S. Government Printing Office, Washington, DC, USA, 10th edition, 1972.

[13] R. Nazikian and L. E. Sharp. CO2 laser scintillation interferometer for the measurement of density fluctuations in plasma confinement devices. *Rev. Sci. Instrum.*, 58(11):2086–2091, 1987.

[14] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, Cambridge, United Kingdom, 7th (expanded) edition, 1999.

[15] S. A. Self. Focusing of spherical Gaussian beams. *Appl. Opt.*, 22(5):658–661, 1983.

[16] S. Coda. *An experimental study of turbulence by phase contrast imaging in the DIII-D tokamak*. PhD thesis, MIT, 1997.

[17] J. C. Rost. PCI low-k response. Unpublished, June 2014.

[18] M. A. Van Zeeland, R. L. Boivin, D. L. Brower, T. N. Carlstrom, J. A. Chavez et al. Conceptual design of the tangentially viewing combined interferometer-polarimeter for ITER density measurements. *Rev. Sci. Instrum.*, 84(4):043501, 2013.

[19] R. V. Bravenec and A. J. Wootton. Effects of limited spatial resolution on fluctuation measurements (invited). *Rev. Sci. Instrum.*, 66(1):802–805, 1995.

3

DESIGN CONSIDERATIONS FOR A HETERODYNE INTERFEROMETER

While Chapter 2 discusses the *optical* foundations for various interferometric methods, most real-world optical diagnostics are complex, integrated systems requiring precise interplay between various components, such as lasers, optics, detectors, and electronics. Optimizing the performance of a given diagnostic requires careful consideration of each component and its role in the measurement. Some of these considerations are generic, and some of them are diagnostic specific.

This chapter examines numerous design considerations that are relevant to heterodyne interferometry. The sections are arranged in roughly sequential order, beginning with the interference signal at the detector and proceeding through successive downstream components until reaching the system's digitizer. In particular, Section 3.1 discusses the geometric effects that affect the magnitude of the heterodyne signal and set the wavenumber response of the interferometer. Section 3.2 explores heterodyne measurements made beyond the saturation intensity of a given detector. Sections 3.3 and 3.4 reveal how phase noise and amplitude noise, respectively, can creep into the interferometer's measurements. Section 3.5 describes demodulation of the heterodyne interference signal and the distortion of the baseband phase signal that results from demodulator imperfections. Section 3.6 discusses the signal quantization that necessarily occurs when generating a digital record. Finally, for ease of reference, Section 3.7 provides a concise summary of the discussed design considerations. These design considerations will be referenced extensively in Chapter 4, which describes the addition of a heterodyne interferometer to the pre-existing phase contrast imaging (PCI) diagnostic on the DIII-D tokamak.

3.1 GEOMETRIC CONSIDERATIONS

Several geometric effects substantially influence the performance of a heterodyne interferometer. Section 3.1.1 provides a minimum threshold on the radii of components in the optical train, while Section 3.1.2 derives the required degree of coalignment between the probe beam and the reference beam. Section 3.1.3 discusses the implications of mismatches between the spatial structures of the probe beam and the refer-

ence beam and develops a criterion for the required level of matching. Section 3.1.4 reveals how the imaging system's magnification M and the detector's size and shape influence the interferometer's wavenumber response. Finally, Section 3.1.5 examines the depth-of-focus effects that become manifest when the detector is displaced from the actual image plane.

3.1.1   *Aperture diffraction*

Diffraction from finite-aperture optics was neglected in Chapter 2's transfer-function derivations. For a propagating Gaussian beam, this neglect of aperture diffraction is a reasonable approximation if

$$a_{\text{eff}} \geqslant \frac{3}{2} w(z), \tag{3.1}$$

for each aperture, where a_{eff} is the effective aperture radius and $w(z)$ is the beam's $1/e$ E radius at the aperture location [1, 2]. For a circular aperture of radius a, the effective aperture radius is simply $a_{\text{eff}} = a$ for a beam propagating along the optical axis (e.g. the unscattered beam from Sec. 2.2.2); however, for a beam located $\rho(z)$ away from the optical axis (e.g. the upscattered or downscattered beam from Sec. 2.2.2), the effective aperture radius is $a_{\text{eff}} = a - |\rho(z)|$.

3.1.2   *Beam coalignment*

For the moment, assume a plane-wave representation for both the reference beam and the unscattered probe beam. Specifically, let the reference beam be given by

$$\boldsymbol{E}_{\text{R}}(\boldsymbol{r}) = E_0 \hat{\boldsymbol{x}} \cdot e^{i[k_0 z - (\omega_0 + \Delta\omega_0)t]}, \tag{3.2}$$

and let the unscattered probe beam be misaligned with the reference beam by angle $\theta \ll 1$ such that, to lowest order in θ, the unscattered probe beam is

$$\boldsymbol{E}_{\text{P}}(\boldsymbol{r}) \approx E_0 \hat{\boldsymbol{x}} \cdot e^{i[k_0(z + \theta x) - \omega_0 t]}. \tag{3.3}$$

The total intensity (averaged over an optical cycle) is then

$$I = \frac{c\varepsilon_0}{2} |\boldsymbol{E}_{\text{R}} + \boldsymbol{E}_{\text{P}}|^2 \approx 2I_0 \left[1 + \cos(\Delta\omega_0 t + k_0 \theta x)\right], \tag{3.4}$$

where $I_0 = c\varepsilon_0 E_0^2/2$ is the corresponding intensity of a single beam. Here, the cosine term corresponds to the interference between the two beams, and the unity term corresponds to the intensity of each individual beam. Optimizing the interference signal requires alignment of the beam polarizations and minimization of the misalignment angle θ. If the interference is measured by a detector with an extent s_x in the x-direction, the misalignment-induced phase $k_0\theta x$ should change by much less than 2π across the detector face; i.e. $|k_0\theta s_x| \ll 2\pi$ or

$$|\theta| \ll \frac{\lambda_0}{s_x} \approx 0.6°, \tag{3.5}$$

where $\lambda_0 = 2\pi/k_0$ is the beam wavelength, and $\lambda_0 = 10.6\,\mu m$ and $s_x = 1\,mm$ have been used for the evaluation. While coalignment constraint (3.5) is relatively easy to satisfy during a system alignment, it does have design implications for CO_2 interferometers that are built for magnetic fusion experiments, which are often characterized by large, pulsed electromagnets whose operation may contort the machine and produce vibrations, potentially destroying the beam coalignment.

3.1.3 *Mismatch between beam spatial structures*

The external reference-beam interferometry derivations in Section 2.4 assumed that the reference beam was exactly matched in both amplitude and spatial structure to the unscattered probe beam. This is obviously an idealization that, at best, can only be approximately met in experiment. This section discusses the geometric effects of such imperfections in beam matching.

The derivation of the heterodyne intensity (2.68) can be easily generalized to account for the geometric effects of unmatched reference and probe beams. Namely, let the image-plane probe radiation be given by (2.59) as

$$E_P(\boldsymbol{r}_J, t) \approx E_{G,P}(\boldsymbol{r}_J, t)e^{i\overline{\phi}}\left[1 + i\widetilde{\phi}(x_J, t)\right], \tag{3.6}$$

and let the corresponding reference beam be given by

$$E_R(\boldsymbol{r}_R, t) = E_{G,R}(\boldsymbol{r}_R, t)e^{-i\Delta\omega_0 t}, \tag{3.7}$$

where $\boldsymbol{r}_J = (x_J, y_J, z_J)$,

$$\boldsymbol{r}_R = \boldsymbol{r}_J + (0, 0, z_R - z_J), \tag{3.8}$$

and $E_{G,j}$ is a Gaussian beam with angular frequency ω_0, waist amplitude $E_{0,j}$, and waist $1/e$ E radius $w_{0,j}$. If $z_R \neq z_J$, the reference beam's waist sits at a different location than that of the unscattered probe beam. Under these circumstances and to first order in $\widetilde{\phi}$, the heterodyne intensity (averaged over an optical cycle) becomes

$$\begin{aligned} I_{het}(\boldsymbol{r}_J, z_R, t) = {} & I_{DC} + I_{AC}\cos(\Delta\omega_0 t + \overline{\phi}_{eff}) \\ & - I_{AC}\widetilde{\phi}(x_J, t)\sin(\Delta\omega_0 t + \overline{\phi}_{eff}), \end{aligned} \tag{3.9}$$

where

$$I_{DC} = I_{G,P}(\boldsymbol{r}_J) + I_{G,R}(\boldsymbol{r}_R), \tag{3.10}$$

$$I_{AC} = 2\sqrt{I_{G,P}(\boldsymbol{r}_J) \cdot I_{G,R}(\boldsymbol{r}_R)} \tag{3.11}$$

are the DC and AC components of the heterodyne intensity, respectively,

$$I_{G,j}(\boldsymbol{r}) = \frac{c\varepsilon_0 |E_{G,j}(\boldsymbol{r})|^2}{2} \tag{3.12}$$

is the intensity profile (averaged over an optical cycle) of Gaussian beam $j \in \{P, R\}$,

$$\overline{\phi}_{eff} = \overline{\phi} + \left[\phi_{G,P}(\boldsymbol{r}_J) - \phi_{G,R}(\boldsymbol{r}_R)\right] \tag{3.13}$$

is the effective bulk phase, and

$$\phi_{G,j}(\boldsymbol{r}) = k_0 z + \frac{k_0 \rho^2}{2R_j(z)} - \psi_j(z) \tag{3.14}$$

is the phase of Gaussian beam $j \in \{P, R\}$ (i.e. $E_{G,j}(\boldsymbol{r}) = |E_{G,j}(\boldsymbol{r})|e^{i\phi_{G,j}(\boldsymbol{r})}$). Note that (3.9) readily reduces to (2.68) if $E_{G,R}(\boldsymbol{r}_R) = E_{G,P}(\boldsymbol{r}_J)$.

It is worth discussing the implications of heterodyne intensity (3.9). First, repeating the derivation between (2.69) and (2.77) with this modified heterodyne intensity, one readily finds that the heterodyne interferometer's wavenumber transfer function (2.77) should be multiplied by the factor $2I_{AC}/(I_{DC} + I_{AC})$. It is then easy to show that the transfer function is *maximized* when $I_{DC} = I_{AC}$, i.e. when the probe beam and reference beam have identical spatial structures and powers. Second, note that the effective bulk phase $\overline{\phi}_{eff}$ is dependent on the geometry of the reference beam and the unscattered probe beam. Specifically, in the context of measuring the plasma-induced bulk phase $\overline{\phi}$, note that

$$\overline{\phi}_{eff}(\rho_J = 0) = \overline{\phi} + k_0(z_J - z_R) - [\psi_P(z_J) - \psi_R(z_R)]. \tag{3.15}$$

If z_J and z_R are fixed, then the beam-geometry contributions to $\overline{\phi}_{eff}(\rho_J = 0)$ constitute an unimportant DC offset that can be removed via baseline subtraction; however, experiments are typically plagued by vibrations, and even small changes to z_J and z_R can make significant time-dependent contributions to $\overline{\phi}_{eff}(\rho_J = 0)$ at CO_2 probe wavelengths. As such, deconvolving the plasma-induced and vibration-induced contributions to $\overline{\phi}_{eff}(\rho_J = 0)$ requires interferometric measurements at two distinct wavelengths (i.e. two-color interferometry) [3]. However, such vibrations occur on slow time-scales (e.g. $f_{vib} \lesssim 5\,\mathrm{kHz}$), and phase measurements at a *single* wavelength are sufficient to quantify plasma-induced phase fluctuations at frequencies above f_{vib} [4]. Finally, note that the beam geometry also imparts a spatially dependent, curvature-induced phase shift

$$
\delta\phi_\kappa(\rho_J) = \overline{\phi}_{eff}(\rho_J) - \overline{\phi}_{eff}(\rho_J = 0)
$$
$$
= \frac{k_0 \rho_J^2}{2} \left[\frac{1}{R_P(z_J)} - \frac{1}{R_R(z_R)} \right],
\tag{3.16}
$$

which can result in signal loss and distortion of the measured wavenumber. To see this, assume that the radiation is interfered on a detector array, as shown in Figure 3.1. As a detector element produces a signal proportional to the average intensity across its face, there will be substantial signal loss if there are large curvature-induced phase shifts across the element's face (i.e. $\delta\phi_\kappa(s_x/2) \gtrsim \pi$ or $\delta\phi_\kappa(s_y/2) \gtrsim \pi$). Further, if there are large curvature-induced phase shifts across the length of the detector array, the spatial structure of the intensity will *not* correspond to the spatial structure of the plasma fluctuation. The latter is the more conservative constraint on the curvature-induced phase shift. Assuming that the detector array shown in Figure 3.1 consists of N_{el} detector elements and that the inter-element spacing is negligible ($\delta_x \ll s_x$), the criterion for negligible curvature-induced phase shifts $\max(\delta\phi_\kappa) = \delta\phi_\kappa(\rho_{J,max}) \ll \pi$ becomes

$$
\max(\delta\phi_\kappa) = \frac{k_0}{8} \left[(N_{el}s_x)^2 + s_y^2 \right] \left| \frac{1}{R_P(z_J)} - \frac{1}{R_R(z_R)} \right| \ll \pi.
\tag{3.17}
$$

3.1.4 *Finite sampling-volume effects*

Practically speaking, detection is always effected via detector elements of *finite* size, with the output of each detector element corresponding to the incident intensity *averaged* over the element's active area. This
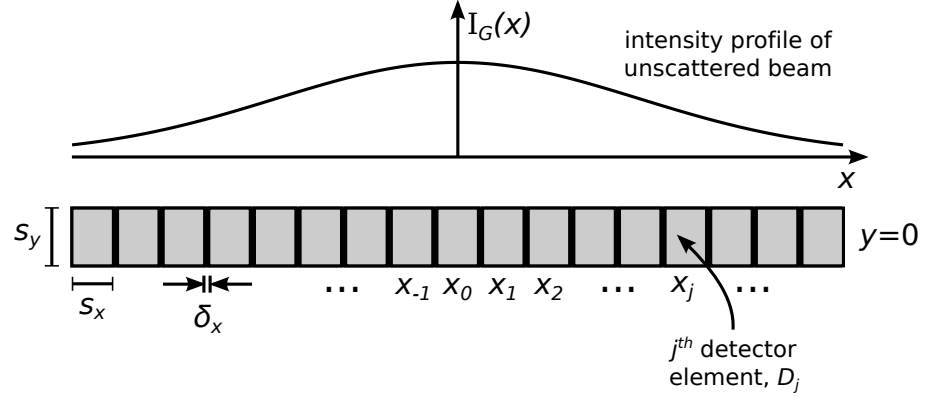
Figure 3.1: The probe radiation and the reference beam are interfered on a detector array. The array consists of numerous detector elements, each of size $s_x \times s_y$ and with interelement spacing δ_x. The unscattered beam is centered on $x = x_0$ and $y = 0$, and its intensity profile varies only weakly over any given element. The finite size of each detector element tends to attenuate short wavelength components of the incident optical signal.

averaging acts as a low-pass filter in the spatial domain and is referred to as the finite sampling-volume effect [5].

Finite sampling-volume effects dictate a heterodyne interferometer's wavenumber response [6]. To see this, assume that measurements are made with the array of rectangular detector elements shown in Figure 3.1 (for circular elements, see Coda's discussion [7, Sec. 3.7]). Let the j^{th} detector element D_j be centered on $x_{\jmath,j}$ and $y_{\jmath} = 0$. Integrating the optical intensity $\widetilde{I}_{IQ}(r_{\jmath}, t)$ corresponding to fluctuations in the baseband signal from (2.75) over the face of detector element D_j yields the corresponding optical power

$$\widetilde{P}_{IQ,j}(t) = \int_{D_j} \widetilde{I}_{IQ}(r_{\jmath}, t) dA$$

$$\approx \frac{2\sqrt{2}}{\pi} I_G(r_{\jmath,j}) s_y \int_{x_{\jmath,j}-s_x/2}^{x_{\jmath,j}+s_x/2} \widetilde{\phi}(x_{\jmath}, t) dx_{\jmath}; \tag{3.18}$$

here, the intensity profile $I_G(r_{\jmath})$ has been assumed to be approximately constant over the face of the detector element. Because (3.18) is linear in $\widetilde{\phi}$, it is suitable to consider a single Fourier mode

$$\widetilde{\phi}(x_{\jmath}, t) = \widetilde{\phi}_0 \cos(k_{\jmath} x_{\jmath} - \omega t) \tag{3.19}$$

for which (3.18) reduces to

$$\widetilde{P}_{IQ,j}(t) = \frac{2\sqrt{2}}{\pi} I_G(r_{\jmath,j}) A \cdot T_{fsv}(k_{\jmath}) \cdot \widetilde{\phi}_0 \cos(k_{\jmath} x_{\jmath} - \omega t), \tag{3.20}$$

where $A = s_x s_y$ is the area of the detector element,

$$T_{fsv}(k_J) \equiv \mathrm{sinc}\left(\frac{k_J}{k_{fsv,J}}\right) \tag{3.21}$$

is the finite sampling-volume transfer function,

$$\mathrm{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{3.22}$$

is the normalized sinc function, and

$$k_{fsv,J} = \frac{2\pi}{s_x} \tag{3.23}$$

is the first zero of $T_{fsv}(k_J)$. Recalling that an object-plane wavenumber k is imaged as $k_J = k/M$ in a magnification-M imaging system, the corresponding object-plane finite sampling-volume wavenumber cutoff is

$$k_{fsv} = \frac{2\pi|M|}{s_x}. \tag{3.24}$$

Now, as in Section 2.4.2, select the central intensity of the unscattered beam at the detector to be $I_G(0) = I_{sat}/4$, where I_{sat} is the detector's linear saturation intensity, such that

$$\frac{\widetilde{P}_{IQ,j}(t)}{I_{sat}A} = \frac{I_G(r_{J,j})}{I_G(0)} \cdot T_{het}(k_J) \cdot \widetilde{\phi}_0 \cos(k_J x_J - \omega t), \tag{3.25}$$

where

$$T_{het}(k_J) = \frac{1}{\sqrt{2}\cdot\pi} \cdot T_{fsv}(k_J) \tag{3.26}$$

is the heterodyne interferometer's wavenumber transfer function. In the limit $s_x \to 0$, $T_{fsv} \to 1$ and the heterodyne interferometer's wavenumber transfer function reduces to (2.77). Thus, finite sampling-volume effects introduce a wavenumber dependence into T_{het}, as shown in Figure 3.2. Note that finite sampling-volume effects introduce similar wavenumber dependencies into the transfer functions of the homodyne interferometer and PCI.

3.1.5 *Depth of focus*

The derivations in Sections 2.3 through 2.5 assumed that the detector sits exactly at the image plane. Empirically, however, uncertainties in
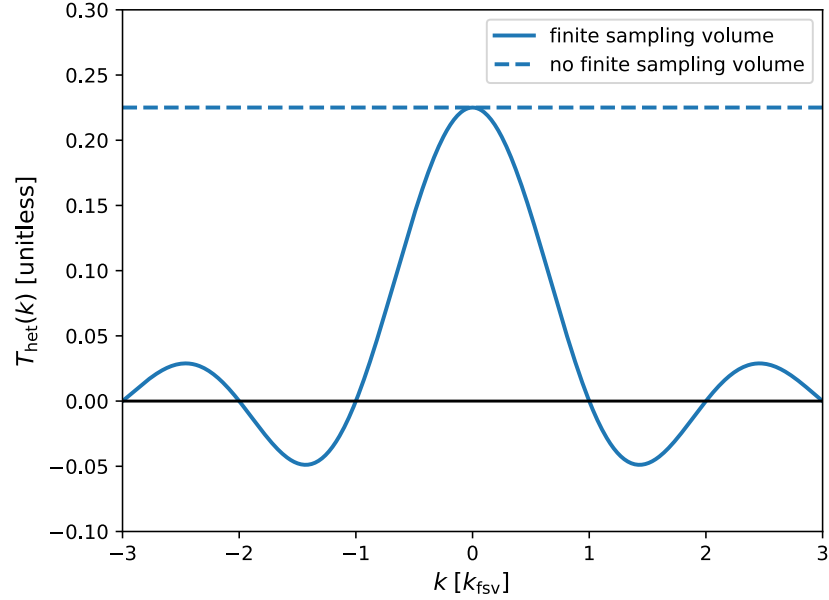
Figure 3.2: The wavenumber transfer function for a heterodyne interferometer with finite sampling-volume effects and without finite sampling-volume effects.

distances and focal lengths produce a corresponding uncertainty in the image-plane location. The axial distance by which the detector location may deviate from the image plane is referred to as the *depth of focus*. Even if all components are perfectly positioned and all focal lengths are equal to their nominal values, the image plane still only maps to the tokamak midplane, and images from points above and below the tokamak midplane will be "out of focus"; such depth-of-field considerations are closely related an imaging system's depth of focus.

In order to investigate a heterodyne interferometer's depth of focus, it is necessary to generalize the derivation of the imaged electric field from Section 2.3. In particular, let the detector be located an axial distance δz_J downstream of the image plane (i.e. $z_{\text{det}} = z_J + \delta z_J$ such that positive δz_J implies that the detector is downstream of the image plane, and negative δz_J implies that the detector is upstream of the image plane). Referencing the detector-plane coordinate transformation (B.14), the imaged field at the detector (2.55) readily generalizes to

$$E(\boldsymbol{r}_{\text{det}}, t) \approx E_G(\boldsymbol{r}_{\text{det}}, t)e^{i\overline{\phi}}\left[1 + ie^{-i\mu}(\widetilde{\phi}_0 \cos\nu')\right], \tag{3.27}$$

where

$$\mu = \left(\frac{k_J^2}{2k_0}\right)\delta z_J, \tag{3.28}$$

$$\nu' = \left\{\left[1 - \frac{\delta z_J}{R(z_{det})}\right]k_J\right\}x_{det} - \omega t, \tag{3.29}$$

$k_J = k/M$ is the image of wavenumber k in a magnification-M imaging system, and x_{det} and δz_J are assumed to be small relative to the radius of curvature $R(z_{det})$ of the probe beam at the detector. Relative to the image-plane field (2.55), the detector-plane field (3.27) has a wavenumber distortion $\{1 - [\delta z_J/R(z_{det})]\}$. Additionally, the detector-plane field has a wavenumber-dependent phase shift μ, which results from the "out-of-focus" interference of the upscattered beam with the downscattered beam; it has been empirically demonstrated that blocking one of these scattered beams eliminates this wavenumber-dependent phase shift [8, Sec. 2.1], albeit at the expense of reducing the amplitude of the fluctuating signal by a factor of two.

Heterodyne detection is implemented by interfering the detector-plane field (3.27) with a frequency-shifted reference beam. In particular, assume a reference beam of the form $E_G(\boldsymbol{r}_{det}, t)e^{-i\Delta\omega_0 t}$ such that intensity (averaged over an optical cycle) becomes

$$\begin{aligned}I_{het}(\boldsymbol{r}_{det}, t) = 2I_G(\boldsymbol{r}_{det})\big[1 &+ \cos(\Delta\omega_0 t + \overline{\phi}) \\ &- \widetilde{\phi}_0\cos\nu'\sin(\Delta\omega_0 t + \overline{\phi} - \mu)\big].\end{aligned} \tag{3.30}$$

Demodulating this heterodyne intensity via the program in Section 2.4, one obtains the equilibrium and fluctuating components of the in-phase and quadrature intensities

$$\overline{I}_I(\boldsymbol{r}_{det}, t) = \frac{2\sqrt{2}}{\pi}I_G(\boldsymbol{r}_{det})\cos\overline{\phi}, \tag{3.31}$$

$$\overline{I}_Q(\boldsymbol{r}_{det}, t) = \frac{2\sqrt{2}}{\pi}I_G(\boldsymbol{r}_{det})\sin\overline{\phi}, \tag{3.32}$$

$$\widetilde{I}_I(\boldsymbol{r}_{det}, t) = -\frac{2\sqrt{2}}{\pi}I_G(\boldsymbol{r}_{det})\sin\left(\overline{\phi} - \mu\right)\cdot\widetilde{\phi}_0\cos\nu', \tag{3.33}$$

$$\widetilde{I}_Q(\boldsymbol{r}_{det}, t) = \frac{2\sqrt{2}}{\pi}I_G(\boldsymbol{r}_{det})\cos\left(\overline{\phi} - \mu\right)\cdot\widetilde{\phi}_0\cos\nu'. \tag{3.34}$$

It is possible to compute the autospectral density of the phase fluctuations by summing the autospectral densities of $\widetilde{I}_I$ and $\widetilde{I}_Q$ and then normalizing by $(\overline{I}_I^2 + \overline{I}_Q^2)$, which fortuitously *eliminates* the μ-dependence. However, it is *not* readily apparent how to extend this method to the

computation of e.g. cross-spectral densities [9, Sec. 5.2] and bispectral densities [10], for which the actual phase signal $\widetilde{\phi}(x, t)$ is needed. For this reason, the phase is often computed via the inverse-tangent method discussed in Section 3.5.1. Unfortunately, this inverse-tangent method has reduced response when $\mu \neq 0$. To see this, first consider the case when $\mu = 0$: $\bar{I}_I$ and $\bar{I}_Q$ map out a circle, and the fluctuation $\widetilde{\phi}_0 \cos \nu'$ produces perturbations $\widetilde{I}_I$ and $\widetilde{I}_Q$ that are always *tangent* to this circle, producing angular deviations in the (I, Q)-plane that can be quantified by the inverse-tangent calculation. Now, if $\mu \neq 0$, the perturbations $\widetilde{I}_I$ and $\widetilde{I}_Q$ are no longer tangent to the circle mapped out by $\bar{I}_I$ and $\bar{I}_Q$; in the extreme that $\mu = (2m + 1)\pi/2$ for integer m, the fluctuations become purely *radial* in the (I, Q)-plane, and the inverse-tangent calculation, sensitive only to angular displacements, fails to detect the fluctuations.

3.2   INTENSITY CONSIDERATIONS

Ideally, a photovoltaic detector produces an output voltage

$$V(t) = \mathcal{R}_0 \cdot I(t), \tag{3.35}$$

where $\mathcal{R}_0$ is the detector responsivity and $I(t)$ is the incident optical intensity. However, every real-world detector has a saturation intensity I_{sat} beyond which the output voltage ceases to be a linear function of the incident optical intensity; that is, the detector responsivity has an intensity dependence $\mathcal{R}(I)$, and the detector voltage can be more generally written as

$$V(t) = \mathcal{R}\left(I(t)\right) \cdot I(t). \tag{3.36}$$

Here, $\mathcal{R}(I)$ is an arbitrary monotonically increasing function of the incident optical intensity I. Despite the potentially nonlinear response, the detector voltage remains periodic in $2\pi/\Delta\omega_0$ and can be expanded in a Fourier series as

$$V(t) = V_0 + \sum_{n=1}^{\infty} V_n \cos\left(n\Delta\omega_0 t + \theta_n\right), \tag{3.37}$$

where V_n and θ_n are the amplitude and phase, respectively, of the n^{th} harmonic. Thus, a nonlinear detector response produces higher-order harmonics in the signal. In general, V_n and θ_n can vary in time, producing sidebands about each harmonic. Provided the bandwidth of these fluctuations is sufficiently low, there will be no spectral overlap
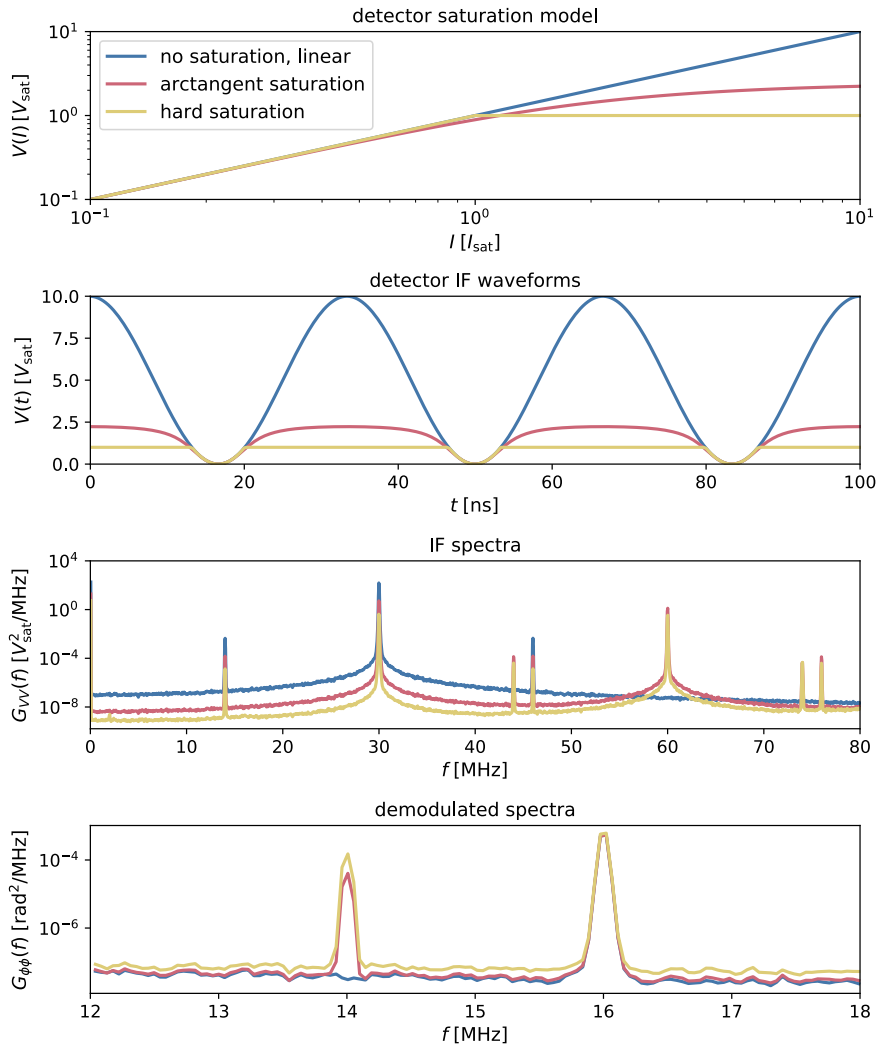
Figure 3.3: Heterodyne detection beyond the saturation intensity with $\Delta\omega_0 = 2\pi \cdot 30\,\mathrm{MHz}$ and $\omega = 2\pi \cdot 16\,\mathrm{MHz}$. (Top panel): Various detector saturation models. The linear model exhibits no saturation, the hard saturation model limits the output voltage to V_{sat} when the incident optical intensity exceeds I_{sat}, and the arctangent saturation model exhibits 1 dB compression when the incident optical intensity is I_{sat}. (2nd panel): The intermediate frequency (IF) waveforms corresponding to each saturation model when $I_{max} = 10\,I_{sat}$. The arctangent and hard saturation models distort the IF waveform, producing numerous higher-order harmonics. (3rd panel): Autospectral densities of the IF waveforms. The IF waveforms all exhibit peaks at the 30 MHz fundamental and its corresponding sidebands at $30\,\mathrm{MHz} \pm 16\,\mathrm{MHz}$. However, the saturated IF waveforms also exhibit peaks at the second harmonic (60 MHz) and its corresponding sidebands ($60\,\mathrm{MHz} \pm 16\,\mathrm{MHz}$). (Bottom panel): Autospectral densities of the demodulated phase. Note that the 16 MHz fluctuation is correctly identified when demodulating all of the IF waveforms. However, the saturated IF waveforms also produce a *spurious* 14 MHz fluctuation, which is attributable to the overlap of the 30 MHz and 60 MHz sidebands.

between the sidebands of adjacent harmonics, and bandpass filtering the detector signal about $\Delta\omega_0$ yields

$$V(t) \approx V_1 \cos[\Delta\omega_0 t + \theta_1(t)] \tag{3.38}$$

with $\theta_1(t) = \phi(t)$, where $\phi(t)$ is the optical phase shift between the plasma and reference arms of the interferometer. However, for fluctuations with sufficiently high bandwidth (e.g. $\omega \sim \Delta\omega_0/2$), the sidebands of adjacent harmonics begin to overlap, potentially corrupting the phase measurement, as demonstrated by the example in Figure 3.3. By processing the bandpass-filtered voltage (3.38) via a procedure similar to that outlined between (2.69) and (2.77), one readily finds that the heterodyne interferometer's wavenumber transfer function (2.77) should be multiplied by the factor $[V_1/V_1(I_{max} = I_{sat})]$, where V_1 is the amplitude of the bandpass-filtered voltage (3.38), and $V_1(I_{max} = I_{sat})$ is the amplitude of the corresponding bandpass-filtered voltage when the system's maximum intensity is scaled to the saturation intensity (the AC and DC fractions of the incident intensity are *not* altered during this scaling; to account for changing the AC and DC fractions, see the discussion in Section 3.1.3).

3.3   PHASE NOISE: SOURCES & EFFECTS

Heterodyne interferometry at 10.6 µm relies on both an optical oscillator (the laser) and a radio-frequency oscillator (usually referred to as the local oscillator (LO)). These oscillators, like all real-world oscillators, exhibit phase noise. The spectral properties and implications of oscillator phase noise are reviewed in Appendix D. Below, Section 3.3.1 shows that a mismatch between the optical path lengths of the probe beam and the reference beam injects the laser's phase noise into the heterodyne interferometer's measurements, while Section 3.3.2 shows that finite coupling time in the Doppler-shifting modulator injects the LO's phase noise into the heterodyne interferometer's measurements.

3.3.1   *Unmatched optical path lengths & laser phase noise*

The external reference-beam interferometry derivations in Section 2.4 assumed that the laser's angular frequency was fixed at its nominal value ω_0. However, the angular frequency of any *real* laser will exhibit small fluctuations in time, much like any other real-world os-

cillator [11, Sec. 1.7]. The electric field of such a Gaussian beam is well-described by

$$E_G(r,t) = E_G(r)e^{-i[\omega_0 t + \phi_{\omega_0}(t)]},$$ (3.39)

where $\phi_{\omega_0}(t)$ is a zero-mean, stationary, random process known as the laser's *phase deviation* whose temporal variation causes the laser's instantaneous angular frequency to wander about its nominal value ω_0.

Now, if the interferometer's probe beam and reference beam traverse different optical path lengths, the laser's phase deviation will inject phase noise into the measured signal. To see this, assume that the optical path length of the probe beam exceeds that of the reference arm by L. Then, if the reference beam impinging on the detector at time t is

$$E_R(r_J,t) = E_G(r_J)e^{-i[(\omega_0 + \Delta\omega_0)t + \phi_{\omega_0}(t)]},$$ (3.40)

the corresponding imaged probe radiation from (2.59) becomes

$$E_P(r_J,t) = E_G(r_J)e^{-i[\omega_0(t-\tau) + \phi_{\omega_0}(t-\tau)]}e^{i\overline{\phi}}\left[1 + i\widetilde{\phi}(x_J,t)\right],$$ (3.41)

where $\tau = L/c$ is the time delay associated with the optical path-length difference L. The resulting intensity (averaged over an optical cycle) of the interfering radiation is

$$I_{het}(r_J,t) \approx 2I_G(r_J)\big[1 + \cos(\Delta\omega_0 t + \overline{\phi}_{eff})$$
$$- \widetilde{\phi}_m(x_J,t)\sin(\Delta\omega_0 t + \overline{\phi}_{eff})\big],$$ (3.42)

where

$$\overline{\phi}_{eff} = \overline{\phi} + \omega_0\tau,$$ (3.43)

$$\widetilde{\phi}_m(x_J,t) = \widetilde{\phi}(x_J,t) + \delta\phi_{\omega_0}(t-\tau,\tau),$$ (3.44)

$$\delta\phi_{\omega_0}(t,\tau) = \phi_{\omega_0}(t+\tau) - \phi_{\omega_0}(t).$$ (3.45)

The quantity $\delta\phi_{\omega_0}(t,\tau)$ is referred to as the "instrumental phase noise", and it is produced by the optical path-length difference L and the laser's phase deviation $\phi_{\omega_0}(t)$. Typically, $|\delta\phi_{\omega_0}(t,\tau)| \ll 1$, and only terms to first order in $\delta\phi_{\omega_0}$ and $\widetilde{\phi}$ were retained in the heterodyne intensity (3.42). Comparing (3.42) with (2.68), one readily sees that the *measured* phase fluctuation is $\widetilde{\phi}_m$ defined in (3.44); that is, the fluctuating signal is contaminated by the instrumental phase noise. The spectral properties of $\delta\phi_{\omega_0}(t,\tau)$ are thoroughly discussed in Appendix D.

As $\delta\phi_{\omega_0}(t, \tau)$ and $\widetilde{\phi}(x_J, t)$ are uncorrelated, the one-sided autospectral density of the measured phase fluctuations is

$$G_{\widetilde{\phi}_m, \widetilde{\phi}_m}(f) = G_{\widetilde{\phi}, \widetilde{\phi}}(f) + 8\sin^2(\pi f \tau)\mathcal{L}_{\omega_0}(f), \qquad (3.46)$$

where $G_{\widetilde{\phi}, \widetilde{\phi}}(f)$ is the *true* one-sided autospectral density of the phase fluctuation $\widetilde{\phi}$ and $\mathcal{L}_{\omega_0}(f)$ is the phase noise of the laser, as defined in Appendix D.

3.3.2   *Modulator's finite coupling time & LO phase noise*

Heterodyne detection is effected by modestly Doppler shifting the reference beam relative to the plasma beam. It is easy to Doppler shift $10.6\,\mu m$ radiation by tens of MHz with a Germanium acousto-optic modulator (AOM). The operation of a typical AOM is sketched in Figure 3.4. A piezo-actuator drives sound waves of angular frequency $\Delta\omega_0$ through the Germanium crystal, and the sound waves act as a diffraction grating that propagates at the crystal's sound speed c_s. When a beam of vacuum wavelength λ_0 impinges upon the crystal at the Bragg angle

$$\theta_B = \frac{\lambda_0 \cdot \Delta\omega_0}{4\pi c_s}, \qquad (3.47)$$

a portion of the beam is deflected and Doppler shifted by angular frequency $\Delta\omega_0$ [12, Sec. 20.1]. The power in the deflected beam is controlled by the intensity of the sound waves.

The coupling of the AOM's drive signal to the deflected beam occurs on the crystal's sound-wave timescale. If the sound waves must propagate a distance d from the piezo-actuator to the AOM's optically active region, the drive signal is coupled to the deflected beam only after time delay $\tau = d/c_s$. The sound speed in Germanium is $c_s = 5400\,m \cdot s^{-1}$ such that a distance $d = 1\,cm$ is accompanied by a time delay $\tau = 1.85\,\mu s$. Note that this is large compared to many other timescales typically considered in interferometry; for example, light propagates through $1\,m$ of air in only $3.33\,ns$, and an RF signal propagates through $1\,m$ of RG-58 coaxial cable (for which the index of refraction is $\sim 3/2$) in only $5\,ns$.

In the presence of LO phase noise, an AOM's finite coupling time can degrade the performance of a heterodyne interferometer. A local oscillator with phase noise is well-described by

$$V_{LO}(t) = V_0 e^{-i[\Delta\omega_0 t + \phi_{\Delta\omega_0}(t)]}, \qquad (3.48)$$
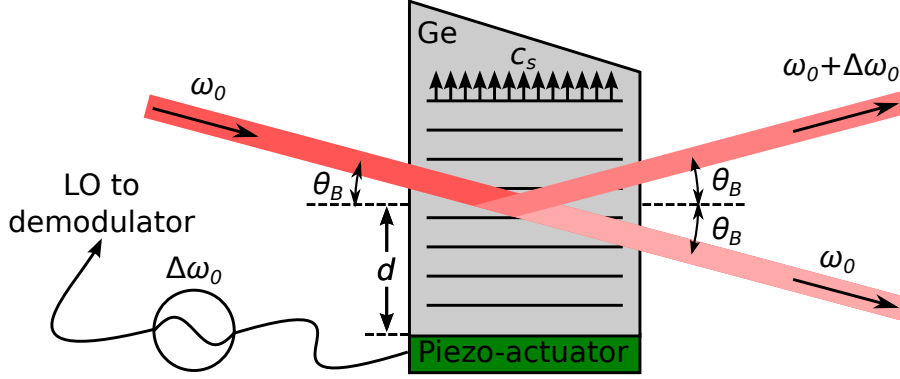
Figure 3.4: An illustration of AOM operation in a heterodyne interferometer. A piezo-actuator drives sound waves of angular frequency $\Delta\omega_0$ through a crystal (usually Germanium for $10.6\,\mu m$ light), and these sound waves deflect and Doppler shift light that is incident upon the crystal at the Bragg angle θ_B. The sound waves propagate from the piezo-actuator to the AOM's optically active region over finite time $\tau = d/c_s$. The RF waveform that drives the piezo-actuator is sampled and used to demodulate the heterodyne interference signal. Note that for simplicity the refraction of the beam as it enters and exits the crystal is *not* depicted.

where $\phi_{\Delta\omega_0}(t)$ is a zero-mean, stationary, random process known as the LO's *phase deviation* whose temporal variation causes the LO's instantaneous angular frequency to wander about its nominal value $\Delta\omega_0$. Then, to account for the AOM's finite coupling time and the LO phase noise, take $\Delta\omega_0 t \to [\Delta\omega_0(t-\tau) + \phi_{\Delta\omega_0}(t-\tau)]$ in the heterodyne intensity (2.68). Neglecting finite sampling-volume effects, the heterodyne output voltage from a given detector element is simply proportional to the local intensity, i.e.

$$
\begin{aligned}
V_{\text{het}}(t) = 2V_0 \big\{ &1 + \cos\big[\Delta\omega_0 t + \overline{\Phi}_{\text{eff}} + \phi_{\Delta\omega_0}(t-\tau)\big] \\
&- \widetilde{\phi}(x_{\mathcal{J}}, t)\sin\big[\Delta\omega_0 t + \overline{\Phi}_{\text{eff}} + \phi_{\Delta\omega_0}(t-\tau)\big]\big\},
\end{aligned}
\tag{3.49}
$$

where $\overline{\Phi}_{\text{eff}} = \overline{\phi} - \Delta\omega_0\tau$. Then, taking inspiration from the demodulated optical intensities (2.69), the demodulated in-phase (V_I) and quadrature (V_Q) voltages are defined as

$$
\begin{aligned}
V_I(t) + i \cdot V_Q(t) &= \frac{1}{V_0}\langle V_{\text{LO}}(t) \cdot V_{\text{het}}(t)\rangle_{\Delta\omega_0} \\
&= V_0 e^{i\overline{\Phi}_{\text{eff}}}\big[1 + i\widetilde{\phi}_m(x_{\mathcal{J}}, t)\big],
\end{aligned}
\tag{3.50}
$$

where

$$
\widetilde{\phi}_m(x_{\mathcal{J}}, t) = \widetilde{\phi}(x_{\mathcal{J}}, t) + \delta\phi_{\Delta\omega_0}(t, -\tau),
\tag{3.51}
$$

$$
\delta\phi_{\Delta\omega_0}(t, \tau) = \phi_{\Delta\omega_0}(t+\tau) - \phi_{\Delta\omega_0}(t).
\tag{3.52}
$$

The quantity $\delta\phi_{\Delta\omega_0}(t,\tau)$ is referred to as the "instrumental phase noise", and it is produced by the modulator's finite coupling time τ and the LO's phase deviation $\phi_{\Delta\omega_0}(t)$. Typically, $|\delta\phi_{\Delta\omega_0}(t,\tau)| \ll 1$, and only terms to first order in $\delta\phi_{\omega_0}$ and $\tilde{\phi}$ were retained in the demodulated voltages (3.50). Comparing (3.50) with the demodulated intensities (2.69), one readily sees that the *measured* phase fluctuation is $\tilde{\phi}_m$ defined in (3.51); that is, the fluctuating signal is contaminated by the instrumental phase noise. The spectral properties of $\delta\phi_{\Delta\omega_0}(t,\tau)$ are thoroughly discussed in Appendix D. As $\delta\phi_{\omega_0}(t,\tau)$ and $\tilde{\phi}(x_J,t)$ are uncorrelated, the one-sided autospectral density of the measured phase fluctuations is

$$G_{\tilde{\phi}_m,\tilde{\phi}_m}(f) = G_{\tilde{\phi},\tilde{\phi}}(f) + 8\sin^2(\pi f\tau)\mathcal{L}_{\Delta\omega_0}(f), \qquad (3.53)$$

where $G_{\tilde{\phi},\tilde{\phi}}(f)$ is the *true* one-sided autospectral density of the phase fluctuation $\tilde{\phi}$ and $\mathcal{L}_{\Delta\omega_0}(f)$ is the phase noise of the LO, as defined in Appendix D.

3.4   AMPLITUDE NOISE: SOURCES & EFFECTS

Detector noise and optical shot noise are omnipresent contributors to amplitude noise in the heterodyne interference signal, while a noisy amplifier can degrade the signal-to-noise ratio. The demodulation of such amplitude noise is throughly discussed by Rakhmanov in [13]. While Rakhmanov does not explicitly consider quadrature heterodyne detection, his results can be naturally applied to quadrature heterodyne detection, as is done below.

3.4.1   *Detector noise*

Real-world detector operation is associated with intrinsic noise. This noise results from, among other things, Johnson thermal noise in the detector and shot noise in the background radiation flux [14]. A detector's noise is often characterized by its noise-equivalent power (NEP): when the power of the incident optical signal is equal to the NEP, the signal-to-noise ratio is unity. Consider optical power $\delta P(t)$ that, when incident upon a detector, produces a signal that emulates the statistical properties of the detector noise (e.g. $\delta P(t)$ is a real-valued, zero-mean,

stationary, random process). Then, the NEP corresponds to the root mean square (RMS) of $\delta P(t)$, and

$$(NEP)^2 = \int_{\Delta f} G_{\delta P, \delta P}(f) df, \tag{3.54}$$

where $G_{\delta P, \delta P}(f)$ is the one-sided autospectral density of $\delta P(t)$ and the integral is performed over the temporal bandwidth Δf of the desired measurement. Note that $G_{\delta P, \delta P}(f)$ depends on both extensive (e.g. element area) and intensive (e.g. element material) properties of the detector. To more easily compare detectors of different sizes and materials, the NEP of a given detector is often parameterized as

$$NEP = \frac{\sqrt{A \cdot \Delta f}}{D^*}, \tag{3.55}$$

where A is the effective area of the detector element, Δf is the temporal bandwidth of the desired measurement, and D^* is the detector's specific detectivity [15]. Note that larger D^* corresponds to increased detector sensitivity. If $G_{\delta P, \delta P}(f)$ is approximately constant over the bandwidth Δf, then equating (3.54) to the square of (3.55) yields

$$G_{\delta P, \delta P}(f) = \frac{A}{(D^*)^2}. \tag{3.56}$$

It is now shown shown that signal demodulation pushes detector noise near the heterodyne angular frequency $\Delta \omega_0$ into the baseband signal. Taking inspiration from the demodulated optical intensities (2.69), define the total NEP contamination of the in-phase (I) and quadrature (Q) signals to be

$$\delta P_{IQ}(t) = \frac{2\sqrt{2}}{\pi} e^{-i\Delta \omega_0 t} \cdot \delta P(t). \tag{3.57}$$

Here, the demodulated noise has *not* yet been low-pass filtered. The autocorrelation function of the demodulated noise is

$$\begin{aligned} R_{\delta P_{IQ}, \delta P_{IQ}}(\tau) &= E\left[\delta P_{IQ}(t) \cdot \delta P_{IQ}^*(t + \tau)\right] \\ &= \frac{8}{\pi^2} \cdot e^{i\Delta \omega_0 \tau} \cdot R_{\delta P, \delta P}(\tau), \end{aligned} \tag{3.58}$$

where z^* indicates the complex conjugate of z and $R_{\delta P, \delta P}(\tau)$ is the autocorrelation function of the NEP. The autospectral density of the demodulated noise is then

$$
\begin{aligned}
S_{\delta P_{IQ}, \delta P_{IQ}}(f) &= \mathcal{F}\left[R_{\delta P_{IQ}, \delta P_{IQ}}(\tau) \right](f) \\
&= \frac{8}{\pi^2} \cdot \mathcal{F}\left[e^{i2\pi\Delta f_0 \tau} \cdot R_{\delta P, \delta P}(\tau) \right](f) \\
&= \frac{8}{\pi^2} \cdot S_{\delta P, \delta P}(f + \Delta f_0),
\end{aligned}
\tag{3.59}
$$

where $\Delta f_0 = \Delta\omega_0 / (2\pi)$ is the heterodyne frequency and $S_{\delta P, \delta P}(f)$ is the autospectral density of the NEP. The demodulated signals are typically low-pass filtered such that only the desired information at $|f| \ll \Delta f_0$ survives

$$
S_{\delta P_{IQ}, \delta P_{IQ}}(f)\big|_{|f| \ll \Delta f_0} \approx \frac{8}{\pi^2} \cdot S_{\delta P, \delta P}(\Delta f_0).
\tag{3.60}
$$

Thus, the NEP at the heterodyne frequency Δf_0 is pushed into the baseband signal via the demodulation process. Note that the autospectral density of the demodulated detector noise (3.60) is in agreement with the literature (e.g. see Rakhmanov's Eq. (47) in [13] with $d_1 = 2/\pi$ for demodulation against the *first harmonic* of a zero-mean square wave with frequency Δf_0 and peak-to-peak amplitude of two; see Section 3.5.2 for the physical significance of demodulation against such a square wave). The corresponding one-sided autospectral density of the demodulated detector noise is

$$
\begin{aligned}
G_{\delta P_{IQ}, \delta P_{IQ}}(f)\big|_{|f| \ll \Delta f_0} &\approx \frac{8}{\pi^2} \cdot G_{\delta P, \delta P}(\Delta f_0). \\
&= \frac{8}{\pi^2} \cdot \frac{A}{(D^*)^2},
\end{aligned}
\tag{3.61}
$$

where the last line follows from (3.56) and the assumption that D^* is the specific detectivity for frequencies f in the neighborhood of the heterodyne frequency Δf_0.

Note that comparing the spectral density of the demodulated detector noise to the spectral density of the measured phase (with the phase having units of radians) requires converting (3.61) from units of $W^2 \cdot Hz^{-1}$ to the corresponding angular equivalent. This unit conversion is accomplished by dividing (3.61) by the *square* of the total demodulated power in the I and Q signals. Assuming that the probe beam and the reference beam vary weakly over the face of the detector element, this squared power is $P_{IQ}^2 \approx (8/\pi^2) P_P P_R$, where P_P is the optical power of the probe beam impinging on the detector, and P_R

is the optical power of the reference beam impinging on the detector. Thus, the one-sided autospectral density of the demodulated detector noise (in angular units of $rad^2 \cdot Hz^{-1}$) is

$$G_{\delta P_{IQ}, \delta P_{IQ}}(f)\Big|_{|f| \ll \Delta f_0} = \frac{A}{P_P P_R (D^*)^2}. \tag{3.62}$$

3.4.2 *Optical shot noise*

The discrete nature of the arriving photons results in shot noise. Well-modeled as a Poisson process, the optical shot noise increases as $N_\gamma^{1/2}$, where N_γ is the number of incident photons. Because the incident optical power (and hence the number of incident photons) in the heterodyne optical signal is modulated as a function of time, the corresponding shot noise is inherently nonstationary. Surprisingly, however, the demodulated shot noise *is* stationary (e.g. see Rakhmanov's Eq. (59) in [13]). Note that Rakhmanov only considers *one* of the demodulated signals, and maximizing the signal-to-noise ratio in the demodulated signal requires careful selection of the local oscillator's phase relative to that of the heterodyne signal (he terms this the "demodulation phase" and represents it via γ). However, by employing quadrature heterodyne detection [3], in which $\gamma_Q = \gamma_I + \pi/2$, the total shot-noise contamination δP_{IQ} of the in-phase (I) and quadrature (Q) signals (with δP_{IQ} as defined in (3.57), but with δP now corresponding to the shot-noise optical power) is *independent* of the demodulation phase, i.e.

$$S_{\delta P_{IQ}, \delta P_{IQ}}(f)\Big|_{|f| \ll \Delta f_0} = \frac{8}{\pi^2} \cdot h f_0 P_0, \tag{3.63}$$

where h is the Planck constant, f_0 is the frequency of the incident photons, and P_0 is the DC optical power impinging upon the detector. The corresponding one-sided autospectral density is

$$G_{\delta P_{IQ}, \delta P_{IQ}}(f)\Big|_{|f| \ll \Delta f_0} = \frac{16}{\pi^2} \cdot h f_0 P_0. \tag{3.64}$$

Rakhmanov notes that (3.64) corresponds to the well-known Schottky formula.

Note that comparing the spectral density of the demodulated optical shot noise to the spectral density of the measured phase (with the phase having units of radians) requires converting (3.64) from units of $W^2 \cdot Hz^{-1}$ to the corresponding angular equivalent. This unit conversion is accomplished by dividing (3.61) by the *square* of the total

demodulated power in the I and Q signals. Assuming that the probe beam and the reference beam vary weakly over the face of the detector element, this squared power is $P_{IQ}^2 \approx (8/\pi^2)P_P P_R$, where P_P is the optical power of the probe beam impinging on the detector, and P_R is the optical power of the reference beam impinging on the detector. Additionally, the DC optical power impinging upon the detector is $P_0 = P_P + P_R$. Thus, the one-sided autospectral density of the demodulated optical shot noise (in angular units of $rad^2 \cdot Hz^{-1}$) is

$$G_{\delta P_{IQ}, \delta P_{IQ}}(f)\big|_{|f| \ll \Delta f_0} = 2hf_0 \left(\frac{P_P + P_R}{P_P P_R} \right). \tag{3.65}$$

3.4.3 Amplifier noise

The noise *factor* F of an RF amplifier is defined as the ratio of the signal-to-noise ratio at the device's input (SNR_{in}) to the signal-to-noise ratio at the device's output (SNR_{out})

$$F = \frac{SNR_{in}}{SNR_{out}}; \tag{3.66}$$

often, the noise factor F is given in terms of the noise *figure* NF [16]

$$NF = 10 \log_{10} F. \tag{3.67}$$

If several amplifiers are cascaded, the total noise factor of the amplifier chain can be computed using the well-known Friis noise-factor formula. Note that the noise factor is only defined in the context of a signal-to-noise ratio, so it is not conventional to write down the corresponding autospectral density of the amplifier noise in absolute units.

3.5 DEMODULATION

The heterodyne interferometer's intermediate frequency (IF) signal must be demodulated in order to recover the baseband phase information. Demodulation is typically described as an analog process in which the IF signal is *mixed* with a local oscillator (LO), but demodulation can also be performed digitally [17, 18] or with non-mixer-based analog electronics [19]. The focus here, however, is on the analog, mixer-based approach. Section 3.5.1 describes ideal analog demodulation. Sections 3.5.2 and 3.5.3 discuss nonideal aspects of analog mixers and demodulators, and Section 3.5.4 analyzes the implications of these imperfections in the context of fluctuation measurements.

3.5.1 *Ideal demodulation*

A typical analog I&Q demodulator consists of a 90° splitter, two double-balanced mixers, and a 0° splitter [20], as shown in Figure 3.5(a). The 90° splitter divides the incident local oscillator (LO) signal

$$\mathrm{LO}(t) = \frac{4\sqrt{2}}{\pi} \cos(\Delta\omega_0 t) \tag{3.68}$$

into an "in-phase" copy of the LO

$$\mathrm{LO}_0(t) = \frac{4}{\pi} \cos(\Delta\omega_0 t) \tag{3.69}$$

and a $\pi/2$ phase-advanced copy of the LO

$$\mathrm{LO}_{\pi/2}(t) = \frac{4}{\pi} \cos\left(\Delta\omega_0 t + \frac{\pi}{2}\right) = -\frac{4}{\pi} \sin(\Delta\omega_0 t). \tag{3.70}$$

Note that the power (i.e. the mean-square value) in the incident LO signal is split evenly between $\mathrm{LO}_0(t)$ and $\mathrm{LO}_{\pi/2}(t)$. Further, the normalization of $\mathrm{LO}_0(t)$ and $\mathrm{LO}_{\pi/2}(t)$ is motivated by and is consistent with the physical processes that occur in a typical ring-diode double-balanced mixer, as discussed in Section 3.5.2. The 0° splitter divides the intermediate frequency (IF) signal

$$\mathrm{IF}(t) = A_{\mathrm{IF}} \cos(\Delta\omega_0 t + \phi) \tag{3.71}$$

into two identical copies of the IF

$$\mathrm{IF}_0(t) = \frac{A_{\mathrm{IF}}}{\sqrt{2}} \cos(\Delta\omega_0 t + \phi). \tag{3.72}$$

Like the LO signal, the power in the incident IF signal is split evenly between the two copies. The signal at the demodulator's in-phase (I) port then corresponds to the product of $\mathrm{IF}_0(t)$ with the in-phase LO signal $\mathrm{LO}_0(t)$:

$$\mathrm{LO}_0(t) \cdot \mathrm{IF}_0(t) = \frac{\sqrt{2}A_{\mathrm{IF}}}{\pi} \left[\cos(\phi) + \cos(2\Delta\omega_0 t + \phi)\right]. \tag{3.73}$$

Such signal multiplication is often referred to as "mixing". Low-pass filtering the signal exiting the demodulator's I port yields the baseband in-phase signal

$$I = \frac{\sqrt{2}A_{\mathrm{IF}}}{\pi} \cos\phi. \tag{3.74}$$

Similar reasoning regarding the product $LO_{\pi/2}(t) \cdot IF_0(t)$ at the demodulator's quadrature (Q) port yields the baseband quadrature signal

$$Q = \frac{\sqrt{2}A_{IF}}{\pi} \sin \phi. \tag{3.75}$$

Note that the total power $I^2 + Q^2$ in the I&Q signals is 4 dB lower than the power in the incident IF signal. This is known as the *conversion loss* of the demodulator; real-world demodulators will have larger conversion losses due to dissipation and other nonideal effects. An *absolute* phase measurement ϕ_m is then obtained via

$$\phi_m = \text{atan2} \left(Q, I \right), \tag{3.76}$$

where $\text{atan2} \left(Q, I \right)$ is the arctangent function of two arguments, which uses the signs of Q and I to correctly determine the quadrant corresponding to the tangent of Q/I. Note that in the ideal case the measured phase is equal to the true phase, i.e. $\phi_m = \phi$, and that ϕ_m is independent of the power in the incident IF signal.

3.5.2 *Nonideal mixing*

In Section 3.5.1, mixing was idealized as the multiplication of the IF signal by a sinusoidal LO signal. However, in practice, more complex processes are used to maximize the mixer's linear dynamic range and minimize its noise figure [21, 22].

For example, a typical ring-diode double-balanced mixer [21] is shown in Figure 3.5(b). The balun transformers T_1 and T_2 isolate the baseband port from the LO and IF ports. Typically, the LO power is $\sim$ 20 dB larger than the IF power. As a result, the LO alone is responsible for biasing the mixer's diodes into conduction. Note that the diodes are not all simultaneously biased into conduction. Instead, when $V_a > V_c$ (neglecting the diode drops), diodes D_1 and D_2 are forced into conduction such that V_b acts as a virtual ground for the IF signal coupled through transformer T_2. Then, when the LO changes sign such that $V_a < V_c$, diodes D_1 and D_2 stop conducting, and diodes D_3 and D_4 begin to conduct, forcing the virtual ground to jump from V_b to V_d. Thus, the ring diode acts as a switch for the polarization of the coupled IF signal, with the switching governed by the sign and frequency of the LO. Low-pass filtering the polarization-modulated IF, of course, yields the desired baseband signal. Note that the diode switching time should be minimized for optimal mod-

(a) analog I&Q demodulator

(b) double-balanced mixer
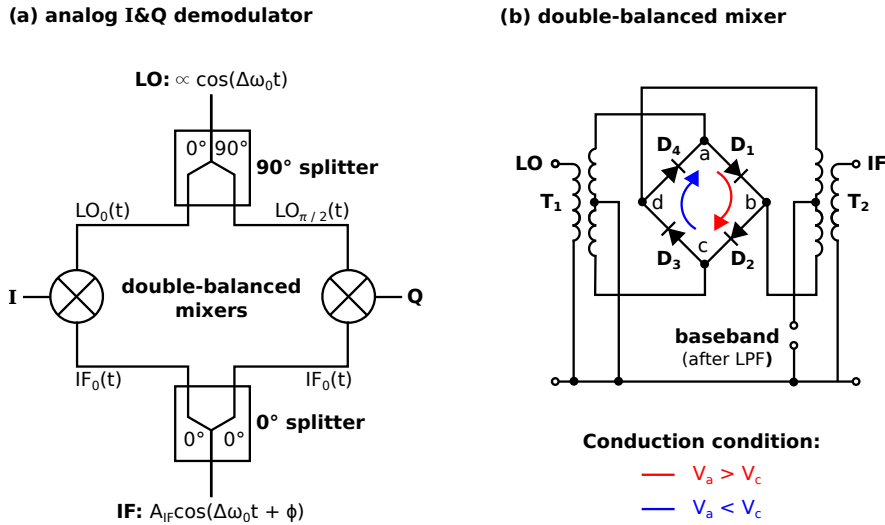
Figure 3.5: (a) A typical analog I&Q demodulator and (b) a typical diode-ring double-balanced mixer.

ulation, explaining why some manufacturers advocate the use of a square, rather than a sinusoidal, LO [23].

This polarization modulation can alter the baseband spectrum. To see this, note that the sign of the in-phase LO signal is simply a zero-mean square wave with even symmetry about the origin, as shown in the lower pane of Figure 3.6. The Fourier series of such a square wave consists of a sum over all of the LO's *odd* harmonics:

$$\text{sgn}\left(\text{LO}_0(t)\right) = \frac{4}{\pi} \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{2n-1} \cos[(2n-1)\Delta\omega_0 t] \tag{3.77}$$

$$= \frac{4}{\pi}\left[\cos(\Delta\omega_0 t) - \frac{1}{3}\cos(3\Delta\omega_0 t) + \cdots\right].$$

Now, if the IF signal is a perfect sinusoid as in (3.71), following Section 3.5.1's program of mixing the LO with the IF and low-pass filtering yields the desired in-phase baseband signal, e.g.

$$I = \left[\text{sgn}\left(\text{LO}_0(t)\right) \cdot \text{IF}_0(t)\right]\Big|_{|\omega| \ll \Delta\omega_0}$$

$$= \frac{\sqrt{2}A_{\text{IF}}}{\pi} \cos\phi \tag{3.78}$$

However, if the path from the mixer's IF port to its baseband port is not wholly linear (and every real-world device exhibits *some* degree of nonlinearity), the IF signal will contain contributions from its higher-
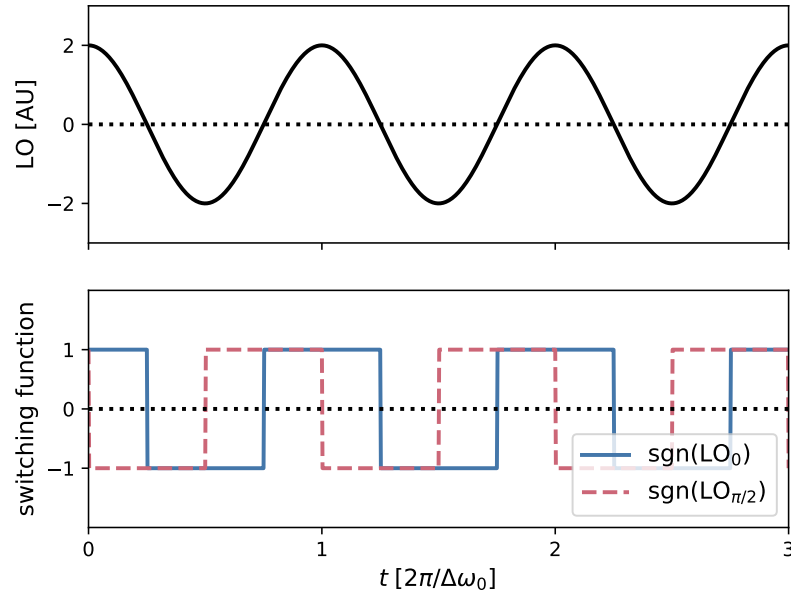
Figure 3.6: The LO switching functions. The top panel displays the incident sinusoidal LO signal, while the bottom panel shows the *sign* of the in-phase (LO_0) and the $\pi/2$ phase-advanced ($LO_{\pi/2}$) copies of the LO signals.

order harmonics. If this nonlinearity depends only on the magnitude of the IF (e.g. double-sided saturation/clipping of the signal), only the odd harmonics of the fundamental will contribute, i.e.

$$IF_0(t) = \frac{1}{\sqrt{2}} \sum_{n=1}^{\infty} A_{2n-1} \cos\left[(2n-1)(\Delta\omega_0 t + \phi)\right], \tag{3.79}$$

where $A_n/\sqrt{2}$ is the Fourier coefficient of the n^{th} harmonic (note that this unconventional normalization facilitates comparison with the previous, ideal definition of $IF_0(t)$ (3.72)). Typically, A_n decreases as n increases, but raising the IF amplitude drives more nonlinear interactions and increases the power in the higher order harmonics relative to the fundamental. Then, following Section 3.5.1's program of mixing the LO with the IF and low-pass filtering yields

$$I = \left[\text{sgn}\left(LO_0(t)\right) \cdot IF_0(t)\right]\Big|_{|\omega| \ll \Delta\omega_0}$$

$$= \frac{\sqrt{2}A_1}{\pi}\left[\cos\phi - \frac{1}{3}\left(\frac{A_3}{A_1}\right)\cos 3\phi + \cdots\right]. \tag{3.80}$$

That is, the higher order harmonics of the IF signal interact with the corresponding higher order harmonics of the LO switching function to produce harmonics in the baseband signal. The coefficient $A_n/(n \cdot A_1)$ gives the *suppression* of the n^{th} harmonic, and it is typically expressed in decibels referenced to the power of the carrier, or dBc, as

$$\text{suppression of } n^{th} \text{ harmonic [dBc]} = 20 \log_{10} \left(\frac{A_n}{n \cdot A_1} \right). \quad (3.81)$$

Noting that $\text{sgn}(\text{LO}_{\pi/2}(t))$ is an inverted, zero-mean square wave with *odd* symmetry about the origin, as shown in the lower pane of Figure 3.6, a similar path of reasoning to that used above shows that the quadrature baseband signal is

$$\begin{aligned} Q &= \left[\text{sgn}\left(\text{LO}_{\pi/2}(t) \right) \cdot \text{IF}_0(t) \right] \Big|_{\omega \ll \Delta\omega_0} \\ &= \frac{\sqrt{2}A_1}{\pi} \left[\sin\phi + \frac{1}{3} \left(\frac{A_3}{A_1} \right) \sin 3\phi + \cdots \right]. \end{aligned} \quad (3.82)$$

Additional distortion of the baseband signal results if the IF power becomes comparable the LO power, say within 10 dB, as the IF signal begins to contribute to the modulation of the diode conduction.

Finally, the diodes of the mixers should be matched as closely as possible. If, for example, diodes D_1 and D_2 have slightly different voltage drops than diodes D_3 and D_4, the virtual grounds at V_b and V_d are not equivalent when referenced to ground, and the resulting baseband signal will have a DC offset. With precision fabrication, however, it is not uncommon for the DC offset to be smaller than 1% of the amplitude of the baseband fundamental harmonic.

3.5.3 *Demodulator imbalances*

In addition to two double-balanced mixers, an analog I&Q demodulator also relies on a 90° splitter and a 0° splitter. Imbalances between any of these components can result in imbalances in the baseband I&Q signals. For example, unequal power division in the splitters produces amplitude imbalances in the baseband I&Q signals, while deviations from the nominal splitter phasings produces phase imbalances in the baseband I&Q signals. As discussed in Section 3.5.2, each double-balanced mixer can produce spectral distortion and DC offsets in the baseband signal; in addition, differences between the two mixers can exacerbate amplitude and phase imbalances in the baseband I&Q sig-

nals. Taken all together, then, the most general form for the baseband I&Q signals is

$$I = I_1 \left\{ \cos \phi - \frac{I_3}{3I_1} \cos (3\phi) + \cdots \right\} + \delta I, \tag{3.83}$$

$$Q = Q_1 \left\{ \sin (\phi + \delta) + \frac{Q_3}{3Q_1} \sin [3 (\phi + \delta)] + \cdots \right\} + \delta Q, \tag{3.84}$$

where I_1 is the amplitude of the in-phase signal's fundamental harmonic, I_3 is the amplitude of the in-phase signal's third harmonic, and δI is the in-phase signal's DC offset. Similar nomenclature applies to the quadrature signal. The phase imbalance of the demodulator is δ. Amplitude imbalances occur when $I_1 \neq Q_1$. The harmonic suppressions are typically comparable, e.g. $|I_3/I_1| \approx |Q_3/Q_1|$. Note that (3.83) and (3.84) generalize previous forms for the I&Q signals [24].

3.5.4 *Effects of demodulator imperfections*

Demodulator imperfections produce systematic errors in the measured phase [24, 25]. Specifically, if the I&Q signals suffer from imbalances and nonlinearities, as shown in (3.83) and (3.84), then the measured phase ϕ_m computed via the inverse tangent formula (3.76) will *not* correspond to the true phase ϕ, i.e.

$$\phi_m = \phi + \delta\phi, \tag{3.85}$$

where $\delta\phi$ is the error in the measured phase. The phase error is a complicated, periodic function of the true phase, i.e. $\delta\phi = \delta\phi(\phi)$ [24]. For small phase fluctuations $|\widetilde{\phi}| \ll 1$ about an equilibrium phase $\overline{\phi}$, the error $\delta\widetilde{\phi}$ in the measured fluctuating phase is simply the *change* in the total phase error between $\overline{\phi}$ and $\overline{\phi} + \widetilde{\phi}$ [25]:

$$\delta\widetilde{\phi} = \delta\phi(\overline{\phi} + \widetilde{\phi}) - \delta\phi(\overline{\phi}) \approx \left[\left. \frac{d(\delta\phi)}{d\phi} \right|_{\overline{\phi}} \right] \widetilde{\phi}. \tag{3.86}$$

Thus, the *relative* error in the measured fluctuating phase is

$$\frac{\delta\widetilde{\phi}}{\widetilde{\phi}} = \left. \frac{d(\delta\phi)}{d\phi} \right|_{\overline{\phi}}. \tag{3.87}$$

Synthetic I&Q Lissajous curves and the relative errors in the measured fluctuating phase that result from various demodulator imperfections are displayed in Figure 3.7. Obviously, each demodulator imperfection
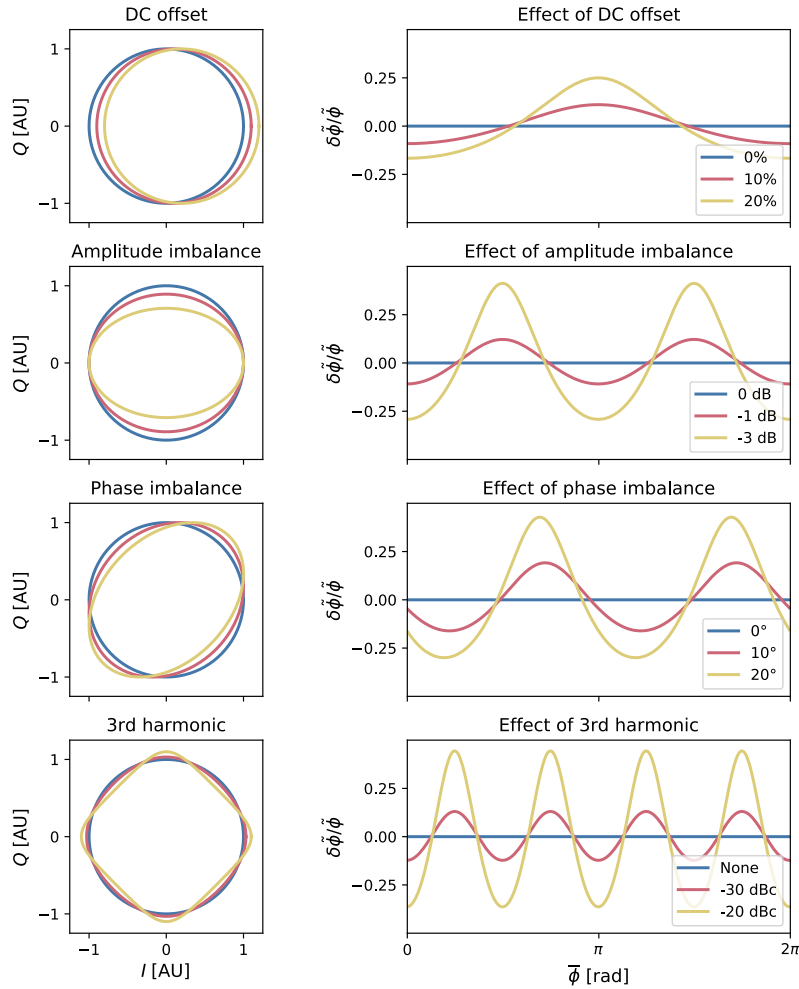
Figure 3.7: Demodulator imperfections produce errors in the measured phase. The left column displays the Lissajous curves that result from plotting synthetic quadrature Q vs. in-phase I signals, while the right column plots the relative error $\delta\widetilde{\phi}/\widetilde{\phi}$ in the measured fluctuating phase as a function of the equilibrium phase $\overline{\phi}$. Each row examines a different demodulator imperfection.

should be minimized in order to minimize the relative error in the measured fluctuating phase.

3.6 QUANTIZATION NOISE

Efficient conversion of an analog signal to a digital record requires quantization of the signal magnitude and temporal sampling of these quantized magnitudes [26]. This analog-to-digital conversion is performed by an instrument known as a digitizer. If a digitizer has bit

depth N_b and a peak-to-peak input-voltage range V_{in}, then its quantum of voltage ΔV is

$$\Delta V = \frac{V_{in}}{2^{N_b} - 1} \approx \frac{V_{in}}{2^{N_b}}, \tag{3.88}$$

where the approximation is well-satisfied for typical digitizer bit depths. At each sampling time, the analog signal's magnitude is approximated by the closest quantized value, whose separation from the true, analog value will be less than or equal to $\Delta V/2$.

In general, then, the quantized signal will differ from the analog signal. This error ϵ is known as quantization noise. The mean square error (i.e. variance) attributable to quantization is simply

$$\overline{\epsilon^2} = \frac{(\Delta V)^2}{12}, \tag{3.89}$$

where ΔV is the digitizer's quantum of voltage, as given by (3.88) [26] [9, Sec. 10.2.4]. For a uniformly sampled record with sampling rate f_s, sufficiently fine quantization ΔV ensures that the quantization noise is *white* [26, Th. 1] [27, Ch. 20]; that is, the one-sided autospectral density of the quantization noise is

$$G_{\epsilon,\epsilon}(f) = \frac{\overline{\epsilon^2}}{f_s/2} = \frac{(\Delta V)^2}{6f_s}, \qquad 0 \leqslant f \leqslant \frac{f_s}{2}. \tag{3.90}$$

In practice, however, aperture error, jitter, and nonlinearities may reduce the effective bit depth by one or two bits [9, Sec. 10.2.4], increasing the realized quantization noise relative to that expected from (3.89) and (3.90).

Quantization noise can be significant when attempting to measure absolute phase fluctuations with a heterodyne interferometer. Recall from the discussion of heterodyne interferometric detection in Section 2.4.2 that measurement of the *absolute* phase requires capturing the full dynamics of the large, slowly varying equilibrium phase $\overline{\phi}$ in addition to measuring the fluctuating phase $\widetilde{\phi}$. Because $\widetilde{\phi} \ll \overline{\phi}$ in typical situations, the fluctuations only occupy a small portion of the digitizer's input-voltage range; that is, fluctuations effectively see a bit depth that is substantially smaller than the digitizer's nominal bit depth. Thus, to minimize the effect of quantization noise, it is absolutely imperative to utilize the full input-voltage range of the digitizer.

The measured phase ϕ_m is computed from the in-phase (I) and quadrature (Q) signals via (3.76). If needed for feedback control, the phase can be calculated in real time with analog or digital electronics.

If the real-time phase is not needed, it is sufficient to digitize the I and Q signals individually; later, offline, the phase can be computed in software. In the offline approach, software may additionally compensate demodulator imperfections (discussed in Section 3.5) prior to computing the phase. Of course, the offline approach subjects both the I and the Q signal to quantization noise. The quantization error ε_I of I and the quantization error ε_Q of Q are uncorrelated, so the total one-sided autospectral density of the quantization noise is simply the sum of the individual one-sided autospectral densities, i.e.

$$G_{\epsilon_I,\epsilon_I}(f) + G_{\epsilon_Q,\epsilon_Q}(f) = \frac{(\Delta V)^2}{3 f_s}, \qquad 0 \leqslant f \leqslant \frac{f_s}{2}. \tag{3.91}$$

Note that comparing the spectral density of the quantization noise to the spectral density of the measured phase (with the phase having units of radians) requires converting (3.91) from units of $V^2 \cdot Hz^{-1}$ to the corresponding angular equivalent. This unit conversion is accomplished by dividing (3.91) by the *squared* radius of the I&Q Lissajous curve. Let the I and Q signals occupy a fraction $\eta_{dyn} \leqslant 1$ of the digitizer's full input-voltage range V_{in} such that the squared radius of the Lissajous curve is $I^2 + Q^2 = (\eta_{dyn} V_{in}/2)^2$, and the total quantization noise (in angular units of $rad^2 \cdot Hz^{-1}$) is

$$G_{\epsilon_I,\epsilon_I}(f) + G_{\epsilon_Q,\epsilon_Q}(f) \approx \frac{1}{3 \left[2^{2(N_b - 1)} \right] \eta_{dyn}^2 f_s}, \quad 0 \leqslant f \leqslant \frac{f_s}{2}, \tag{3.92}$$

where the definition of the voltage quantum (3.88) has been utilized.

3.7 SUMMARY

For ease of reference, this section provides a concise summary of the design considerations discussed in this chapter. If more details are desired regarding any given design consideration, the summary also points to the appropriate section in the text.

3.7.1 *Geometric considerations*

The measurements of a heterodyne interferometer are influenced by several geometric effects, which are summarized below:

- *Aperture diffraction* can alter the propagation of the unscattered and scattered Gaussian beams that formed the basis of the anal-

ysis in Chapter 2. Aperture diffraction of a Gaussian beam is minimal if

$$a_{\text{eff}} \geqslant \frac{3}{2} w(z),\tag{3.93}$$

for each aperture, where a_{eff} is the effective aperture radius and $w(z)$ is the beam's $1/e$ E radius at the aperture location. For a beam located transverse distance $\rho(z)$ away from the optical axis, a circular aperture of radius a has an effective aperture radius $a_{\text{eff}} = a - |\rho(z)|$. See Section 3.1.1 for more details.

- *Beam coalignment* affects the size of the interference signal. If the reference beam and the unscattered probe beam are misaligned by angle θ, then obtaining a finite interference signal from a detector element of length s_x requires that

$$|\theta| \ll \frac{\lambda_0}{s_x} \approx 0.6°,\tag{3.94}$$

where $\lambda_0 = 2\pi/k_0$ is the beam wavelength, and the typical values $\lambda_0 = 10.6\,\mu\text{m}$ and $s_x = 1\,\text{mm}$ have been used for the evaluation. See Section 3.1.2 for more details.

- *Differences in the radii of curvature* between the reference beam and the probe beam can alter the interference pattern at the detector, potentially decreasing the size of the interference signal and distorting the imaged wavenumbers. These detrimental curvature effects are negligible if

$$\max(\delta\phi_\kappa) = \frac{k_0}{8} \left[(N_{\text{el}}s_x)^2 + s_y^2\right] \left|\frac{1}{R_P(z_J)} - \frac{1}{R_R(z_R)}\right| \ll \pi.\tag{3.95}$$

where $R_P(z_J)$ is the radius of curvature of the probe beam at the detector location, $R_R(z_R)$ is the radius of curvature of the reference beam at the detector location, k_0 is the wavenumber of the probe radiation, s_x and s_y are the linear dimensions of a single detector element in the x- and y-dimensions, respectively, and N_{el} is the number of such detector elements (arranged in the x-direction with negligible inter-element spacing). See Section 3.1.3 for more details.

- The *finite sampling-volume effect* introduces a wavenumber dependence into the transfer function of a heterodyne interferometer. See Section 3.7.2 for a summary of this effect; see Section 3.1.4 for additional details.

- The *depth of focus* is the axial distance by which the detector location may deviate from the image plane. If δz_J is the axial displacement of the detector from the image plane (i.e. $z_{det} = z_J + \delta z_J$), a heterodyne interferometer will measure wavenumbers k_{meas} that are biased away from their true values k by

$$k_{meas} = \left[1 - \frac{\delta z_J}{R(z_{det})}\right] k,\tag{3.96}$$

where $R(z_{det})$ is the radius of curvature of the probe beam at the detector. Additionally, the "out-of-focus" interference of the upscattered and downscattered beams produces a wavenumber-dependent phase shift

$$\mu = \left(\frac{k^2}{2M^2 k_0}\right)\delta z_J,\tag{3.97}$$

where M is the magnification of the imaging system and k_0 is the wavenumber of the probe radiation. Autospectral densities of the phase fluctuations can be computed independent of μ, but it is less clear how to compute e.g. cross-spectral densities or bispectral densities in a μ-independent manner. An imaging system should be designed such that reasonable uncertainties in δz_J produce negligible wavenumber distortion (i.e. $|\delta z_J / R(z_{det})| \ll 1$) and negligible wavenumber-dependent phase shifts (i.e. $|\mu| \ll 1$). See Section 3.1.5 for more details.

3.7.2 *Wavenumber transfer function*

The heterodyne interferometer's basic wavenumber transfer function (derived in Chapter 2) is modified by several design decisions, becoming

$$T_{het}(k) = \frac{1}{\sqrt{2}\pi} \cdot \frac{2I_{AC}}{I_{DC} + I_{AC}} \cdot \frac{V_1}{V_1(I_{max} = I_{sat})} \cdot T_{fsv}(k),\tag{3.98}$$

where

- The first term $1/(\sqrt{2}\cdot\pi)$ corresponds to the heterodyne interferometer's basic wavenumber transfer function (2.77). See Section 2.4.2 for more details.

- The second term $2I_{AC}/(I_{DC} + I_{AC})$ specifies the ratio of peak-to-peak AC intensity to the maximum intensity in the heterodyne interference signal, where the DC intensity I_{DC} is defined

by (3.10) and the AC intensity I_{AC} is defined by (3.11). This modification to the transfer function is maximized when $I_{DC} = I_{AC}$, i.e. when there is full-depth modulation of the heterodyne intensity, which occurs when the probe beam and reference beam have identical spatial structures and powers. See Section 3.1.3 for more details.

- The third term $V_1/V_1(I_{max} = I_{sat})$ specifies the ratio of the realized amplitude (in volts) V_1 of the IF signal's fundamental harmonic relative to the corresponding amplitude $V_1(I_{max} = I_{sat})$ when the maximum incident intensity I_{max} is scaled to the detector's saturation intensity I_{sat} (the AC and DC fractions of the incident intensity are *not* altered during this scaling; to account for changing the AC and DC fractions, see the transfer function's second term, $2I_{AC}/(I_{DC} + I_{AC})$). Thus, operation above the saturation intensity *may* increase the heterodyne interferometer's wavenumber transfer function (i.e. if $V_1 > V_1(I_{max} = I_{sat})$, where V_1 depends sensitively on the detector's saturation physics). Similarly, operating below the saturation intensity will decrease the heterodyne interferometer's wavenumber transfer function. (i.e. $V_1 < V_1(I_{max} = I_{sat})$). See Section 3.2 for more details.

- The fourth and final term $T_{fsv}(k)$ quantifies the effect of sampling the interfering radiation field with detector elements of *finite* size. The output of each detector element corresponds to the incident intensity *averaged* over the element's active area, and this averaging acts as a low-pass filter in the spatial domain. This is referred to as the *finite sampling-volume effect*, and a square detector of linear dimension s_x placed in the image plane of a magnification M imaging system has the corresponding transfer function

$$T_{fsv}(k) = \text{sinc}\left(\frac{k}{k_{fsv}}\right), \tag{3.99}$$

where

$$\text{sinc}(x) = \frac{\sin(\pi x)}{\pi x} \tag{3.100}$$

is the normalized sinc function, and

$$k_{fsv} = \frac{2\pi |M|}{s_x} \tag{3.101}$$

is the first zero of $T_{fsv}(k)$. Wavenumbers $|k| \lesssim k_{fsv}$ are measurable, while wavenumbers $|k| \gtrsim k_{fsv}$ are not measurable. See Section 3.1.4 for more details. The finite sampling-volume effect is exploited in Chapter 4 to provide an overlap in the wavenumber sensitivities of DIII-D's pre-existing phase contrast imaging (PCI) system and its newly installed heterodyne interferometer.

3.7.3 *The measured phase*

An absolute phase measurement ϕ_m is computed from a heterodyne interferometer's demodulated in-phase (I) and quadrature (Q) signals via

$$\phi_m = \text{atan2}\,(Q, I), \tag{3.102}$$

where $\text{atan2}\,(Q, I)$ is the arctangent function of two arguments, which uses the signs of Q and I to correctly determine the quadrant corresponding to the tangent of Q/I. Demodulator imperfections produce a systematic error $\delta\phi$ in the measured phase, i.e.

$$\phi_m = \phi + \delta\phi, \tag{3.103}$$

where ϕ is the true phase. This systematic error is a complicated, periodic function of the true phase, i.e. $\delta\phi = \delta\phi(\phi)$. For small phase fluctuations $|\widetilde{\phi}| \ll 1$ about an equilibrium phase $\overline{\phi}$, the *relative* error in the measured fluctuating phase is

$$\frac{\widetilde{\delta\phi}}{\widetilde{\phi}} = \left.\frac{d(\delta\phi)}{d\phi}\right|_{\overline{\phi}}. \tag{3.104}$$

See Section 3.5 for more details.

3.7.4 *Noise sources & their spectral densities*

The phase measurements of a heterodyne interferometer can be corrupted by oscillator phase noise, optical shot noise, electrical noise, and quantization noise. The origins and the spectral densities of these noise sources are summarized below:

- *Laser phase noise* $\mathcal{L}_{\omega_0}(f)$ injects noise into a heterodyne interferometer's measurements. The one-sided autospectral density of this noise is

$$G(f) = 8\sin^2(\pi f\tau)\mathcal{L}_{\omega_0}(f), \qquad f \geqslant 0, \tag{3.105}$$

where $\tau = L/c$ is the time delay associated with the optical path-length difference L between the interferometer's probe arm and reference arm. See Section 3.3.1 for more details.

- *Local oscillator (LO) phase noise* $\mathcal{L}_{\Delta\omega_0}(f)$ injects noise into a heterodyne interferometer's measurements. The one-sided autospectral density of this noise is

$$G(f) = 8\sin^2(\pi f\tau)\mathcal{L}_{\Delta\omega_0}(f), \qquad f \geqslant 0, \tag{3.106}$$

where τ is the time required to couple the LO signal to the interferometer's Doppler-shifted beam. See Section 3.3.2 for more details.

- *Detector noise* near the intermediate frequency Δf_0 is demodulated and contaminates a heterodyne interferometer's measurements. The one-sided autospectral density of the demodulated detector noise (in angular units of $rad^2 \cdot Hz^{-1}$) is

$$G(f) = \frac{A}{P_P P_R (D^*)^2}, \qquad 0 \leqslant f \ll \Delta f_0, \tag{3.107}$$

where A is the area of the detector element, P_P is the optical power of the probe beam impinging on the detector element, P_R is the optical power of the reference beam impinging on the detector element, and D^* is the specific detectivity of the detector near Δf_0. Here, it is assumed that the probe beam and the reference beam vary weakly over the face of the detector element. See Section 3.4.1 for more details.

- *Optical shot noise* is demodulated and contaminates a heterodyne interferometer's measurements. The one-sided autospectral density of the demodulated optical shot noise (in angular units of $rad^2 \cdot Hz^{-1}$) is

$$G(f) = 2hf_0 \left(\frac{P_P + P_R}{P_P P_R}\right), \qquad 0 \leqslant f \ll \Delta f_0, \tag{3.108}$$

where h is the Planck constant, f_0 is the frequency of the incident photons, P_P is the optical power of the probe beam impinging on the detector element, P_R is the optical power of the reference beam impinging on the detector element, and Δf_0 is the intermediate frequency of the heterodyne signal. Here, it is assumed that the probe beam and the reference beam vary weakly over

the face of the detector element. See Section 3.4.2 for more de-
tails.

- *Amplifier noise* is characterized via a noise figure (NF), which
  quantifies the degradation in signal-to-noise ratio (SNR) after
  passing through the amplifier. See Section 3.4.3 for more details.

- *Quantization noise* is produced by the creation of a digital record.
  Assuming that the I and Q signals of a heterodyne interferom-
  eter are individually digitized with bit depth N_b and sample
  rate f_s (in Hz), the one-sided autospectral density of the total
  quantization noise (in angular units of $rad^2 \cdot Hz^{-1}$) is

$$G(f) \approx \frac{1}{3\left[2^{2(N_b - 1)}\right] \eta_{dyn}^2 f_s}, \qquad 0 \leqslant f \leqslant \frac{f_s}{2}, \qquad (3.109)$$

where $\eta_{dyn} \leqslant 1$ is the fraction of the digitizers's peak-to-peak
input-voltage range V_{in} occupied by the I and Q signals, i.e.

$$\eta_{dyn} = \frac{\max(I) - \min(I)}{V_{in}} = \frac{\max(Q) - \min(Q)}{V_{in}}. \qquad (3.110)$$

See Section 3.6 for more details.

[1] J. P. Campbell and L. G. DeShazer. Near fields of truncated-Gaussian apertures. *J. Opt. Soc. Am.*, 59(11):1427–1429, 1969.

[2] J. C. Rost. Personal communication, Mar 2014.

[3] T. N. Carlstrom, D. R. Ahlgren and J. Crosbie. Real-time, vibration-compensated CO_2 interferometer operation on the DIII-D tokamak. *Rev. Sci. Instrum.*, 59(7):1063–1066, 1988.

[4] M. A. Van Zeeland, G. J. Kramer, R. Nazikian, H. L. Berk, T. N. Carlstrom et al. Alfvén eigenmode observations on DIII-D via two-colour CO_2 interferometry. *Plasma Phys. Control. Fusion*, 47(9):L31, 2005.

[5] R. V. Bravenec and A. J. Wootton. Effects of limited spatial resolution on fluctuation measurements (invited). *Rev. Sci. Instrum.*, 66(1):802–805, 1995.

[6] E. M. Davis, J. C. Rost, M. Porkolab, A. Marinoni and M. A. Van Zeeland. A phase contrast imaging–interferometer system for detection of multiscale electron density fluctuations on DIII-D. *Rev. Sci. Instrum.*, 87(11):11E117, 2016.

[7] S. Coda. *An experimental study of turbulence by phase contrast imaging in the DIII-D tokamak*. PhD thesis, MIT, 1997.

[8] J. R. Dorris. *Phase contrast imaging measurements and modeling of short wavelength turbulence in the DIII-D tokamak*. PhD thesis, MIT, 2010.

[9] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 4th edition, 2010.

[10] Y. C. Kim and E. J. Powers. Digital bispectral analysis and its applications to nonlinear wave interactions. *IEEE Trans. Plasma Sci.*, 7(2):120–131, 1979.

[11] A. E. Siegman. *Lasers*. University Science Books, Sausalito, CA, USA, 1986.

[12] B. E. A. Saleh and M. C. Teich. *Fundamentals of Photonics*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 1st edition, 1991.

[13] M. Rakhmanov. Demodulation of intensity and shot noise in the optical heterodyne detection of laser interferometers for gravitational waves. *Appl. Opt.*, 40(36):6596–6605, 2001.

[14] Hamamatsu. *Characteristics and use of infrared detectors (technical information SD-12)*, Nov. 2004.

[15] R. C. Jones. Proposal of the detectivity D^{**} for detectors limited by radiation noise. *J. Opt. Soc. Am.*, 50(11):1058–1059, 1960.

[16] Mini-Circuits. *Amplifier terms defined (AN-60-038)*, Apr. 2015. Rev. B.

[17] M. A. Van Zeeland, R. L. Boivin, T. N. Carlstrom and T. M. Deterly. CO_2 laser polarimeter for Faraday rotation measurements in the DIII-D tokamak. *Rev. Sci. Instrum.*, 79(10):10E719, 2008.

[18] A. Mlynek, G. Pautasso, M. Maraschek, H. Eixenberger and the ASDEX Upgrade Team. Infrared interferometry with submicrosecond time resolution in massive gas injection experiments on ASDEX Upgrade. *Fusion Sci. Tech.*, 61(4):290–300, 2012.

[19] A. Mlynek, H. Faugel, H. Eixenberger, G. Pautasso and G. Sellmair. A simple and versatile phase detector for heterodyne interferometers. *Rev. Sci. Instrum.*, 88(2):023504, 2017.

[20] Mini-Circuits. *Modulators (AN-MOD11-2)*, Sept. 1999.

[21] Analog Devices. *Mixers and modulators tutorial (MT-080)*, Oct. 2008. Rev. 0.

[22] J. Bryant. Multipliers vs. modulators. *Analog Dialogue*, 47(06):1–2, 2013.

[23] Mini-Circuits. *Frequently asked questions about mixers (AN00-011)*, Apr. 2015. Rev. B.

[24] M. A. Van Zeeland and T. N. Carlstrom. Phase error correction method for a vibration compensated interferometer. *Rev. Sci. Instrum.*, 75(10):3423–3425, 2004.

[25] C. P. Kasten. Two-color interferometry as a fluctuation diagnostic on Alcator C-Mod. Master's thesis, MIT, 2013.

[26] W. R. Bennett. Spectra of quantized signals. *Bell Sys. Tech. J.*, 27(3):446–472, July 1948.

[27] B. Widrow and I. Kollár. *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*. Cambridge University Press, Cambridge, UK, 2008.

4

IMPLEMENTATION OF A COMBINED PCI-INTERFEROMETER ON DIII-D

The optical foundations in Chapter 2 revealed that heterodyne interferometry and PCI have complementary wavenumber capabilities, motivating the addition of a heterodyne interferometer to DIII-D's preexisting PCI system. With the intent of constructing a heterodyne interferometer, Chapter 3 extensively discussed the relevant design considerations. Heavily referencing both Chapters 2 and 3, this chapter details the physical implementation of the heterodyne interferometer and its integration into DIII-D's pre-existing PCI system.

Below, Section 4.1 briefly reviews optical-diagnostic access on DIII-D, and Section 4.2 discusses the relevant hardware and capabilities of the pre-existing PCI system. Section 4.3 then examines the heterodyne-interferometer optical layout, which sets the interferometer's wavenumber response and influences the interferometer's sensitivity to vibration-induced misalignment and small uncertainties in component placement. Next, Section 4.4 develops a simple model to investigate the distribution of the laser's finite optical power between the PCI and heterodyne-interferometer systems. While the PCI and heterodyne interferometer share several components, such as the CO_2 laser, the beam-delivery and beam-collection optics, and the digitizer, numerous other components are exclusively dedicated to the operation of the heterodyne interferometer; Section 4.5 details this dedicated hardware. Section 4.6 then describes the digital signal processing that is typically applied to the heterodyne interferometer's digital records prior to performing spectral analysis. Section 4.7 provides a detailed examination of the debilitating local-oscillator (LO) phase noise that initially plagued the heterodyne interferometer's measurements — identifying this noise source was singularly the most time-consuming and difficult aspect of the diagnostic development pursued in this work. Ultimately, this noise was eliminated with the procurement and installation of a low phase-noise LO. Section 4.7 also characterizes the predicted and measured noise in the heterodyne interferometer (after the LO upgrade) and compares the resulting noise floor to typical plasma-fluctuation spectra. Finally, Section 4.8 empirically verifies the predicted response of the heterodyne interferometer and cross-calibrates

the PCI against the absolute phase measurements of the heterodyne interferometer.

The robust response of the combined PCI-interferometer across the calibration wavenumber range demonstrates the complementary nature of the interferometer and PCI measurements and confirms the ability of the combined system to simultaneously monitor low-k and high-k instabilities.

4.1   OPTICAL-DIAGNOSTIC ACCESS ON DIII-D

DIII-D provides optical access to its plasmas through a number of ports, as indicated in Figure 4.1. The ports are labeled according to their toroidal positions and their sightlines, and an experimentalist should have at least a rough familiarity with these conventions. The toroidal location of a port is given in degrees clockwise from "machine north" when viewing the machine from above (note that machine north does *not* correspond to geographic or magnetic north). The angular separation of adjacent toroidal ports is 15°. Port sightlines can be vertical or radial. Ports with vertical (V) sightlines are labeled sequentially in terms of increasing major radius, with V1 having the smallest major radius and V3 having the largest major radius. Radial ports (R) have sightlines that are roughly aligned with the plasma's minor radius, and they are labeled according to their positions relative to the plasma midplane: R0 sits at the plasma midplane, R+1 and R+2 are the first and second ports *above* the plasma midplane, respectively, and R-1 and R-2 are the first and second ports *below* the plasma midplane, respectively.

4.2   DIII-D'S PRE-EXISTING PCI SYSTEM

The DIII-D PCI system is thoroughly described elsewhere [1, 2], but the system components of relevance to this work are briefly summarized below for completeness.

4.2.1   *CO_2 laser*

The PCI CO_2 laser [3, Sec. 3.3] has been in service since the inception of the DIII-D PCI project in the mid-1990s. Lasing occurs via high-voltage, DC excitation inside a 2 m-long, sealed-off glass tube to produce a TEM_{00} (Gaussian) mode with linear polarization. The beam waist occurs at the output coupler, and the corresponding 1/e

(a) **DIII-D top-down view**   (b) **DIII-D cross section**
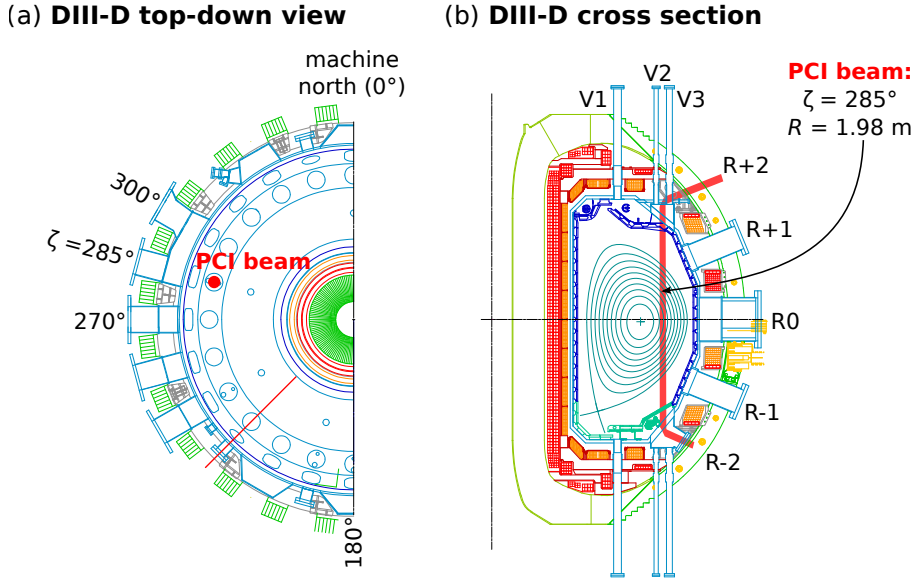
Figure 4.1: (a) View of DIII-D from above, indicating the toroidal-labeling convention. (b) View of DIII-D cross section, indicating the labeling convention for vertical (V) and radial (R) sightlines. The PCI beam enters the vessel through the 285° R+2 port, propagates vertically downwards through the plasma at a major radius of R = 1.98 m, and exits the vessel through the 285° R-2 port.

E radius is $w_0 = 1.25\,\text{mm}$. Historically, the laser power has been relatively constant at $\mathcal{P}_s = 14\,\text{W}$, but the power has declined in the last year. The manufacturer's specifications state that the laser has a "short-term" (0.1 s) peak-to-peak frequency variation $\lesssim 300\,\text{kHz}$ and a "long-term" (10^3 s) peak-to-peak frequency variation $\lesssim 3\,\text{MHz}$; unfortunately, without knowledge of the corresponding bandwidth, it is impossible to characterize the laser's phase noise. A ceiling on the laser phase noise is empirically established in Section 4.7.2.

4.2.2 *System geometry*

The system is currently configured in the "Phase II" geometry [1], with the probe beam propagating vertically downwards from the 285° R+2 port to the 285° R-2 port. The beam center sits at major radius R = 1.98 m. Both the toroidal and major radial positions of the PCI beam are shown in Figure 4.1. Because it is a line-integrated measurement, only fluctuations propagating perpendicular to the beam path are detected, as fluctuations propagating parallel to the beam path are effectively averaged out of the signal. Additionally, electrostatic drift-wave turbulence (e.g. ITG, ETG) tends to be field-aligned such that $k_\perp \gg k_\parallel$ [4, 5], where the $\perp$ and $\parallel$ subscripts are used here to

indicate orientations that are perpendicular to and parallel to the local magnetic field, respectively. To lowest order, then, electrostatic fluctuations propagate perpendicular to a tokamak's toroidal field. PCI's vertical beam path and the field-aligned constraint of electrostatic turbulence imply that PCI is sensitive to fluctuations with $k_Z = 0$ and finite major-radial wavenumber k_R. (For a more exacting discussion of the wavevectors that PCI is sensitive to, see Section 5.4.3). Thus, PCI's 32-element, 1-dimensional detector array is oriented in the image plane such that each detector element corresponds to a unique major radius in the plasma. In some situations, spatially filtering "masks" [1, 2, 6] or 2-dimensional detector arrays [7, 8] can be used to localize measurements by exploiting the spatial variation in the magnetic field's orientation along the beam path. These localization techniques typically work best for high-k measurements.

4.2.3 *Spatial bandwidth*

Several critical wavenumbers were defined in Section 2.5 that characterize the spatial bandwidth of a PCI system. The goal of this section is to evaluate each of these wavenumbers with the relevant parameters from the DIII-D PCI system.

PCI's low-k cutoff, k_g, is physically constrained by the free-space diffraction of the in-vessel probe beam. This constraint was derived both by examining the far-field overlap of the scattered and unscattered beams as in (2.80) and by matching the focal-plane spot size of the unscattered beam with the width of the phase-plate groove as in (2.88). Both approaches yield the identical constraint that $k_g^{min} = 2/w_0$, where w_0 is the $1/e$ E radius of the in-vessel probe beam. The maximum value of w_0 is set by the aperture diffraction criterion (3.93). After beam expansion and collimation, the 4"-diameter 285° R+2 port window is the smallest aperture seen by the beam prior to its interaction with the plasma. To satisfy the aperture diffraction criterion (3.93) with this 4"-diameter window, the expansion optics are configured to produce a collimated beam with $w_0 = 4/3" \approx 3.4\,\mathrm{cm}$, corresponding to

$$k_g^{min} \approx 0.6\,\mathrm{cm}^{-1}. \tag{4.1}$$

PCI's *realized* low-k cutoff, however, is typically $\sim 2-3\times$ larger than the diffraction-limited k_g^{min}. This occurs if the width of the phase-plate groove is oversized relative to the unscattered beam's focal-plane spot size. In such situations the realized low-k cutoff is given by

(2.86). Despite sacrificing some of the system's low-k range, operating with an oversized phase-groove width can be advantageous on large, vibration-prone fusion experiments. To see this, recall that the phase groove typically reflects only a fraction $\eta < 1$ of the incident unscattered beam power (the forward-facing surface of the DIII-D phase-plate groove is uncoated ZnSe, which has $\eta = 0.17$ at $10.6\,\mu m$); if vibration-induced misalignments push the unscattered beam out of the phase groove, there will be large power modulations on the PCI detector that are *not* attributable to plasma fluctuations and that push the detector beyond its saturation limits. While DIII-D's PCI system has an elaborate feedback control system that dynamically centers the unscattered beam on the phase-plate groove [3, Sec. 3.5], operating with an oversized phase-groove width gives the feedback system some leeway. As such, the groove width of the DIII-D phase plate is $d = 1\,mm$, and the probe radiation is focused onto the phase plate by an off-axis parabolic mirror of focal length $f = 80.7''$ such that the realized low-k cutoff (2.86) is

$$k_g \approx 1.5\,cm^{-1}, \tag{4.2}$$

approximately $2.5\times$ larger than the diffraction-limited minimum in (4.1). Now, for a $2\,T$ magnetic field and a $1\,keV$ temperature typical of DIII-D's pedestal, a deuteron has a gyroradius $\rho_i \approx 0.3\,cm$; assuming that the PCI beam and the local flux surface intersect at an angle $\alpha \sim 45°$, (4.2) corresponds to detection of fluctuations from the pedestal with $k_\theta \rho_i \gtrsim 0.25$. The higher-temperature core has larger ρ_i but also typically has smaller α, so the corresponding $k_\theta \rho_i$ cutoff in the core depends on the details of the temperature profile and the magnetic equilibrium.

PCI's high-k limits are dictated by finite collection-optic size and system magnification. The DIII-D phase plate has a diameter $D = 2''$ such that the phase plate's high-k cutoff (2.89) is

$$k_D \approx 75\,cm^{-1}. \tag{4.3}$$

Although the $5''$ diameter $285°$ R-2 exit window and the subsequent $12''$ diameter steering mirrors are large enough to accommodate beams scattered from fluctuations with wavenumbers $|k| \sim k_D$, apertures in the imaging optics only allow beams scattered from fluctuations with $|k| \lesssim 30\,cm^{-1}$ to reach the PCI detector unclipped. Upon reaching the detector, the measured PCI signal is subject to finite sampling-volume effects, which result from spatial averaging over the face of a detector element. The PCI's 32-element, 1-dimensional detector array

has elements of height $1\,\mathrm{mm}$ and width $s_x^{\mathrm{pci}} = 0.5\,\mathrm{mm}$. Coupled with the system's magnification $|M^{\mathrm{pci}}| = 0.45$, the finite sampling-volume cutoff (3.24) of the PCI is

$$k_{\mathrm{fsv}}^{\mathrm{pci}} \approx 55\,\mathrm{cm}^{-1}. \qquad (4.4)$$

Further, the center-to-center element spacing of adjacent detector elements is $\Delta x = 0.55\,\mathrm{mm}$ such that the corresponding Nyquist wavenumber is

$$k_{\mathrm{Ny}} = \frac{\pi |M^{\mathrm{pci}}|}{\Delta x} \approx 25\,\mathrm{cm}^{-1}; \qquad (4.5)$$

the measured wavenumber is aliased if $|k| > k_{\mathrm{Ny}}$.

4.2.4 *Temporal bandwidth*

PCI's detector (MCT-16-32, Infrared Associates; Stuart, FL USA) and its associated preamplifiers (also through Infrared Associates) are the dominant constraint on the system's temporal bandwidth. The HgCdTe detector array operates in the photoconductive regime and is cooled by liquid nitrogen; the liquid-nitrogen cooling reduces noise and boosts the response such that the detector-preamplifier combination has a specific detectivity $D^* \approx 2 \times 10^{10}\,\mathrm{cm}\sqrt{\mathrm{Hz}}/\mathrm{W}$ and a $500\,\mathrm{V} \cdot \mathrm{W}^{-1}$ responsivity to incident $10.6\,\mu\mathrm{m}$ light (note that the detector also has a saturation intensity $I_{\mathrm{sat}} \sim 1\,\mathrm{mW} \cdot \mathrm{mm}^{-2}$). The benefits of cooling, however, come at the expense of reduced bandwidth: the detector-preamplifier combination has a high-frequency, 2-pole cutoff at $\sim 800\,\mathrm{kHz}$ [9]. As the DC PCI signal is of little interest, the detector-preamplifier combination also has a low-frequency, 1-pole cutoff at $\sim 2\,\mathrm{kHz}$ [9].

The components downstream of the detector and preamplifiers have a small impact on the system bandwidth, but they are briefly summarized below for completeness. The Variable Gain and Filter (VGAF) circuits [2, Sec. 3.3.3] are located immediately downstream of the preamplifiers and have a low-frequency, 1-pole cutoff that can be easily switched between $10\,\mathrm{kHz}$ and $100\,\mathrm{kHz}$; the VGAFs are typically operated in the $10\,\mathrm{kHz}$ configuration. Following the VGAFs, fiber optic links (732 T/R-2.5-33k, Analog Modules; Longwood, FL USA) with an analog bandwidth of DC to $10\,\mathrm{MHz}$ transmit the signal from the DIII-D pit to the annex for digitization.

4.2.5 *Digitizer*

Upon reaching the annex, the PCI signals are digitized with two ACQ216CPCI boards (D-tAcq Solutions Ltd.; Glasgow, Scotland UK). The digitizer boards have bit depth $N_b = 14$, 8 V peak-to-peak input voltage range, and input impedance $Z_{in} = 50\,\Omega$. The digitizer sampling rate is typically $f_s = 4$ mega-samples/second (MSPS).

4.3 OPTICAL LAYOUT OF HETERODYNE INTERFEROMETER

The optical layout of a heterodyne interferometer sets the interferometer's wavenumber response, governs the injection of laser phase noise into the interferometer's measurements, and influences the interferometer's sensitivity to vibration-induced misalignment and small uncertainties in component placement. Having reviewed the spatial bandwidth of the pre-existing PCI system in Section 4.2.3, a complementary spatial bandwidth is selected for the heterodyne interferometer in Section 4.3.1, and the resulting imaging-system requirements are discussed in 4.3.2. Section 4.3.3 demonstrates that a carefully designed imaging system and the PCI focal-plane feedback system can synergistically act to dynamically maintain the coalignment of the interferometer beams in the presence of machine vibrations. Finally, Sections 4.3.4 and 4.3.5 discuss the generation of the interferometer's reference beam and probe beam and examine the sensitivity of the resulting interference signal to small uncertainties in the placement of the imaging optics.

4.3.1 *Desired spatial bandwidth*

Because of its external reference beam, a heterodyne interferometer has a minimum detectable wavenumber of $k = 0$. This ability to detect fluctuations below the PCI low-k cutoff (4.2) was one of the primary motivations for the addition of a heterodyne interferometer to DIII-D's pre-existing PCI system. Assuming negligible aperture diffraction (criterion (3.93)), finite sampling-volume effects govern the wavenumber dependence of a heterodyne interferometer's transfer function (3.98), producing the finite sampling-volume cutoff (3.101). Through careful selection of both the linear size s_x of the detector element and the magnification M of the imaging system, the finite sampling-volume cutoff can be engineered to provide any desired degree of overlap with the PCI wavenumber domain. Producing a reasonable degree of overlap with the PCI and ensuring that the finite sampling-volume
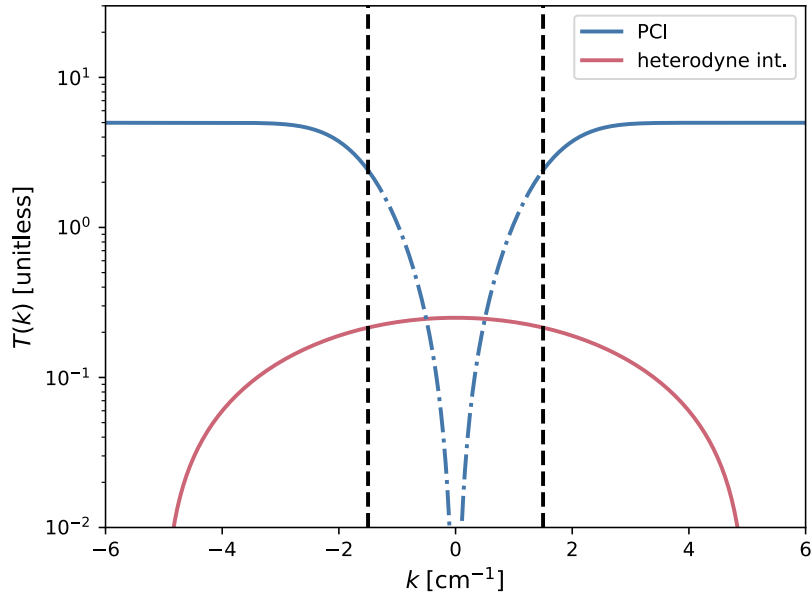
Figure 4.2: The design-point wavenumber transfer function of the heterodyne interferometer relative to that of the pre-existing PCI system. The heterodyne interferometer can detect fluctuations from $k = 0$ up to its finite sampling-volume cutoff (4.6). The vertical, dashed lines indicate the low-k cutoff of the PCI phase-plate groove, k_g from (4.2); the phase-plate high-k cutoff (4.3) and the PCI finite sampling-volume cutoff (4.4) are sufficiently large to negligibly affect the PCI transfer function over the depicted wavenumber domain. The reflectivity of the PCI phase groove is $\eta = 0.17$. Because the PCI transfer function is not defined for $|k| \lesssim k_g$, the low-k PCI amplitude response (2.99) is indicated by the dash-dot curves instead. The depicted transfer functions should be compared to those without finite sampling-volume effects in Figure 2.6.

effect (rather than aperture diffraction) governs the high-k response leads naturally to the design point for the heterodyne interferometer's finite sampling-volume cutoff

$$k_{fsv}^{het} = 5 \, cm^{-1}. \tag{4.6}$$

Figure 4.2 displays the desired wavenumber transfer function of the heterodyne interferometer relative to that of the pre-existing PCI system.

4.3.2 *Detector-element size*

The heterodyne interferometer's design-point finite sampling-volume cutoff (4.6) fixes the ratio of the imaging system's (absolute) magnifi-

cation $|M|$ to the detector element's linear size s_x. Commercially available detector elements for use at $10.6\,\mu m$ range in linear size from fractions of a millimeter to several millimeters [10]. Smaller linear element size s_x eases both the beam-coalignment constraint (3.94) and the curvature-matching constraint (3.95). The variation of detector noise and optical shot noise with s_x (but fixed $|M|/s_x$) can also be examined. To proceed, assume that the interferometer image plane sits in the corresponding Rayleigh range (i.e. $|z_\mathcal{J}| \ll z_R$) such that $w(z_\mathcal{J}) \approx |M|w_\mathcal{O}$, where $w(z_\mathcal{J})$ is the $1/e$ E radius of the probe beam at the image plane, and $w_\mathcal{O}$ is the $1/e$ E radius of the probe beam at the object plane. For any given beam power, the probe-beam optical power P_P impinging on the detector element is a function of the element size relative to the beam size, i.e. $P_P = P_P(s_x/w(z_\mathcal{J})) \approx P_P(s_x/(|M|w_\mathcal{O}))$; because $|M|/s_x$ is fixed by the design-point finite sampling-volume cutoff (4.6), P_P is also approximately fixed. The optical power in the reference beam impinging on the detector element, however, is independent of the magnification and proportional to the element area A, i.e. $P_R \propto A$. Thus, the one-sided autospectral density of the demodulated detector noise (3.107) is *independent* of the element size. The one-sided autospectral density of the demodulated optical shot noise (3.108) is, in general, a function of the element size; in the extreme $P_R \gg P_P$, the spectral density is *independent* of the element size; in the opposite extreme $P_R \ll P_P$, the spectral density varies inversely with element area A. In light of the above considerations, a square detector element of intermediate linear size

$$s_x = 1\,mm \tag{4.7}$$

was chosen. Referencing the definition of the finite sampling-volume cutoff (3.101) and the interferometer's corresponding design-point finite sampling-volume cutoff (4.6), it naturally follows that the interferometer imaging system must have magnification

$$|M| = 0.08. \tag{4.8}$$

4.3.3 *Beam coalignment in presence of machine vibrations*

DIII-D's PCI system has an elaborate feedback control system that dynamically centers the unscattered beam on the phase-plate groove [3, Sec. 3.5]. Coda has extensively characterized the effect of this feedback system on the lateral position of the PCI image [3, Sec. 3.5(f)]. This section extends Coda's analysis by examining the effect of the feedback

system on the misalignment angle θ between the unscattered probe beam and the image-plane optical axis. It readily follows from this discussion that, for certain imaging systems, the PCI feedback system can dynamically maintain the beam coalignment (constraint (3.94)) of a heterodyne interferometer in the presence of machine vibrations.

First, consider the situation without feedback. Let the probe beam propagate a distance l from the plasma midplane to a focusing optic with focal length f, and let the focused beam come to a waist a distance l' from the focusing optic. Now, let a single mirror, located a distance d upstream from the focusing optic, be tilted away from its nominal position (e.g. from machine vibrations), deflecting the unscattered beam by angle β_1 relative to the nominal optical axis. The resulting focal-plane position and orientation of this beam are incorrect (i.e. the beam does not lie on the optical axis, but it should). This misalignment is depicted by the solid line in Figure 4.3(a). In a perfectly aligned system, the focal-plane position and orientation of this beam would correspond to the "effective" beam path depicted by the dashed line in Figure 4.3(a). The imaging optics image the object-plane ray corresponding to this effective beam path.

Imaging systems are discussed in Appendix B, but the relevant details are briefly summarized here. The symmetry axis of a Gaussian beam behaves as a ray in the geometric-optics sense, where a ray is fully described by its transverse distance ρ to the optical axis and its angular orientation θ relative to the optical axis. Ray propagation through a magnification-M imaging system is well-characterized by an ABCD ray matrix of the form (B.3). Specifically,

$$\rho_{\mathcal{I}} = M\rho_{\mathcal{O}}, \tag{4.9}$$

$$\theta_{\mathcal{I}} = \frac{\theta_{\mathcal{O}}}{M} + C\rho_{\mathcal{O}}, \tag{4.10}$$

where $\mathcal{O}$ indicates object-plane quantities, $\mathcal{I}$ indicates image-plane quantities, and C is constant determined by the particulars of the imaging system. In a perfectly aligned system, the unscattered beam's object-plane ray ($\rho_{\mathcal{O}} = 0$ and $\theta_{\mathcal{O}} = 0$) is correctly imaged as $\rho_{\mathcal{I}} = 0$ and $\theta_{\mathcal{I}} = 0$. However, in a *misaligned* system, the unscattered beam's *effective* object-plane ray ($\rho_{\mathcal{O}} \neq 0$ and $\theta_{\mathcal{O}} \neq 0$) is imaged as $\rho_{\mathcal{I}} \neq 0$ and $\theta_{\mathcal{I}} \neq 0$. Thus, without feedback, a tilted mirror produces errors in both the image-plane position and orientation of the unscattered beam.

Now, add a steering mirror a distance d' downstream of the focusing optic. This steering mirror is tasked with compensating for the tilted mirror by returning the unscattered beam's focal-plane position to the nominal optical axis, as depicted in Figure 4.3(b). In the
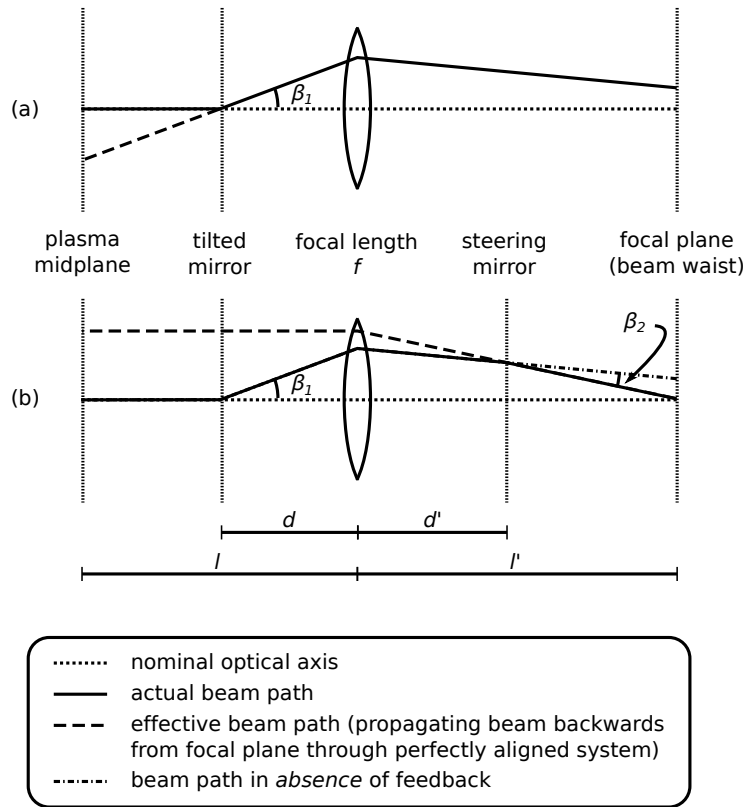
Figure 4.3: Effect of a tilted mirror (a) without focal-plane feedback and (b) with focal-plane feedback on the actual and effective paths of a Gaussian beam. An imaging system images the object-plane ray corresponding to the effective beam path. Without focal-plane feedback, both the transverse distance and the angular orientation of the effective object-plane ray are incorrect, producing errors in the corresponding image-plane quantities. With focal-plane feedback, however, the angular orientation of the effective object-plane ray is corrected (i.e. the ray is parallel to the optical axis, albeit laterally displaced); then, if the imaging system is engineered such that $C = 0$ in the $ABCD$ ray matrix, the feedback will dynamically maintain the correct angular orientation of the corresponding image-plane ray.

geometric-optics limit, a ray passing through the intersection of the optical axis and the focal plane corresponds to a collimated beam ($\theta = 0$) upstream of the focusing optic. Thus, in a perfectly aligned system, the focal-plane position and orientation of the feedback-compensated beam would correspond to the "effective" beam path, depicted by the dashed line in Figure 4.3(b). Thus, upstream of the focusing optic, the effective beam path is displaced from but *parallel* to the nominal optical axis. As is the case without feedback, the imaging optics image the object-plane ray ($\rho_{\mathcal{O}} \neq 0$, $\theta_{\mathcal{O}} = 0$) corresponding to this effective beam path, i.e. $\rho_{\mathcal{I}} = M\rho_{\mathcal{O}}$ and $\theta_{\mathcal{I}} = C\rho_{\mathcal{O}}$. Thus, for imaging sys-

tems with $C = 0$, the PCI feedback system will dynamically maintain the correct angular orientation of the image-plane unscattered beam (i.e. $\theta_J = 0$); an obvious application is satisfying the heterodyne interferometer's coalignment constraint (3.94) in the presence of machine vibrations. Of course, the lateral image location shifts in accordance with the discussion by Coda [3, Sec. 3.5(f)].

4.3.4 *Reference-beam generation*

A heterodyne interferometer interferes the imaged probe radiation with a frequency-shifted reference beam to make an absolute phase measurement, as discussed in Section 2.4.2. It is easy to Doppler shift 10.6 μm radiation by tens of MHz with an acousto-optic modulator (AOM). The operation of a typical AOM is described in in Section 3.3.2.

The heterodyne interferometer's reference beam is generated with a Gooch & Housego (Ilminster, UK) 37027-5 Germanium AOM, pictured in Figure 4.4. The resonant frequency of the AOM's piezo-actuator is 27.12 MHz, but the AOM's deflection efficiency varies negligibly between 25 MHz and 30 MHz. For reasons discussed in Section 4.7.1, the AOM is operated at $\Delta f_0 = 30$ MHz. At this frequency, the deflected beam is separated from the undeflected beam by $2\theta_B = 59$ mrad, where θ_B is the Bragg angle from (3.47). Deflection efficiency scales roughly linearly with RF power, and deflection efficiencies in excess of 75% can be obtained at the maximum-rated RF power of 30 W. The RF power is CW such that the AOM simply deflects and Doppler shifts a fraction of the incident optical beam. (In contrast, rapidly varying the RF power modulates the optical powers in the deflected and undeflected beams, but such modulation is *not* desirable in a heterodyne interferometer). The AOM's static optical insertion loss is measured to be ∼ 10%, and the incident optical intensity should be limited to $\leqslant 5\,W \cdot mm^{-2}$ to avoid thermal lensing of the beam. The diameter of the piezo-driven acoustic beam is ∼ 5 mm. Satisfying both the AOM's peak-intensity constraint and the aperture-diffraction constraint (3.93, with the acoustic beam providing an "effective aperture") requires placement of the AOM ∼ 15" downstream of the laser's beam waist.

Due to space constraints, the reference-beam path length is *not* matched to the probe-beam path length. This path-length discrepancy (∼ 10 m) injects the laser's phase noise into the heterodyne interferometer's measurements, as quantified by (3.105). However, as shown in Section 4.7.2, the laser's phase noise negligibly contributes to the heterodyne interferometer's noise floor.
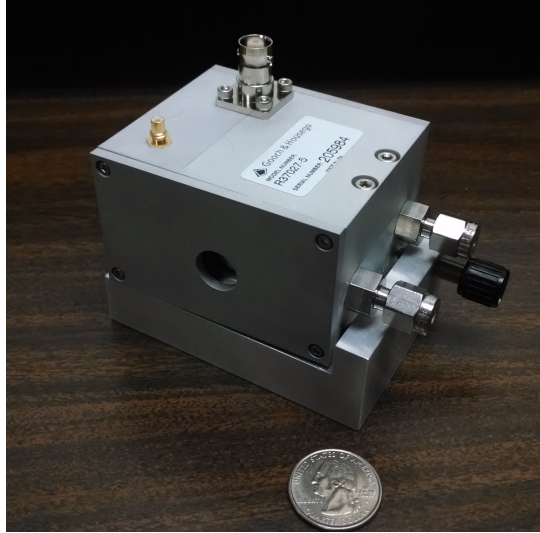
Figure 4.4: The heterodyne interferometer's Germanium acousto-optic modulator (AOM). The AOM deflects and Doppler shifts ($\Delta f_0 = 30\,\text{MHz}$) a fraction of the incident CO_2 beam ($f_0 = 28.3\,\text{THz}$) to produce a reference beam. The 30 MHz RF signal is coupled to the AOM through the upper-left BNC jack, and the resulting sound waves propagate through the Germanium crystal from left to right. The CO_2 beam enters through the aperture on the AOM's face and exits through a similar aperture on the AOM's rear. The twin Swagelock fittings on the AOM's right provide an inlet and outlet for water cooling, while the lower-left SMC jack connects to a normally closed, creep-action PEPI N thermostat that provides a thermal interlock to the AOM's RF driver. The AOM is mounted on a "Bragg mount", which allows easy optimization of the beam's angle of incidence (via the mount's adjustment screw, seen in black at the middle-right).

Allowing such a "modest" path-length discrepancy substantially simplifies the optical design of the reference arm. Following its generation at the AOM, the reference beam is directed to the interferometer detector with four broadband metallic mirrors (ER.2 protected-silver coated from Newport Corporation; Irvine, CA, USA) and is combined with the probe beam via a 50% reflective, polarization-independent ZnSe beam splitter (II-VI Infrared; Saxonburg, PA, USA). No lenses are used to condition the reference beam. However, a multiple-order $10.6\,\mu\text{m}$ half-waveplate (also from II-VI Infrared) sits between the AOM and the polarization-independent beam splitter. As the probe beam is *not* confined to a single plane, out-of-plane mirror reflections can rotate the probe-beam polarization; the half-waveplate allows one to easily align the polarization of the reference beam with that of the probe beam, thereby maximizing the interference signal. From source to detector, the reference beam propagates a total distance 59-3/8".
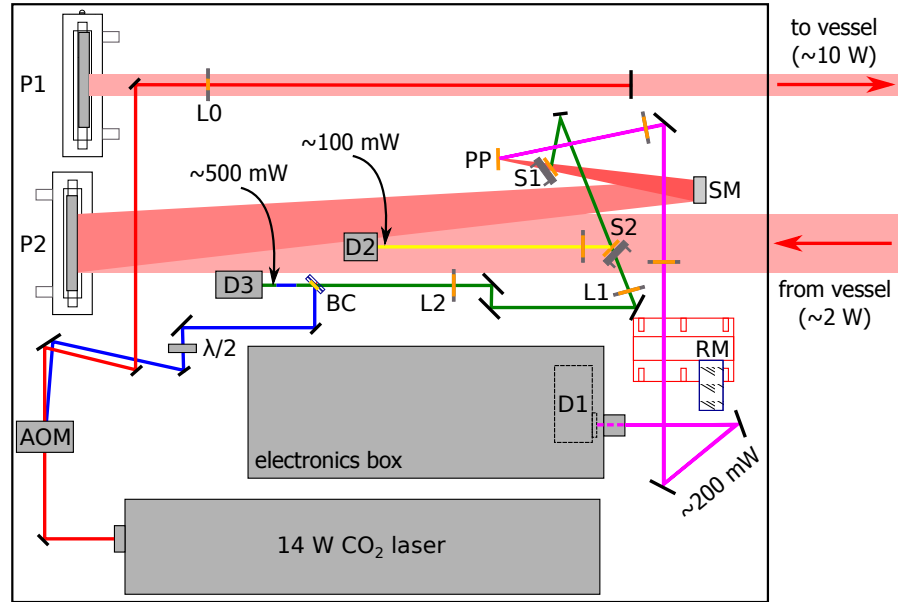
Figure 4.5: Combined PCI-interferometer optical table. Salmon: shared probe beam; green: interferometer plasma arm; blue: interferometer reference arm; fuchsia: PCI beam; yellow: feedback beam. AOM: Ge acousto-optic modulator. BC: beam combiner. D1: LN_2-cooled PCI detector array; D2: room-temperature quadrant detector for beam steering; D3: thermoelectrically-cooled interferometer detector. L0, L1, L2: ZnSe lenses. P1, P2: off-axis parabolic mirrors. PP: phase plate. RM: rotatable mask. S1, S2: ZnSe beam splitters. SM: steering mirrors. $\lambda/2$: half-wave plate.

4.3.5   *Probe-beam generation & imaging*

A schematic of the combined-PCI interferometer optical table is shown in Figure 4.5. The interferometer and PCI *share* the in-vessel probe beam. The only modification to the beam-generation optics was the addition of the AOM, discussed extensively in Section 4.3.4. The AOM insertion loss and the diversion of some of the laser power to the reference arm reduce the power in the probe beam relative to the PCI-only configuration.

As in the PCI system, the probe arm of the heterodyne interferometer is configured to *image* the probe radiation from the tokamak midplane. The heterodyne interferometer and PCI imaging optics share the focusing $f = 80.7''$ off-axis parabolic mirror and the feedback steering mirrors [3, Sec. 3.5]. Because the phase plate's spatial filtering produces the PCI's low-k cutoff, a 2'' diameter S-polarization ZnSe splitter (II-VI Infrared; Saxonburg, PA, USA) located 5-1/8'' upstream of the phase plate diverts a fraction of the probe radiation to dedicated heterodyne-interferometer optics. (The selection of the splitter reflectivity is discussed in Section 4.4.5. This beam splitter also diverts a

portion of the unscattered probe beam to the feedback system's quadrant detector; the heterodyne-interferometer probe beam and the feedback beam are separated with an additional 2" diameter S-polarization ZnSe splitter). Four broadband metallic mirrors (ER.2 protected-silver coated from Newport Corporation; Irvine, CA, USA) direct the probe radiation through the remainder of the heterodyne interferometer's imaging optics, which consist of two plano-convex ZnSe lenses (also from II-VI Infrared). The first lens L1 has a 7.5" focal length and a 2" diameter, and the second lens L2 has a 7.5" focal length and a 1.5" diameter. Then, under the constraint of imaging the plasma midplane, the remaining design parameters are the distance $d_{P2,L1}$ between the focusing off-axis parabolic mirror and L1 and the distance $d_{L1,L2}$ between L1 and L2.

The sensitivity of the heterodyne-interferometer optical layout to $d_{P2,L1}$ and $d_{L1,L2}$ is shown in Figure 4.6, and the resulting implications are discussed here in detail. The *top-left* panel of Figure 4.6 displays the distance from L2 to the image plane $\mathfrak{I}$. The image-plane location is a strong function of $d_{L1,L2}$ and a weak function of $d_{P2,L2}$. The *top-right* panel of Figure 4.6 displays the imaging system's magnification M, with the design point in accord with (4.8). The magnification is a strong function of $d_{L1,L2}$. The *middle-left* panel of Figure 4.6 displays the C parameter of the imaging system's ABCD ray matrix. As discussed in Section 4.3.3, the PCI feedback system will dynamically maintain the coalignment of the heterodyne interferometer's beams if C = 0. The C parameter is a strong function of both $d_{P2,L1}$ and $d_{L1,L2}$, and experimental uncertainties in distances and focal lengths make it difficult to enforce C = 0. (Indeed, the amplitude of the interference signal from the realized system *does* vary with vibrations, indicating that C $\neq$ 0 and/or the action of the feedback system is not sufficient to maintain coalignment). The *middle-right* panel of Figure 4.6 displays the maximum curvature-induced phase shift (3.95) for a single, $1\,\mathrm{mm}^2$ square detector. The design point satisfies $\max(\delta\phi_\kappa) \ll 1$ such that curvature-induced phase shifts minimally reduce the interference power and negligibly distort the imaged wavenumbers. The *bottom-left* panel of Figure 4.6 displays the axial distance z from the waist of the probe beam to the image plane (this is the usual axial parameter used to characterize Gaussian beams, with $z > 0$ indicating that the waist is upstream of the image plane and $z < 0$ indicating that the waist is downstream of the image plane). There is only a very narrow strip where the image plane and the beam waist coincide (i.e. where $z = 0$); note that this strip closely resembles the C = 0 curve of the middle-left panel. This is no coincidence: as discussed in the
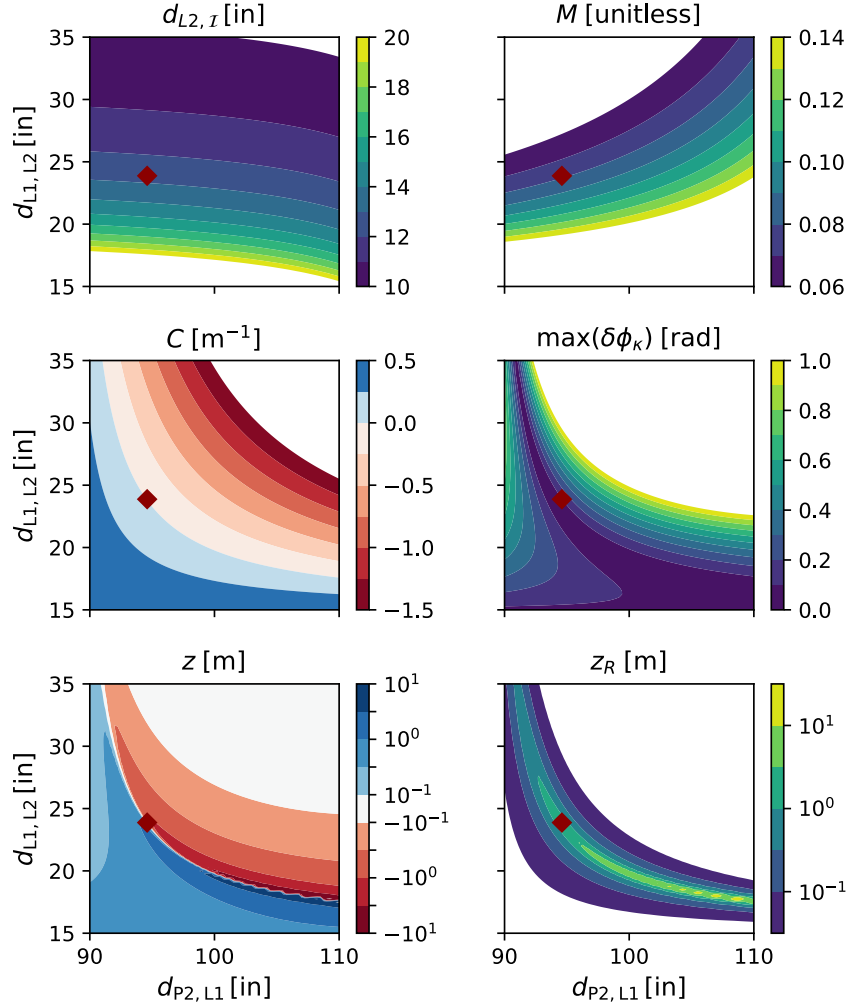
Figure 4.6: Sensitivity of heterodyne-interferometer optical design to the placement of lenses L1 and L2. The distance between the focusing off-axis parabolic mirror P2 and the first lens L1 of the imaging system is denoted by $d_{P2,L1}$ (along the x-axis), and the distance between L1 and the second lens L2 is denoted by $d_{L1,L2}$ (along the y-axis). The design point is marked with a burgundy diamond.

text surrounding (B.9), the image plane and the beam waist coincide only if $|C| \ll 1/|Mq_{\mathcal{O}}|$, where $q_{\mathcal{O}}$ is the object-plane complex beam parameter and M is the imaging-system magnification; the $10.6\,\mu m$ PCI probe beam, with in-vessel $1/e$ E radius $w_0 \approx 3.4\,cm$, has $q_{\mathcal{O}} \approx 340\,m$. Thus, for moderate magnifications, the large size of $q_{\mathcal{O}}$ couples the $C = 0$ and the $z = 0$ curves to each other. Finally, the *bottom right* panel of Figure 4.6 displays the Rayleigh range z_R of the image-plane probe beam. The discussion of detector noise and optical shot noise in
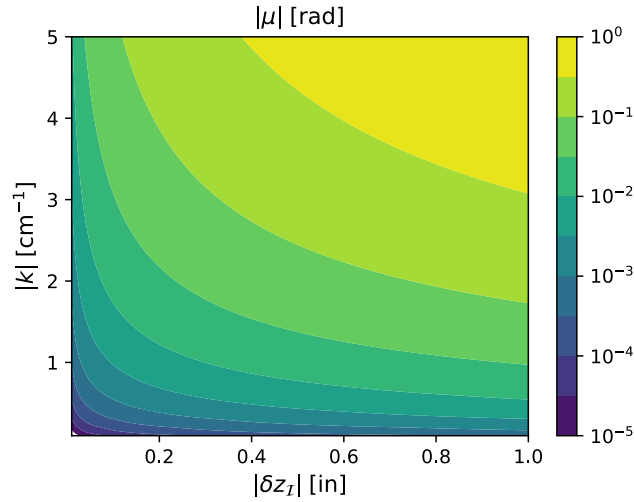
Figure 4.7: The wavenumber-dependent phase shift μ that occurs when the upscattered and downscattered $10.6\,\mu m$ beams interfere a distance $\delta z_\mathcal{J}$ away from the image plane of an $M = 0.08$ imaging system; $|\mu| \ll 1$ is desirable.

Section 4.3.2 assumed that the image plane sits in the Rayleigh range (i.e. $|z| \lesssim z_R$), and the design point satisfies this assumption. Further, the depth-of-focus criterion (3.96) states that there will be minimal wavenumber distortion if the detector is $|\delta z_\mathcal{J}| \ll R(z_{\text{det}})$ within the image plane. As a Gaussian beam has minimum radius of curvature $\min(R(z)) = 2z_R \leqslant R(z_{\text{det}})$, this depth-of-focus criterion simply reduces to $|\delta z_\mathcal{J}| \ll 2z_R$. Within the neighborhood of the design point, $z_R \geqslant 10\,\text{cm}$ such that wavenumber distortion is minimal as long as the detector is well within $\pm 20\,\text{cm} \approx \pm 8''$ of the image plane, which is relatively easy to satisfy experimentally.

The imaging system's depth of focus is additionally constrained by the "out-of-focus" interference of the upscattered and downscattered beams, which produces the wavenumber-dependent phase shift μ from (3.97). Because a heterodyne interferometer exhibits nulls in its response when $\mu = (2m+1)\pi/2$ for integer m, it is desirable to have $|\mu| \ll 1$ over the system's full wavenumber range. Figure 4.7 displays $|\mu|$ for $10.6\,\mu m$ probe radiation and the design-point magnification $M = 0.08$ from (4.8). Clearly, the power loss from this "out-of-focus" interference effect will be small for $|k| \lesssim 5\,\text{cm}^{-1}$ if the detector is within $\pm 0.5''$ of the image plane.

Axial profiles of the unscattered beam and a scattered beam in the design-point imaging system are shown in Figure 4.8. Because scattering from density fluctuations with larger wavenumbers produces
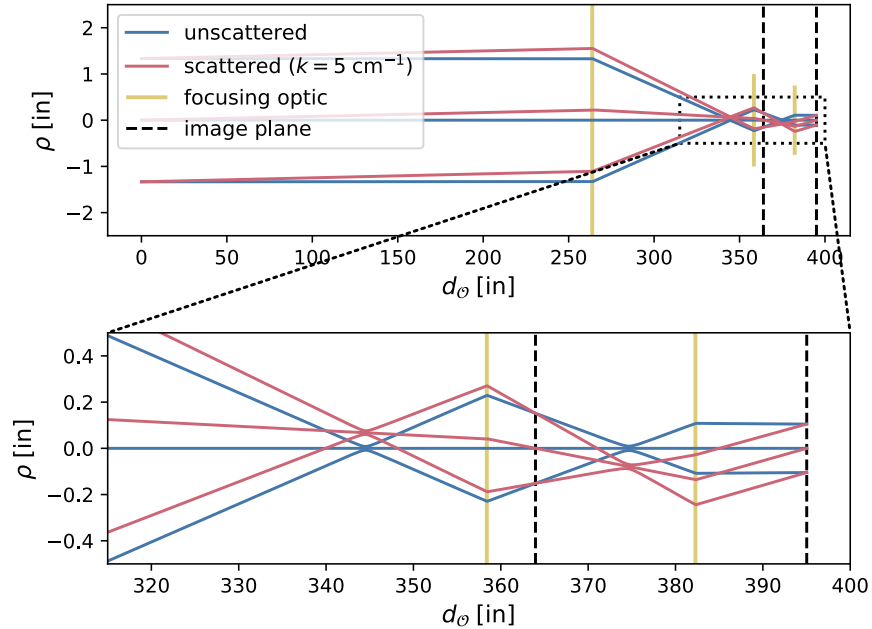
Figure 4.8: Axial beam profiles in interferometer probe arm as a function of distance d_O from the tokamak midplane (i.e. the object plane). For a given beam, the central ray corresponds to the symmetry axis of the beam, while the other two rays correspond to the $1/e$ E radius of the beam. The interferometer detector sits at the second image plane, approximately 395″ downstream of the tokamak midplane.

larger transverse deviations from the optical axis, the scattered beam plotted in Figure 4.8 corresponds to a $k = 5\,\mathrm{cm}^{-1}$ fluctuation, which is the maximum measurable wavenumber (4.6) of the interferometer. Note that the diameter of each focusing optic is accurately depicted (if an optic touches the horizontal axes, the optical diameter exceeds the plotted dimensions) and that the aperture diffraction criterion (3.93) is met for both the scattered and unscattered beams at each focusing optic. For simplicity, planar mirrors are not depicted, but they similarly satisfy the aperture diffraction criterion. The probe radiation is combined with the reference beam and interfered on a detector located at the second image plane, approximately 395″ downstream of the tokamak midplane.

4.4   DISTRIBUTION OF OPTICAL POWER

Having identified a suitable heterodyne-interferometer optical layout in Section 4.3, the finite power of the $10.6\,\mu\mathrm{m}$ laser must be distributed between the PCI and interferometer systems. Section 4.4.1 reviews the

experimentally relevant relations between a Gaussian beam's global properties (e.g. total beam power) and local properties (e.g. peak, on-axis intensity). Sections 4.4.2 through 4.4.4 develop approximate expressions for the total power in the beam(s) at each detector location; significantly, two design parameters, η_R and η_P, control the power distribution in the combined PCI-interferometer. Finally, Section 4.4.5 identifies a suitable operational point in (η_R, η_P)-space.

4.4.1 *Gaussian-beam intensity & power*

The Gaussian-beam electric field (2.33) produces the optical intensity (averaged over an optical cycle)

$$I(\rho, z) = \frac{c\varepsilon_0}{2}|E_G(r)|^2 = I(0, z)\exp\left[\frac{-2\rho^2}{w(z)^2}\right], \tag{4.11}$$

where

$$I(0, z) = \frac{c\varepsilon_0}{2}\left[\frac{E_0 w_0}{w(z)}\right]^2 \tag{4.12}$$

is the peak, on-axis intensity of the Gaussian beam at axial distance z from the beam waist. The optical power within ρ_0 of the symmetry axis is

$$P(\rho \leqslant \rho_0) = \mathcal{P}\left\{1 - \exp\left[\frac{-2\rho_0^2}{w(z)^2}\right]\right\}, \tag{4.13}$$

where

$$\mathcal{P} = \frac{\pi[w(z)]^2}{2} \cdot I(0, z) \tag{4.14}$$

is the *total* optical power in the beam. Similarly, the optical power within a square of side length s and centered on the optical axis is

$$P = \mathcal{P}\left[\mathrm{erf}\left(\frac{s}{\sqrt{2}w(z)}\right)\right]^2, \tag{4.15}$$

where $\mathcal{P}$ is again the total optical power (4.14) in the beam and

$$\mathrm{erf}(z) = \frac{2}{\sqrt{\pi}}\int_0^z e^{-t^2}\,\mathrm{d}t \tag{4.16}$$

is the error function.

4.4.2 *PCI beam*

Prior to the heterodyne-interferometer upgrade, the total beam power at the PCI detector was $\mathcal{P}_{\mathrm{pci},0} \sim 300\,\mathrm{mW}$. The placement of the AOM in the expansion optics produces a static $\sim 10\%$ optical insertion loss, and driving the AOM with RF power deflects an additional fraction η_R of the beam power to the reference arm of the interferometer. Further, upstream of the phase plate, a fraction η_P of the beam power is diverted to the probe arm of the interferometer. Thus, after the upgrade, the total beam power at the PCI detector is

$$\mathcal{P}_{\mathrm{pci}} = 0.9(1-\eta_R)(1-\eta_P)\mathcal{P}_{\mathrm{pci},0}. \tag{4.17}$$

To prevent substantial performance degradation of the PCI, the criterion $\mathcal{P}_{\mathrm{pci}}/\mathcal{P}_{\mathrm{pci},0} \gtrsim 2/3$ is enforced.

4.4.3 *Interferometer probe beam*

Upstream of the phase plate, a fraction η_P of the beam power is diverted to the probe arm of the interferometer. Recall that the phase-plate groove is uncoated ZnSe, which reflects only 17% of the unscattered probe beam. Additionally, 50% of the probe beam's power is lost at the beam combiner that combines the probe beam and the reference beam. Thus, the total probe-beam power $\mathcal{P}_P$ at the heterodyne-interferometer detector is given through power conservation as

$$\mathcal{P}_{\mathrm{pci}} + \left(\frac{0.17}{0.5}\right)\mathcal{P}_P = \mathcal{P}_{\mathrm{pci},0}. \tag{4.18}$$

(Of course, some of the beam power is also diverted to the feedback system's quadrant detector, but the power in this feedback beam was negligibly altered by the upgrade). Using the probe-beam optical layout discussed in Section 4.3.5, the 1/e E radius of the probe beam at the heterodyne-interferometer detector is

$$w_P = 2.7\,\mathrm{mm}. \tag{4.19}$$

4.4.4 *Interferometer reference beam*

The AOM produces a static $\sim 10\%$ optical insertion loss and deflects a fraction η_R of the remaining beam power to the reference arm of the interferometer. Additionally, 50% of the reference beam's power is lost at the beam combiner that combines the probe beam and the reference

beam. Thus, the total reference-beam power $\mathcal{P}_R$ at the heterodyne-interferometer detector is

$$\mathcal{P}_R = \left(\frac{0.9}{2}\right) \eta_R \mathcal{P}_S, \tag{4.20}$$

where $\mathcal{P}_S \approx 14\,W$ is the total beam power at the laser source. Using the reference-beam optical layout discussed in Section 4.3.4, the $1/e\,E$ radius of the probe beam at the heterodyne-interferometer detector is

$$w_R = 4.3\,\text{mm}. \tag{4.21}$$

4.4.5   *Constraining the distribution of optical power*

Suitably distributing the optical power in the combined PCI and interferometer system is an exercise in constrained optimization. The objective function is the optical power in the heterodyne-interference signal, while the constraints are the operational limits of the interferometer detector and the minimum required power in the PCI beam ($\mathcal{P}_{pci}/\mathcal{P}_{pci,0} \gtrsim 2/3$). Because the PCI constraint is somewhat "soft", a graphical exploration of the design space is appropriate; such a graphical exploration is shown in Figure 4.9. The *top-left* panel of Figure 4.9 displays the AC-to-DC multiplicative factor $2I_{AC}/(I_{DC} + I_{AC})$ from the heterodyne-interferometer transfer function (3.98). The AC-to-DC multiplicative factor attains a maximum value of unity when the AC and DC intensities of the heterodyne interference signal are equal (i.e. when $I_{AC} = I_{DC}$). Clearly, the AC-to-DC multiplicative factor is a strong function of η_R, and the design point achieves a reasonable balance between the AC and DC components of the heterodyne signal. A large value of $2I_{AC}/(I_{DC} + I_{AC})$ alone does *not* guarantee optimal performance; the normalizing peak intensity $I_{DC} + I_{AC}$ should also be maximized. The *top-right* panel of Figure 4.9 displays this peak intensity. The photovoltaic HgCdTe detectors that are often employed for $10.6\,\mu\text{m}$ heterodyne interferometry typically have a saturation intensity $I_{sat} = 100\,\text{mW}\cdot\text{mm}^{-2}$ and a damage intensity $I_{dam} = 1\,W\cdot\text{mm}^{-2}$. Note that $I_{DC} + I_{AC}$ is a strong function of η_R and that $I_{DC} + I_{AC} \sim I_{sat}$ at the design point. Depending on the saturation physics of the detector, operating above I_{sat} (but below I_{dam}) may improve the system performance, as quantified by the $V_1/V_1(I_{max} = I_{sat})$ multiplicative factor in the heterodyne-interferometer transfer function (3.98). However, as the *bottom-left* panel of Figure 4.9 shows, the constraint $\mathcal{P}_{pci} \gtrsim (2/3)\mathcal{P}_{pci,0} \sim 200\,\text{mW}$ prevents operating the heterodyne interferometer with intensities above I_{sat}. This inability to
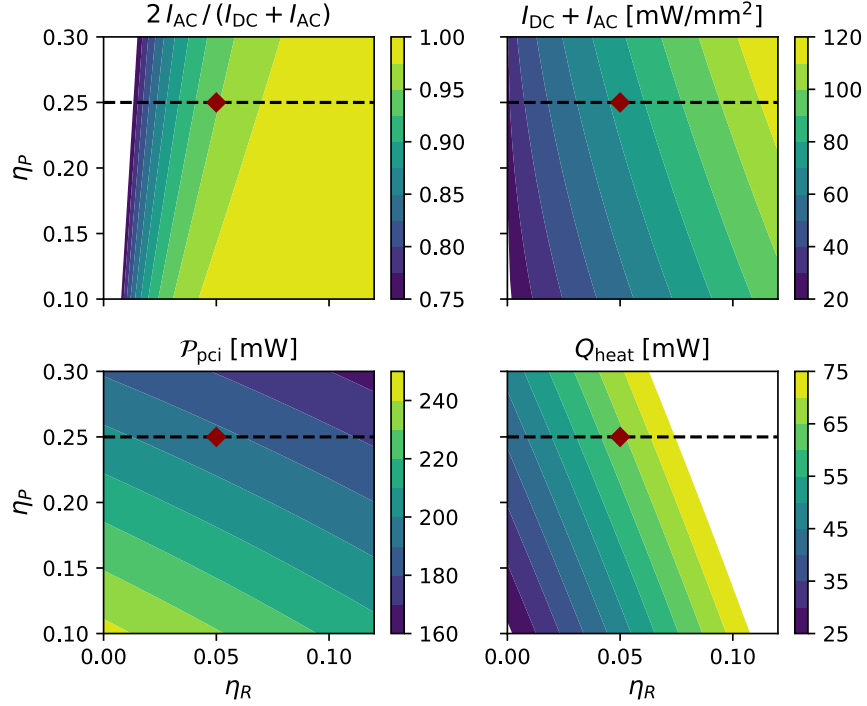
Figure 4.9: Constraints on the distribution of optical power in the combined PCI-interferometer. The AOM deflection efficiency, which governs the fraction η_R of beam power diverted to the interferometer reference arm, is displayed along the x-axis. The fraction η_P of probe-beam power (immediately upstream of the phase plate) diverted to the interferometer probe arm is displayed along the y-axis. All intensities correspond to peak, on-axis values. The design point is marked with a burgundy diamond; note that $\eta_P = 0.25$ is fixed for the installed splitter, but varying the AOM RF power readily changes η_R, allowing easy movement along the horizontal dashed line. See text for discussion.

operate above I_{sat} can be compensated by actively cooling the detector to improve its sensitivity. The interferometer described in this thesis uses two-stage thermoelectric cooling to cool the detector to $T \sim 230\,K$, boosting the specific detectivity D^* by an order of magnitude; the cooling capacity is limited to $Q_{heat} \lesssim 75\,mW$. Because the total power in the probe beam and the reference beam greatly exceeds this cooling capacity, a circular aperture of radius $a = 0.75\,mm$ is positioned immediately upstream of the detector; the power of a Gaussian beam passing through such an aperture is given by (4.13) such that the heat that must be dissipated reduces to

$$Q_{heat} = \mathcal{P}_P \left[1 - \exp\left(\frac{-2a^2}{w_P^2} \right) \right] + \mathcal{P}_R \left[1 - \exp\left(\frac{-2a^2}{w_R^2} \right) \right]. \quad (4.22)$$

The *bottom-right* panel of Figure 4.9 displays Q_{heat}. Now, in practice, η_P is fixed by the reflectivity of the ZnSe splitter located upstream of the phase plate (with 25% reflectivity being a common, off-the-shelf value); then, during an alignment, η_R is increased slowly until the cooling capacity of the detector is reached.

4.5 DEDICATED HETERODYNE-INTERFEROMETER HARDWARE

Utilizing a shared in-vessel probe beam implies that the PCI and the heterodyne interferometer share several components, including the CO_2 laser, the beam-expansion optics, and the beam-delivery and beam-collection optics. The two systems also share a digitizer. Numerous other components, however, are exclusively dedicated to the operation of the heterodyne interferometer. Some of these components have been discussed in previous sections (i.e. the reference-beam optics in Section 4.3.4 and the probe-beam imaging optics in Section 4.3.5). This section details the remainder of the heterodyne interferometer's dedicated hardware.

4.5.1 *Oven-controlled crystal oscillator (OCXO)*

Placing a crystal oscillator (XO) in a constant-temperature thermal bath substantially improves the XO stability, dramatically reducing the XO phase noise. For technical reasons, ovens operated tens of degrees above the ambient temperature provide the most robust constant-temperature thermal baths. The stability of such an oven-controlled crystal oscillator (OCXO) is typically enhanced by a factor $\geqslant 10^3$ relative to that of a comparable XO [11, Table 7.2].

The local oscillator (LO) of the heterodyne interferometer is derived from an OCXO. The OCXO was procured from Wenzel Associates, Inc. (Austin, TX, USA) and is enclosed in their Sprinter packaging. The OCXO has a 5-minute warm-up time, during which it draws a maximum power of 5 W; after warming up, it draws 2.2 W. The OCXO output is a 13 dBm sinusoidal wave at a frequency of 30 MHz (i.e. in the notation of Chapters 2 and 3, $\Delta\omega_0 = 2\pi \cdot 30\,\mathrm{MHz}$). The OCXO phase noise is

$$\mathcal{L}_{\Delta\omega_0}(f) = -165\ \mathrm{dBc}, \qquad f \geqslant 10\ \mathrm{kHz}. \tag{4.23}$$

The conversion between dBc and $\mathrm{rad}^2 \cdot \mathrm{Hz}^{-1}$ is discussed in Appendix D.

A custom, rack-mounted module was built to house the OCXO. The module housing protects the OCXO from rough handling and encloses
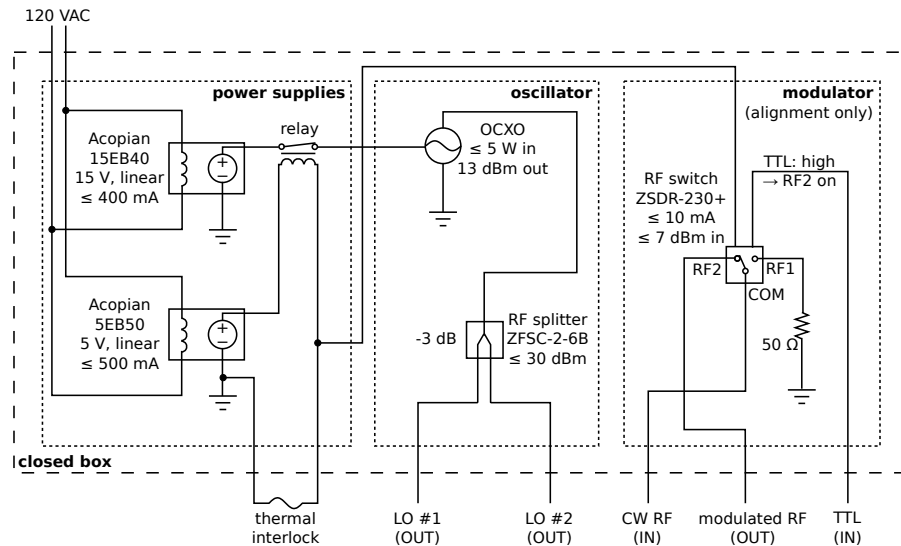
Figure 4.10: Schematic for OCXO rack-mounted module.

additional electronics that facilitate the OCXO integration into the rest
of the heterodyne-interferometer system. The OCXO module is in-
stalled in the PCI rack located in the DIII-D annex. A schematic of this
module is shown in Figure 4.10. A 15 V linear-regulated power supply
(15EB40, Acopian Technical Co.; Easton, PA, USA) powers the OCXO,
while a 5 V linear-regulated power supply (5EB50, Acopian) powers
the AOM's thermal-regulation and RF-modulation electronics. Linear-
regulated power supplies are less noisy than either switching-mode
or unregulated power supplies, so it was deemed essential to power
the OCXO with a linear-regulated power supply. To prevent overheat-
ing the AOM (e.g. due to a water-cooling failure), a normally closed,
creep-action thermostat (Pepi Model N, Portage Electric Products, Inc;
North Canton, OH, USA) monitors the AOM temperature; if the AOM
temperature exceeds 30 °C, the thermostat opens. Now, assuming the
AOM temperature is less than 30 °C, connecting the thermostat of the
AOM to the thermal interlock of the OCXO module produces a current
that closes the normally open, non-latching relay that sits between the
15 V linear-regulated power supply and the OCXO, powering on the
OCXO. The OCXO signal is evenly split via a Mini-Circuits (Brooklyn,
NY, USA) ZFSC-2-6B RF splitter to produce two LO signals; one of the
LO signals is routed to the AOM RF driver (described in Section 4.5.2)
located in the DIII-D pit, and the other LO signal is routed to the
interferometer demodulation electronics (described in Section 4.5.6)
located in the PCI rack of the DIII-D annex. A Mini-Circuits FTB-1-1
balun transformer breaks the ground loop between the OCXO module

and the AOM RF driver. A Mini-Circuits ZSDR-230+ RF switch also enables amplitude modulation of the LO signal that is routed to the AOM RF driver; this capability is only used during system alignment to "chop" the reference and probe beams, allowing each beam to be separately centered on the interferometer's AC-coupled detector. Finally, a switch at the rear of the module allows one to easily toggle the module grounding between the electrical mains (i.e. for bench testing) and the PCI rack (i.e. for operations).

4.5.2 *AOM RF driver*

The LO signal must be amplified to several watts in order to drive the AOM. An ENI (Rochester, NY, USA) 3100L performs this amplification. (Note that ENI no longer exists but that the closely related E&I currently manufactures comparable amplifiers). The 3100L provides broadband power amplification between 250 kHz to 105 MHz. The 3100L has 50 dB gain and a rated power output above 100 W (as the AOM is typically driven at ~ 1 W, the 3100L greatly exceeds the drive requirements; note that the 3100L is on long-term loan from the DIII-D RF group, though, and they had no suitable, lower-power amplifiers). The harmonic distortion of the 3100L is ⩽ −25 dBc, and its noise figure is ⩽ 10 dB. Self-demodulation tests (methodology shown schematically in Figure 4.17) conducted with and without the 3100L demonstrate that the 3100L does *not* add appreciable noise to the interferometer measurements. The 3100L draws 1.1 kW of 120 VAC wall power; a 2.4 kVA Topaz isolation transformer isolates the 3100L from high-frequency wall-power noise. At a hefty 70 lbs, the 3100L is mounted in the PCI rack adjacent to the PCI optical table in the DIII-D pit; the Topaz isolation transformer is mounted to the pit wall behind the PCI rack. The AOM drive level is adjusted by appropriately attenuating the LO signal just upstream of the 3100L input; this "just-in-time" attenuation minimizes the effect of electrical pickup during the ~ 60 m signal transit from the annex to the pit (relative to e.g. attenuating the signal in the annex and then sending the signal to the pit).

4.5.3 *Detector*

Detection of 10.6 μm light is typically effected with Mercury Cadmium Telluride (HgCdTe), an alloy of HgTe and CdTe [10]. The ratio of HgTe to CdTe is tuned to provide optimal detection at the device's intended operational wavelength and temperature. Cooling the HgCdTe typically increases the cut-off wavelength, increases optical responsivity,
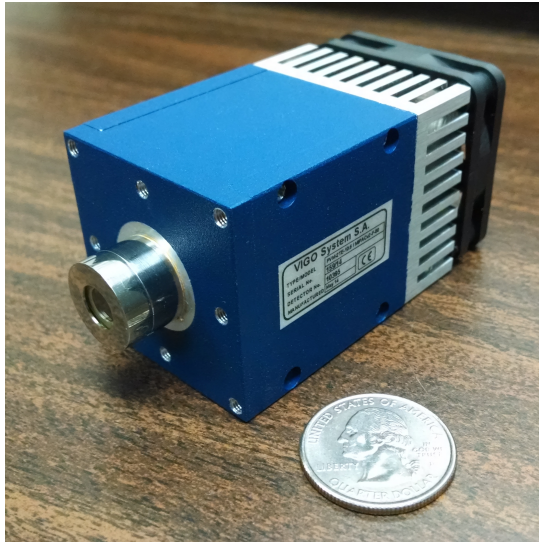
Figure 4.11: The heterodyne interferometer's detector module. The detector element and thermoelectric (TE) cooler are enclosed within the TO-8 silver "can" at middle left. The blue housing encloses additional TE-cooling components and the detector preamplifier. A heat sink and small fan are located opposite of the detector face. The female SMA connector (for detector output) and the receptacle for a 9-pin LEMO cable (for communication with the module controller) sit on the unpictured side of the detector module.

and decreases device noise, but often at the expense of device speed. Device noise and speed also depend on whether the HgCdTe detector is photovoltaic (PV) or photoconductive (PC); in particular, PC detectors exhibit $1/f$ noise and are often slower than their PV counterparts.

Resolving the temporal dynamics of the heterodyne interferometer's 30 MHz signal requires a sufficiently fast detector. This immediately precludes the use of a liquid-nitrogen cooled detector, as is used for the PCI. Room-temperature HgCdTe detectors are sufficiently fast but also exhibit relatively low sensitivity. Modest thermoelectric (TE) cooling, however, can engender some of the positive aspects of cooling while maintaining sufficient time resolution.

The 30 MHz heterodyne-interference signal is measured with a PVM-2TE-10.6 detector element mounted on a MIPACv2 detection module, both procured from VIGO System S.A. (Ozarow Mazowiecki, Poland). The multiple-junction PV element provides superior speed (3 dB high-frequency cutoff ~ 50 MHz) and sensitivity ($D^* \sim 1.9 \times 10^8$ cm$\sqrt{\text{Hz}}/W$) relative to a PC element or a single-junction PV element. The element is square with side length $s_x = 1$ mm, in accordance with (4.7). The element has linear saturation intensity $I_{sat} = 100$ mW $\cdot$ mm^{-2} and damage intensity $I_{dam} = 1$ W $\cdot$ mm^{-2}. The element and TE cooler are en-

closed in TO-8 packaging with a BaF_2 window; care should be taken not to scratch the soft BaF_2 window or to subject the system to mechanical shocks, which may damage the TE cooler. The TE cooler is capable of maintaining the detector element at a constant 235 K temperature for incident DC optical fluxes $\lesssim 75$ mW. The MIPACv2 detection module houses additional TE-cooling components as well as the detector preamplifier. The preamplifier is AC-coupled with a cut-on frequency of 1 kHz and a cut-off frequency of 50 MHz, and it is capable of producing a ± 1 V output-voltage swing into a 50 Ω load. The MIPACv2 detection module is powered by VIGO's STCC-04 controller. The specific detectivity of the element and the detection module taken as an integrated unit is

$$D^* = 5.35 \times 10^7 \text{ cm}\sqrt{\text{Hz}}/\text{W}. \tag{4.24}$$

The detector module is pictured in Figure 4.11.

4.5.4 *Coaxial cables to DIII-D annex*

RG-58 coaxial cables transmit signals between the interferometer components located in the DIII-D annex and the DIII-D pit. Initially installed for PCI signal transmission but abandoned following an upgrade to fiber-optic links [2, Sec. 3.3.3], these cables were reclaimed for the interferometer. The 16 cables sit bundled beneath the PCI optical table and connect to channels C27 through C42 of panel 5B in the DIII-D annex. The single-transit propagation time of each cable was measured to be 315 ns by launching a square wave of modest frequency (tens of kHz) down the *unterminated* cable and halving the observed time delay between the forward and reflected waves. As RG-58 has an index of refraction $\sim 3/2$, this corresponds to a cable length of 63 m. DC signals are negligibly attenuated along this cable length, but AC signals are subject to the skin effect [11, Sec. H.1.4]; the measured attenuation at 30 MHz is 5.9 dB. The DC signal between the AOM thermostat and the thermal interlock of the OCXO module travels along one of these coaxial cables. A second coaxial cable transmits one of the 30 MHz LO signals from the OCXO module to the AOM RF driver, and a third coaxial cable transmits the 30 MHz heterodyne-interference signal from the interferometer detector to the signal-conditioning RF amplifiers.
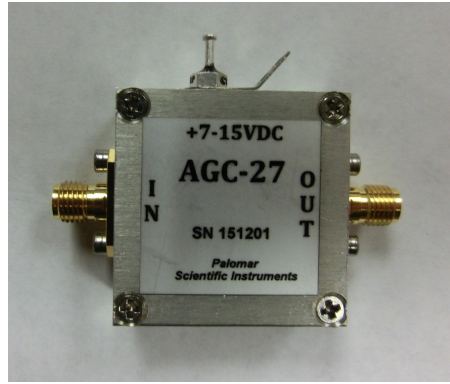
Figure 4.12: Automatic gain-control (AGC) amplifier.

4.5.5   *Signal-conditioning RF amplifiers*

Despite feedback stabilization of the beam coalignment (discussed in Sections 4.3.3 and 4.3.5), vibration-induced misalignment can still substantially reduce the amplitude of the heterodyne-interference signal, dramatically degrading the signal-to-noise ratio of the interferometer. These amplitude variations can be compensated with an automatic gain-control (AGC) amplifier, which monitors the amplitude of the input signal and dynamically adjusts its gain to maintain a constant-amplitude output signal.

An AGC amplifier was graciously provided *pro bono* by Palomar Scientific Instruments (San Marcos, CA, USA). The AGC operates over a frequency range from 20 MHz to 100 MHz. The AGC produces a -6 dBm output signal for an input signal between -24 dBm and $+6$ dBm; the input damage limit is 20 dBm. The input and output impedances are both 50 Ω, and the input and output ports are female SMA. The AGC draws 160 mA from a DC power supply that can sit anywhere between 7 V and 15 V; however, higher DC voltages produce significant heat dissipation, and operation at the upper limit of 15 V should be avoided, if possible. The AGC is pictured in Figure 4.12. If the heterodyne interference signal falls below the lower limit of the AGC input range (e.g. due to long-term drift of the system alignment), a Mini-Circuits (Brooklyn, NY, USA) ZFL-500B amplifier is added immediately upstream of the AGC.

These signal-conditioning RF amplifiers are located in the PCI rack of the DIII-D annex. A 12 V linear-regulated power supply (12EB40, Acopian Technical Co.; Easton, PA, USA) powers the amplifiers. A Mini-Circuits FTB-1-1 balun transformer breaks the ground loop be-
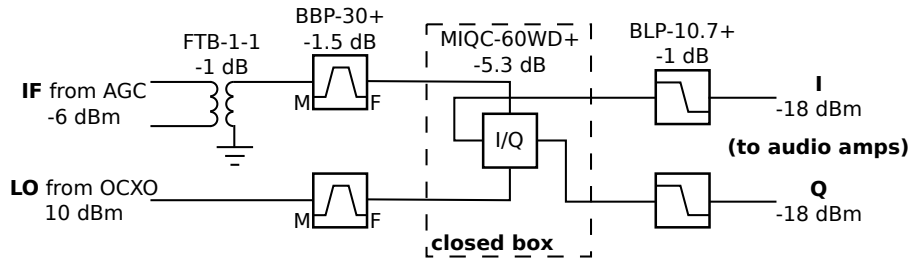
Figure 4.13: Schematic for demodulation electronics. Note that the polarity of the bandpass filters is explicitly noted as M (male) and F (female). The OCXO and the electronics downstream of the I&Q demodulator are all grounded to the PCI rack in the DIII-D annex; the IF signal conditioning amplifiers are *not* grounded to the rack, so a balun transformer breaks the ground loop. Measurements along any point in the circuit should be into $50\,\Omega$.

tween the interferometer detector and the signal-conditioning RF amplifiers.

4.5.6 *Demodulation electronics*

The theory of ideal and real-world demodulation is discussed in Section 3.5, and readers are encouraged to review that section if context for the hardware described below is desired. A schematic of the demodulation electronics is shown in Figure 4.13.

Prior to demodulation, the LO and IF signals are each bandpass filtered with a Mini-Circuits (Brooklyn, NY, USA) BBP-30+. The BBP-30+ has $\leqslant 1.5\,\text{dB}$ insertion loss over its 27 MHz to 33 MHz passband. Bandpass filtering the 30 MHz IF suppresses higher-order harmonics and out-of-band noise, both of which would degrade the heterodyne interferometer's phase measurement. Although Section 3.5.2 discusses the potential benefits of demodulating against a square LO, no such LO was available, so the LO is instead bandpass filtered to ensure that the IF is demodulated against a sinusoid. The BBP-30+ is *directional*, so care should be taken to install the filter with the correct orientation. Following the bandpass filter, a Mini-Circuits FTB-1-1 breaks the ground loop between the IF signal-conditioning amplifiers and the downstream electronics. The resulting $-8.5\,\text{dBm}$ IF and $8.5\,\text{dBm}$ LO are now ready to be demodulated.

Demodulation is performed with a Mini-Circuits MIQC-60WD+ analog I&Q demodulator. The MIQC-60WD+ can demodulate IF signals with frequencies between 20 MHz and 60 MHz, while the bandwidth of the resulting I and Q signals may span from DC up to 5 MHz. To minimize parasitic capacitances, the MIQC-60WD+ is soldered "dead-

bug" style [12] to a large grounding plane, and every electrical connection is soldered directly to its corresponding pin; the demodulator, wiring, and grounding plane are enclosed in a protective box with external BNC jacks. The MIQC-60WD+ conversion loss is 5.3 dB. The amplitude imbalance is $\leqslant 0.6$ dB, the phase imbalance $\leqslant 5°$, and the DC offset is typically ~ 1 mV. The 3$^{\text{rd}}$- and 5$^{\text{th}}$-order harmonic suppressions were measured to be 53 dB and 64 dB, respectively.

The signals exiting the I and Q ports of the demodulator must be low-pass filtered to remove the 60 MHz "sum" components that result from mixing the LO and IF. Low-pass filtering is performed with a pair of Mini-Circuits BLP-10.7+ filters, which have $\leqslant 1$ dB insertion loss over their DC to 11 MHz passband. The resulting I and Q signals are each -18 dBm (i.e. 28.2 mV RMS, 80 mV peak-to-peak, into 50 Ω).

4.5.7 *Audio amplifiers*

The demodulated I and Q signals are 80 mV peak-to-peak, while the digitizer input voltage range is 8 V peak-to-peak. Thus, direct digitization of the I and Q signals corresponds to a fractional use $\eta_{\text{dyn}} = 0.01$ of the digitizer's dynamic range. When $\eta_{\text{dyn}} = 0.01$, the autospectral density of the corresponding quantization noise (3.109) is 10^4 *larger* than when $\eta_{\text{dyn}} = 1$. Figure 4.20 compares the quantization noise for $\eta_{\text{dyn}} \approx 1$ to typical plasma fluctuations — clearly, $\eta_{\text{dyn}} = 0.01$ would produce unacceptably large quantization noise that would mask all but the largest plasma fluctuations. To minimize quantization noise, then, the I and Q signals must be amplified to $\eta_{\text{dyn}} \approx 1$ prior to digitization.

A pair of custom "audio amplifiers" were built to perform the desired amplification of the I and Q signals. The "audio" qualifier indicates that these amplifiers are lower bandwidth than the upstream RF amplifiers discussed in Section 4.5.5 (the BBP-30+ bandpass filters restrict the bandwidth of the I and Q signals to $\lesssim 3$ MHz, and plasma fluctuations above the interferometer noise floor are almost always $\lesssim 1$ MHz). The required amplitude gain of the amplifier is $G \approx 100$. A schematic for a single channel of the amplifier is shown in Figure 4.14. The amplifier consists of two stages of high-gain, low-noise amplification followed by a 50 Ω line driver, which is required to drive the digitizer's 50 Ω input impedance. Each high-gain, low-noise amplification stage utilizes an OP37 bipolar operational amplifier (Analog Devices; Norwood, MA, USA) in an inverting configuration. The OP37 is *decompensated* [11, Sec. 4.9], sacrificing low-gain stability for higher bandwidth (gain-bandwidth product of 63 MHz),
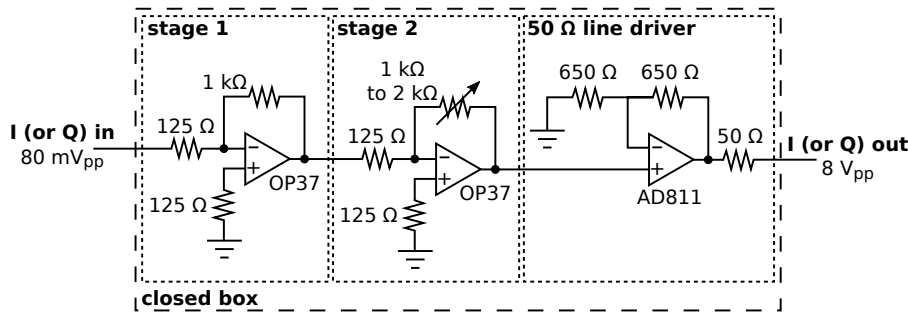
Figure 4.14: Audio-amplifier schematic. The amplifier consists of two high-gain, low-noise amplification stages followed by a 50 Ω line driver, which is required to drive the digitizer's 50 Ω input impedance.

and it is stable for closed-loop gains $\geqslant 5$. The first stage has a fixed gain $|G_1| = 8$, and the second stage has a variable gain $8 \leqslant |G_2| \leqslant 16$ that can be easily adjusted via a potentiometer (allowing quick "on-the-fly" optimization of the amplifier gain). The inverting and non-inverting inputs of each OP37 see roughly balanced impedances, minimizing the non-ideal effect of input bias current [11, Sec. 4.4.2.E]. The 50 Ω line driver utilizes an AD811 current-feedback operational amplifier (also from Analog Devices) in a non-inverting configuration with an amplitude gain of two. The output impedance of each audio amplifier is $Z_{\text{out}} = 50\,\Omega$ such that the total gain when driving a 50 Ω load is $64 \leqslant G \leqslant 128$. The measured bandwidth of each audio amplifier is in excess of $2\,\text{MHz}$. The input impedance of each audio amplifier is $Z_{\text{in}} = 125\,\Omega$. Thus, there is a slight impedance mismatch between the audio amplifiers and the 50 Ω demodulation components; this impedance mismatch decreases the signal-transmission efficiency, but more nefarious transmission-line effects are negligible for the $\lesssim 1\,\text{MHz}$ signals in the $\sim 0.5\,\text{m}$ coaxial cables connecting the demodulation electronics to the audio amplifiers. Impedance matching could be improved, for example, by utilizing a non-inverting amplifier in the first stage (effectively infinite input impedance) coupled with a parallel 50 Ω resistor to ground at the amplifier input. The power-supply rails of each operational amplifier are bypassed with $1\,\mu\text{F}$ ceramic capacitors to prevent coupling high-frequency power-supply noise into the amplifier [11, Sec. 4.2.7]. Feedback-loop parasitic capacitances, which can reduce the high-frequency, closed-loop gain and excite instability, are minimized via "deadbug" construction [12], as shown in Figure 4.15. Although deadbug construction is ideal for prototyping, any future increase to the number of interferometer channels would call for a printed-circuit-board (PCB) construction. The au-
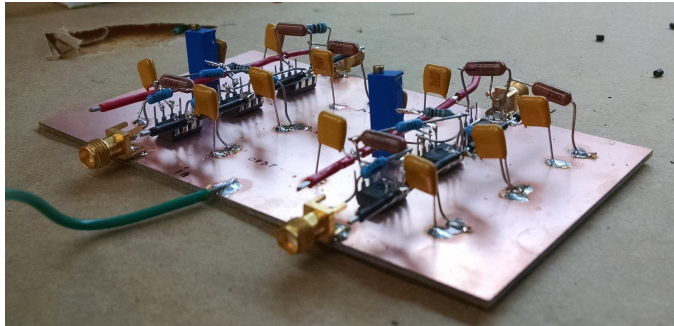
Figure 4.15: The "deadbug" construction of the audio amplifiers minimizes parasitic capacitances that can produce instability. The operational amplifiers sit upside down, resembling "dead bugs", and are held in place by their solder connections to other components in the circuit. An unetched printed circuit board serves as the grounding plane. The amplifiers are packaged within a protective box (not pictured here).

dio amplifiers share 15 V power supplies with the PCI fiber-optic-link receivers in the PCI rack of the DIII-D annex.

4.5.8 *Anti-aliasing filters*

Passive anti-aliasing filters (J3715-840K-50-720B, TTE Filters; Arcade, NY, USA) sit immediately upstream of the digitizer. Each filter has $50\,\Omega$ input and output impedance, a DC to $840\,\text{kHz}$ passband, and $3\,\text{dB}$ frequency of just over $1\,\text{MHz}$. These filters are reclaimed PCI components, but future procurement of filters with a larger passband ($\lesssim 2\,\text{MHz}$ for typical $f_s = 4$ MSPS sampling rates) would improve the interferometer's bandwidth.

4.6 DATA PREPARATION

The interferometer-measured phase ϕ_m is computed from the in-phase I and quadrature Q signals via the two-argument arctangent function in (3.102). Before computing the measured phase ϕ_m, however, the ellipticity of the I and Q signals is compensated, as described in Section 4.6.1. Then, before computing the autospectral density, a zero-delay, finite-impulse-response, high-pass filter is applied to the measured phase ϕ_m, as described in Section 4.6.2.
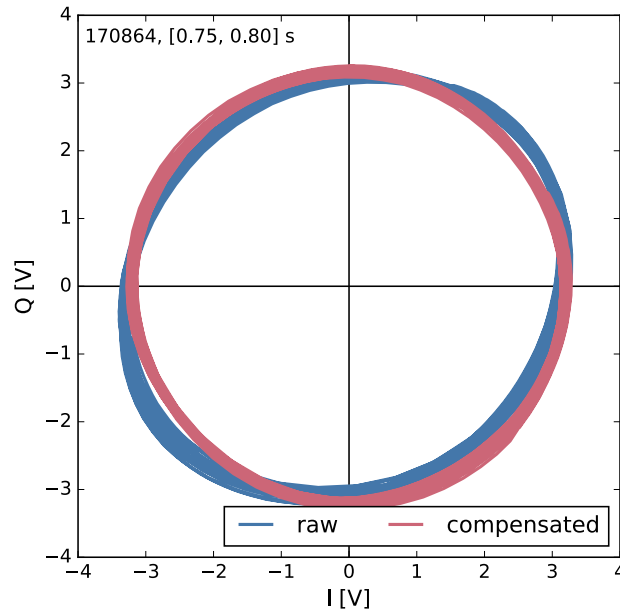
Figure 4.16: An example of the ellipticity compensation applied to the I and Q signals prior to computing the measured phase ϕ_m.

4.6.1 *Ellipticity compensation of I & Q signals*

Demodulator imperfections produce a relative error in the measured fluctuating phase, as described in Section 3.5.4. Because the IF power entering the demodulator is sufficiently small, demodulator nonlinearities are negligible, and the I and Q signals possess negligible higher-order harmonics (i.e. $|I_3| \ll I_1$, $|Q_3| \ll Q_1$, etc.). Thus, a Lissajous figure of Q from (3.84) vs. I from (3.83) is an ellipse. This ellipse is fitted with a direct, efficient, least-squares algorithm [13, 14]. The resulting fit is then used to compensate any DC offsets, amplitude imbalance, and phase imbalance in the I and Q signals, minimizing the relative error (3.87) in the measured fluctuating phase. Figure 4.16 displays an example of raw and compensated I and Q signals.

4.6.2 *High-pass filtering the measured phase ϕ_m*

Vibrations contaminate the low-frequency components ($f \lesssim 10\,\text{kHz}$) of the measured phase ϕ_m. At the lowest frequencies, vibrational contributions to ϕ_m are orders-of-magnitude larger ($\gtrsim 100\,\text{dB}$) than the corresponding plasma contribution. Spectral estimates are often com-

puted after time-history tapering with a Hanning window and with frequency resolution $\sim 1\,\text{kHz}$. While the Hanning window provides at least 32 dB side-lobe suppression [15, Sec. 11.5.2.1], the vibrational component to ϕ_m is so large that substantial spectral leakage can still occur. To minimize spectral leakage, then, the measured phase is high-pass filtered in software to remove the large vibrational contributions. A non-causal, type-I, finite-impulse-response (FIR) filter has the desirable property that it produces zero delay for all frequencies [16][17, Sec. 5.7.3]. Such a filter can be easily designed via the Kaiser window method [17, Sec. 7.5.3][18]. Typically, a Kaiser-designed, high-pass filter with -120 dB ripple, 10 kHz cut-on frequency, and 5 kHz transition-region width performs sufficiently; for the typical $f_s = 4$ MSPS digitization rate, such a filter has a length of 12489 points (corresponding to $\sim 3\,\text{ms}$), and only points in the filtered phase that are free from the filter's boundary effects are included in subsequent analysis (that is, $\sim 1.5\,\text{ms}$ at the beginning and the end of the record are lost to the filter's boundary effects).

4.7   NOISE IN HETERODYNE INTERFEROMETER

The sensitivity of the heterodyne interferometer is set by the noise in the system. Initially, the heterodyne interferometer was plagued by enormous noise that obscured all but the strongest coherent plasma fluctuations. Through substantial effort, this noise was found to be wholly attributable to LO phase noise and the finite coupling time of the AOM. The theory of noise generation via this mechanism was discussed in Section 3.3.2. Below, Section 4.7.1 describes a simple, all-electrical means of characterizing LO phase noise that, in hindsight, would have saved a great deal of time and heartache. This section also discusses compensation of LO phase noise with a delay line, an approach that was empirically found to suffer from electrical pickup during DIII-D operations. Ultimately, the heterodyne-interferometer noise issues were resolved by procuring the OCXO described in Section 4.5.1 and eliminating the pickup-prone delay line. Having solved the LO phase-noise issues, Section 4.7.2 proceeds with a spectral characterization of the predicted and measured sources of noise in the heterodyne interferometer and compares the resulting noise floor to typical plasma-fluctuation spectra.
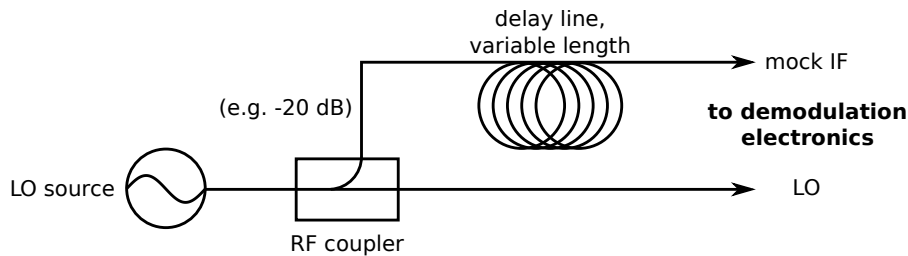
Figure 4.17: LO self-demodulation schematic for investigation of LO phase noise. A small fraction of the LO is coupled and delayed to produce a mock IF. The mock IF is demodulated and digitized, and the autospectral density of the measured phase is compared to the corresponding autospectral density with zero delay.

4.7.1 *Quantification & delay-line compensation of LO phase noise*

The finite coupling time of an AOM injects LO phase noise into the measurements of a heterodyne interferometer. This noise mechanism was thoroughly discussed in Section 3.3.2, and it can easily be the dominant source of noise in a heterodyne interferometer, potentially producing severe constraints on the system sensitivity. Thus, if LO phase noise is not known *a priori*, it should be measured. Often, the standard deviation of the oscillator frequency is specified (for example, the oscillator initially procured for this project was specified to be 0.01% quartz-stabilized); however, without knowledge of the corresponding spectrum, the specification of the standard deviation is insufficient to characterize the resulting heterodyne-interferometer noise.

The injection of LO phase noise into the heterodyne interferometer can be empirically characterized by demodulating a delayed copy of the LO, as depicted schematically in Figure 4.17. A Mini-Circuits (Brooklyn, NY, USA) ZFBDC20-61HP+ bi-directional coupler provides adequate coupling for $\sim 30\,\text{MHz}$, moderate-power ($\lesssim 25\,W$) RF signals. The delay line should provide a delay comparable to the AOM coupling time. To lowest order, the AOM coupling time can be estimated as the time required for sound waves in the AOM Germanium crystal (sound speed $c_s = 5400\,\text{m}\cdot\text{s}^{-1}$) to propagate from the AOM RF port (assumed to be adjacent to the corresponding piezo-actuator) to the AOM optical axis; for the Gooch & Housego AOM discussed in Section 4.3.4, this distance is $d \approx 1.5\,\text{cm}$, corresponding to a coupling time $\tau \approx 2.8\,\mu\text{s}$ and an RG-58 coaxial-cable (index of refraction $\sim 3/2$) delay line of length $L \approx 560\,\text{m}$. The mock IF signal may need to be amplified to combat attenuation along the delay line, but amplifier noise should be minimized to prevent corrupting the phase-noise measurement; Mini-Circuits ZFL-500B amplifiers were found to
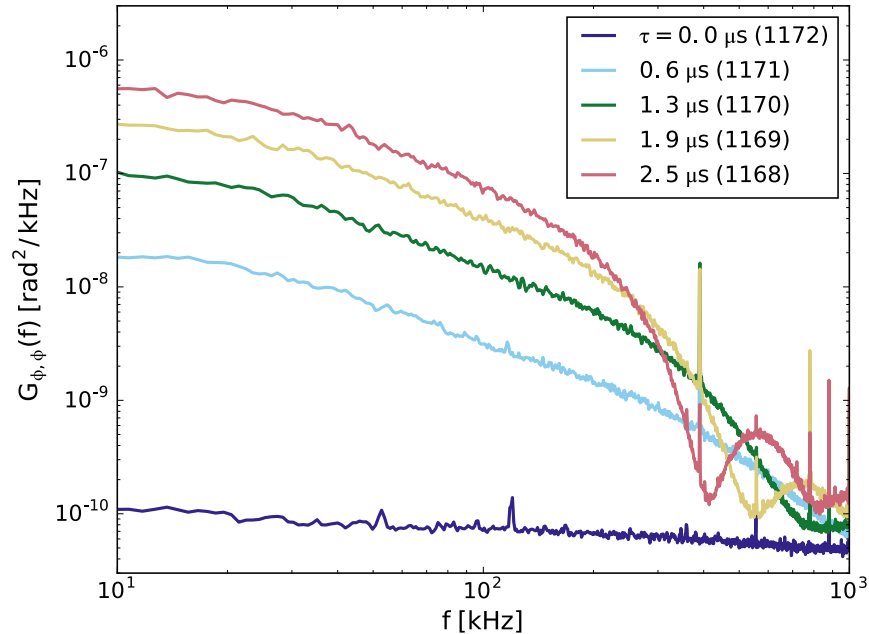
Figure 4.18: Self-demodulation of a 0.1% quartz-stabilized LO gauges the injection of LO phase noise into a heterodyne interferometer with various AOM coupling times τ (the 4-digit numbers in parentheses indicate the corresponding diagnostic test shot). If the injection of LO phase noise were negligible, the autospectral density $G_{\phi,\phi}(f)$ would be independent of the delay τ; clearly, the injection of phase noise with *this* LO is significant for finite τ. The coupling time (4.25) of the AOM used in this work is $\tau \approx 2.5\,\mu s$.

provide acceptable low-noise amplification. Additionally, because the bulk phase difference between the mock IF and the LO is fixed by the delay-line length, the demodulated I and Q signals do *not* trace out a full Lissajous ellipse; for this reason, the ellipticity compensation discussed in Section 4.6.1 should *not* be performed prior to computing the measured phase ϕ_m. Finally, the autospectral density of the measured phase should be estimated. If this autospectral density exceeds the corresponding autospectral density with zero delay, the injection of LO phase noise into the heterodyne interferometer will be significant. Results from such an LO self-demodulation test with the 0.01% quartz-stabilized oscillator initially procured for this project are shown in Figure 4.18; clearly, the phase noise of this oscillator is unacceptably large.

The injection of LO phase noise into the heterodyne interferometer can be minimized by delaying the LO by the AOM coupling time such that the LO and IF are equally delayed (i.e. the time delay between the LO and IF vanishes; effectively, $\tau \rightarrow 0$). The rough estimate

of the AOM coupling time in the previous paragraph can be refined by noting that the sine-squared weighting of the autospectral density corresponding to the injected LO phase noise (3.106) exhibits "ringing", with nulls at frequencies that are integer multiples of $1/\tau$. Such nulls were observed in the interferometer-measured spectrum at integer multiples of $f \approx 400 \, \text{kHz}$, implying an AOM coupling time of

$$\tau \approx 2.5 \, \mu\text{s}, \tag{4.25}$$

in reasonable agreement with the rough estimate from the previous paragraph. Delaying the LO by the AOM coupling time (4.25) can be easily accomplished with an $L \approx 500 \, \text{m}$ RG-58 coaxial-cable delay line. Such a delay line was constructed by stringing together eight of the pit-to-annex coaxial cables discussed in Section 4.5.4 and amplifying the signal to appropriately compensate for attenuation. As expected, the delay line eliminated the injection of LO phase noise into the heterodyne interferometer, and a whole host of coherent and broadband plasma fluctuations were subsequently measured with the system.

However, it was suspected that the use of a 500 m delay line increased the system's susceptibility to electrical pickup, particularly during DIII-D operations. To quantify the electrical pickup attributable to the delay line, the LO self-demodulation tests schematically depicted in Figure 4.17 were repeated with equal (but extended) LO and IF cable lengths during DIII-D operations. Because a non-zero time delay between the LO and IF injects LO phase noise that greatly exceeds the delay-line electrical pickup, the LO and IF cable lengths were extended by equal amounts, with the total additional cable length acting as a proxy for the delay-line cable length L. The results of this test are shown in Figure 4.19. Clearly, even a 250 m delay line suffers from electrical pickup during DIII-D operations; the broadband electrical pickup between $\sim 100 \, \text{kHz}$ and $\sim 400 \, \text{kHz}$ is particularly constraining, as many plasma fluctuations of interest occupy this frequency range.

The above noise issues were resolved by procuring the OCXO described in Section 4.5.1. The OCXO phase noise $\mathcal{L}_{\Delta\omega_0}(f)$ is sufficiently small that the LO phase noise injected into the heterodyne interferometer (3.106) is negligible for *any* AOM coupling time τ, allowing removal of the pickup-prone delay line.
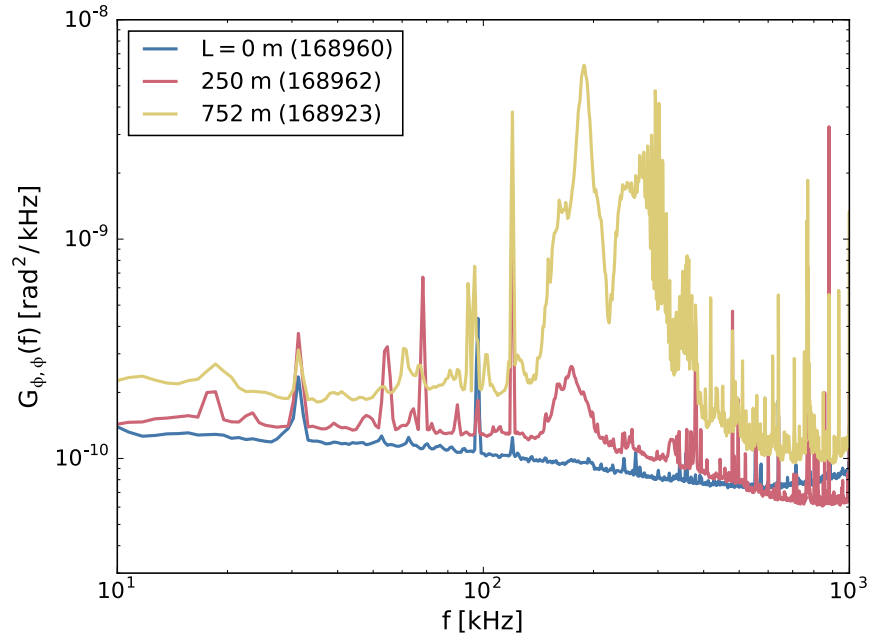
Figure 4.19: Delay-line electrical pickup during DIII-D operations. Delay-line length L is indicated (the 6-digit numbers in parentheses indicate the corresponding DIII-D shot). Clearly, even a 250 m delay line suffers from electrical pickup during DIII-D operations. With the 2.5 µs AOM coupling time (4.25), compensating for poor LO stability requires a 500 m delay line. The delay line can be eliminated (L = 0) when a more stable LO is used.

4.7.2 *Spectral characterization of heterodyne-interferometer noise*

Having resolved the debilitating noise issues described in Section 4.7.1, the resulting noise floor of the heterodyne interferometer should be quantified. By comparing noise measurements to theoretical predictions, the dominant source of noise can be identified and the potential for future sensitivity improvements can be evaluated. Figure 4.20 spectrally characterizes both predicted and measured noise sources and compares them to typical plasma-fluctuation spectra.

Noise predictions establish a baseline for the expected performance of the heterodyne interferometer. Recall that Chapter 3 extensively discusses heterodyne-interferometer noise sources and their corresponding autospectral densities, with a concise summary provided in Section 3.7.4. The dashed, annotated lines in Figure 4.20 indicate the predicted (P) autospectral densities from these noise sources in the heterodyne-interferometer implementation described in this chapter. Clearly, the demodulated detector noise (3.107) is predicted to be the dominant source of noise. Note that the demodulated detector noise
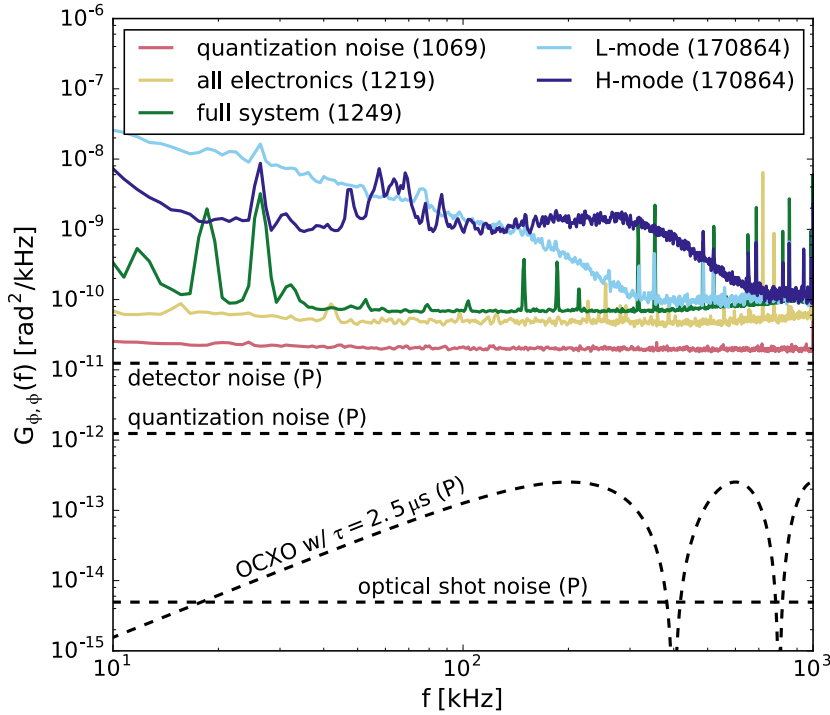
Figure 4.20: Spectral characterization of heterodyne-interferometer noise and comparison to typical plasma-fluctuation spectra. Dashed lines indicate predicted (P) quantities, while colored traces correspond to measured quantities with shot numbers in parentheses (4-digit numbers indicate a diagnostic test shot, and 6-digit numbers indicate a DIII-D shot). Importantly, the heterodyne-interferometer noise floor (i.e. the "full system" trace) sits at least an order of magnitude below typical plasma fluctuations. Vibrations dominate the heterodyne-interferometer spectrum for $f \lesssim 10\,\text{kHz}$.

depends on the probe-beam and reference-beam optical powers impinging on the detector element, P_P and P_R, respectively. Referencing (4.15), the probe-beam optical power impinging on the square detector element of the heterodyne interferometer is

$$P_P = \mathcal{P}_P \left[\text{erf}\left(\frac{s_x}{\sqrt{2}w_P} \right) \right]^2 \approx 25\,\text{mW}, \tag{4.26}$$

where $s_x = 1\,\text{mm}$ from (4.7), $w_P = 2.7\,\text{mm}$ from (4.19), and $\mathcal{P}_P = 315\,\text{mW}$ from (4.18) and the design point indicated in Figure 4.9. Simi-

larly the reference-beam optical power impinging on the square detector element of the heterodyne interferometer is

$$P_R = \mathcal{P}_R \left[\text{erf} \left(\frac{s_x}{\sqrt{2}w_R} \right) \right]^2 \approx 10\,\text{mW}, \tag{4.27}$$

where $s_x = 1\,\text{mm}$ from (4.7), $w_R = 4.3\,\text{mm}$ from (4.21), and $\mathcal{P}_R = 315\,\text{mW}$ from (4.20) and the design point indicated in Figure 4.9. The demodulated detector noise also depends on the detector-element area $A = s_x^2 = 1\,\text{mm}^2$ and the detector-module specific detectivity D^* from (4.24). The predicted quantization noise (3.109) corresponds to a bit depth $N_b = 14$ and a sample rate $f_s = 4\,\text{MSPS}$, as described in Section 4.2.5, and an assumed use of the digitizer's full dynamic range ($\eta_{dyn} = 1$). The OCXO phase noise $\mathcal{L}_{\Delta\omega_0}(f)$ from (4.23) is injected into the interferometer by the $\tau = 2.5\,\mu\text{s}$ AOM coupling time from (4.25); because the OCXO has very low phase noise, it makes a negligible contribution to the heterodyne-interferometer noise. The demodulated optical shot noise (3.108) is easily evaluated using the values of P_P from (4.26) and P_R from (4.27), and it also makes a negligible contribution to the heterodyne-interferometer noise. Unfortunately, the laser phase noise $\mathcal{L}_{\omega_0}(f)$ is *not* known, so the injection of laser phase noise (3.105) into the heterodyne interferometer cannot be predicted (recall that the probe beam and the reference beam do *not* have matched optical path lengths); thus, quantifying the effect of laser phase noise requires empirical investigation.

While the noise predictions establish a baseline, the *measured* noise floor actually sets the sensitivity of the heterodyne interferometer. Several of the colored traces in Figure 4.20 display the noise contributions from various system components. The system quantization noise was measured with the I and Q channels of the digitizer *open*, and equivalent results were obtained when closing both channels with resistors matching the $50\,\Omega$ input impedance of each digitizer channel. The normalization of the quantization noise assumes full use of the digitizer's dynamic range ($\eta_{dyn} = 1$); utilizing less of the dynamic range increases the quantization noise (e.g. $\eta_{dyn} \sim 0.8$ increases the quantization noise by $\sim 50\%$). Interestingly, the measured quantization noise exceeds the predicted quantization noise by more than an order of magnitude. While it is not uncommon for aperture error, jitter, and nonlinearities to reduce the effective bit depth by one or two bits (increasing the quantization noise by a factor of two to four) [15, Sec. 10.2.4], this discrepancy between the measured and predicted quantization noises is quite a bit larger. The origin of this anomalously high quantization noise is not understood, and it was not investigated any further in

this work. The system electrical noise was measured during an LO self-demodulation test (shown schematically in Figure 4.17) with the OCXO, a $\tau = 2.5\,\mu s$ delay between the LO and mock IF in accordance with the AOM coupling time (4.25), the appropriate LO and IF powers, and all the remaining electronics (filters, demodulator, audio amplifiers, etc.) configured as they are during operations. The electronics produce an approximately threefold increase in the system noise relative to the quantization noise. Finally, the full-system noise was measured by interfering the probe beam and the reference beam on the detector in the absence of plasma. Relative to the all-electrical noise measurement, the full-system test additionally includes the effects of laser phase noise, detector noise, and vibrations. Clearly, for frequencies $f \gtrsim 10\,kHz$, the full-system, broadband noise is only marginally larger than the electrical noise (the coherent peaks at $f \lesssim 30\,kHz$ may correspond to vibrations or laser phase noise; however, without either implementing two-color detection to perform vibration subtraction [19], or matching the probe-beam and reference-beam optical path lengths to remove the injection of laser phase noise, it is impossible to determine the origin of these low-frequency coherent peaks). Note that the full-system broadband noise sets the heterodyne-interferometer noise floor at

$$\min\left[G_{\phi,\phi}(f)\right] \approx 7 \times 10^{-11}\ rad^2/kHz, \tag{4.28}$$

corresponding to a sensitivity to line-integrated electron-density fluctuations of $3 \times 10^{14}\ m^{-2}/\sqrt{kHz}$. For context, autospectral densities for typical plasma-density fluctuations in L-mode and H-mode are also shown in Figure 4.20. Importantly, the fluctuations are one to two orders of magnitude larger than the noise floor across a large fraction of the heterodyne interferometer's temporal bandwidth.

4.8 CALIBRATION OF COMBINED PCI-INTERFEROMETER

It is always desirable to perform an end-to-end calibration of a system to ensure the system is operating as desired. The combined PCI-interferometer can be end-to-end calibrated by injecting sound waves into the shared probe beam to produce well-known refractive-index fluctuations that are detectable by both systems. Below, Section 4.8.1 describes the sound-wave verification of the heterodyne interferometer's amplitude and wavenumber response. Section 4.8.2 then discusses the cross-calibration of the PCI against the absolute phase measurements of the heterodyne interferometer. The robust response of

the combined system across the surveyed wavenumber range demonstrates the complementary nature of the interferometer and PCI measurements and confirms the ability of the combined system to simultaneously monitor low-k and high-k instabilities.

4.8.1   *Verification of heterodyne-interferometer response*

As discussed in Section 2.4.2, an ideal heterodyne interferometer possesses a flat wavenumber response and measures absolute phase fluctuations, independent of any calibration factors. However, as discussed in Section 3.1.4, detection is always effected via finite-size detector elements, which effectively low-pass filter the spatial content of the signal; in particular, a square detector element weights the wavenumber transfer function of an ideal heterodyne-interferometer with a normalized sinc function, as shown in (3.99). The first zero of this normalized-sinc weighting sets the nominal high-k cutoff of the system, and the design point for the heterodyne interferometer described in this work is $5\,\mathrm{cm}^{-1}$, as specified in (4.6). Of course, inadvertent clipping of the probe beam would alter the wavenumber response. Further, as demonstrated in Section 3.5.4, demodulator imperfections produce relative errors in the measured phase fluctuations. For these reasons, it is desirable to verify the amplitude and wavenumber response of the heterodyne interferometer.

Sound waves are ideal tools for verifying the wavenumber response of an interferometer. Sensitive to the index of refraction N, interferometers can detect the fluctuations in N corresponding to sound-wave pressure perturbations. Further, the well-known sound-wave dispersion relation allows one to infer the sound-wave wavenumber k from the measured sound-wave frequency f. Often, the sound-wave frequency is repeatedly swept in time, enabling verification of the interferometer response throughout its full wavenumber domain within a single diagnostic test shot. However, quantitatively comparing the predicted and the measured interferometer responses requires an accurate description of the sound-wave pressure fluctuations. The sound waves used throughout this work are characterized in detail with an absolutely calibrated microphone in Appendix E.

The response of the heterodyne interferometer to the sound waves can be *predicted* from the measurements of the sound-wave pressure perturbations. In particular, (E.11) reveals that the phase shift $\widetilde{\phi}_{CO_2}$
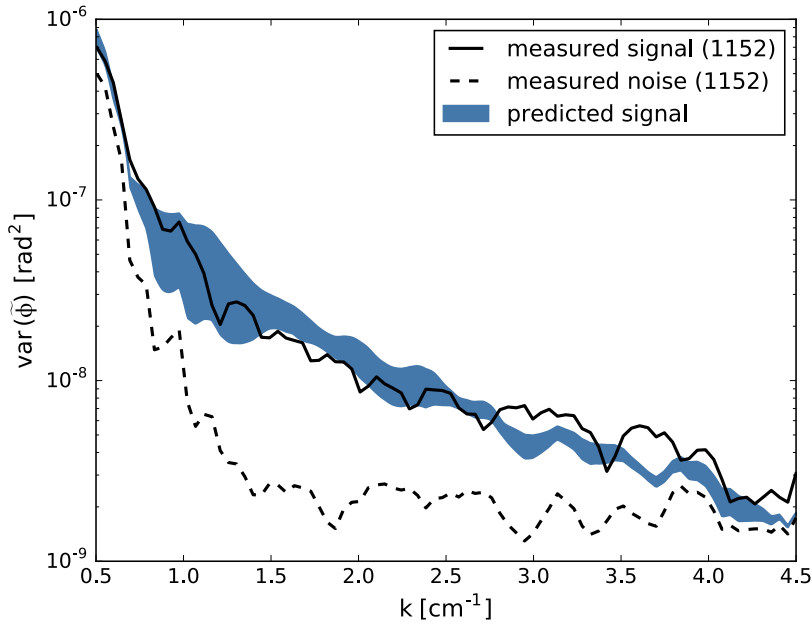
Figure 4.21: Sound-wave verification of heterodyne-interferometer response. In diagnostic test shot 1152, a speaker was centered beneath the interferometer probe beam, and swept-frequency sound waves were driven orthogonally through the beam. The interferometer-measured variance of the resulting phase fluctuations is indicated by the solid black trace. The shaded blue region indicates the predicted variance (4.31), where $\text{var}(\widetilde{\phi}_{CO_2})$ is computed from Appendix E's measurements with a calibrated microphone of the sound waves' absolute pressure perturbations. Note that the variance predicted here differs from that in Figure E.6 due to the inclusion of the interferometer's finite sampling-volume effects and noise. Further, the apparent increase in noise at low-k actually corresponds to low-f vibrational noise; that is, low-k *plasma* fluctuations at higher frequencies will not be plagued by this noise.

imparted to a $10.6\,\mu\text{m}$ CO_2 probe beam by a sound-wave pressure perturbation $\widetilde{p}$ is

$$\widetilde{\phi}_{CO_2}\,[\text{rad}] = (1.1 \times 10^{-5}\,\text{cm}^{-1}) \int (\widetilde{p}\,[\text{Pa}])\,dl, \qquad (4.29)$$

where the differential path length dl must have units of centimeters. However, to make a true "apples-to-apples" comparison with the interferometer measurements, finite sampling-volume effects and the system noise should be included such that the predicted phase fluctuation at each wavenumber k becomes

$$\widetilde{\phi}(k) = T_{\text{fsv}}(k) \cdot \widetilde{\phi}_{CO_2}(k) + \widetilde{\phi}_{\text{noise}}(f), \qquad (4.30)$$

where $T_{fsv}(k)$ is the finite sampling-volume transfer function (3.99) with k_{fsv}^{int} from (4.6), and $\widetilde{\phi}_{noise}(f)$ is the system noise. (Here, the noise is written as an explicit function of frequency f to emphasize that the noise is frequency dependent and *not* wavenumber dependent; however, because the sound-wave wavenumber is inferred from the sound-wave frequency, this frequency-dependent noise has the *appearance* of a wavenumber dependence in the ensuing figures). Because the sound-wave phase fluctuations and the system noise are uncorrelated, the predicted variance of the sound-wave phase fluctuations is simply

$$\text{var}\left[\widetilde{\phi}(k)\right] = T_{fsv}^2(k) \cdot \text{var}\left[\widetilde{\phi}_{CO_2}(k)\right] + \text{var}\left[\widetilde{\phi}_{noise}(f)\right]. \qquad (4.31)$$

Figure 4.21 compares the predicted variance (4.31) to the interferometer-measured variance, indicating very good agreement.

4.8.2 *Cross-calibration of PCI*

As discussed in Section 2.5.6, PCI does *not* make an absolute measurement of phase fluctuations. Instead, quantifying the absolute magnitude of the PCI-measured phase fluctuations requires a calibration factor, Φ_{pci}, which is a function of the probe-beam power, the alignment quality at the phase plate, the detector responsivity, and the detector-to-digitizer electrical throughput. Of course, this calibration factor can be estimated from specified or measured system parameters. However, an independent and quite possibly more accurate approach is to cross-calibrate the PCI measurements against the absolute phase measurements of the heterodyne interferometer.

As the current heterodyne-interferometer implementation only utilizes a single detector element, cross-calibrating the PCI requires a sound-wave test shot, such as the test shot used in Section 4.8.1 to verify the heterodyne-interferometer response. The sound waves should be moderate-k, well away from the cutoffs of both the PCI and the heterodyne interferometer, and the sound-wave signal measured by both systems should sit robustly above the corresponding noise floor; for the work described here, this corresponds to sound waves with wavenumbers between $2\,\text{cm}^{-1}$ and $3.5\,\text{cm}^{-1}$. Because the heterodyne interferometer and the PCI share the probe beam, these moderate-k sound waves should produce identical signals on both systems, less any finite sampling-volume effects. Equating the measured sound-wave variance from both systems and appropriately accounting for

finite sampling-volume effects allows determination of the PCI calibration constant as

$$\Phi_{\text{pci}} = \frac{T_{\text{fsv}}^{\text{het}}(k)}{T_{\text{fsv}}^{\text{pci}}(k)} \sqrt{\frac{\text{var}\left[\widetilde{\phi}_{\text{het}}(k)\right]}{\text{var}\left[\widetilde{y}_{\text{pci}}(k)\right]}} = 1.3 \times 10^{-6} \text{ rad/bit}, \qquad (4.32)$$

where $T_{\text{fsv}}^{\text{het}}(k)$ is the heterodyne-interferometer finite sampling-volume transfer function (3.99) with $k_{\text{fsv}}^{\text{int}}$ from (4.6), $T_{\text{fsv}}^{\text{pci}}(k)$ is the PCI finite sampling-volume transfer function (3.99) with $k_{\text{fsv}}^{\text{pci}}$ from (4.4), and $\widetilde{y}_{\text{pci}}(k)$ is the raw PCI signal (in bits); the numerical value corresponds to evaluation of this expression for the DIII-D PCI system, and this calibration factor is used throughout the remainder of this work (i.e. $\widetilde{\phi}_{\text{pci}}(t) = \Phi_{\text{pci}} \cdot \widetilde{y}_{\text{pci}}(t)$; additionally, note that 4 is idiosyncratically the least-significant bit of the DIII-D PCI digitizer but that the numerical value in (4.32) corresponds to a unitary bit). Figure 4.22 displays the cross-calibration, indicating very good agreement for sound-wave wavenumbers between $2\,\text{cm}^{-1}$ and $3.5\,\text{cm}^{-1}$.

Upgrading the heterodyne-interferometer detector from a single element to a multi-element array would allow comparable cross-calibration of the PCI with plasma fluctuations rather than sound waves. To see this, note that the cross-calibration requires accounting for the wavenumber transfer functions of both the PCI and the heterodyne interferometer. Sound waves are currently used for the cross-calibration because their wavenumber can be unambiguously inferred from their measured frequency via the sound-wave dispersion relation. However, a detector array can measure the wavenumber directly, allowing cross-calibration with arbitrary (albeit moderate-k) plasma fluctuations. Such a capability would allow robust and accurate evaluation of the PCI calibration on a shot-to-shot and an intra-shot basis.
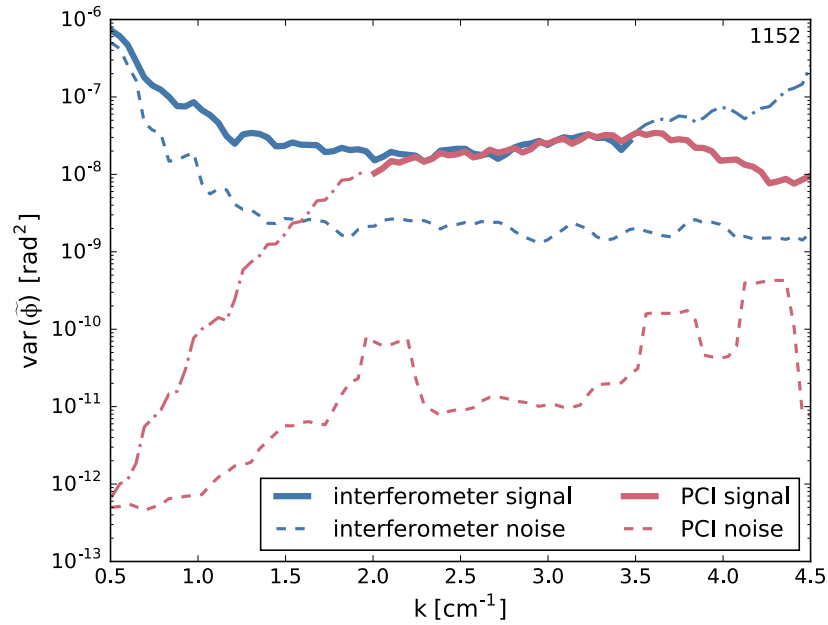
Figure 4.22: The PCI can be very accurately calibrated against the absolute measurements of the heterodyne interferometer. The results are demonstrated here for diagnostic test shot 1152. The interferometer signal corresponds to the solid black trace in Figure 4.21 *after* compensating for the attenuation of the high-k signal by finite sampling-volume effects; the corresponding dash-dot line for $k \geqslant 3.5\,\mathrm{cm}^{-1}$ is dominated by noise, so the finite sampling-volume compensation fails to produce accurate values of the sound-wave variance in this region. The dash-dot line for $k \leqslant 2\,\mathrm{cm}^{-1}$ corresponds to the cutoff PCI signal (note that the signal enters cutoff somewhat above the nominal low-k cutoff of $1.5\,\mathrm{cm}^{-1}$ from (4.2)). The robust response of the combined system across the surveyed wavenumber range demonstrates the complementary nature of the interferometer and PCI measurements and confirms the ability of the combined system to simultaneously monitor low-k and high-k instabilities. Also shown are the system noise floors; the enhanced signal-to-noise ratio of the PCI relative to the heterodyne interferometer is in order-of-magnitude agreement with the transfer-function predictions in Figure 2.6.

BIBLIOGRAPHY

[1] J. R. Dorris, J. C. Rost and M. Porkolab. Localized measurement of short wavelength plasma fluctuations with the DIII-D phase contrast imaging diagnostic. *Rev. Sci. Instrum.*, 80(2):023503, 2009.

[2] J. R. Dorris. *Phase contrast imaging measurements and modeling of short wavelength turbulence in the DIII-D tokamak*. PhD thesis, MIT, 2010.

[3] S. Coda. *An experimental study of turbulence by phase contrast imaging in the DIII-D tokamak*. PhD thesis, MIT, 1997.

[4] W. Horton. Drift waves and transport. *Rev. Mod. Phys.*, 71(3):735–778, 1999.

[5] G. R. Tynan, A. Fujisawa and G. McKee. A review of experimental drift turbulence studies. *Plasma Phys. Control. Fusion*, 51(11):113001, 2009.

[6] L. Lin, E. M. Edlund, M. Porkolab, Y. Lin and S. J. Wukitch. Vertical localization of phase contrast imaging diagnostic in Alcator C-Mod. *Rev. Sci. Instrum.*, 77(10):10E918, 2006.

[7] A. L. Sanin, K. Tanaka, L. N. Vyacheslavov, K. Kawahata and T. Akiyama. Two-dimensional phase contrast interferometer for fluctuations study on LHD. *Rev. Sci. Instrum.*, 75(10):3439–3441, 2004.

[8] K. Tanaka, S. Coda, M. Yoshida, H. Sasao, Y. Kawano et al. Design of tangential viewing phase contrast imaging for turbulence measurements in JT-60SA. *Rev. Sci. Instrum.*, 87(11):11E118, 2016.

[9] J. C. Rost. DIII–D PCI detector calibration check. Unpublished, March 2015.

[10] VIGO System, S.A. *Infrared detectors and related electronic devices*, 1 2014. VS-14-01-27MB.

[11] P. Horowitz and W. Hill. *The Art of Electronics*. Cambridge University Press, New York, NY, USA, 3rd edition, 2015.

[12] J. Williams. *High Speed Amplifier Techniques: A Designer's Companion for Wideband Circuitry (AN47)*. Linear Technology, Aug. 1991. Rev. A.

[13] A. Fitzgibbon, M. Pilu and R. B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. Pattern Analysis & Machine Intelligence*, 21(5):476–480, 1999.

[14] Nicky van Foreest. *fit_ellipse*, Feb. 2017. https://github.com/ndvanforeest/fit_ellipse.

[15] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 4th edition, 2010.

[16] J. F. Kaiser and W. A. Reed. Data smoothing using low-pass digital filters. *Rev. Sci. Instrum.*, 48(11):1447–1457, 1977.

[17] A. V. Oppenheim and R. W. Schafer. *Discrete-Time Signal Processing*. Pearson, Upper Saddle River, NJ, USA, 3rd edition, 2010.

[18] The Scipy community. *scipy.signal.kaiser (v0.19.1)*, June 2017.

[19] T. N. Carlstrom, D. R. Ahlgren and J. Crosbie. Real-time, vibration-compensated CO_2 interferometer operation on the DIII-D tokamak. *Rev. Sci. Instrum.*, 59(7):1063–1066, 1988.

5

MULTISCALE TURBULENCE MEASUREMENTS

The development of a first-principles understanding of turbulent transport in a tokamak is a long-standing goal of the fusion community. Such a model would allow accurate design and optimization of a tokamak. Before such a model is accepted, however, it must be validated against experimental measurements. Ideally, this validation should be multi-tiered, with the model accurately and robustly reproducing experimental observations at all scales, from macroscopic parameters such as the turbulent heat flux to the details of the turbulent spectrum (e.g. power spectral densities, cross phases, nonlinear couplings, etc.).

Along these lines, this chapter presents measurements of electron-density turbulence made over a wide spatiotemporal bandwidth during a recent DIII-D experiment. Section 5.1 begins with a brief overview of the drift-wave turbulence often observed in tokamak plasmas, emphasizing recent results from realistic multiscale simulations [1, 2, 3, 4, 5]. Section 5.2 then describes the recent DIII-D experiment, which was designed to probe the multiscale nature of plasma turbulence. As such, this experiment presents an ideal opportunity for multiscale turbulence investigations with the combined PCI-interferometer. Next, Section 5.3 details the measurements from the combined PCI-interferometer. Numerous turbulent branches are observed. In particular, the interferometer measures a low-k electromagnetic mode driven unstable by collisionality, properties consistent with the micro-tearing mode (MTM), and the PCI measures a wavenumber spectrum that exhibits distinct flattening when increasing the electron-scale turbulent drive relative to the ion-scale turbulent drive, reminiscent of results from realistic multiscale gyrokinetic simulations [3]. Finally, to aid the interpretation of these measurements, linear-stability analysis and quasilinear-transport modeling are performed with the TGLF code in Section 5.4, and qualitative agreement with the PCI-measured wavenumber spectrum is obtained.

5.1 OVERVIEW OF DRIFT-WAVE TURBULENCE

The radial transport of particles, heat, and momentum in a tokamak plasma is often larger than that predicted by collisional (i.e. neoclassical) theory. There is strong theoretical and experimental evidence that

this "anomalous" transport results from drift-wave turbulence driven by the free energy in the plasma gradients [6, 7]. This turbulence can be characterized by its driving gradients, its spatiotemporal characteristics, the role of trapped particles, and whether or not it is electrostatic (as opposed to electromagnetic).

The plasma gradients are often quantified in terms of scale lengths. The scale length of quantity x is defined as $L_x = x/|\nabla x| = |\nabla \ln x|^{-1}$. Smaller values of L_x indicate more rapid spatial variation (i.e. steeper gradient) of quantity x. Thus, the free energy in profile x increases with the inverse scale length $1/L_x$. For theoretical analysis, it is conventional to normalize L_x to a relevant length in the system under study; below, scale lengths are normalized to the tokamak minor radius a, which is appropriate for studies of radial transport. Thus, the drive for instability is quantified by the normalized inverse scale length a/L_x. The ratio of density scale length to temperature scale length $\eta_j = L_{n_j}/L_{T_j}$ for species j is an additional dimensionless parameter that is often used for instability characterization.

Due to the relatively large size of its eddies, ion-scale ($k_\theta \rho_s < 1$) turbulence is often considered to be the most detrimental to confinement. Here, k_θ is a typical poloidal wavenumber of the turbulent mode, $\rho_s = c_s/\Omega_i$ is the ion gyroradius evaluated at the ion sound speed, $c_s = (T_e/m_i)^{1/2}$ is the ion sound speed, and $\Omega_i = eB/m_i$ is the ion angular cyclotron frequency. The temporal ordering of ion-scale turbulence is usually expressed as $k_\| v_{ti} \ll \omega \ll k_\| v_{te}$ [8, Sec. 8.3], where ω is the angular frequency of the instability, $k_\|$ is the instability wavenumber parallel to the magnetic field, and v_{tj} is the thermal speed of species j; thus, the electrons respond rapidly (approximately adiabatically) to the electrostatic-potential fluctuations of the instability. Typically, trapped-ion dynamics make insignificant contributions to ion-scale drift-wave turbulence in tokamak plasmas [6, Sec. IV.E]; passing-particle dynamics, however, are significant and produce the η_i mode, which is driven linearly unstable above a critical value of η_i ($\eta_i^{crit} \sim 1$) [8, Sec. 8.3] and peaks about $k_\theta \rho_s \sim 0.3$ [7, Sec. 1.2.1]. For sufficiently flat density profiles ($a/L_{n_i} \lesssim 1$, which is valid across most of the radial domain for the plasmas considered in this work), the toroidal η_i mode has a critical value $\eta_i^{crit} = \eta_i^{crit}(L_{T_i})$, independent of L_{n_i}; for this reason, the η_i mode is also often referred to as the ion temperature gradient (ITG) mode [8, Sec.8.3], driven more unstable by larger values of a/L_{T_i}.

Despite having much smaller eddy size than its ion-scale counterparts, electron-scale ($k_\theta \rho_s > 1$) turbulence may still make large contributions to plasma transport. Of particular importance is the electron

temperature gradient (ETG) mode. The physics of the ETG is comparable to that of the ITG, with the roles of the ions and electrons reversed. (Here, the ion response is approximately adiabatic because the electrostatic potential fluctuations occur on spatial scales much smaller than the ion gyroradius [9, Sec. 2.3.4.2]). As the name suggests, the ETG is driven more unstable by larger values of a/L_{T_e} [8, Sec. 8.3]. Relative to the ITG, the ETG spatial scales in a $T_e \sim T_i$ deuterium plasma are reduced by the square root of the electron-to-ion mass ratio $(m_e/m_i)^{1/2} \approx 60$ such that the ETG peaks at $k_\theta \rho_s \sim 20$. While electron-scale simulations have long shown that the ETG may be capable of forming radially elongated "streamers" [10] capable of driving experimentally relevant electron heat transport [11], it was unknown whether or not such streamers would survive in the presence of ion-scale turbulence, as self-consistently and simultaneously simulating both ion- and electron-scale turbulence with realistic mass ratios was computationally intractable until very recently.

By exploiting the latest advances in high-performance computing, Howard *et al.* have performed multiscale gyrokinetic simulations with realistic mass ratios for both L-mode [1, 2, 3] and H-mode plasmas [4, 5]. Indeed, under certain realistic conditions, ETG streamers are predicted to coexist with the ITG and to drive empirically relevant levels of electron heat flux [1]. Local and non-local energy cascades between the ion and electron scales are also predicted [2], indicating the importance of cross-scale coupling. This cross-scale coupling becomes stronger as the ETG drive a/L_{T_e} increases relative to the ITG drive a/L_{T_i} [3]. Further, flattening of the wavenumber spectrum is predicted to be a tell-tale signature of this enhanced cross-scale coupling [3].

Seeking empirical validation of these predictions, Howard *et al.* recently performed an experiment at DIII-D to alter the local a/L_{T_e} and a/L_{T_i}. The reference discharge for this experiment, DIII-D shot 153523, was selected because multiscale gyrokinetic simulations indicate that its turbulent transport is intrinsically multiscale in nature [5]. The experimental conditions are discussed in Section 5.2, and the combined PCI-interferometer measurements of the resulting electron-density fluctuations are presented in Section 5.3. To aid the interpretation of these measurements, linear-stability analysis and quasilinear-transport modeling is performed in Section 5.4. While the plasmas simulated in the above multiscale work were ITG-ETG dominant, Sections 5.3 and 5.4 indicate the presence of additional modes; the basic properties of candidate modes consistent with these observations are briefly reviewed below.

The linear-stability analysis performed in Section 5.4 indicates the presence of a mid-k ($0.5 \lesssim k_\theta \rho_s \lesssim 5$) electron mode. The mode is destabilized with increasing a/L_{T_e} (and decreasing a/L_{T_i}). It is conceivable that this mode is the trapped electron mode (TEM). In contrast to the passing-particle dynamics of the ITG and ETG, trapped-electron dynamics destabilize the TEM [8, Sec. 8.4]. Trapped particles effectively average out their parallel velocities over a bounce period, but they also spend most of their time in regions of bad magnetic curvature, producing stability characteristics not unlike those seen in a magnetic mirror [8, Sec. 8.4]. Of course, trapped-particle dynamics will only be important if trapped particles are able to execute one or more bounce orbits before being detrapped by collisions. Thus, the TEM requires that most of the trapped electrons are in the banana collisionality regime ($\nu_e^* < 1$) [9, Sec. 2.3.4.3]. The spatiotemporal aspects of the TEM are characterized by $k_\theta \rho_s \sim 1$ and $\omega \sim \omega_{*e}$ [9, Sec. 2.3.1], where $\omega_{*e} = -k_\theta T_e (dn_e/dr)/(eBn_e)$ is the electron diamagnetic frequency [8, 8.2] and the other parameters were introduced in the above ITG discussion. (Note that ω_{*e} is derived from considerations of the adiabatic electron response along the magnetic field, which corresponds to the parallel electron temperature $T_{e,||}$; however, in an isotropic temperature distribution, $T_{e,||} = T_{e,\perp} = T_e$). Both a/L_{T_e} and a/L_{n_e} provide free energy to the TEM [9, Sec. 2.3.1].

Finally, Section 5.3 identifies a low-k ($k_R < 1.5\,\text{cm}^{-1}$) electromagnetic mode driven unstable by collisionality. These properties are consistent with the micro-tearing mode (MTM) [8, Sec. 8.5], which was predicted to be marginally unstable for $k_y \rho_s \sim 0.2$ in this experiment's reference discharge [5]. While the ITG, ETG, and TEM are electrostatic, the MTM is electromagnetic and can produce small-scale magnetic islands. Relative to conventional, macroscopic tearing modes, the MTM has short wavelength and high poloidal mode number m [8, Sec. 8.5]. Tearing-mode linear stability is often quantified by Δ', a parameter derived from resistive MHD [8, Sec. 7.3]; for the large poloidal mode number of the MTM, the resulting Δ' predicts stability. However, resistive MHD is an inadequate model for the MTM, and both kinetic and nonlinear effects can compete with Δ' to destabilize the MTM in the presence of finite η_e (i.e. finite electron temperature gradient). Importantly, via kinetic effects, the MTM is unstable at "moderate" ν_{ei}/ω_{*e}, where ν_{ei} is the electron-ion collision frequency and ω_{*e} is the electron diamagnetic frequency [8, Sec. 8.5].

5.2 EXPERIMENTAL CONDITIONS

The experiment was run in H-mode with the ITER-similar shape (i.e. aspect ratio, elongation, and triangularity all closely matched to those of the ITER-baseline scenario [8, Sec. 13.5 & 13.6]). The on-axis toroidal field $B_T = -1.7\,T$ and plasma current $I_p = 1.3\,MA$ produced $q_{95} = 3.15$, where q_{95} is the average value of the safety factor q [8, Sec. 3.4] over the surface that encloses 95% of the poloidal flux within the last-closed flux surface. Neutral beam injection (NBI) [8, Sec. 5.3-5.5] was performed with feedback to maintain $\beta_N = 1.9$, where β_N is the normalized plasma pressure [8, Sec. 6.18]. In order to suppress core MHD, an average NBI torque of approximately $1.5\,N \cdot m$ was injected into the plasma; note that this is approximately four times larger than the projected ITER-equivalent torque [12]. In order to alter the local electron-scale and ion-scale drives, the electron cyclotron resonance heating (ECH) [8, Sec. 5.10] location was scanned between $\rho = 0.5$ and $\rho = 0.8$, where ρ is the square root of the normalized toroidal flux (which scales as r/a, with r being the minor-radial coordinate and a being the minor radius of the plasma). Intra-shot scans of the ECH location were plagued with core MHD, so only shot-to-shot, MHD-free scans of the ECH location are considered here. The line-averaged density was $\overline{n}_e = 5.2 \times 10^{19}\,m^{-3}$. Impurities were removed from the plasma by both large and small edge localized modes (ELMs) [8, Sec. 7.17]. The time histories of several actuators and plasma parameters are shown in Figure 5.1. Note that multiscale gyrokinetic simulations of this experiment's reference discharge, DIII-D shot 153523 with ECH at $\rho = 0.5$, indicate that the turbulent transport is intrinsically multiscale in nature [5].

Equilibrium profiles were obtained by averaging over $200\,ms$, as indicated by the shaded regions in Figure 5.1. Magnetic equilibria were reconstructed with the EFIT code [13] and were constrained to match the total plasma pressure and motional Stark effect (MSE) measurements of the local magnetic pitch angle. Electron densities and temperatures were measured via Thomson scattering, while ion densities and temperatures were inferred from charge exchange recombination (CER) measurements of C^{6+}, the dominant impurity in DIII-D. The radial electric field was computed by invoking force balance on C^{6+}. To minimize the impact of ELMs on the profile fits, only measurements falling within the last 50% – 99% of each inter-ELM window were included in the fitting. The fitted profiles and their corresponding gradients or normalized inverse scale lengths are shown in Figure 5.2. While it may seem counterintuitive that the central electron tempera-
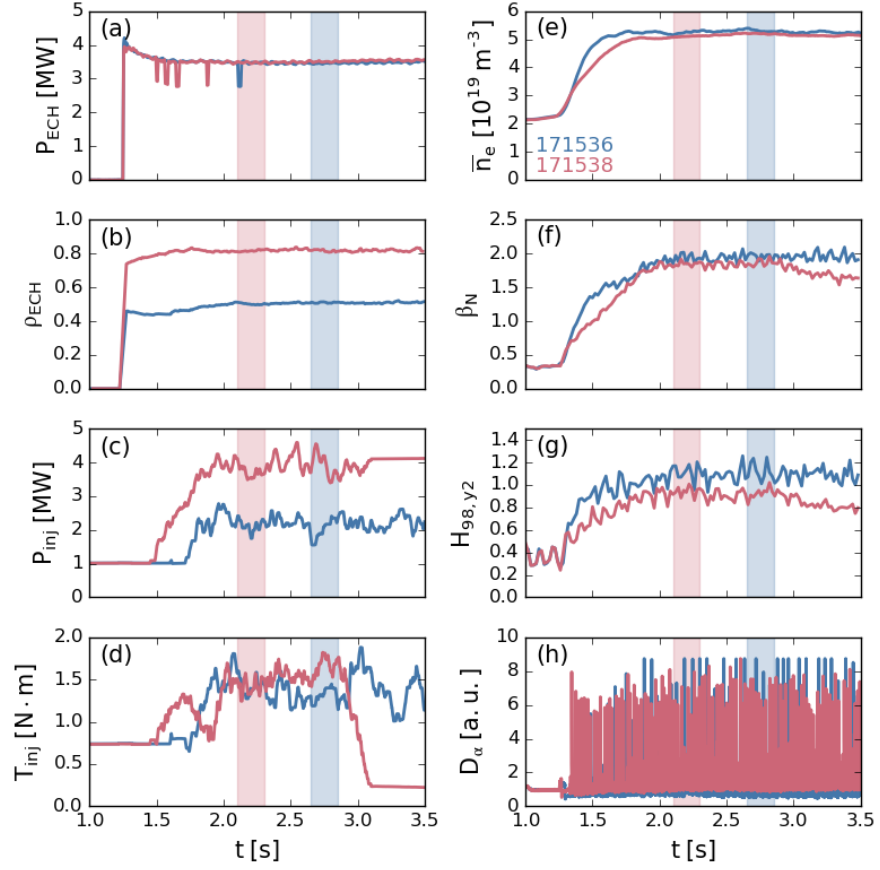
Figure 5.1: Time histories of various actuators and plasma parameters: (a) electron cyclotron resonance heating (ECH) power P_{ECH}, (b) ECH location ρ_{ECH}, (c) neutral beam injected (NBI) power P_{inj}, (d) NBI torque T_{inj}, (e) line-averaged density $\bar{n}_e$, (f) normalized plasma pressure β_N, (g) confinement quality $H_{98,y2}$, and (h) divertor D_α light, indicating the presence of large and small edge localized modes (ELMs).

ture $T_e(0)$ increases when moving ECH from $\rho = 0.5$ to $\rho = 0.8$, maintaining constant β_N requires increased NBI heating (see Figure 5.1(c)), which enhances the NBI electron heating density $q_{e,NBI}$ across the full plasma profile, as shown in Figure 5.3(b). The 1σ uncertainties in the profile fits, indicated by the shaded bands in Figure 5.2, were quantified by performing separate fits to 100 distinct data sets generated via Monte Carlo variation of the measurements about their uncertainties. Clearly, moving ECH from $\rho = 0.5$ to $\rho = 0.8$ produces large changes in the electron-scale and ion-scale drives, a/L_{T_e} and a/L_{T_i}, respectively, in the region of the plasma accessible to the PCI probe beam ($R = 1.98\,\mathrm{m}$). Using these profiles, power-balance analysis was performed with the ONETWO code [14], with NUBEAM [15] calcu-
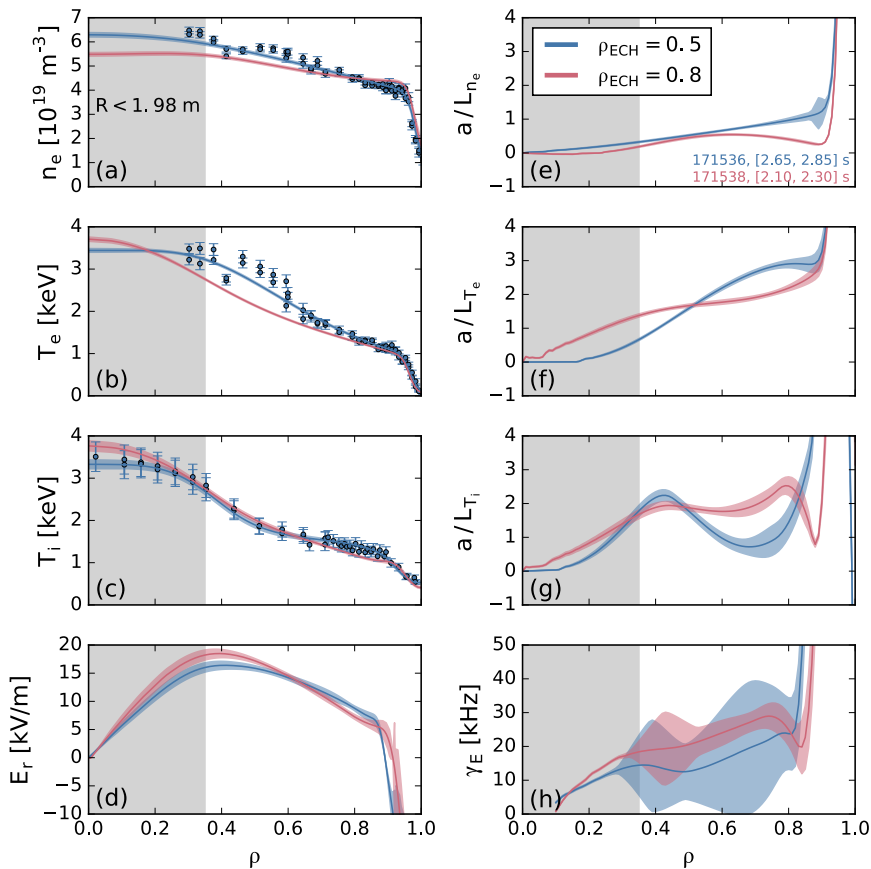
Figure 5.2: Profiles, normalized inverse scale lengths, and $E \times B$ shearing rate: (a) electron density n_e, (b) electron temperature T_e, (c) deuterium temperature T_i, (d) radial electric field E_r along the outboard midplane, (e) normalized inverse n_e scale length a/L_{n_e}, (f) normalized inverse T_e scale length a/L_{T_e}, (g) normalized inverse T_i scale length a/L_{T_i}, (h) $E \times B$ shearing rate γ_E. The shaded bands indicate the 1σ uncertainties in the profiles, as determined by performing separate fits to 100 distinct data sets generated via Monte Carlo variation of the measurements about their uncertainties. Representative measurements and their uncertainties are indicated for a 10 ms window from a single shot. The relatively large uncertainty on γ_E is dominated by uncertainty in the curvature of the T_i profile.

lations for NBI heating and torque and TORAY [16] calculations for ECH; the resulting loop voltages, stored energies, and neutron rates match their measured values to within ±5%.

5.3 COMBINED PCI-INTERFEROMETER MEASUREMENTS

The experiment described in Section 5.2 presents an ideal opportunity for multiscale turbulence investigations with the combined PCI-
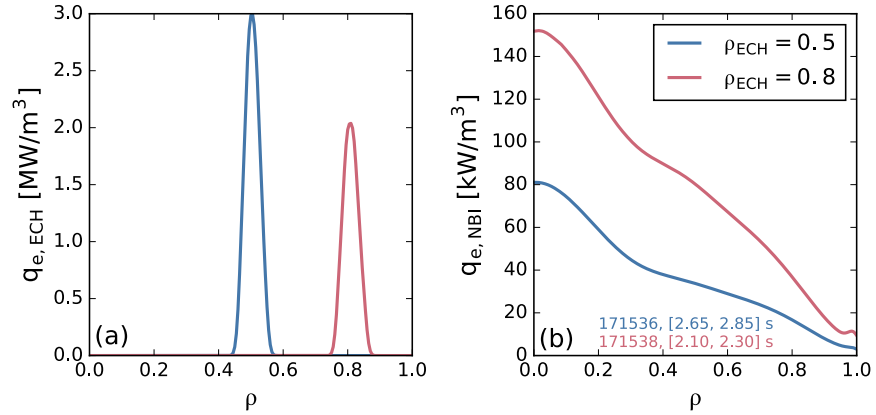
Figure 5.3: (a) ECH electron heating density $q_{e,ECH}$ and (b) NBI electron heating density $q_{e,NBI}$. When moving ECH from $\rho = 0.5$ to $\rho = 0.8$, maintaining constant β_N requires increased NBI heating (see Figure 5.1(c)), which enhances the NBI electron heating density $q_{e,NBI}$ across the full plasma profile. The radiated power and ohmic heating negligibly change between the two discharges.

interferometer. Below, Section 5.3.1 discusses the automated filtering of transient bursts attributable to edge localized modes (ELMs) that would otherwise bias spectral estimates of the background turbulence. Section 5.3.2 then compares the interferometer and PCI frequency spectra for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$; interestingly, the interferometer identifies a novel turbulent branch with properties consistent with a micro-tearing mode (MTM). Next, Section 5.3.3 presents the PCI frequency-wavenumber spectra, which reveal the presence of several distinct turbulent branches. Finally, Section 5.3.4 demonstrates that the PCI wavenumber spectrum distinctly flattens with $\rho_{ECH} = 0.5$ relative to that with $\rho_{ECH} = 0.8$, which is reminiscent of results from realistic multiscale gyrokinetic simulations.

5.3.1 *ELM filtering*

Edge localized modes (ELMs) expel impurities from the plasma but will also present severe challenges to plasma-facing components in future reactors [8, Sec. 7.17]. Because of their virulence and their bursty nature, ELMs produce strong spiking in the interferometer and PCI measurements, whitening the measured spectra [Sec. 10.3.2.3][17]. Additionally, the temperature and density profiles relax during an ELM, altering the turbulent drives in the plasma edge. In order to accurately estimate the spectrum of the background turbulence, then, the ELM
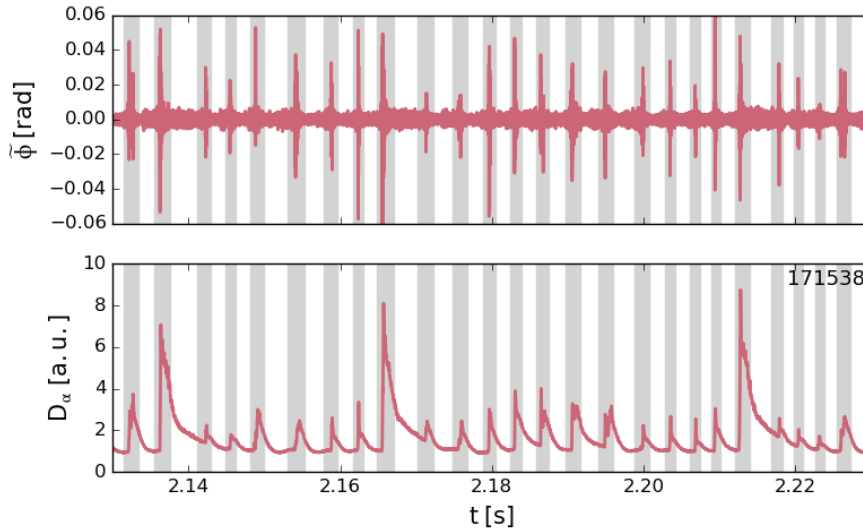
Figure 5.4: Edge localized modes (ELMs) must be removed from the PCI and interferometer measurements prior to spectral analysis of the background turbulence. (Upper panel): The interferometer-measured fluctuating phase $\widetilde{\phi}$, with large, ELM-induced spiking. (Lower panel): Divertor D_α emission, indicating the presence of large Type I ELMs as well as smaller ELMs. Windows *excluded* from spectral analysis are shown in gray. The DIII-D shot number is shown in the upper right of the lower panel.

contributions to the interferometer and PCI measurements must be removed.

In this work, ELMs are simply and automatically detected using measurements from the interferometer. After the high-pass filtering described in Section 4.6.2, the interferometer-measured fluctuating phase $\widetilde{\phi}$ is a zero-mean, random process, as shown in the upper panel of Figure 5.4. Large, intermittent spikes pepper $\widetilde{\phi}(t)$ during ELMy H-mode, and the lower panel of Figure 5.4 indicates that these spikes are well correlated with ELM-induced D_α emission in the divertor. While the D_α emission following large Type I ELMs exhibits a relatively slow decay, the interferometer-measured $\widetilde{\phi}$ returns to stationarity much more rapidly. Thus, it is desirable to identify stationary inter-ELM windows from the interferometer measurements rather than the D_α emission. Points in the interferometer-measured $\widetilde{\phi}$ exceeding $3\times$ the RMS value are identified as ELMs, and successive ELMs are required to be separated by at least a 0.5 ms "debouncing time" (spikes separated by less than the debouncing time are classified as belonging to the same ELM). Subsequent spectral analysis is then performed using only the 20% – 80% inter-ELM windows of the interferometer and PCI measurements. Figure 5.4 shows the windows *excluded* from spectral analysis in gray.

5.3.2 *Frequency spectra*

One-sided autospectral densities estimates $G_{\phi,\phi}(f)$ of the phase fluctuations are calculated using the methodology described in Section F.1. The interferometer and PCI signals from the shaded windows in Figure 5.1 are split into realizations of 1024 points (corresponding to roughly 250 µs) resulting in a frequency resolution of approximately 4 kHz in the spectral estimates. As described in Section 5.3.1, only realizations falling within 20% – 80% of each inter-ELM window are included in the ensemble averaging; the exact number of realizations N_r included in each ensemble depends on the details of the ELM dynamics, but $N_r \sim 450$ for the shots considered here, corresponding to a relative random error in $G_{\phi,\phi}(f)$ of approximately 5%. A Hanning window is applied to each realization prior to computation of its fast Fourier transform (FFT). To simplify inter-ELM bookkeeping, adjacent realizations have zero overlap. As described in Section 4.6.2, the interferometer and PCI phase signals are high-pass filtered prior to spectral analysis, and no further detrending is performed. The resulting spectral estimates for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$ are shown in Figure 5.5. The corresponding noise floors are estimated from 50 ms of data prior to plasma breakdown; the knee in the PCI noise floor at approximately 500 kHz corresponds to the roll-off in the temporal bandwidth of the PCI detector and its preamplifiers. As expected theoretically (see Figure 2.6) and observed empirically in sound-wave calibrations (see Figure 4.22), the PCI is more sensitive than the heterodyne interferometer.

Interestingly, the autospectral density of the heterodyne interferometer indicates the presence of a distinct broadband fluctuation with a central frequency $f_0 \sim 450$ kHz and a bandwidth $\Delta f \sim 300$ kHz. The toroidally separated V2 interferometer corroborates the presence of this fluctuation, but the fluctuation is only vaguely coherent between the two interferometers, with magnitude-squared coherences $\gamma_{xy}^2(f) \leqslant 0.1$. The fluctuation is larger than the corresponding PCI-measured fluctuations in this frequency range, but it is absent from the autospectral density of the PCI; this indicates that the fluctuation wavenumber is smaller than the PCI low-k cutoff (4.2). Further, this fluctuation has a magnetic component, as shown by the beige inset to Figure 5.5(a). The red curve in the beige inset corresponds to the autospectral density (in arbitrary units) of the poloidal magnetic-field fluctuations measured by a high-frequency magnetic probe (b5) [18] during the $\rho_{ECH} = 0.8$ window, while the dashed line corresponds to the noise floor, as estimated from 50 ms of data prior to plasma break-
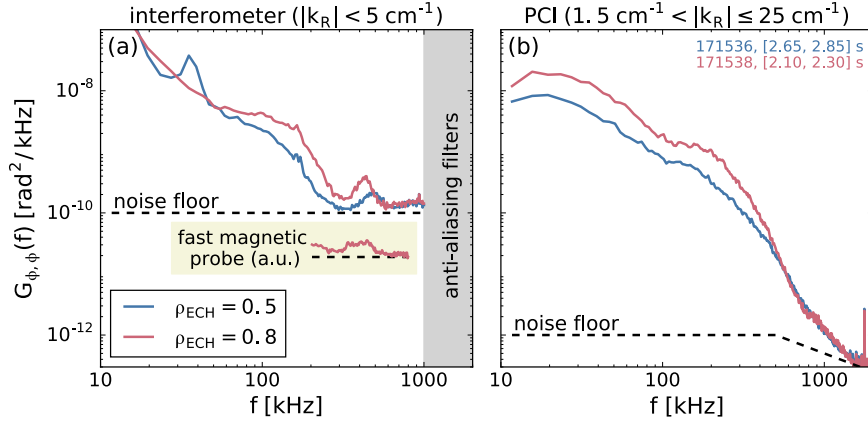
Figure 5.5: One-sided autospectral-density estimates $G_{\phi,\phi}(f)$ from (a) the heterodyne interferometer and (b) the PCI. The low-frequency ($f \lesssim 300\,\text{kHz}$) dynamics are comparable, with both the interferometer and the PCI measuring a $2-4\times$ increase in fluctuation power when moving from $\rho_{\text{ECH}} = 0.5$ to $\rho_{\text{ECH}} = 0.8$. The higher-frequency ($f \gtrsim 300\,\text{kHz}$) dynamics, however, are distinct. In particular, the interferometer-measured broadband fluctuation between $300\,\text{kHz}$ and $600\,\text{kHz}$ is low-k (due to its absence in the PCI spectrum), electromagnetic (as seen by the beige inset to (a)), and driven by collisionality (as shown in Figure 5.7), suggesting that it may be a micro-tearing mode (MTM). The subtle break in slope at $f \sim 800\,\text{kHz}$ in the PCI spectrum when $\rho_{\text{ECH}} = 0.8$ is shown to be a distinct turbulent branch in Figure 5.9.

down. The magnetic autospectral density is estimated in a manner consistent with those of the interferometer and PCI (realization length of approximately $250\,\mu\text{s}$ with zero overlap between adjacent realizations, application of Hanning window to each realization prior to FFT computation, ensemble averaging only over realizations falling within $20\% - 80\%$ of each inter-ELM window, and a total number of realizations $N_r \sim 450$). The autospectral density of the interferometer also indicates that the power in this fluctuation increases when moving from $\rho_{\text{ECH}} = 0.5$ to $\rho_{\text{ECH}} = 0.8$. To prove that this is a robust trend, the total power in this fluctuation is computed for stationary windows from 7 distinct shots with $\rho_{\text{ECH}} = 0.5$ and from 7 distinct shots with $\rho_{\text{ECH}} = 0.8$, each of which are nominally identical to the corresponding discharges shown in Figures 5.1 and 5.2. The total fluctuation power is quantified as

$$\text{var}(\widetilde{\phi}) = \int_{300\,\text{kHz}}^{600\,\text{kHz}} \left[G_{\phi,\phi}^{\text{int}}(f) - \text{N.F.} \right] df, \tag{5.1}$$

where $G_{\phi,\phi}^{\text{int}}(f)$ is the autospectral density of the heterodyne interferometer and $\text{N.F.} = 10^{-10}\,\text{rad}^2 \cdot \text{kHz}^{-1}$ is the corresponding noise floor.
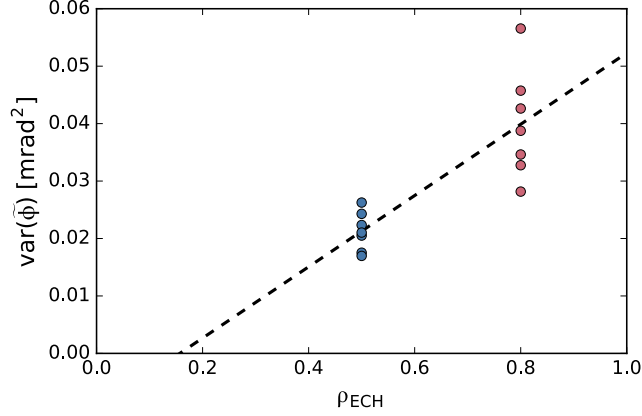
Figure 5.6: Power $\mathrm{var}(\widetilde{\phi})$ in interferometer-measured low-k, mid-f, electromagnetic turbulence vs. ECH location ρ_{ECH}. Powers are estimated via (5.1). Different points correspond to distinct, nominally identical shots from the same experimental run day; the dashed line is the linear least-squares fit to the points, indicating an approximate doubling in fluctuation power when moving from $\rho_{\mathrm{ECH}} = 0.5$ to $\rho_{\mathrm{ECH}} = 0.8$.

Figure 5.6 plots $\mathrm{var}(\widetilde{\phi})$ versus ρ_{ECH}; a linear least-squares fit indicates that the fluctuation power approximately doubles when moving from $\rho_{\mathrm{ECH}} = 0.5$ to $\rho_{\mathrm{ECH}} = 0.8$. Further, Figure 5.7 shows that the electron-ion collisionality ν_{ei} in the region of the plasma accessible to the PCI probe beam (R = 1.98 m) also roughly doubles when moving from $\rho_{\mathrm{ECH}} = 0.5$ to $\rho_{\mathrm{ECH}} = 0.8$. Thus, this fluctuation is low-k, electromagnetic, and driven by collisionality, suggesting that this fluctuation may be a micro-tearing mode (MTM) [8, Sec. 8.5][19], which is predicted to be marginally unstable in this experiment's reference discharge [5].

It is also interesting to examine the low-frequency (f $\lesssim$ 300 kHz) fluctuations measured by both the interferometer and the PCI. Both systems measure a $2 - 4\times$ increase in the low-frequency fluctuation power when moving from $\rho_{\mathrm{ECH}} = 0.5$ to $\rho_{\mathrm{ECH}} = 0.8$, which may be responsible for the slightly reduced confinement in Figure 5.1(g). The spatial content of these fluctuations can be roughly characterized via the magnitude-squared coherence $\gamma_{xy}^2(f)$ between the interferometer (x) and PCI (y) measurements. As discussed in Section F.1, $\gamma_{xy}^2(f)$ is bounded between 0 and 1, and it quantifies the linear correlation between signals x and y, with $\gamma_{xy}^2(f) = 0$ indicating 0% correlation at frequency f and $\gamma_{xy}^2(f) = 1$ indicating 100% correlation at frequency f. Thus, $\gamma_{xy}^2(f)$ characterizes the fraction of fluctuation power sitting in the mid-k overlap of the interferometer and PCI [17, Sec. 5.2.6]. As with the auto-spectral density estimates, the interferometer and
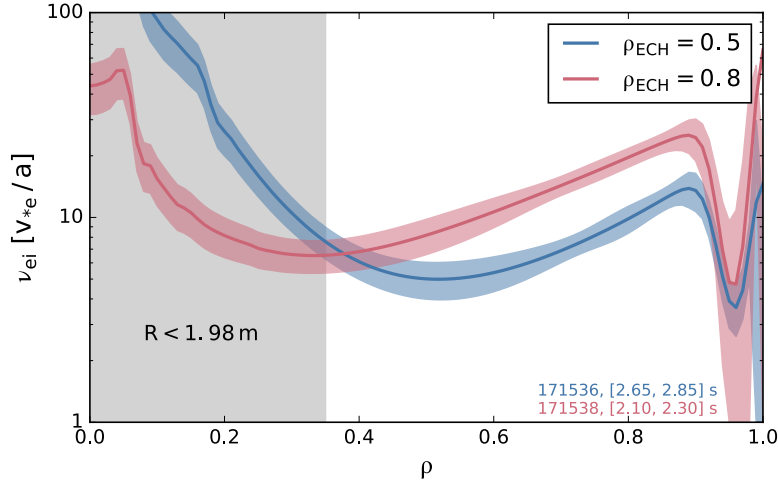
Figure 5.7: The electron-ion collisionality ν_{ei} in the region of the plasma accessible to the PCI probe beam ($R = 1.98$ m) roughly doubles when moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$. The collisionality is computed from the profiles in Figure 5.2, and the shaded bands indicate the 1σ uncertainties. The collisionality normalization is v_{*e}/a, where v_{*e} is the electron diamagnetic velocity and a is the minor radius of the plasma; note that the angular electron diamagnetic frequency $\omega_{*e} \propto v_{*e}/a$ [8, Sec. 8.2] and that micro-tearing mode (MTM) linear-stability calculations are performed in expansions of ν_{ei}/ω_{*e} [8, Sec. 8.5][19]. Unfortunately, without a wavenumber measurement, the more theoretically relevant ω_{*e} normalization cannot be performed. Both discharges are in the banana regime ($\nu_{ei}^* < 1$) [8, Sec. 4.6] for $\rho \lesssim 0.95$.

PCI signals are split into realizations of 1024 points with zero overlap between adjacent realizations, a Hanning window is applied to each realization prior to computing its FFT, the ensemble average is performed only over realizations falling within $20\% - 80\%$ of each inter-ELM window, and the total number of realizations is $N_r \sim 450$. The resulting $\gamma_{xy}^2(f)$ estimates are shown in Figure 5.8. From the definition of the lab-frame phase velocity $v_{ph} = 2\pi f/k$, where f is the lab-frame frequency and k is the wavenumber, it can be seen that the lowest frequencies in a particular fluctuation branch are associated with the lowest wavenumbers, and the highest frequencies in a particular fluctuation branch are associated with the highest wavenumbers. Thus, the low values of $\gamma_{xy}^2(f)$ for $f \lesssim 100$ kHz in Figure 5.8 can be interpreted as the interferometer measuring substantial power in low-k fluctuations that sit below the PCI's low-k cutoff (4.2). Similarly, the low values of $\gamma_{xy}^2(f)$ for $f \gtrsim 300$ kHz correspond to the PCI measuring substantial power in high-k fluctuations that sit above the interferometer's high-k cutoff (4.6) or below the interferometer's noise floor. Between 100 kHz and 300 kHz, $0.25 \lesssim \gamma_{xy}^2 \lesssim 0.5$, indicating between 25% and 50% of
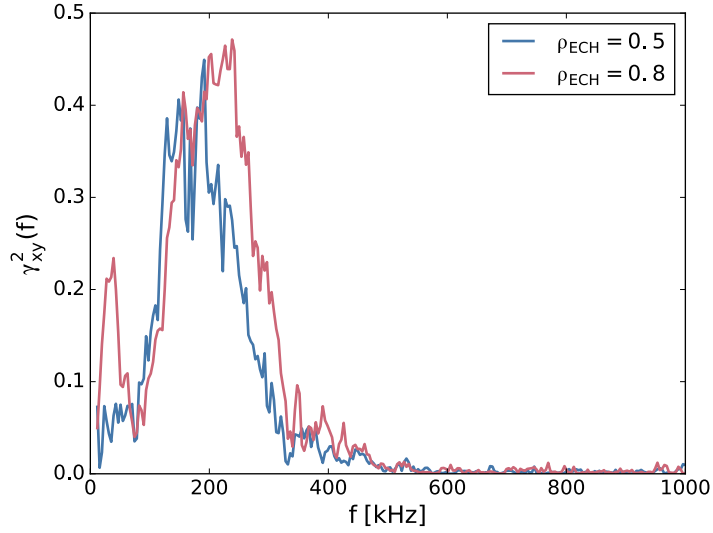
Figure 5.8: The magnitude-squared coherence $\gamma^2_{xy}(f)$ between the interferometer and the PCI measurements characterizes the fraction of total fluctuation power sitting in the mid-k overlap of the interferometer and the PCI. The increase in $\gamma^2_{xy}(f)$ between 200 kHz and 300 kHz when moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$ suggests a wavenumber downshift in these fluctuations, which is corroborated by the increase in the lab-frame phase velocity shown in Figure 5.9.

the total fluctuation power sits in the mid-k overlap of the interferometer and the PCI. Further, moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$ produces a substantial increase in $\gamma^2_{xy}(f)$ for $200 \, \text{kHz} \lesssim f \lesssim 300 \, \text{kHz}$, suggesting that the wavenumbers of these fluctuations decrease when moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$; this is corroborated by the increase in the lab-frame phase velocity shown in Figure 5.9.

5.3.3 *Frequency-wavenumber spectra*

As the PCI measurements are made with a multi-element detector array, the spatial content of the PCI signal can also be characterized. The two-dimensional autospectral density $S_{\phi,\phi}(k, f)$ simultaneously quantifies the spatial and temporal content of a signal. (Unfortunately, as the interferometer measurements are currently made with a single-element detector, $S_{\phi,\phi}(k, f)$ cannot be estimated from the interferometer measurements; the only spatial knowledge about the interferometer signal is that the wavenumber magnitudes |k| sit below the interferometer's high-k cutoff (4.6)).

The PCI $S_{\phi,\phi}(k, f)$ is estimated using the hybrid non-parametric-in-time, parametric-in-space technique discussed in Section F.3.3. Specif-

ically, the hybrid autocorrelation function $\widetilde{R}_{\phi,\phi}(\delta,f)$ is estimated non-parametrically over the shaded windows in Figure 5.1 via (F.22) using the same spectral-estimation parameters as those in Section 5.3.2 (i.e. realization length of 1024 points with zero overlap between adjacent realizations, application of Hanning window to each realization prior to FFT computation, ensemble averaging only over realizations falling within 20% – 80% of each inter-ELM window, and a total number of realizations $N_r \sim 450$). Then, a parametric $p = 6$ Burg autoregression (AR) in the spatial lag δ estimates the autospectral density $S_{\widetilde{R},\widetilde{R}}(k,f)$ of $\widetilde{R}(\delta,f)$, and the autospectral density $S_{\phi,\phi}(k,f)$ of the phase fluctuations is computed via (F.28). The Burg AR is evaluated on a uniformly spaced, 1000-element wavenumber grid. Often, an AR model is referred to as an "all-pole model" because the only frequency dependence in the spectral estimate appears in the denominator; this all-pole feature of an AR model allows fitting very sharp spectral features and substantially improves the wavenumber resolution of the sparsely sampled (in space) PCI data (see Figure F.3 for a comparison to a conventional Fourier-in-space estimate).

The PCI $S_{\phi,\phi}(k,f)$ for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$ are shown in Figure 5.9. The spectra indicate the existence of numerous turbulent branches, each of which may respond differently when ρ_{ECH} is varied. Two of these branches are discussed here. First, perhaps one of the most surprising qualitative effects of moving $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$ is the excitation of a new, distinct turbulent mode at $k_R \sim -5\,cm^{-1}$ with a central frequency $f_0 \sim 1\,MHz$ and a bandwidth $\Delta f \sim 400\,kHz$. This mode corresponds to the subtle break in slope at $f \sim 800\,kHz$ in the PCI $G_{\phi,\phi}(f)$ in Figure 5.9. This mode straddles the current spatiotemporal limits of the interferometer and sits well below the interferometer noise floor; thus, this mode is invisible to the interferometer. The mode is robustly observed in all 7 of the steady-state $\rho_{ECH} = 0.8$ windows from this experimental run day, and it is robustly absent from all 7 of the corresponding $\rho_{ECH} = 0.5$ windows. It should be noted that the PCI observes a similar mode during ELM-free H-mode and wide-pedestal QH-mode [20] as well as during NBI-only ELMy H-mode (see Figure F.3). Second, and perhaps of more relevance to multiscale studies, is the turbulent branch bounded by the annotating dashed white lines in Figure 5.9. When $\rho_{ECH} = 0.5$, this turbulent branch has a lab-frame phase velocity $v_{ph} = 5.6\,km \cdot s^{-1}$ with wavenumbers and frequencies extending to $14\,cm^{-1}$ and $1250\,kHz$, respectively; moving to $\rho_{ECH} = 0.8$ increases the lab-frame phase velocity to $v_{ph} = 6.5\,km \cdot s^{-1}$ but reduces the spatiotemporal bandwidth to $7.2\,cm^{-1}$ and $750\,kHz$. These observations are generic to all 7 of
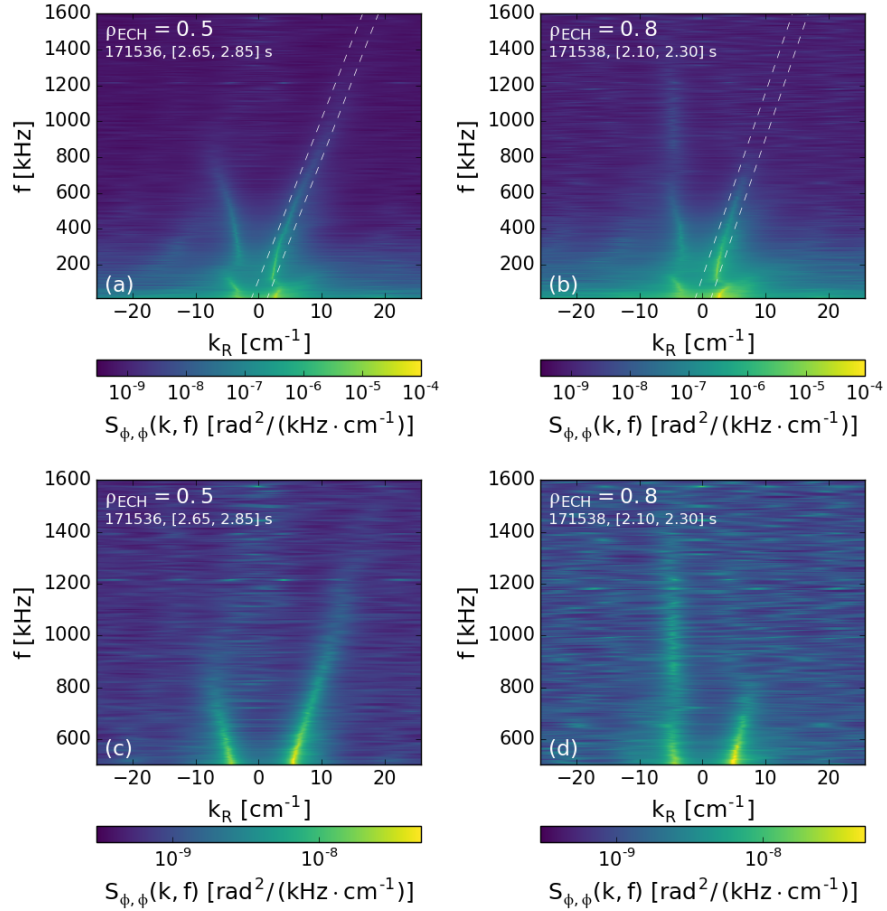
Figure 5.9: PCI two-dimensional autospectral densities $S_{\phi,\phi}(k, f)$ for (a) $\rho_{ECH} = 0.5$ and (b) $\rho_{ECH} = 0.8$. To better resolve the high-k and high-f features, (c) and (d) display the same spectra from (a) and (b), respectively, but only for $f \geqslant 500\,\text{kHz}$ and with an altered colorscale. Moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$ increases the lab-frame phase velocity but decreases the spatiotemporal bandwidth of the broadband turbulence bounded by the annotating dashed white lines and also excites a new turbulent branch at $k \sim -5\,\text{cm}^{-1}$ and $f \sim 1\,\text{MHz}$.

the steady-state $\rho_{ECH} = 0.5$ windows and all 7 of the steady-state $\rho_{ECH} = 0.8$ windows from this experimental run day. Note that for a given frequency f, an increase in phase velocity corresponds to a wavenumber downshift such that the observed change in v_{ph} with ρ_{ECH} confirms the wavenumber-downshift speculation from Figure 5.8. This turbulent branch will be investigated further in Section 5.3.4.
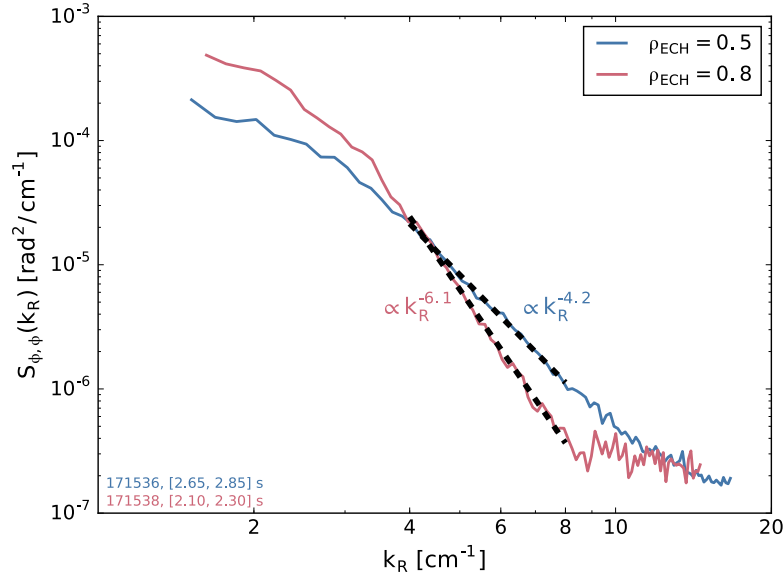
Figure 5.10: PCI wavenumber autospectral densities $S_{\phi,\phi}(k)$ obtained by integrating $S_{\phi,\phi}(k, f)$ in frequency between the bounds denoted by the annotating dashed white lines in Figure 5.9. Here, the black dashed lines indicate least-squares power-law fits to $S_{\phi,\phi}(k)$. The flattening (i.e. slower decay) of $S_{\phi,\phi}(k)$ with $\rho_{ECH} = 0.5$ relative to that with $\rho_{ECH} = 0.8$ is reminiscent of results from multiscale gyrokinetic simulations in Alcator C-Mod, but quantitative comparisons await completion of the corresponding multiscale gyrokinetic simulations in DIII-D.

5.3.4 *Wavenumber spectra*

Because simulations are finely sampled in space but often have limited temporal histories, experimental wavenumber spectra are particularly valuable for validating theoretical models and computational predictions. The PCI wavenumber autospectral density $S_{\phi,\phi}(k)$ corresponding to a given turbulent branch is computed by integrating the two-dimensional autospectral density $S_{\phi,\phi}(k, f)$ from Section 5.3.3 over the temporal bandwidth Δf of the turbulent branch, i.e.

$$S_{\phi,\phi}(k) = \int_{f_0(k)-(\Delta f/2)}^{f_0(k)+(\Delta f/2)} S_{\phi,\phi}(k, f) df, \tag{5.2}$$

where $f_0(k) = k v_{ph}/(2\pi)$ is the central frequency of the turbulent branch at wavenumber k, and v_{ph} is the lab-frame phase velocity of the turbulent branch. Note that restricting the integration domain in this manner minimizes contributions to $S_{\phi,\phi}(k)$ from noise outside of the branch's bandwidth. For example, integrating between the bounds denoted by the annotating dashed white lines in Figure 5.9 produces

the $S_{\phi,\phi}(k)$ displayed in Figure 5.10. Note that the $\rho_{ECH} = 0.5$ spectrum is substantially flattened (i.e. decays more slowly) relative to the $\rho_{ECH} = 0.8$ spectrum. As shown in Figure 5.2, relative to $\rho_{ECH} = 0.8$, $\rho_{ECH} = 0.5$ corresponds to increased a/L_{T_e} and decreased a/L_{T_i} over the majority of the plasma accessible to the PCI probe beam ($R = 1.98\,m$). This spectral flattening with enhanced electron-scale drive relative to ion-scale drive is reminiscent of results from multiscale gyrokinetic simulations [3] of an ion cyclotron resonance heated (ICRH) [8, Sec. 5.8] L-mode discharge in Alcator C-Mod [8, Sec. 11.5]. While the only existing, realistic, multiscale gyrokinetic simulations of DIII-D correspond to the reference discharge for the multiscale experiment discussed here, the $E \times B$ shearing rate γ_E was varied rather than a/L_{T_e} or a/L_{T_i} [5]. Thus, the spectral flattening observed in Figure 5.10 with increased a/L_{T_e} and decreased a/L_{T_i} eagerly awaits quantitative comparison to corresponding multiscale gyrokinetic simulations, which are expected to be completed by Howard *et al.* in roughly the next six months.

5.4 TGLF MODELING

To aid the interpretation of the combined PCI-interferometer measurements described in Section 5.3, linear-stability analysis and quasilinear-transport modeling were performed with the TGLF code [1]. Below, Section 5.4.1 provides a brief overview of the TGLF code, and Section 5.4.2 presents a global overview of the TGLF-predicted linear stability for $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$. In order to facilitate comparisons between theory and measurement, Section 5.4.3 then derives the relationship between the theoretically relevant field-aligned wavenumber k_y and the PCI-measured major-radial wavenumber k_R. Next, in an attempt to localize analysis of the TGLF results, Section 5.4.4 compares the advection of the TGLF-predicted instabilities by the $E \times B$ velocity to the PCI-measured phase velocities; the comparison suggests that the PCI-measured branch of interest is localized to $0.6 \leqslant \rho \leqslant 0.65$. Finally, Section 5.4.5 compares the TGLF-predicted electron-density fluctuation spectrum to the corresponding PCI measurements, finding qualitative agreement.

---

1 TGLF simulations were graciously performed by Dr. Alessandro Marinoni, but the subsequent analysis is the author's own.

5.4.1 *TGLF overview*

TGLF is a gyro-Landau fluid (GLF) model that captures the dynamics of both trapped (T) and passing particles in shaped geometry with finite aspect ratio and $E \times B$ shear [21, 22]. A GLF model consists of velocity-moment equations of the gyroaveraged kinetic equation with moment closures selected to retain important kinetic effects, such as Landau damping. Due to their reduced dimensionality, GLF models are computationally less expensive than full gyrokinetic simulations. TGLF can solve for the linear eigenmodes of trapped ion and electron modes (TIM and TEM, respectively), ion and electron temperature gradient (ITG and ETG, respectively) modes, and kinetic ballooning modes (KBM). Unfortunately, the default eigenfunction basis of four Hermite polynomials is typically insufficient to resolve micro-tearing modes (MTMs) [23], so there are no attempts to simulate the MTM in this work. TGLF additionally uses its eigenmodes to predict quasi-linear transport using a saturation model fit to results from nonlinear gyrokinetic simulations [22], the most recent of which include realistic multiscale physics [24].

5.4.2 *Global overview of predicted linear stability*

Using the TGLF_scan module in the OMFIT integrated modeling framework [25], TGLF simulations were run to investigate the change in linear stability when moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$. The simulations span $0.1 \leqslant k_y \rho_s \leqslant 24$ and $0.3 \leqslant \rho \leqslant 0.9$ (only $\rho \gtrsim 0.35$ is accessible to the PCI probe beam, which propagates vertically through the plasma at major radius $R = 1.98 \, \mathrm{m}$). The equilibrium profiles shown in Figure 5.2 were used as input to TGLF. Figure 5.11 displays the resulting linear growth rates γ and plasma-frame phase velocities $v_{ph} = \omega/k_y$ as a function of radial position ρ and normalized fluctuation wavenumber $k_y \rho_s$. The predicted change in linear stability when moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$ is largely in accord with intuition derived from the changes to the normalized inverse scale lengths shown in Figure 5.2(e)-(g). Perhaps the most substantial change is that $\rho_{ECH} = 0.5$ destabilizes a continuum of mid-k ($0.5 \lesssim k_y \rho_s \lesssim 5$) electron modes in the outer region ($\rho \gtrsim 0.6$) of the plasma. Additionally, in the outer region of the plasma ($\rho \gtrsim 0.6$), the low-k ion modes ($k_y \rho_s \lesssim 0.3$) are marginally suppressed and the high-k electron modes ($k_y \rho_s \gtrsim 10$) are marginally enhanced relative to $\rho_{ECH} = 0.8$.
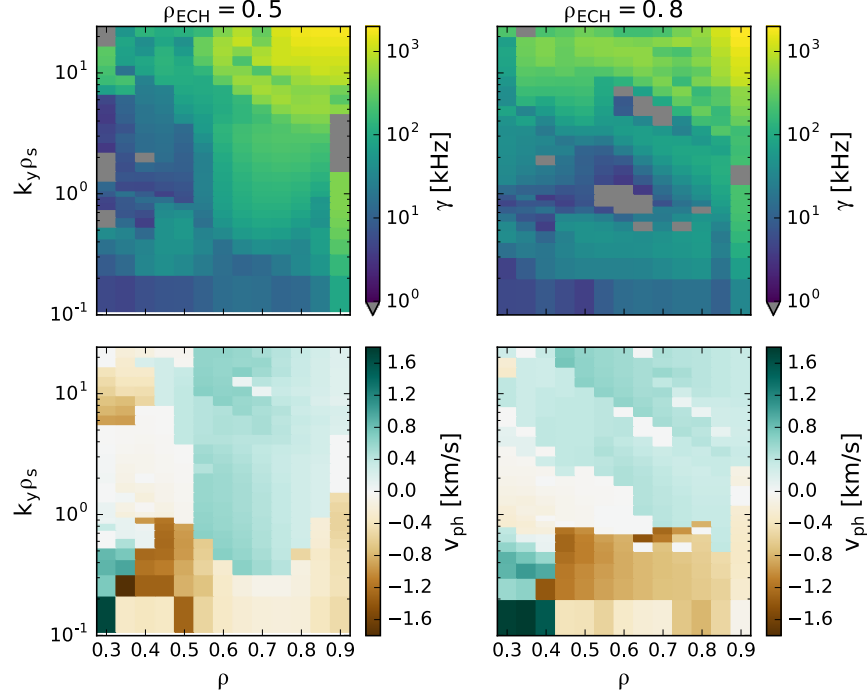
Figure 5.11: Global overview of TGLF-predicted linear stability for the profiles shown in Figure 5.2. Linear growth rates γ and plasma-frame phase velocities $v_{ph} = \omega/k_y$ are plotted as a function of radial position ρ and normalized fluctuation wavenumber $k_y \rho_s$ for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$. Electron modes have $v_{ph} > 0$, and ion modes have $v_{ph} < 0$. The predicted change in linear stability when moving from $\rho_{ECH} = 0.5$ to $\rho_{ECH} = 0.8$ is largely in accord with intuition derived from the changes to the normalized inverse scale lengths shown in Figure 5.2(e)-(g). Growth rates can be compared to the $E \times B$ shearing rates γ_E in Figure 5.2(h).

5.4.3 *Relation between field-aligned & PCI-measured wavevectors*

As transport parallel to the magnetic field $\boldsymbol{B}$ is much more rapid than transport across the magnetic field, drift-wave turbulent fluctuations have a dominant, central wavevector $\boldsymbol{k}_0$ that satisfies the field-aligned constraint $\boldsymbol{k}_0 \cdot \boldsymbol{B} \approx 0$. For electrostatic fluctuations, the magnetic field $\boldsymbol{B}$ is well-represented by the equilibrium field $\boldsymbol{B}_0$ such that the field-aligned constraint reduces to $\boldsymbol{k}_0 \cdot \boldsymbol{B}_0 \approx 0$, requiring

$$\boldsymbol{k}_0 = k_{\rho,0} \hat{\boldsymbol{\rho}} + k_{\theta,0} \left[\hat{\boldsymbol{\theta}} - \left(\frac{B_{\theta,0}}{B_{\zeta,0}} \right) \hat{\boldsymbol{\zeta}} \right], \tag{5.3}$$

where $k_{\rho,0}$ is the dominant radial wavenumber of the turbulence (note that velocity shear produces finite $k_{\rho,0}$ [26]), $k_{\theta,0}$ is the dominant poloidal wavenumber of the turbulence, $B_{\theta,0}$ is the equilibrium poloidal
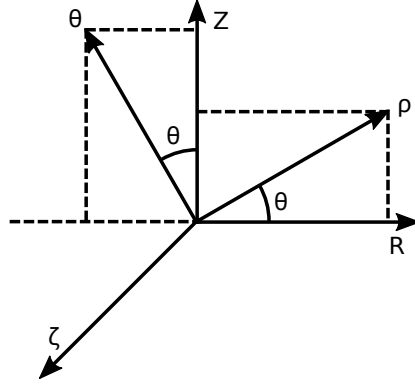
Figure 5.12: Measurement coordinate system. Here, R is the major-radial direction, Z is the lab-frame vertical direction, ζ is the toroidal angle, θ is the poloidal angle, and ρ is a flux-surface label corresponding to the square root of the normalized toroidal magnetic-field flux.

magnetic field, $B_{\zeta,0}$ is the equilibrium toroidal magnetic field, and the (ρ, θ, ζ) coordinate system is defined in Figure 5.12. Analytic theory and computation are often performed on field-aligned coordinate systems in which the z-coordinate is along $\boldsymbol{B}_0$, the x-coordinate is in the radial direction, and the y-coordinate is perpendicular to $\boldsymbol{B}_0$ but within a flux surface. In such a field-aligned coordinate system, $\boldsymbol{k}_0 - k_{\rho,0}\hat{\boldsymbol{\rho}}$ is aligned with the y-coordinate such that one is led naturally to define

$$\boldsymbol{k}_y = k_{\theta,0} \left[\hat{\boldsymbol{\theta}} - \left(\frac{B_{\theta,0}}{B_{\zeta,0}} \right) \hat{\boldsymbol{\zeta}} \right]. \tag{5.4}$$

The wavenumber k_y (or its dimensionless equivalent, $k_y\rho_s$) is an important parameter in the classification and understanding of turbulent fluctuations.

To make contact with theory, then, it is important to understand the relation between k_y and the PCI-measured wavevectors. As a line-integrated measurement, PCI is sensitive to fluctuations with wavevectors $\boldsymbol{k}_{\mathrm{pci}}$ that are perpendicular to the beam path. Thus, for the vertical beam path of DIII-D's PCI,

$$\boldsymbol{k}_{\mathrm{pci}} \cdot \hat{\boldsymbol{Z}} = 0. \tag{5.5}$$

As with $\boldsymbol{k}_0$, $\boldsymbol{k}_{\mathrm{pci}}$ is also field-aligned such that (for electrostatic fluctuations)

$$\boldsymbol{k}_{\mathrm{pci}} \cdot \boldsymbol{B}_0 \approx 0. \tag{5.6}$$

In general, $k_{\mathrm{pci}} \neq k_0$. Instead, it is suitable to define

$$k_{\mathrm{pci}} = k_0 + \delta k. \tag{5.7}$$

The PCI will measure finite signal only if there exists a δk within the spatial bandwidth Δk of the turbulence such that the resulting k_{pci} from (5.7) satisfies both (5.5) and (5.6). (Note that the above reasoning also holds for any other vertically line-integrated measurement of field-aligned turbulence, such as those made by the heterodyne interferometer constructed in this work). Decomposing δk in the (ρ, θ, ζ) coordinate system defined in Figure 5.12 such that $\delta k = \delta k_\rho \hat{\boldsymbol{\rho}} + \delta k_\theta \hat{\boldsymbol{\theta}} + \delta k_\zeta \hat{\boldsymbol{\zeta}}$, constraint (5.5) requires

$$\delta k_\rho = -\delta k_{\rho,0} - (k_{\theta,0} + \delta k_\theta) \cot \theta, \tag{5.8}$$

and constraint (5.6) requires

$$\delta k_\zeta = -\left(\frac{B_{\theta,0}}{B_{\zeta,0}}\right) \delta k_\theta. \tag{5.9}$$

Using (5.8) and (5.9), k_{pci} from (5.7) reduces to

$$k_{\mathrm{pci}} = -\left(\frac{k_{\theta,0} + \delta k_\theta}{\sin \theta}\right) \left[\hat{\boldsymbol{R}} + \left(\frac{B_{\theta,0} \sin \theta}{B_{\zeta,0}}\right) \hat{\boldsymbol{\zeta}}\right], \tag{5.10}$$

where $\hat{\boldsymbol{R}}$ is the unit vector in the major-radial direction and is related to $\hat{\boldsymbol{\rho}}$ and $\hat{\boldsymbol{\theta}}$ as defined in Figure 5.12. Because $|B_{\theta,0}| \ll |B_{\zeta,0}|$, the PCI is sensitive to fluctuations with wavevectors that are predominantly in the major-radial direction. If measurements are made with a one-dimensional detector array, only a projection of k_{pci} can be reconstructed. The one-dimensional detector array of the DIII-D PCI is nominally aligned in the major-radial direction such that the major-radial wavenumber k_R of k_{pci} can be reconstructed, i.e.

$$k_R = k_{\mathrm{pci}} \cdot \hat{\boldsymbol{R}} = -\left(\frac{k_{\theta,0} + \delta k_\theta}{\sin \theta}\right). \tag{5.11}$$

In subsequent numerical evaluations, it is assumed that $|\delta k_\theta| \ll |k_{\theta,0}|$ in order to relate the field-aligned k_y from TGLF to the PCI-measured major radial wavenumber k_R. Figure 5.13 displays profiles of $|k_y \rho_s|$ for various k_R of interest.
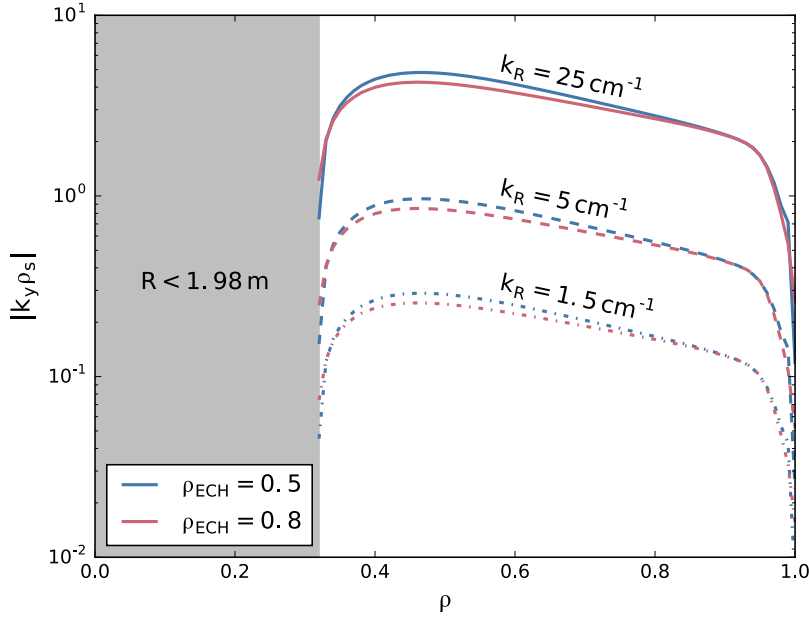
Figure 5.13: Profiles of $|k_y \rho_s|$ vs. radial coordinate ρ for the PCI low-k cutoff $k_R = 1.5\,\mathrm{cm}^{-1}$, the interferometer high-k cutoff $k_R = 5\,\mathrm{cm}^{-1}$, and the PCI high-k cutoff $k_R = 25\,\mathrm{cm}^{-1}$ in the $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$ discharges shown in Figure 5.2. The gray region ($\rho \lesssim 0.35$) is inaccessible to the PCI probe beam, which propagates vertically through the plasma at major radius $R = 1.98\,\mathrm{m}$. The roll-off in $|k_y \rho_s|$ for $\rho \lesssim 0.4$ is attributable to the $\sin\theta$ major-radial projection, while the roll-off for $\rho \gtrsim 0.95$ is attributable to the edge pedestal.

5.4.4 *Comparison to PCI-measured phase velocities*

In an attempt to localize analysis of the TGLF results, the advection of the TGLF-predicted instabilities by the $E \times B$ velocity can be compared to the PCI-measured phase velocities. The PCI-measured phase velocity is

$$v_{ph}^{pci} = \frac{\omega_{pci}}{k_R} = \frac{\boldsymbol{k}_{pci} \cdot \boldsymbol{v}}{k_R}, \tag{5.12}$$

where ω_{pci} is the PCI-measured angular frequency, $\boldsymbol{k}_{pci}$ is the wavevector (5.10), $\boldsymbol{v}$ is the lab-frame velocity of the fluctuation, and k_R is the reconstructed major-radial wavenumber (5.11). By field-aligned constraint (5.6), $\boldsymbol{k}_{pci} \cdot \boldsymbol{v} = \boldsymbol{k}_{pci} \cdot \boldsymbol{v}_\perp$, where $\boldsymbol{v}_\perp$ is the velocity perpendicular to the equilibrium magnetic field. The perpendicular velocity $\boldsymbol{v}_\perp$ is simply the sum of the $E \times B$ velocity $\boldsymbol{v}_E$ and the plasma-frame phase velocity v_{ph}^{plasma} of the fluctuation; here, the plasma-frame

phase velocity is approximated by that of the dominant wavevector, i.e. $\boldsymbol{v}_{\text{ph}}^{\text{plasma}} \approx v_{\text{ph}}^{\text{plasma}} \hat{\boldsymbol{k}}_0$, where $\hat{\boldsymbol{k}}_0 = \boldsymbol{k}_0/|\boldsymbol{k}_0|$ and $\boldsymbol{k}_0$ is the dominant wavevector defined in (5.3). Thus, $\boldsymbol{v}_\perp \approx \boldsymbol{v}_{\text{E}} + v_{\text{ph}}^{\text{plasma}} \hat{\boldsymbol{k}}_0$, and

$$\omega_{\text{pci}} = \boldsymbol{k}_{\text{pci}} \cdot \boldsymbol{v}_\perp \approx \boldsymbol{k}_{\text{pci}} \cdot \boldsymbol{v}_{\text{E}} + (\boldsymbol{k}_{\text{pci}} \cdot \hat{\boldsymbol{k}}_0) v_{\text{ph}}^{\text{plasma}}. \tag{5.13}$$

Now, the electrostatic potential $\varphi = \varphi(\rho)$ is a flux function such that the corresponding electric field is $\boldsymbol{E} = -\nabla\varphi = E_r(\rho, \theta)\hat{\boldsymbol{\rho}}$. The resulting $E \times B$ velocity is

$$\boldsymbol{v}_{\text{E}} = \frac{\boldsymbol{E} \times \boldsymbol{B}_0}{B_0^2} = \frac{E_r(\rho, \theta)}{B_0^2} \left(B_\theta \hat{\boldsymbol{\zeta}} - B_\zeta \hat{\boldsymbol{\theta}} \right). \tag{5.14}$$

Using $\boldsymbol{k}_{\text{pci}}$ from (5.10), $\boldsymbol{v}_{\text{E}}$ from (5.14), $\boldsymbol{k}_0$ from (5.3), ω_{pci} from (5.13), k_R from (5.11), and a bit of algebra, the PCI-measured phase velocity (5.12) readily reduces to

$$v_{\text{ph}}^{\text{pci}} = \left[\frac{E_r(\rho, \theta)}{B_0} - v_{\text{ph}}^{\text{plasma}} \right] \left(\frac{B_0}{B_{\zeta,0}} \right) \sin\theta. \tag{5.15}$$

The sign convention for the plasma-frame phase velocity is such that $v_{\text{ph}}^{\text{plasma}} > 0$ for electron modes and $v_{\text{ph}}^{\text{plasma}} < 0$ for ion modes. For physical intuition regarding this sign convention, consider a region of the plasma with $E_r > 0$; here, electron modes propagate against the $E \times B$ direction (decreasing $v_{\text{ph}}^{\text{pci}}$ relative to the $E \times B$ contribution alone), and ion modes propagate in the $E \times B$ direction (increasing $v_{\text{ph}}^{\text{pci}}$ relative to the $E \times B$ contribution alone). The $B_0/B_{\zeta,0}$ multiplicative enhancement to $v_{\text{ph}}^{\text{pci}}$ results from $|k_R| < |\boldsymbol{k}_{\text{pci}}|$ when the magnetic field is not solely toroidal, and the $\sin\theta$ term provides the major-radial projection required by the constraint (5.5); note that $v_{\text{ph}}^{\text{pci}}$ changes sign about the midplane ($\theta = 0$) due to this major-radial projection. As the spatially filtering mask [27, 28] was not used in this experiment, the PCI measurements cannot be localized to above or below the midplane, and only the magnitude of $v_{\text{ph}}^{\text{pci}}$ will be considered here. Figure 5.14 compares the magnitude of the predicted PCI phase velocity from (5.15) to the measured PCI phase velocities from the annotated turbulent branches in Figure 5.9 and infers that this turbulence is localized to $0.6 \leqslant \rho \leqslant 0.65$.

As a brief aside, it should be emphasized that the radial electric field E_r in (5.15) is *not* a flux function. To see this, recall that the corresponding electrostatic potential *is* a flux function, i.e. $\varphi = \varphi(\rho) =$
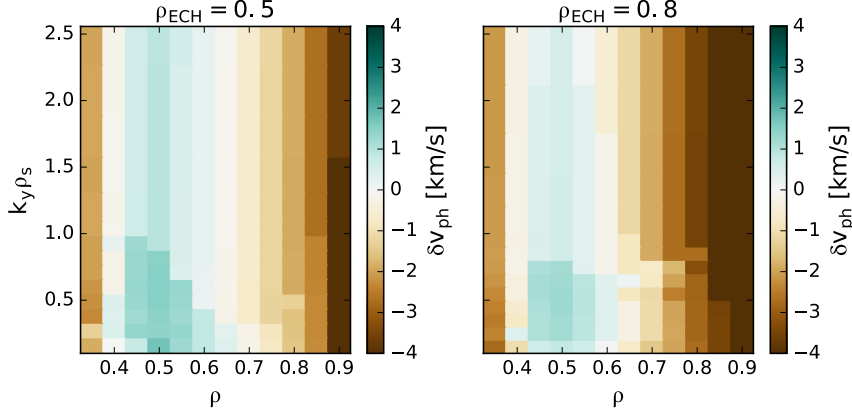
Figure 5.14: Comparison of predicted and measured PCI phase velocities. Here, δv_{ph} is the difference between the predicted phase velocity and the measured phase velocity. The predicted phase velocity (5.15) is computed using the equilibrium profiles shown in Figure 5.2 and the TGLF-predicted plasma-frame phase velocites shown in Figure 5.11. The measured phase velocities correspond to the annotated turbulent branches in Figure 5.9. The wavenumber range $k_y \rho_s$ considered here is restricted to $k_R \leqslant 15 \, cm^{-1}$, the maximum PCI-measured wavenumber in Fig 5.9. The minimum discrepancy between the predicted and measured PCI phase velocities occurs between at $0.6 \leqslant \rho \leqslant 0.65$, suggesting that the annotated turbulent branches in Figure 5.9 are localized to $0.6 \leqslant \rho \leqslant 0.65$.

$\varphi(\psi)$, where ψ is the flux-surface label corresponding to the poloidal magnetic-field flux per radian. Now,

$$
\begin{aligned}
E_r(\rho, \theta) &= -\left(\frac{\partial \varphi}{\partial r}\right) \\
&= -\left(\frac{d\varphi}{d\psi}\right)\left(\frac{\partial \psi}{\partial r}\right) \\
&= -\left(\frac{d\varphi}{d\psi}\right)(RB_\theta),
\end{aligned} \tag{5.16}
$$

where the last line follows from the definition of ψ as the poloidal magnetic-field flux per radian. The derivative $d\varphi/d\psi$ is a flux function because φ is a flux function, and this implies that $E_r/(RB_\theta)$ is also a flux function. Thus, the radial electric field at any point within the last closed flux surface can be computed from the radial electric field along the outboard midplane (where $\theta = 0$) as follows

$$
\left.\frac{E_r}{RB_\theta}\right|_{\rho,\theta} = \left.\frac{E_r}{RB_\theta}\right|_{\rho,\theta=0}. \tag{5.17}
$$

5.4.5   *Quantitative local results*

A global overview of the TGLF-predicted linear stability was presented in Section 5.4.2. In Section 5.4.4, the annotated turbulent branches from Figure 5.9 were localized to $0.6 \leqslant \rho \leqslant 0.65$. Thus, it is now reasonable to perform a more quantitative assessment of the corresponding local TGLF predictions. Below, TGLF predictions are shown for $\rho = 0.6$, but comparable predictions are made at $\rho = 0.65$.

Figure 5.15 displays the TGLF-predicted linear stability at radial location $\rho = 0.6$ for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$. Both operational scenarios are predicted to destabilize ion and electron modes across multiple spatiotemporal scales, but $\rho_{ECH} = 0.5$ is predicted to destabilize a continuum of mid-k ($0.5 \lesssim k_y \rho_s \lesssim 5$) electron modes that bridge the gap between low-k ($k_y \rho_s \lesssim 0.5$) ion modes and high-k ($k_y \rho_s \gtrsim 5$) electron modes, potentially facilitating cross-scale coupling. Further, relative to $\rho_{ECH} = 0.8$, the $\rho_{ECH} = 0.5$ ion-mode growth rates are substantially reduced (becoming comparable to the $E \times B$ shearing rate shown in Figure 5.2(h)), and the high-k electron-mode growth rates are marginally enhanced, both of which may also facilitate cross-scale coupling.

Howard *et al.* suggest a linear-stability "rule of thumb" for gauging the importance of cross-scale coupling [3]. The rule of thumb attempts to quantify the relative "strength" of electron-scale to ion-scale turbulence by examining the ratio of the maximum electron-scale growth rate γ_{high-k} to the maximum ion-scale growth rate γ_{low-k}, with larger values of $\gamma_{high-k}/\gamma_{low-k}$ corresponding to increased cross-scale coupling. Howard *et al.* constrain their search for γ_{high-k} to $2 \leqslant k_y \rho_s \leqslant 48$ and for γ_{low-k} to $0.25 \leqslant k_y \rho_s \leqslant 0.75$. For consistency, this same selection criterion is used below. For $\rho_{ECH} = 0.8$, Figure 5.15 shows that $\gamma_{high-k} = 540 \, \text{kHz}$ at $k_y \rho_s = 14$ and $\gamma_{low-k} = 43 \, \text{kHz}$ at $k_y \rho_s = 0.4$ such that $\gamma_{high-k}/\gamma_{low-k} = 12.6$; with such a small value, the rule of thumb suggests that cross-scale coupling is insignificant. For $\rho_{ECH} = 0.5$, Figure 5.15 shows that $\gamma_{high-k} = 720 \, \text{kHz}$ at $k_y \rho_s = 14$ and $\gamma_{low-k} = 63 \, \text{kHz}$ at $k_y \rho_s = 0.7$ such that $\gamma_{high-k}/\gamma_{low-k} = 11.4$, less than the $\rho_{ECH} = 0.8$ case. It should be noted, however, that the distinct mid-k electron branch in Figure 5.15 is *absent* from the plasmas used to develop the rule of thumb; further, the above γ_{low-k} for $\rho_{ECH} = 0.5$ corresponds to this mid-k electron branch. If γ_{low-k} is instead restricted to an ion mode, $\gamma_{low-k} = 17 \, \text{kHz}$ at $k_y \rho_s = 0.35$ such that $\gamma_{high-k}/\gamma_{low-k} = 42$; cross-scale coupling can be important in ELMy H-mode plasmas with this $\gamma_{high-k}/\gamma_{low-k}$ [4].
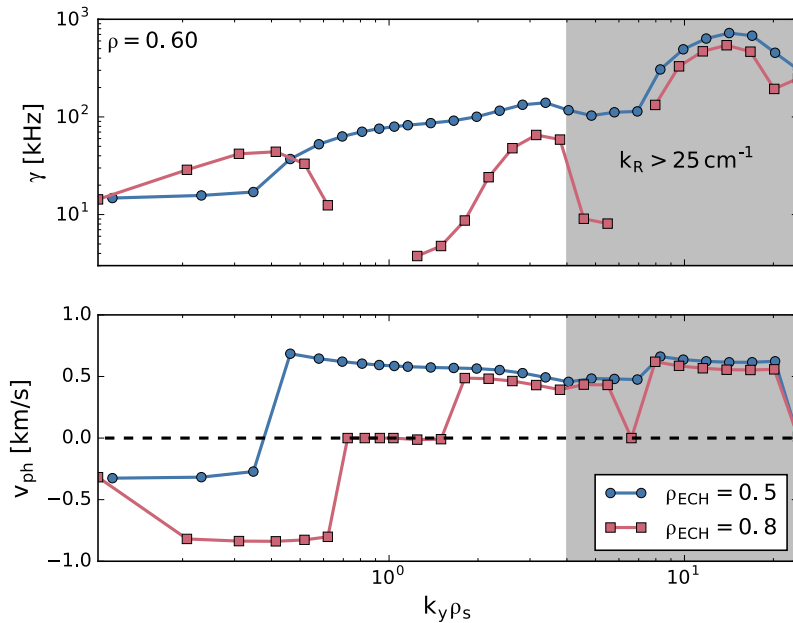
Figure 5.15: TGLF-predicted linear stability at $\rho = 0.6$ for the profiles shown in Figure 5.2. Linear growth rates γ and plasma-frame phase velocities $v_{ph} = \omega/k_y$ are plotted as a function of normalized fluctuation wavenumber $k_y \rho_s$ for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$. Electron modes have $v_{ph} > 0$, and ion modes have $v_{ph} < 0$. The gray region exceeds the spatial bandwidth of the combined PCI-interferometer developed in this work. Growth rates can be compared to the E × B shearing rates γ_E in Figure 5.2(h).

In addition to assessing linear stability, TGLF can predict the electron-density fluctuation spectrum using quasilinear transport and a model for the nonlinear saturation of the turbulence. This saturation model is fit to results from nonlinear gyrokinetic simulations [22], the most recent of which include realistic multiscale physics [24]. The multiscale saturation model is referred to as SAT1. The density fluctuation spectrum predicted by TGLF-SAT1 at $\rho = 0.6$ is shown Figure 5.16. Within the PCI sensitivity and wavenumber domain, the predicted spectra are qualitatively consistent with the PCI-measured spectra in Figure 5.10. Specifically, the $\rho_{ECH} = 0.5$ spectrum is noticeably flatter (i.e. decays more slowly) than the $\rho_{ECH} = 0.8$ spectrum, which may be indicative of enhanced cross-scale coupling [3]. Note that the PCI spectrum also depends on $k_{\rho,0}$ and Δk, as discussed in Section 5.4.3, so quantitative comparison with the PCI measurements requires a full-physics simulation (i.e. beyond the abilities of a reduced model like TGLF) and a synthetic diagnostic [29]. It should also be emphasized that the TGLF-SAT1 model is fit to just a handful [24] of realistic multiscale gyroki-
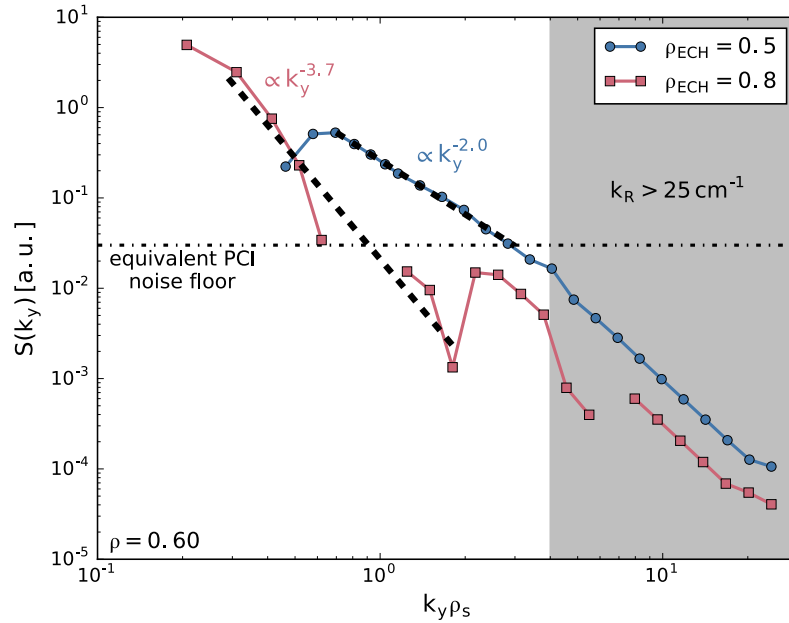
Figure 5.16: Electron-density fluctuation spectra predicted by TGLF-SAT1 at radial location $\rho = 0.6$ for $\rho_{ECH} = 0.5$ and $\rho_{ECH} = 0.8$. The gray region exceeds the spatial bandwidth of the combined PCI-interferometer developed in this work. The equivalent PCI noise floor, indicated by the horizontal dash-dot line, is inferred from the noise floor in Figure 5.10; that is, for $\rho_{ECH} = 0.5$ the equivalent noise floor intersects $S(k_y)$ at $k_y \rho_s \approx 2.7$, and $k_y \rho_s \approx 2.7 \leftrightarrow k_R \approx 15\,\mathrm{cm}^{-1}$ at $\rho = 0.6$, which is where the corresponding $S(k_R)$ intersects the noise floor in Figure 5.10. Within the PCI sensitivity and wavenumber domain, the predicted spectra are qualitatively consistent with the PCI-measured spectra in Figure 5.10. The black dashed lines indicate least-squares power-law fits to the predicted $S(k_R)$. Note that the PCI spectrum also depends on $k_{\rho,0}$ and Δk, as discussed in Section 5.4.3, so quantitative comparison with the PCI measurements requires a full-physics simulation and a synthetic diagnostic.

netic simulations corresponding to Alcator C-Mod L-mode plasmas, so the model may improve when additional multiscale simulations are completed and the TGLF-SAT1 model is recalibrated.

[1] N. T. Howard, C. Holland, A. E. White, M. Greenwald and J. Candy. Synergistic cross-scale coupling of turbulence in a tokamak plasma. *Phys. Plasmas*, 21(11):112510, 2014.

[2] 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. *Nucl. Fusion*, 56(1):014004, 2016.

[3] N. T. Howard, C. Holland, A. E. White, M. Greenwald, J. Candy et al. Multi-scale gyrokinetic simulations: comparison with experiment and implications for predicting turbulence and transport. *Phys. Plasmas*, 23(5):056109, 2016.

[4] N T Howard, C Holland, A E White, M Greenwald, P Rodriguez-Fernandez et al. Multi-scale gyrokinetic simulations of an Alcator C-Mod, ELM-y H-mode plasma. *Plasma Phys. Control. Fusion*, 60(1):014034, 2018.

[5] C. Holland, N.T. Howard and B.A. Grierson. Gyrokinetic predictions of multiscale transport in a DIII-D ITER baseline discharge. *Nucl. Fusion*, 57(6):066043, 2017.

[6] W. Horton. Drift waves and transport. *Rev. Mod. Phys.*, 71(3):735–778, 1999.

[7] G. R. Tynan, A. Fujisawa and G. McKee. A review of experimental drift turbulence studies. *Plasma Phys. Control. Fusion*, 51(11):113001, 2009.

[8] J. Wesson. *Tokamaks*. Oxford University Press, Oxford, UK, 4th edition, 2011.

[9] International Atomic Energy Agency. *Fusion Physics*. International Atomic Energy Agency, Vienna, 2012.

[10] W. Dorland, F. Jenko, M. Kotschenreuther and B. N. Rogers. Electron temperature gradient turbulence. *Phys. Rev. Lett.*, 85(26):5579–5582, 2000.

[11] F. Jenko and W. Dorland. Prediction of significant tokamak turbulence at electron gyroradius scales. *Phys. Rev. Lett.*, 89(22):225001, 2002.

[12] A.M. Garofalo, W.M. Solomon, J.-K. Park, K.H. Burrell, J.C. DeBoo et al. Advances towards QH-mode viability for ELM-stable operation in ITER. *Nucl. Fusion*, 51(8):083018, 2011.

[13] L. L. Lao, H. E. St. John, Q. Peng, J. R. Ferron, E. J. Strait et al. MHD equilibrium reconstruction in the DIII-D tokamak. *Fusion Sci. Tech.*, 48(2):968–977, 2005.

[14] W. W. Pfeiffer, R. H. Davidson, R. L. Miller and R. Waltz. ONETWO: A computer code for modeling plasma transport in tokamaks. Technical Report GA-A16178, General Atomics, 2009.

[15] Alexei Pankin, Douglas McCune, Robert Andre, Glenn Bateman and Arnold Kritz. The tokamak Monte Carlo fast ion module NUBEAM in the National Transport Code Collaboration library. *Comp. Phys. Comm.*, 159(3):157 – 184, 2004.

[16] K. Matsuda. Ray tracing study of the electron cyclotron current drive in DIII-D using 60 GHz. *IEEE Trans. Plasma Sci.*, 17(1):6–11, 1989.

[17] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 4th edition, 2010.

[18] E. J. Strait. Magnetic diagnostic system of the DIII-D tokamak. *Rev. Sci. Instrum.*, 77(2):023502, 2006.

[19] J. F. Drake and Y. C. Lee. Kinetic theory of tearing instabilities. *Phys. Fluids*, 20(8):1341–1353, 1977.

[20] J. C. Rost. Personal communication.

[21] G. M. Staebler, J. E. Kinsey and R. E. Waltz. Gyro-Landau fluid equations for trapped and passing particles. *Phys. Plasmas*, 12(10):102508, 2005.

[22] G. M. Staebler, J. E. Kinsey and R. E. Waltz. A theory-based transport model with comprehensive physics. *Phys. Plasmas*, 14(5):055909, 2007.

[23] G. M. Staebler. Personal communication.

[24] G.M. Staebler, N.T. Howard, J. Candy and C. Holland. A model of the saturation of coupled electron and ion scale gyrokinetic turbulence. *Nucl. Fusion*, 57(6):066046, 2017.

[25] O. Meneghini, S.P. Smith, L.L. Lao, O. Izacard, Q. Ren et al. Integrated modeling applications for tokamak experiments with OMFIT. *Nucl. Fusion*, 55(8):083008, 2015.

[26] J. C. Rost, M. Porkolab, J. Dorris and K. H. Burrell. Short wavelength turbulence generated by shear in the quiescent H-mode edge on DIII–D. *Phys. Plasmas*, 21(6):062306, 2014.

[27] J. R. Dorris, J. C. Rost and M. Porkolab. Localized measurement of short wavelength plasma fluctuations with the DIII-D phase contrast imaging diagnostic. *Rev. Sci. Instrum.*, 80(2):023503, 2009.

[28] J. R. Dorris. *Phase contrast imaging measurements and modeling of short wavelength turbulence in the DIII-D tokamak*. PhD thesis, MIT, 2010.

[29] J. C. Rost, L. Lin and M. Porkolab. Development of a synthetic phase contrast imaging diagnostic. *Phys. Plasmas*, 17(6):062506, 2010.

6

CORRELATION OF DIII-D'S TOROIDALLY SEPARATED INTERFEROMETERS

The toroidal structure of an MHD mode can have important implications for the mode's stability and its interaction with the surrounding plasma. A mode's toroidal structure is typically characterized by its toroidal mode number n. Historically, measurement of toroidal (and poloidal) mode numbers with external magnetic probes has provided rich insight into the physics governing numerous operational regimes and stability limits. However, core-localized MHD produces weak signals outside of the plasma volume, making measurement of the corresponding mode numbers via external magnetic probes difficult or impossible. Recently, measurements from toroidally separated electron cyclotron emission imaging (ECEI) systems on the KSTAR tokamak have identified mode numbers of sawteeth [1], demonstrating the utility of using more exotic measurements to probe the structure of core-localized MHD.

This chapter describes the correlation of toroidally separated interferometers to measure toroidal mode numbers. To the author's knowledge, this is the first such implementation in a tokamak. As the interferometers are capable of probing the plasma core, their correlation allows direct measurement of the toroidal structure of core-localized modes — a long-sought after and first-of-its-kind measurement at DIII-D. Below, Section 6.1 reviews the two-point correlation technique, which allows inference of a propagating wave's spatial structure from measurements made at two distinct spatial locations. Section 6.2 then examines the geometry of the interferometers and develops a formula for the measured toroidal mode number. Section 6.3 describes the careful efforts to eliminate timebase discrepancies between the two interferometer systems, validates the interferometer-measured toroidal mode numbers against those measured by external magnetic probes, and discusses the effect of the interferometers' radial offset. Finally, Section 6.4 provides an encouraging proof of principle regarding the ability of the correlated interferometers to diagnose core-localized MHD.

6.1   TWO-POINT CORRELATION

Two-point correlation allows inference of a propagating wave's spatial structure from measurements made at two distinct spatial locations. Below, Section 6.1.1 details the two-point correlation technique. Section 6.1.2 then discusses the aliasing of large wavenumbers and establishes a measurement's Nyquist wavenumber, below which wavenumbers are *not* aliased. Finally, Section 6.1.3 converts the results of the previous sections into their toroidal-mode-number equivalents, as is conventional for fluctuation characterization in a tokamak.

6.1.1   *Wavenumber measurement via two-point correlation*

Consider a 1-dimensional, coherent density fluctuation with amplitude $\widetilde{n}_0$, wavenumber k, and angular frequency ω

$$\widetilde{n}(z, t) = \widetilde{n}_0 e^{i(kz - \omega t)}. \tag{6.1}$$

Imagine that the density is measured at two locations separated by distance Δz, producing two time series, $x(t)$ and $y(t)$, given as

$$x(t) = \widetilde{n}(z_0, t) \tag{6.2}$$

$$y(t) = \widetilde{n}(z_0 + \Delta z, t) = x(t) e^{ik\Delta z}. \tag{6.3}$$

The cross phase α_{xy} between these two time series is

$$\alpha_{xy} = \arg\left[x^*(t) \cdot y(t)\right] \tag{6.4}$$

$$= k\Delta z, \tag{6.5}$$

where x^* is the complex conjugate of x. Thus, a measured wavenumber k_{meas} can be inferred from the cross phase via

$$k_{meas} = \frac{\alpha_{xy}}{\Delta z}. \tag{6.6}$$

The cross phase α_{xy} of $x(t)$ and $y(t)$ can be readily estimated via the non-parametric, FFT-based spectral-estimation techniques discussed in Appendix F.1. Note that such cross-phase estimates are frequency-resolved, allowing simultaneous characterization of multiple fluctuations at distinct frequencies.

6.1.2 *Aliasing & the Nyquist wavenumber*

The measured wavenumber (6.6) may be *aliased*. To see this, note that the cross-phase estimate α_{xy} is only unique modulo 2π (that is, 0 is equivalent to $\pm 2\pi$, $\pm 4\pi$, etc.). If the fluctuation wavenumber k is sufficiently large (i.e. if it exceeds the so-called Nyquist wavenumber k_{Ny}), this cross-phase ambiguity aliases the measured wavenumber (6.6) away from the true wavenumber.

   In the most general case, the wave's propagation direction is not known *a priori*, and both positive and negative wavenumbers should be considered (i.e. $-\pi < \alpha_{xy} \leqslant \pi$). This cross-phase domain yields a Nyquist wavenumber

$$k_{Ny} = \frac{\pi}{\Delta z}, \qquad \text{for } \textit{unknown} \text{ propagation direction.} \tag{6.7}$$

Wavenumber measurements (6.6) from fluctuations with $|k| > k_{Ny}$ are aliased, while measurements from fluctuations with $|k| \leqslant k_{Ny}$ are not aliased. Note that (6.7) is equivalent to the famed Nyquist frequency: making the transformations $\Delta z \to 1/f_s$ for temporal sampling rate f_s and $k \to 2\pi f$, (6.7) readily becomes $f_{Ny} = f_s/2$. Thus, $1/\Delta z$ can be thought of as the spatial sampling rate, with a larger sampling rate (i.e. smaller Δz) allowing un-aliased measurements of larger wavenumbers.

   Now, if the wave's propagation direction is known *a priori* (for example, if the wave propagation is dominated by advection, and the fluid velocity is well-diagnosed), only a single polarity of wavenumbers need to be considered. For concreteness, positive wavenumbers are considered below (i.e. $0 \leqslant \alpha_{xy} < 2\pi$). This cross-phase domain yields a Nyquist wavenumber

$$k_{Ny} = \frac{2\pi}{\Delta z}, \qquad \text{for } \textit{known} \text{ propagation direction.} \tag{6.8}$$

6.1.3 *Measurement of toroidal mode numbers*

Fluctuations in a torus are often characterized by their toroidal n and poloidal m mode numbers, both of which are constrained to be integers by the toroidal and poloidal periodicities, respectively, of the torus. For a torus with major radius R, the toroidal mode number n is related to the toroidal wavenumber k_ζ as

$$k_\zeta = \frac{n}{R}, \tag{6.9}$$

and the toroidal angular separation $\Delta\zeta$ is related to the spatial separation Δz as

$$\Delta\zeta = \frac{\Delta z}{R}. \tag{6.10}$$

Using these definitions, the above formulas for the measured wavenumber and the Nyquist wavenumber can be rewritten in terms of toroidal mode numbers as follows:

$$n_{\text{meas}} = \frac{\alpha_{xy}}{\Delta\zeta}, \tag{6.11}$$

$$n_{\text{Ny}} = \begin{cases} \frac{\pi}{\Delta\zeta} & \text{for } \textit{unknown} \text{ propagation direction} \\ \frac{2\pi}{\Delta\zeta} & \text{for } \textit{known} \text{ propagation direction} \end{cases}. \tag{6.12}$$

6.2 TOROIDAL CORRELATION OF INTERFEROMETERS

This section applies two-point correlation to the measurement of toroidal mode numbers with DIII-D's toroidally separated, heterodyne CO_2 interferometers. Section 6.2.1 describes the geometry of the interferometer probe beams, which establishes the Nyquist toroidal mode number. Sections 6.2.2 and 6.2.3 then examine the relationship between the interferometer-measured phase fluctuations and the toroidal mode number. As the interferometers are capable of probing the plasma core, their correlation allows direct measurement of the toroidal structure of core-localized modes — a long-sought after and first-of-its-kind measurement at DIII-D.

6.2.1 *DIII-D's interferometers*

DIII-D's multichannel, two-color, heterodyne CO_2 interferometer measures both the equilibrium [2] and fluctuating [3, 4] components of the line-integrated electron density. Each channel is configured as a Michelson interferometer, with each probe beam making a double-pass through the plasma. The three vertical chords pass through the V1, V2, and V3 ports at a toroidal location of 240°, while the radial chord passes through the R0 port at a toroidal location of 225°. Of particular relevance to this work is the V2 interferometer, which has a major-radial location R = 1.94 m and is shown in Figure 6.1. The viewing geometry influences an interferometer's sensitivity to various MHD instabilities; for example, vertical chords are more sensitive to
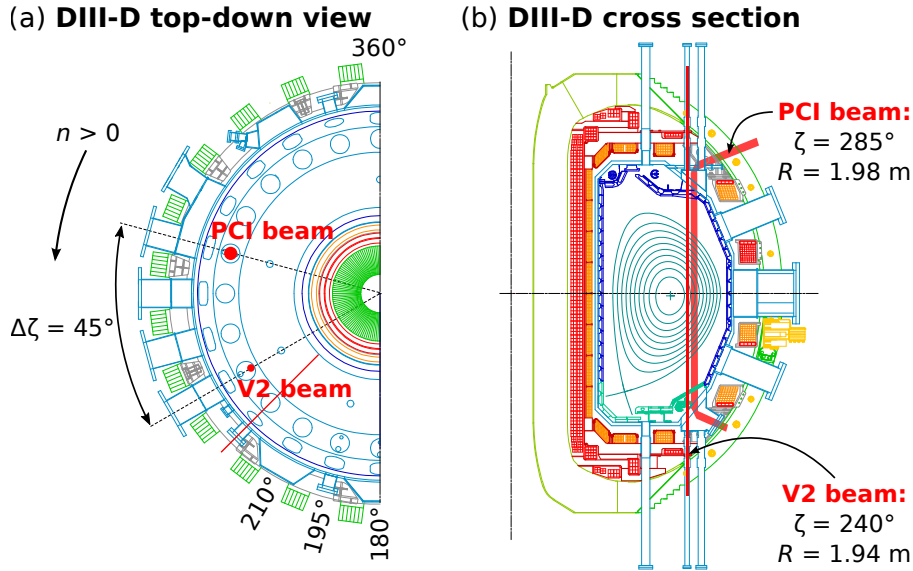
Figure 6.1: (a) A top-down view of the DIII-D vessel. The toroidal location of the V2 interferometer beam is $\zeta = 240°$, and the toroidal location of the PCI beam is $\zeta = 285°$. The DIII-D sign convention for toroidal mode numbers is that $n > 0$ for modes that rotate counterclockwise when viewing the torus from above, as this corresponds to the direction of dominant torque injection. (b) A poloidal cross section of the DIII-D vessel. The major-radial location of the V2 interferometer beam is $R = 1.94 \, \text{m}$, and the major-radial location of the PCI beam is $R = 1.98 \, \text{m}$.

toroidal Alfvén eigenmodes (TAEs), while radial chords are more sensitive to reversed-shear Alfvén eigenmodes (RSAEs) [3].

The addition of a heterodyne-interferometer channel to DIII-D's pre-existing phase contrast imaging (PCI) system is discussed extensively in Chapter 4 and elsewhere [5], but the details of relevance to the toroidal-correlation measurement are briefly reviewed here. The PCI probe beam sits at a toroidal location of 285° and has a major-radial location $R = 1.98 \, \text{m}$. The location of the PCI beampath relative to that of the V2 interferometer is shown in Figure 6.1.

The geometry of the V2 and PCI interferometers has consequences for the toroidal-correlation measurement. The interferometers are toroidally spaced by $\Delta\zeta = 45°$ such that the Nyquist toroidal mode number (6.12) becomes

$$n_{Ny} = \begin{cases} 4 & \text{for } \textit{unknown} \text{ propagation direction} \\ 8 & \text{for } \textit{known} \text{ propagation direction} \end{cases} . \qquad (6.13)$$

The DIII-D sign convention for toroidal mode numbers is indicated in Figure 6.1; that is, $n > 0$ for modes that rotate counterclockwise when viewing the torus from above, as this corresponds to the direction of

dominant torque injection. Consistency with this sign convention requires that $x(t)$ corresponds to the V2 interferometer signal and that $y(t)$ corresponds to the PCI interferometer signal when computing the toroidal mode number via (6.11). Finally, the V2 and PCI interferometer beam paths have a slight radial offset ($\Delta R = 4 \, \text{cm}$ with $R_{v2} = 1.94 \, \text{m}$ and $R_{pci} = 1.98 \, \text{m}$); the consequences of this offset are discussed in Section 6.3.4.

For completeness, it should also be mentioned that a three-chord, radially viewing Faraday-effect polarimeter-interferometer was recently installed on DIII-D's 285° R0 port [6]. While diagnosing core-localized magnetic fluctuations is the primary impetus for this installation, the system also measures line-integrated electron-density fluctuations. Presumably, these measurements could be correlated with those from the 225° R0 CO_2 heterodyne interferometer. However, these two systems do not have phase-locked sampling rates (the necessity of which is discussed in Section 6.3.1), so no attempts to correlate their interferometric measurements are performed in this work.

6.2.2 *Interferometer-measured phase fluctuations*

Consider a tokamak electron density fluctuation with complex amplitude $\tilde{n}_0(r)$, toroidal mode number n, poloidal mode number m, and angular frequency ω

$$\tilde{n}_e(\boldsymbol{r}, t) = \tilde{n}_0(r) e^{i(n\zeta + m\theta - \omega t)}. \tag{6.14}$$

The phase fluctuation (2.31) imparted to a CO_2 probe beam propagating *vertically* through $\tilde{n}_e(\boldsymbol{r}, t)$ becomes

$$\begin{aligned} \tilde{\phi} &= -r_e \lambda_0 \int \tilde{n}_e(\boldsymbol{r}, t) \, dz. \\ &= \Phi e^{i(n\zeta - \omega t)}, \end{aligned} \tag{6.15}$$

where

$$\Phi = -r_e \lambda_0 \int \tilde{n}_0(r) e^{im\theta} \, dz \tag{6.16}$$

is a complex-valued function of the beam's major-radial location and the radial and poloidal mode structure. For a given mode, the V2 and PCI interferometers see the same radial and poloidal mode structure such that Φ effectively reduces to a one-dimensional function $\Phi = \Phi(R)$. Further, Φ can be written explicitly as a complex value $\Phi =$

$|\Phi|e^{i\sigma}$. Thus, the phase fluctuation (6.15) can be written alternatively as

$$\widetilde{\phi} = |\Phi(R)|e^{i[n\zeta - \omega t + \sigma(R)]}, \tag{6.17}$$

where the dependence on the major radius R of the beam has been noted explicitly.

6.2.3 *Interferometer-measured toroidal mode number*

Define $x(t)$ and $y(t)$ to be the V2-measured and PCI-measured phase fluctuations, respectively; i.e.

$$x(t) = \widetilde{\phi}_{v2}(t) = |\Phi(R_{v2})|e^{i[n\zeta_{v2} - \omega t + \sigma(R_{v2})]}, \tag{6.18}$$

$$y(t) = \widetilde{\phi}_{pci}(t) = |\Phi(R_{pci})|e^{i[n\zeta_{pci} - \omega t + \sigma(R_{pci})]}. \tag{6.19}$$

Then the cross phase (6.4) becomes

$$\alpha_{xy} = n\Delta\zeta + \left[\sigma(R_{pci}) - \sigma(R_{v2})\right], \tag{6.20}$$

where $\Delta\zeta = \zeta_{pci} - \zeta_{v2} = 45°$ is the toroidal separation between the interferometer beams. The cross phase α_{xy} of $x(t)$ and $y(t)$ can be readily estimated via the non-parametric, FFT-based spectral-estimation techniques discussed in Appendix F.1, but $\sigma(R_{pci})$ and $\sigma(R_{v2})$ are not typically known *a priori*. Nonetheless, the measured toroidal mode number n_{meas} is defined as

$$n_{meas} = \frac{\alpha_{xy}}{\Delta\zeta}, \qquad \text{where } \Delta\zeta = 45°. \tag{6.21}$$

However, if $\sigma(R_{pci}) \neq \sigma(R_{v2})$, the measured mode number will not be equal to the true mode number (i.e. $n_{meas} \neq n$); this is discussed further in Section 6.3.4. Further, if the true mode number exceeds the Nyquist mode number (6.13), the measured mode number will be aliased away from the true mode number (i.e. $n_{meas} \neq n$).

6.3 IMPLEMENTATION DETAILS AND NON-IDEAL EFFECTS

This section discusses various implementation details and non-ideal effects regarding the toroidal correlation of the V2 and PCI interferometers. Specifically, Section 6.3.1 describes the modifications to the V2 and PCI digitizers that now enable phase-locked measurements. Section 6.3.2 unveils the presence of a deleterious "trigger offset" between the phase-locked systems but also demonstrates an easy and ro-

bust methodology for estimating and compensating this offset in post-processing software. Section 6.3.3 then compares the interferometer-measured toroidal mode numbers with those measured by DIII-D's array of midplane magnetic probes, typically resulting in good agreement for large, robust tearing modes. Finally, Section 6.3.4 discusses how the $\Delta R = 4\,cm$ major-radial offset between the V2 and PCI interferometers can bias the measured toroidal mode numbers; eliminating this radial offset is not possible with the current port allocations, but future work to empirically or computationally account for the radial and poloidal mode structures may aid the interpretation of the interferometer-measured toroidal mode numbers.

6.3.1 *Phase-locked sampling*

Extracting useful information from the correlation of two measurements requires that the sampling rates of the two measurements are *phase-locked*. If the sampling rates are *not* phase-locked, slippage between the sampling times will result in artificial evolution of the measured phase.

Digitizing the two signals on a common digitizer is the simplest method for ensuring phase-locked sampling. Unfortunately, such an approach is not suitable for the V2 and PCI interferometer signals. The V2 interferometer's $\Delta f_0 = 40\,MHz$ intermediate-frequency signal is demodulated using an all-digital technique that mandates use of a sampling rate $f_s = (4/3)\Delta f_0$ [7], and the baseband I and Q signals never exist in analog form. In contrast, as described in Chapter 4 and [5], the PCI interferometer's $\Delta f_0 = 30\,MHz$ intermediate-frequency signal is demodulated with analog electronics, and the analog baseband I and Q signals are then digitized on two channels of a D-tAcq ACQ216 CPCI board. Because the V2 interferometer's baseband I and Q signals never exist in analog form, it is not possible to digitize the V2 I and Q signals on the digitizer used by the PCI interferometer. Further, because of the intermediate-frequency mismatch between the V2 and PCI interferometers, it is also not possible digitally demodulate and digitize the PCI interferometer's 30 MHz intermediate-frequency signal using the V2 interferometer's 40 MHz digital demodulation system. An alternative approach is to directly digitize both intermediate-frequency signals with a high-bandwidth digitizer and demodulate both signals in software, as has been done elsewhere [8]. While DIII-D's ion cyclotron emission (ICE) digitizer has a 200 MSPS sample rate, channels on the ICE digitizer are not consistently available.

As sharing a common digitizer is not possible, phase-locked sampling between the V2 and PCI interferometers requires that the digitizers of both systems derive their clocks from a common source. The V2 interferometer derives its clock from a 320 MHz oven-controlled crystal oscillator (OCXO), and its digital demodulation system has two auxiliary outputs that can be programmed to output phase-locked LVCMOS signals at 320 MHz/N, where $N \in \{1, 2, 3, ..., 32\}$. Both outputs are currently programmed with divisor $N = 20$ such that they each provide a 16 MHz clock signal. One of these 16 MHz clock signals is routed via an RG-58 coaxial cable from the V2 digital demodulation hardware in the first row of the DIII-D annex to the PCI digitizer in the third row of the DIII-D annex.

The PCI digitizer typically samples at $f_s = 4$ MSPS. Thus, division of the V2 16 MHz clock by four (in hardware or software) yields a clock appropriate for typical sampling. Each board of the PCI digitizer can accept an external clock through the front panel's single-pin LEMO CLK input. The CLK signal passes through an optocoupler with a bandwidth ~ 10 MHz [9], but in-house tests have demonstrated that the input clock frequency can actually exceed 16 MHz. As a result, the 16 MHz signal is directly connected to the front-panel LEMO CLK input of the "master" digitizer board ("board 8"), and the necessary division (i.e. divide by four to achieve typical 4 MSPS sampling rate) is performed within the digitizer board, which is cleaner, simpler, and more easily extensible than performing the division in hardware. The resulting 4 MHz clock is routed to the "slave" digitizer board ("board 7") via the PXI backplane. The phase-locked sampling of the V2 and PCI systems has been in place since June 2016.

6.3.2 *Estimating & compensating the "trigger offset"*

As discussed in Appendix G, phase-locked digitizers can still suffer from a so-called "trigger offset", as defined by (G.7). A finite trigger offset results (a) when a digitizer triggers at a time δt later than its nominal trigger time and (b) when the sampling rate deviates from the nominal sampling rate and the nominal trigger times of both digitizers differ.

A finite trigger offset biases mode-number measurements. For example, Figure 6.2(a) displays the toroidal mode numbers of two coherent modes as measured by the uncompensated interferometers. Note that the measured mode number of each mode (unphysically) evolves with the mode frequency, which is consistent (see (G.9)) with the presence of a finite trigger offset between the interferometer measurements.

Referencing (G.11), the trigger offset can be estimated from the measured mode-number evolution during a linear frequency sweep as

$$\delta t_{\text{trig}} = \frac{d\alpha_{xy}}{d\omega} \approx \frac{\Delta\alpha_{xy}}{\Delta\omega} \approx 30\,\mu s, \tag{6.22}$$

where $\Delta\alpha_{xy} = \Delta\zeta \cdot \Delta n_{\text{meas}}$ is the measured phase-angle change and $\Delta\zeta = \pi/4$ is the angular separation of the interferometers (consistent with (6.21)); here, the numerical value results from $\Delta n_{\text{meas}} \approx 5$ and $\Delta\omega \approx 2\pi \cdot 20\,\text{kHz}$, as indicated by the red circular annotations to the lower-frequency mode in Figure 6.2(a). Note that $\delta t_{\text{trig}} > 0$ implies that the actual sampling times of the PCI interferometer lag (i.e. occur later than) those of the V2 interferometer, as shown schematically in Figure G.1. Applying standard techniques [10, Sec. 4.5] in post-processing software to eliminate this 30 μs timebase discrepancy dramatically reduces the unphysical mode-number evolution observed in Figure 6.2(a). Scanning δt_{trig} about 30 μs shows that the unphysical mode-number evolution is minimized for

$$\delta t_{\text{trig}} = 32.5\,\mu s. \tag{6.23}$$

Figure 6.2(b) displays the interferometer-measured toroidal mode numbers after compensating for trigger offset (6.23); note that the mode numbers are constant in time, consistent with physical expectations. It should be emphasized that the only difference between Figure 6.2(a) and Figure 6.2(b) is the compensation of trigger offset (6.23).

Trigger offset (6.23) has been robustly stable since phase-locking the V2 and PCI digitizers in June 2016, a timescale currently exceeding one year. It is interesting to examine the physical origin of this trigger offset to understand when it might change. Referencing the definition of δt_{trig} in (G.7), there are three distinct mechanisms contributing to the trigger offset. The first effect is the difference between the trigger-time discrepancies δt_j of each digitizer; in a properly functioning digitizer, this is discrepancy is typically smaller than one or two sample times (i.e. $\lesssim 1\,\mu s$ for the digitizers considered here), making it an unlikely source of the relatively large trigger offset (6.23). The second effect is the digitizers' normalized sampling-rate deviation $\overline{\delta F_j}$ (phase-locked digitizers have *equal* normalized sampling-rate deviations; see (G.5)). Because the trigger offset (6.23) does *not* appear to vary in time, temporal fluctuations in $\overline{\delta F_j}$ cannot be responsible; however, $\overline{\delta F_j}$ will be constant and finite if the sampling rate has a constant offset from the nominal sampling rate. The third effect is the weighting of $\overline{\delta F_j}$ by the difference in nominal trigger times. Now, the V2 digitizer nomi-
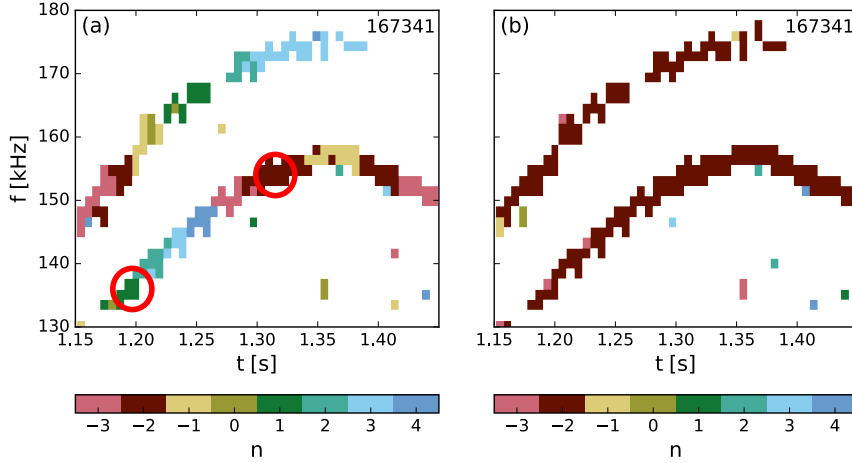
Figure 6.2: Interferometer-measured toroidal mode numbers (a) *without* compensation for the trigger offset and (b) *with* compensation for the trigger offset. In (a), the finite trigger offset (6.23) results in the unphysical evolution of the measured toroidal mode numbers. The trigger offset is roughly estimated from the measured mode-number evolution between the red circular annotations to the lower-frequency mode ($\Delta f \approx 20\,\mathrm{kHz}$ and $\Delta n \approx 5$; note that Δn_meas refers to the total number of mode numbers evolved through, i.e. it is path-dependent and is *not* simply the difference between the final and initial n_meas). In (b), the digital records of both interferometers have been synchronized by compensating for the finite trigger offset (6.23), and the measured toroidal mode numbers are constant in time, consistent with physical expectations. The 6-digit number in the upper right-hand corner indicates the DIII-D shot.

nally triggers at $t_\mathrm{v2}^\mathrm{nom}[0] \approx -1.5\,\mathrm{s}$ relative to plasma breakdown, while the PCI digitizer nominally triggers at $t_\mathrm{pci}^\mathrm{nom}[0] = -0.05\,\mathrm{s}$ relative to plasma breakdown. Solving for the $\overline{\delta F_j}$ in (G.7) and neglecting the small trigger-time discrepancies δt_j yields

$$\overline{\delta F_j} \approx \frac{\delta t_\mathrm{trig}}{t_\mathrm{pci}^\mathrm{nom}[0] - t_\mathrm{v2}^\mathrm{nom}[0]} \approx 2.2 \times 10^{-5}, \tag{6.24}$$

which is not an unreasonable precision for the digitizer sampling rates. Thus, digitizer offset (6.23) appears to result from the combined effects of finite sampling-rate precision (i.e. nonzero $\overline{\delta F_j}$) and the difference between the nominal trigger times of the V2 and PCI digitizers. Importantly, if the nominal trigger times of the V2 or PCI digitizers are altered in the future, the trigger offset will also change and will need to be estimated from new measurements.
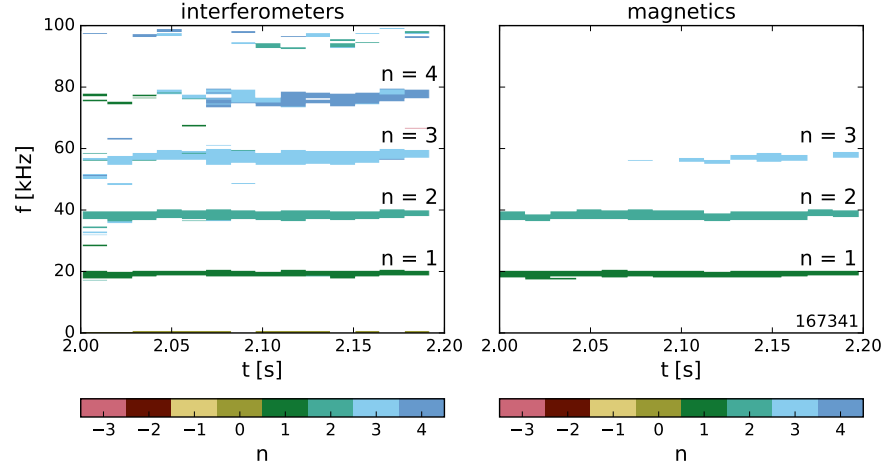
Figure 6.3: A comparison of interferometric (left) and magnetic (right) measurements of toroidal mode numbers. Here, cross-spectral densities are estimated by averaging over $\approx 14\,\mathrm{ms}$ "ensembles", each of which consists of 10 "realizations"; adjacent realizations have 50% overlap, and a Hanning window is applied to each realization prior to performing any spectral computations. Only points with magnitude-squared coherence $\gamma_{xy}^2 \geqslant 0.5$ are displayed in the interferometer spectrum, while only points with coefficient of determination $R^2 \geqslant 0.9$ are displayed in the magnetic spectrum. See Appendix F for background regarding spectral estimation. The 6-digit number in the lower right-hand corner indicates the DIII-D shot.

6.3.3 *Validation against magnetic measurements*

After compensating for the trigger offset (6.23), the interferometer-measured toroidal mode numbers can be validated against those measured by the DIII-D midplane magnetic probes [11]. The magnetic probes measure the poloidal magnetic-field perturbation $\widetilde{B}_\theta$, as eddy currents in the wall tend to reinforce $\widetilde{B}_\theta$ (in contrast, eddy currents tend to shield the radial perturbations $\widetilde{B}_r$). While toroidal mode numbers can be computed via two-point correlation of a single magnetic-probe pair, a more robust estimate can be made by correlating each unique magnetic-probe pair and then least-squares fitting the resulting cross-phase estimates vs. the corresponding toroidal separations to a linear model; all of the magnetics-measured toroidal mode numbers presented in this work are computed via this more robust estimation method. Figure 6.3 compares the interferometer-measured and the magnetics-measured toroidal mode numbers; both measurements indicate the presence of an $n = 1$ tearing mode and several higher order harmonics. The good agreement between the interferometric and magnetic measurements is typical for large, robust tearing modes.
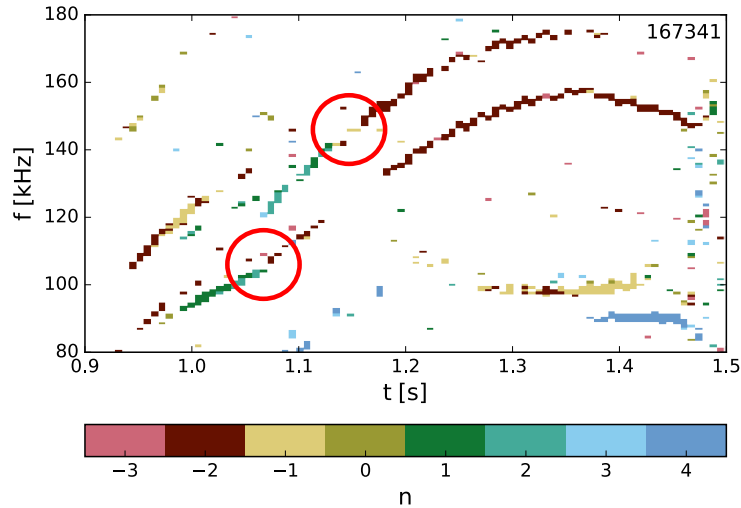
Figure 6.4: When the radial mode structure evolves, the $\Delta R = 4\,cm$ major-radial offset between the V2 and PCI interferometers can produce unphysical "jumping" of the measured toroidal mode number, indicated here with red circular annotations. The 6-digit number in the upper right-hand corner indicates the DIII-D shot.

6.3.4 *Effect of major-radial offset*

As discussed in Section 6.2.1, the beams of the V2 and PCI interferometers have a $\Delta R = 4\,cm$ major-radial offset. This major-radial offset can bias the measured cross phase (6.20) between the two interferometer signals, which produces a corresponding bias in the interferometer-measured toroidal mode number (6.21). The cross-phase bias in (6.20) is $\sigma(R_{pci}) - \sigma(R_{v2})$, where $\sigma(R) = \arg[\Phi(R)]$ and $\Phi(R)$ is proportional to the line-integrated radial and poloidal mode structure at major radius R, as defined in (6.16). Now, if the mode structure is up-down symmetric about the midplane, Φ is real, and, modulo 2π, $\sigma(R_{pci}) - \sigma(R_{v2}) \in \{0, \pi\}$. The toroidal mode number is correctly identified when $\sigma(R_{pci}) - \sigma(R_{v2}) = 0$ but is aliased (and thus incorrectly identified) when $\sigma(R_{pci}) - \sigma(R_{v2}) = \pi$. If the radial mode structure evolves such that $\sigma(R_{pci}) - \sigma(R_{v2}) = 0 \rightarrow \pi$ or vice versa, the measured mode number will "jump"; an example of such mode-number "jumping" is shown in Figure 6.4. If the mode structure is *not* up-down symmetric about the midplane, however, Φ is complex, and the bias to the measured mode number becomes continuous, rather than discrete. Fast-ion shearing of Alfvén eigenmodes [12], for example, produces such up-down asymmetries.

Currently, there is not a tested, robust method for correcting the bias introduced by the major-radial offset of the interferometers (other

than reducing the offset, i.e. $\Delta R \to 0$, which is not possible with current port allocations). Although external magnetic probes can corroborate the interferometer-measured mode numbers in some cases, such as in Figure 6.3, they are blind to core-localized fluctuations. It may be possible to account for the radial and poloidal mode structure via measurements from microwave imaging reflectometry (MIR) [13] or electron cyclotron emission imaging (ECEI) [14] or via predictions from ideal MHD (e.g. NOVA [15, 16]) or hybrid MHD-gyrofluid (e.g. TAEFL [17, 18]) codes, but such efforts were deemed beyond the scope of this work.

6.4 DIAGNOSIS OF CORE-LOCALIZED MHD

Directly measuring the toroidal mode numbers of core-localized MHD was one of the primary motivations for the addition of a heterodyne-interferometer channel to DIII-D's pre-existing PCI system. As discussed in Section 6.3.4, however, the $\Delta R = 4\,\text{cm}$ major-radial offset between the V2 and PCI beams introduces ambiguities in the interferometer-measured mode numbers. Further, because core-localized MHD is, by definition, invisible to external magnetic probes, the interferometer-measured mode numbers cannot be corroborated by magnetics. Although the interferometer-measured mode numbers may currently be biased away from their true values, the mere presence of core-localized MHD in the interferometer signals should still be considered an encouraging proof of principle, where future work to eliminate the major-radial offset would allow direct and accurate mode-number measurements of core-localized MHD.

Figure 6.5 provides such a proof of principle. The top panel displays the fluctuating poloidal magnetic-field spectrum measured by a high-frequency magnetic probe [11], while the second panel displays the toroidal mode numbers measured by the midplane array of lower-bandwidth magnetic probes [11]. The third panel shows the interferometer-measured toroidal mode numbers, indicating very good agreement with the low-bandwidth magnetic measurements; however, between 1.8 s and 2.2 s, the interferometers measure a burst of fluctuations that are invisible to the low- and high-bandwidth magnetic probes, indicating that the modes are core-localized. The modes are triggered when the on-axis safety factor q_0 drops below unity. Unfortunately, additional observations of core-localized MHD proved to be surprisingly rare, and no systematic behavior could be identified linking the few-and-far-between occurrences.
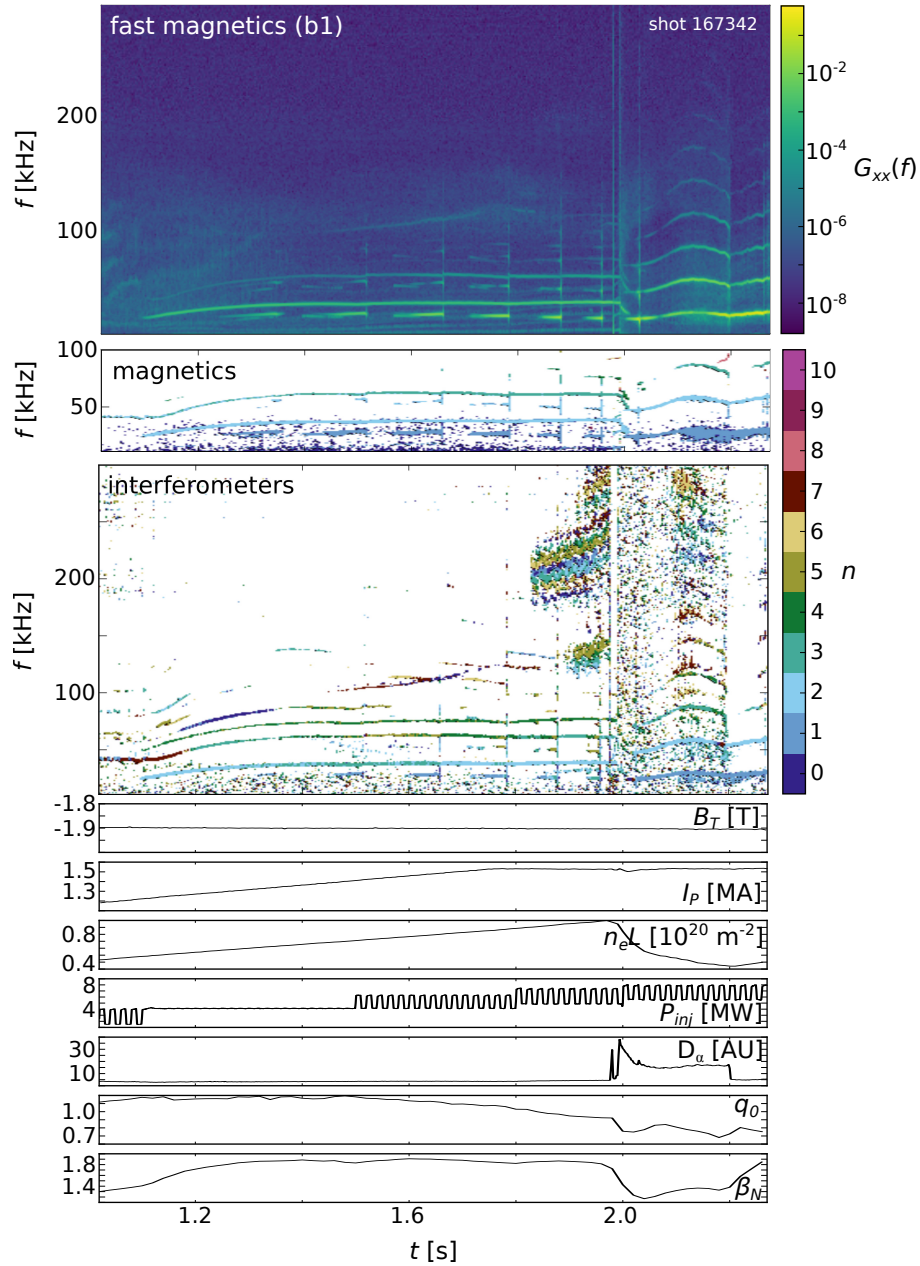
Figure 6.5: Between 1.8 s and 2.2 s, the correlated interferometers (3rd panel) measure a burst of fluctuations that are *invisible* to magnetics (1st and 2nd panels). This suggests that the modes are *core-localized* and that the correlated interferometers are indeed capable of measuring core-localized MHD. Various plasma parameters of interest are shown in the lower panels; the modes are triggered when the on-axis safety factor q_0 drops below unity.

[1] G. H. Choe, G. S. Yun, Y. Nam, W. Lee, H. K. Park et al. Dynamics of multiple flux tubes in sawtoothing KSTAR plasmas heated by electron cyclotron waves: I. Experimental analysis of the tube structure. *Nucl. Fusion*, 55(1):013015, 2015.

[2] T. N. Carlstrom, D. R. Ahlgren and J. Crosbie. Real-time, vibration-compensated CO_2 interferometer operation on the DIII-D tokamak. *Rev. Sci. Instrum.*, 59(7):1063–1066, 1988.

[3] M. A. Van Zeeland, G. J. Kramer, R. Nazikian, H. L. Berk, T. N. Carlstrom et al. Alfvén eigenmode observations on DIII-D via two-colour CO_2 interferometry. *Plasma Phys. Control. Fusion*, 47(9):L31, 2005.

[4] D. C. Pace, C. S. Collins, B. Crowley, B. A. Grierson, W. W. Heidbrink et al. Control of power, torque, and instability drive using in-shot variable neutral beam energy in tokamaks. *Nucl. Fusion*, 57(1):014001, 2017.

[5] E. M. Davis, J. C. Rost, M. Porkolab, A. Marinoni and M. A. Van Zeeland. A phase contrast imaging–interferometer system for detection of multiscale electron density fluctuations on DIII-D. *Rev. Sci. Instrum.*, 87(11):11E117, 2016.

[6] J. Chen, W. X. Ding, D. L. Brower, D. Finkenthal, C. Muscatello et al. Faraday-effect polarimeter diagnostic for internal magnetic field fluctuation measurements in DIII-D. *Rev. Sci. Instrum.*, 87(11):11E108, 2016.

[7] M. A. Van Zeeland, R. L. Boivin, T. N. Carlstrom and T. M. Deterly. CO_2 laser polarimeter for Faraday rotation measurements in the DIII-D tokamak. *Rev. Sci. Instrum.*, 79(10):10E719, 2008.

[8] A. Mlynek, G. Pautasso, M. Maraschek, H. Eixenberger and the ASDEX Upgrade Team. Infrared interferometry with submicrosecond time resolution in massive gas injection experiments on ASDEX Upgrade. *Fusion Sci. Tech.*, 61(4):290–300, 2012.

[9] P. Milne. Personal communication, May 2016.

[10] A. V. Oppenheim and R. W. Schafer. *Discrete-Time Signal Processing*. Pearson, Upper Saddle River, NJ, USA, 3rd edition, 2010.

[11] E. J. Strait. Magnetic diagnostic system of the DIII-D tokamak. *Rev. Sci. Instrum.*, 77(2):023502, 2006.

[12] B. J. Tobias, I. G. J. Classen, C. W. Domier, W. W. Heidbrink, N. C. Luhmann et al. Fast ion induced shearing of 2D Alfvén eigenmodes measured by electron cyclotron emission imaging. *Phys. Rev. Lett.*, 106:075003, 2011.

[13] C. M. Muscatello, C. W. Domier, X. Hu, G. J. Kramer, N. C. Luhmann Jr. et al. Technical overview of the millimeter-wave imaging reflectometer on the DIII-D tokamak (invited). *Rev. Sci. Instrum.*, 85(11):11D702, 2014.

[14] B. Tobias, C. W. Domier, T. Liang, X. Kong, L. Yu et al. Commissioning of electron cyclotron emission imaging instrument on the DIII-D tokamak and first data. *Rev. Sci. Instrum.*, 81(10):10D928, 2010.

[15] C.Z Cheng and M.S Chance. NOVA: A nonvariational code for solving the MHD stability of axisymmetric toroidal plasmas. *J. Comp. Phys.*, 71(1):124 – 146, 1987.

[16] C.Z. Cheng. Kinetic extensions of magnetohydrodynamics for axisymmetric toroidal plasmas. *Phys. Rep.*, 211(1):1 – 51, 1992.

[17] D. A. Spong, B. A. Carreras and C. L. Hedrick. Linearized gyrofluid model of the alpha-destabilized toroidal Alfvén eigenmode with continuum damping effects. *Phys. Fluids B*, 4(10):3316–3328, 1992.

[18] D A Spong, B A Carreras, C L Hedrick, N Dominguez, L A Charlton et al. Alpha destabilization of the TAE mode using a reduced gyrofluid model with Landau closure. *Phys. Scr.*, 45(2):159, 1992.

7

CONCLUSIONS & FUTURE WORK

7.1 SUMMARY & CONCLUSIONS

The work described in this thesis can be summarized as follows:

- Chapter 2 discusses the theory of optical interferometric methods in the context of measuring tokamak plasma-density fluctuations. The laser-plasma interaction is quantified via Fraunhofer scalar-diffraction theory, and the resulting diffracted field is imaged onto a square-law detector. Interfering the imaged field with a known reference field produces measurable intensity fluctuations; the specification of this reference field defines the interferometric method. Details of two particular interferometric methods — external-reference-beam interferometry and phase contrast imaging (PCI) — are discussed, with an emphasis on their sensitivity to fluctuations and their spatiotemporal bandwidths. Significantly, while PCI can measure fluctuations more sensitively than an external-reference-beam interferometer, PCI suffers from a low-k cutoff; an external-reference-beam interferometer does *not* suffer from such a low-k cutoff.

- Chapter 3 considers the design of an external-reference-beam, heterodyne interferometer (hereafter referred to as a heterodyne interferometer). A criterion for satisfactory wavefront matching between the probe beam and the reference beam is developed, and finite-sampling-volume effects are shown to constrain the heterodyne interferometer's spatial bandwidth. The effects of phase noise, amplitude noise, and digitizer bit noise are each discussed in the context of the heterodyne interferometer's signal-to-noise ratio, and the systematic errors resulting from imperfect demodulation of the heterodyne interference signal are quantified.

- Chapter 4 details the addition of a heterodyne interferometer to the pre-existing PCI system on the DIII-D tokamak. Both systems operate simultaneously, sharing a single $10.6\,\mu m$ probe beam through the plasma. Optical-diagnostic access on DIII-D and the capabilities of the pre-existing PCI system are briefly reviewed. Referencing the design considerations in Chapter 3

and adopting the philosophy that the pre-existing PCI system should be minimally perturbed, the optical layout for the heterodyne interferometer is developed; the magnification of the interferometer's imaging system is selected such that the spatial bandwidths of the PCI and interferometer have a mid-k overlap. The design, procurement, and installation of the new optical and electrical components required to make the heterodyne interferometric measurement are summarized. Of note is the interferometer's radio-frequency local oscillator: the phase noise of a crystal oscillator (XO) was empirically found to be *too large* to make meaningful fluctuation measurements in most tokamak plasmas, but the substantially lower phase noise of an oven-controlled crystal oscillator (OCXO) allows measurements of a whole zoo of coherent and broadband plasma fluctuations. The interferometer response and the multiscale capabilities of the combined PCI-interferometer are empirically verified via sound-wave calibrations. Specifically, the PCI is shown to measure high-k ($1.5\,\mathrm{cm^{-1}} < |k_R| \leqslant 25\,\mathrm{cm^{-1}}$) fluctuations with sensitivity $3 \times 10^{13}\ \mathrm{m^{-2}}/\sqrt{\mathrm{kHz}}$, while the interferometer simultaneously measures low-k ($|k_R| < 5\,\mathrm{cm^{-1}}$) fluctuations with sensitivity $3 \times 10^{14}\ \mathrm{m^{-2}}/\sqrt{\mathrm{kHz}}$. Both systems have temporal bandwidths in excess of $1\,\mathrm{MHz}$.

• Chapter 5 demonstrates the multiscale capabilities of the combined PCI-interferometer. During a recent DIII-D experiment, the location of electron cyclotron resonance heating (ECH) was moved from $\rho_{\mathrm{ECH}} = 0.5$ to $\rho_{ECH} = 0.8$, altering the local a/L_{T_e} and a/L_{T_i} in an attempt to change the coupling between the electron-scale and ion-scale turbulence. As such, this experiment presents an ideal opportunity for multiscale turbulence investigations with the combined PCI-interferometer. Numerous turbulent branches are observed. In particular, the interferometer measures a low-k electromagnetic mode driven unstable by collisionality, properties consistent with the micro-tearing mode (MTM), and the PCI measures a wavenumber spectrum that exhibits distinct flattening when a/L_{T_e} is increased relative to a/L_{T_i}, reminiscent of results from realistic multiscale gyrokinetic simulations [1]. To aid the interpretation of these measurements, linear-stability analysis and quasilinear-transport modeling are performed with the gyro-Landau fluid code TGLF, and qualitative agreement with the PCI-measured wavenumber spectrum is obtained.

- Chapter 6 discusses the correlation of the newly installed interferometer with DIII-D's toroidally separated, pre-existing V2 interferometer. Capable of probing the core plasma, the interferometers are shown to be sensitive to core-localized fluctuations that are *invisible* to external magnetic probes. The chapter begins with a brief review of the two-point correlation technique and shows how toroidal mode numbers can be extracted from a pair of toroidally separated measurements. Meaningful correlation requires that the two measurements share the same timebase. The digitizers of both interferometers were modified to phase lock their clocks, and a residual "trigger offset" was measured and is compensated in software. Where comparisons can be made with magnetic probes, the interferometer-measured toroidal mode numbers are in good agreement. Currently, there is not a tested, robust method for correcting the bias introduced by the 4 cm major-radial offset between the interferometer beam centers, which unfortunately limits the deployment of this system for physics studies of core-localized MHD.

7.2 FUTURE WORK

The combined PCI-interferometer developed in this work has a clear application in the burgeoning study of multiscale turbulence and cross-scale coupling, which may be significant in the reactor relevant $T_e \approx T_i$ regime. In roughly the next six months, Howard *et al.* expects to complete realistic multiscale gyrokinetic simulations for the experiment described in Chapter 5. It will be very interesting to see if the predicted wavenumber spectrum matches the PCI-measured wavenumber spectrum. It should be noted that a synthetic PCI diagnostic already exists for the interpretation of such gyrokinetic simulations [2]. Small modifications to the synthetic PCI should also allow a synthetic interferometer diagnostic. Previous multiscale simulations predict significant local and non-local energy cascades between the ion and electron scales [1], so it is desirable to investigate such coupling empirically. With its large spatiotemporal bandwidth, the combined PCI-interferometer may be ideally suited for measurement of such coupling, which may be suitably quantified by the bicoherence [3] between various channels of the system or some other suitable measure of nonlinear processes. (Note that the author has performed preliminary bispectral analysis of the measurements discussed in Chapter 5; interestingly, the $|k_R| \sim 5\,cm^{-1}$ and $f \sim 1\,MHz$ mode observed in Figure 5.9(b) has an exceptionally large autobicoherence).

The interferometer-measured, low-k, electromagnetic modes that are destabilized by collisionality are also deserving of further study. The properties of this mode are consistent with the micro-tearing mode, which was predicted to be marginally unstable in the multiscale experiment's reference discharge [4]. Unfortunately, TGLF's default eigenfunction basis of four Hermite polynomials is typically insufficient to resolve micro-tearing modes (MTMs) [5], so there are no attempts to simulate the MTM in this work. However, it may be conceivable that increasing the number of Hermite polynomials will allow identification of the MTM in TGLF. Alternatively, linear simulations with the gyrokinetic code GYRO [6] could be pursued. (Note that the reference-discharge simulations indicating marginal MTM instability were performed with GYRO). Experimentally, it is desirable to map out the parametric dependencies of this mode, particularly its response to the plasma β and collisionality. If dedicated experiments cannot be performed, it should be noted that the relevant experimental conditions (i.e. ITER-baseline scenario) are fairly typical at DIII-D, and a fair amount may still be learned by "piggybacking" on other experiments.

With regards to the combined PCI-interferometer, the most substantial improvement to the system would be upgrading the heterodyne-interferometer detector from a single element to a multi-element array. This would allow reconstruction of k_R from the interferometer measurements, enabling estimates of frequency-wavenumber spectra $S_{\phi,\phi}(k, f)$ and wavenumber spectra $S_{\phi,\phi}(k)$ much like with the PCI. This capability is desirable for several reasons. First and foremost, interferometric measurements across a multi-element array would allow accurate estimates of $S_{\phi,\phi}(k)$ below the PCI low-k cutoff (4.2), which may have important implications for validation of spectral-flattening predictions from multiscale gyrokinetic predictions. Further, as discussed in Section 4.8.2, interferometric measurements across a multi-element array would allow robust and accurate cross-calibration of the PCI on a shot-to-shot and an intra-shot basis. Note that each additional detector element would require its own set of electronics (e.g. signal conditioning RF amplifiers, demodulation electronics, and audio amplifiers) and two additional digitizer channels (to digitize both the in-phase I and quadrature Q signals). While the "deadbug" circuit construction utilized in this work is ideal for prototyping, any future increase to the number of interferometer channels would call for a printed-circuit-board (PCB) construction of the electronics. Thus, increasing the number of interferometer channels is not a small undertaking.

A simpler, cheaper, and faster performance improvement would be the procurement of anti-aliasing filters with a higher cutoff frequency. The current anti-aliasing filters limit the bandwidth of the interferometer to approximately 1 MHz, but the upstream components have bandwidths in excess of 2 MHz. Thus, new anti-aliasing filters could, quite literally overnight, nearly double the temporal bandwidth of the interferometer.

Finally, there is not currently a tested, robust method for correcting the bias introduced by the major-radial offset of the toroidally correlated interferometers (other than reducing the offset, i.e. $\Delta R \rightarrow 0$, which is not possible with current port allocations). It may be possible to account for the radial and poloidal mode structure via e.g. measurements from microwave imaging reflectometry (MIR) [7] or electron cyclotron emission imaging (ECEI) [8]. A fully functional system would, for example, enable the detection and identification of Alfvén eigenmode toroidal mode numbers and, consequently, quantification of their damping rates, as has been done in JET [9].

Looking towards ITER and other next-step devices, the combined PCI-interferometer may allow sensitive, high spatiotemporal bandwidth measurements of multiscale turbulence. The diagnostic development pursued in this thesis proves that heterodyne-interferometric detection and PCI detection can be simultaneously implemented using a shared probe beam and a shared set of ports. The addition of PCI detection to e.g. the ITER interferometer [10], however, is not without its challenges. For example, the ITER interferometer employs a Michelson configuration, with the beam making a second pass through the plasma after bouncing off of a retroreflector inside the vacuum vessel. In contrast, at least to the author's knowledge, all previous PCI implementations have employed a Mach-Zehnder configuration, with the probe beam making a single pass through the plasma. In principle, PCI can use a Michelson configuration, but the double pass and retroreflector may complicate interpretation of the measurements, particularly if attempting to localize the measurements with a spatially filtering mask [11, 12, 13] or with 2-dimensional detector arrays [14, 15].

Regardless, the spatial bandwidth of a PCI system that shares its probe beam with the ITER interferometer can be considered. The $1/e$ E waist of the ITER interferometer's 10.6 μm probe beam is $w_0 \approx$ 8 mm [10] such that a PCI system using this probe beam would have a diffraction-limited low-k cutoff (2.88) of 2.5 cm^{-1}. (However, recall that the DIII-D PCI is operated two to three times above the diffraction limit to give some leeway to the PCI feedback system). Of course, as demonstrated in this thesis, simultaneous heterodyne-interferometric

and PCI detection can obviate the PCI's low-k cutoff. A somewhat larger problem, however, may be the limited collection volumes and long path lengths between the vacuum vessel and the detector, which may impose severe constraints on the high-k cutoff of a $10.6\,\mu m$ PCI or interferometer. One potential solution is to use a smaller probe wavelength (i.e. larger k_0) to decrease the scattering angle θ_m from (A.14). As the burning plasma regime will be predominantly electron heated (i.e. via fusion alpha particles slowing down on electrons), it is extremely important that both high-k electron turbulence and any cross-scale coupling with low-k ion turbulence is accurately diagnosed and understood.

BIBLIOGRAPHY

[1] N. T. Howard, C. Holland, A. E. White, M. Greenwald, J. Candy et al. Multi-scale gyrokinetic simulations: comparison with experiment and implications for predicting turbulence and transport. *Phys. Plasmas*, 23(5):056109, 2016.

[2] J. C. Rost, L. Lin and M. Porkolab. Development of a synthetic phase contrast imaging diagnostic. *Phys. Plasmas*, 17(6):062506, 2010.

[3] Y. C. Kim and E. J. Powers. Digital bispectral analysis and its applications to nonlinear wave interactions. *IEEE Trans. Plasma Sci.*, 7(2):120–131, 1979.

[4] C. Holland, N.T. Howard and B.A. Grierson. Gyrokinetic predictions of multiscale transport in a DIII-D ITER baseline discharge. *Nucl. Fusion*, 57(6):066043, 2017.

[5] G. M. Staebler. Personal communication.

[6] J. Candy and R.E. Waltz. An Eulerian gyrokinetic-Maxwell solver. *J. Comp. Phys.*, 186(2):545 – 581, 2003.

[7] C. M. Muscatello, C. W. Domier, X. Hu, G. J. Kramer, N. C. Luhmann Jr. et al. Technical overview of the millimeter-wave imaging reflectometer on the DIII-D tokamak (invited). *Rev. Sci. Instrum.*, 85(11):11D702, 2014.

[8] B. Tobias, C. W. Domier, T. Liang, X. Kong, L. Yu et al. Commissioning of electron cyclotron emission imaging instrument on the DIII-D tokamak and first data. *Rev. Sci. Instrum.*, 81(10):10D928, 2010.

[9] V. Aslanyan, M. Porkolab, P. Woskov, P. Puglia, W. Pires de Sa et al. Progress and first measurements from the upgraded Alfvén eigenmode active diagnostic on JET. *Proc. European Phys. Soc.*, page P1.105, 2017.

[10] M. A. Van Zeeland, R. L. Boivin, D. L. Brower, T. N. Carlstrom, J. A. Chavez et al. Conceptual design of the tangentially viewing combined interferometer-polarimeter for ITER density measurements. *Rev. Sci. Instrum.*, 84(4):043501, 2013.

[11] J. R. Dorris, J. C. Rost and M. Porkolab. Localized measurement of short wavelength plasma fluctuations with the DIII-D phase contrast imaging diagnostic. *Rev. Sci. Instrum.*, 80(2):023503, 2009.

[12] J. R. Dorris. *Phase contrast imaging measurements and modeling of short wavelength turbulence in the DIII-D tokamak*. PhD thesis, MIT, 2010.

[13] L. Lin, E. M. Edlund, M. Porkolab, Y. Lin and S. J. Wukitch. Vertical localization of phase contrast imaging diagnostic in Alcator C-Mod. *Rev. Sci. Instrum.*, 77(10):10E918, 2006.

[14] A. L. Sanin, K. Tanaka, L. N. Vyacheslavov, K. Kawahata and T. Akiyama. Two-dimensional phase contrast interferometer for fluctuations study on LHD. *Rev. Sci. Instrum.*, 75(10):3439–3441, 2004.

[15] K. Tanaka, S. Coda, M. Yoshida, H. Sasao, Y. Kawano et al. Design of tangential viewing phase contrast imaging for turbulence measurements in JT-60SA. *Rev. Sci. Instrum.*, 87(11):11E118, 2016.

A

DIFFRACTION OF A GAUSSIAN PROBE BEAM

A.1 KIRCHHOFF SCALAR-DIFFRACTION THEORY

As discussed in the text surrounding (2.25), a CO_2 probe beam in a tokamak plasma propagates as a transverse electromagnetic wave with near-constant polarization (with any small changes to beam polarization being of little practical interest to the present work). Thus, it is suitable to pursue a *scalar* theory of the beam's interaction with the plasma. Below, Kirchoff's scalar-diffraction theory is summarized.

A monochromatic scalar wave $U(r)e^{-i\omega t}$ in vacuum satisfies the Helmholtz equation

$$(\nabla^2 + k^2)U = 0, \tag{A.1}$$

where $k = \omega/c$. The Helmholtz-Kirchhoff integral theorem states that the field at a point P is

$$U(P) = \frac{1}{4\pi} \int_S \left[U \frac{\partial}{\partial n} \left(\frac{e^{iks}}{s} \right) - \frac{e^{iks}}{s} \frac{\partial U}{\partial n} \right] dS, \tag{A.2}$$

where S is an arbitrary surface that encloses P, s is the vector from point P to differential area element dS, n is the *inward*-pointing normal of surface S, and U is assumed to be differentiable to second order within and on S [1, Sec. 8.3]. The relevant geometry is sketched in Figure A.1(a).
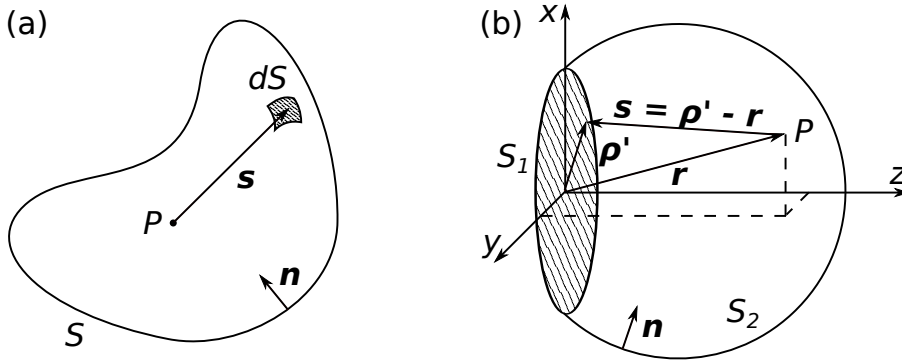
Figure A.1: Geometries for Kirchhoff scalar-diffraction calculations. When S_1 is imaged by an optical system, S_1 is referred to as the *object plane*.

To proceed with the diffraction calculation, assume that the incident waves propagate in the $+z$-direction and adopt the surface drawn in Figure A.1(b). That is, $S = S_1 + S_2$, where S_1 is a circle in the (x, y)-plane, and S_2 is a spherical segment centered on the optical axis. When S_1 is imaged by an optical system, S_1 is referred to as the *object plane*. Now, assume that the incident waves were "turned on" at some finite time in the past, and take the radius of S_2 to be large enough such that none of the diffracted waves have had sufficient time to reach S_2, i.e. $U \equiv 0$ on S_2. (Of course, strictly speaking, the source's finite turn-on time requires relaxation of the monochromatic assumption. Finite turn-on time does not preclude a pseudo-monochromatic source, however, and such a source is assumed hereafter). Thus, the integral over S_2 vanishes, and the diffraction calculation reduces to an integral over S_1.

Now, evaluation of the Helmholtz-Kirchhoff integral (A.2) requires knowledge of U on S_1. For free-space propagation, S_1 is an imagined (rather than a physical) surface that does not impede the propagation of the incident wave $U^{(i)}$ (that is, $U = U^{(i)}$ and $\partial U/\partial n = \partial U^{(i)}/\partial n$ on S_1). If, however, S_1 contains opaque obstacles, the free-space propagation conditions are no longer valid; instead, the Kirchhoff boundary conditions can be adopted:

$$\text{surfaces of clear aperture:} \quad U = U^{(i)}, \quad \frac{\partial U}{\partial n} = \frac{\partial U^{(i)}}{\partial n},$$
$$\text{opaque surfaces:} \quad U = 0, \quad \frac{\partial U}{\partial n} = 0.$$

While these boundary conditions are adequate for the current application, it should be noted that they are not physical for points that are very close to the boundaries of the opaque obstacles.

A.2   FREE-SPACE DIFFRACTION OF A GAUSSIAN BEAM

This section demonstrates that the Fraunhofer diffraction formalism gives the correct form for a free-space Gaussian beam in the far-field limit, and it also lays the groundwork for examining the diffraction of a Gaussian beam from plasma-density fluctuations. Note that the Gaussian-beam definition provided in Section 2.2.1 is used throughout the remainder of this appendix.

Assume that the incident Gaussian beam has a waist at S_1, and take the radius of S_1 to be much larger than the beam waist w_0 such that the domain of integration effectively extends over the whole (x, y)-plane. For free-space propagation, S_1 does not perturb the Gaussian

beam; thus, $E(\mathbf{r}, t) = E^{(i)}(\mathbf{r}, t) = E_G(\mathbf{r})e^{-i\omega_0 t}$, where $E_G(\mathbf{r})$ is the Gaussian beam's spatial dependence, as defined by (2.33). Now, in the far-field ($k_0 s \gg 1$) and paraxial ($\mathbf{s} \approx -z\hat{\mathbf{z}}$) approximations

$$\frac{e^{ik_0 s}}{s}\bigg|_{S_1} \approx \frac{e^{ik_0 s}}{z},$$

$$\frac{\partial}{\partial n}\left(\frac{e^{ik_0 s}}{s}\right)\bigg|_{S_1} \approx -ik_0\left(\frac{e^{ik_0 s}}{z}\right).$$

The s-dependence in the phase arguments has been retained, as it is the mechanism responsible for diffraction, but the s-dependence in the amplitude has been dropped as it only gives rise to negligible variations in the amplitude of the diffracted wave. Relative to a spherical wave, a Gaussian beam has several additional z-dependencies; however, at the beam's waist

$$\frac{\partial w(z)}{\partial z}\bigg|_{\text{waist}} \equiv 0,$$

$$\frac{\partial}{\partial z}\left[\frac{1}{R(z)}\right]\bigg|_{\text{waist}} = \frac{1}{z_R^2},$$

$$\frac{\partial \psi(z)}{\partial z}\bigg|_{\text{waist}} = \frac{1}{z_R}.$$

Then, if the beam's Rayleigh range is much greater than the probe wavelength ($k_0 z_R \gg 1$) and the relevant transverse dimensions are much less than the Rayleigh range ($w_0 \ll z_R$), the Gaussian beam at S_1 satisfies

$$E_G(\mathbf{r}')\big|_{S_1} \approx E_0 e^{-(\rho'/w_0)^2},$$

$$\frac{\partial E_G(\mathbf{r}')}{\partial n}\bigg|_{S_1} \approx ik_0\left[E_0 e^{-(\rho'/w_0)^2}\right].$$

Note that the CO_2 laser beams ($k_0 \approx 2\pi \times 10^5\,\text{m}^{-1}$) that probe tokamak plasmas often have $z_R \gg 10\,\text{m}$ such that $k_0 z_R \gg 1$ and $w_0 \ll z_R$ (the transverse dimensions are constrained by the machine size such that $w_0 \ll 1\,\text{m}$) are very well-satisfied.

Substituting the above expressions for the incident waves and their surface-normal derivatives into the Helmholtz-Kirchhoff integral (A.2) and simplifying yields

$$E(\mathbf{r}) \approx \frac{-iE_0}{\lambda_0 z}\int_{S_1} e^{-(\rho'/w_0)^2}e^{ik_0 s}\,dS. \tag{A.3}$$

To proceed further, s must be approximated:

$$
\begin{aligned}
s &= |\boldsymbol{\rho}' - \boldsymbol{r}| \\
&= \left[r^2 - 2(x'x + y'y) + (x'^2 + y'^2) \right]^{1/2} \\
&\approx r - \frac{x'x + y'y}{r},
\end{aligned}
\tag{A.4}
$$

where only terms linear in (x'/r) and (y'/r) have been retained. This is known as the Fraunhofer limit, and it is valid in the far-field $z \gg z_R$ [1, Sec. 8.3.3]. Under the Fraunhofer limit the diffraction integral (A.3) becomes

$$
E(\boldsymbol{r}) \approx \frac{-iE_0}{\lambda_0 z} e^{ik_0 r} D_x D_y,
\tag{A.5}
$$

where

$$
\begin{aligned}
D_x &= \int_{-\infty}^{\infty} e^{-(x'/w_0)^2} e^{-ik_0 x' x/r} dx' \\
&= \mathcal{F}\left[e^{-(x'/w_0)^2} \right] (k_0 x/r) \\
&= \sqrt{\pi} w_0 e^{-(k_0 w_0 x/2r)^2}
\end{aligned}
\tag{A.6}
$$
$$
\tag{A.7}
$$

gives the diffraction pattern in the x-direction, and the integral has been easily evaluated by noting that it is simply the Fourier transform of a Gaussian. The diffraction pattern in the y-direction, D_y, is similarly determined. Note that the $e^{-ik_0 x' x/r} = e^{-ik_0 x' \sin\theta}$ term in (A.6) is the typical geometric phase factor that results from path-length differences between points on surface S_1 and the field point $\boldsymbol{r}$, as shown in Figure A.2. Substituting (A.7) into (A.5) yields

$$
E(\boldsymbol{r}) \approx -iE_0 \left(\frac{z_R}{z} \right) e^{-(k_0 w_0 \rho/2r)^2} e^{ik_0 r}.
\tag{A.8}
$$

Is this consistent with the expected far-field representation of a Gaussian beam? Yes! To see this, note that in the far-field ($z \gg z_R$)

$$
\begin{aligned}
\frac{z_R}{z} &\approx \frac{w_0}{w(z)}, \\
\frac{k_0 w_0 \rho}{2r} &\approx \frac{\rho}{w(z)}, \\
r &\approx z + \frac{\rho^2}{2R(z)}, \\
-i = e^{-i\pi/2} &\approx e^{-i\psi(z)},
\end{aligned}
$$
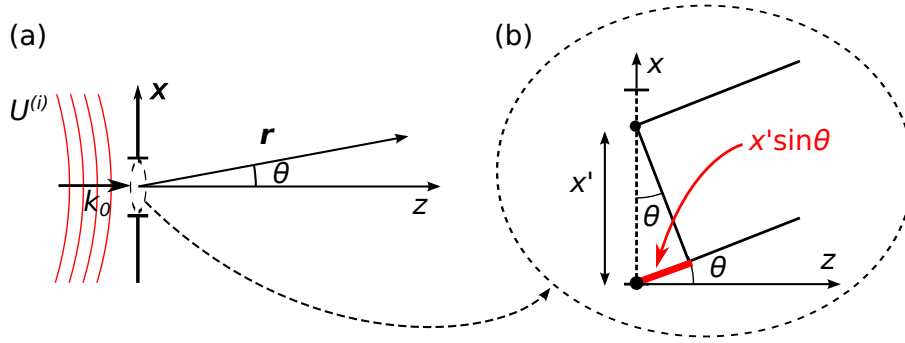
Figure A.2: (a) Typical Fraunhofer diffraction geometry and (b) a close up that displays the path-length difference $x' \sin \theta$ between a wave emanating from the origin and a wave emanating from x'

such that the diffracted field in the Fraunhofer limit (A.8) can be cast in the form of a typical Gaussian beam as expressed in (2.33), i.e. $E(\boldsymbol{r}) = E_G(\boldsymbol{r})$ for $z \gg z_R$. Of course, when considering free-space propagation, $E(\boldsymbol{r}) \equiv E_G(\boldsymbol{r})$ for $0 \leqslant z < \infty$, but the above work *proves* that the Fraunhofer diffraction formalism gives the correct results under the appropriate limits; it also lays the groundwork for examining the diffraction of a phase-modulated Gaussian beam.

A.3 DIFFRACTION FROM PLASMA-DENSITY FLUCTUATIONS

Now, allow a Gaussian CO_2 probe beam to pass through a tokamak plasma. The beam acquires the plasma-induced phase delay $\phi(\boldsymbol{\rho}', t)$ given by (2.30), where $\boldsymbol{\rho}'$ corresponds to the beam's transverse dimensions. Explicitly dividing ϕ into bulk $\overline{\phi}(t)$ and spatially varying $\widetilde{\phi}(\boldsymbol{\rho}', t)$ components, the plasma-induced phase delay becomes

$$\phi(\boldsymbol{\rho}', t) = \overline{\phi}(t) + \widetilde{\phi}(\boldsymbol{\rho}', t). \tag{A.9}$$

Typically, $\widetilde{\phi}$ varies on much faster time scales than $\overline{\phi}$, but this is not required. The spatial variation of the plasma-induced phase delay contributes to the diffraction of the incident Gaussian probe beam, and the remainder of this section uses scalar-diffraction theory to determine the diffracted field in the Fraunhofer limit; the near-field form consistent with the computed Fraunhofer field is then inferred. (The object planes of the imaging systems relevant to this work sit in the near field, so computation of the imaged field requires knowledge of the near-field diffraction pattern).

The response functions of the diagnostics investigated in Sections 2.4 and 2.5 are shown to be linear in their regimes of relevance, so it is suf-

ficient to examine phase fluctuations $\widetilde{\phi}$ consisting of a single Fourier mode

$$\widetilde{\phi}(\boldsymbol{\rho}', t) = \widetilde{\phi}_0 \cos(kx' - \omega t). \tag{A.10}$$

As a CO_2 beam's optical cycles ($\omega_0 = 2\pi \cdot 28.3\,\text{THz}$) occur much more rapidly than the temporal evolution of the plasma ($\omega \lesssim 2\pi \cdot 1\,\text{GHz}$), the problem can be treated adiabatically by solving for the beam propagation at each instant in time during the relatively slow evolution of ϕ. Then, following the formalism pioneered by Raman and Nath [2, 3], this plasma-induced phase delay makes an additional phase contribution to the diffraction integral (A.6) such that the diffraction pattern in the x-direction is given as

$$
\begin{aligned}
D_x &= \int_{-\infty}^{\infty} e^{-(x'/w_0)^2} e^{-ik_0 x x'/r} e^{i\phi(x',t)}\,dx' \\
&= e^{i\overline{\phi}} \int_{-\infty}^{\infty} e^{-(x'/w_0)^2} e^{-ik_0 x x'/r} e^{i\widetilde{\phi}_0 \cos(kx'-\omega t)}\,dx' \\
&= e^{i\overline{\phi}} \int_{-\infty}^{\infty} e^{-(x'/w_0)^2} e^{-ik_0 x x'/r} \\
&\qquad \times \left\{ \sum_{m=-\infty}^{\infty} i^m \left[J_m(\widetilde{\phi}_0) \right] e^{im(kx'-\omega t)} \right\} dx' \\
&= e^{i\overline{\phi}} \sum_{m=-\infty}^{\infty} i^m \left[J_m(\widetilde{\phi}_0) \right] e^{-im\omega t} \\
&\qquad \times \int_{-\infty}^{\infty} e^{-(x'/w_0)^2} e^{-i\left(\frac{k_0 x}{r} - mk\right)x'}\,dx' \\
&= \sqrt{\pi}w_0 e^{i\overline{\phi}} \sum_{m=-\infty}^{\infty} \left\{ i^m \left[J_m(\widetilde{\phi}_0) \right] e^{-im\omega t} \right. \\
&\qquad\qquad\qquad \left. \times e^{-\left[\frac{w_0}{2}\left(\frac{k_0 x}{r}-mk\right)\right]^2} \right\},
\end{aligned}
\tag{A.11}
$$

where the bracketed expression in the third equality follows from application of the well-known Jacobi-Anger expansion and J_m is the m^{th} Bessel function of the first kind. Noting that $E(\boldsymbol{r}, t) = E(\boldsymbol{r})e^{-i\omega_0 t}$, substitution of (A.11) into (A.5) yields

$$
\begin{aligned}
E(\boldsymbol{r}, t) &\approx \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0) \right] e^{-im\omega t} e^{-\left[\frac{w_0}{2}\left(\frac{k_0 x}{r}-mk\right)\right]^2} \\
&\quad \times e^{i\overline{\phi}} \left[-iE_0 \left(\frac{z_R}{z}\right) e^{-(k_0 w_0 y/2r)^2} e^{i(k_0 r - \omega_0 t)} \right].
\end{aligned}
\tag{A.12}
$$
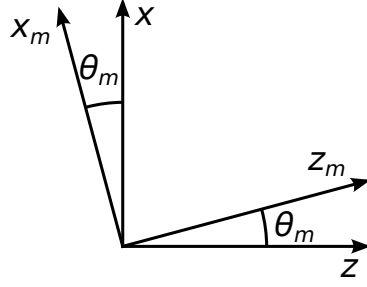
Figure A.3: Coordinate transformation for interpretation of the diffraction pattern of a phase-modulated Gaussian beam

Here, only the $|m| \leqslant 1$ terms have been retained because $|J_m(\widetilde{\phi}_0)| \sim \widetilde{\phi}_0^{|m|}$ for experimentally relevant values of $\widetilde{\phi}_0 \ll 1$. (The complete small-argument, asymptotic form for J_m is discussed in Section 2.3.3). The effect of higher order terms can be easily investigated by e.g. including the $m = \pm 2$ terms etc.

To put (A.12) in a more familiar form, consider the coordinate transformation from the lab-frame coordinate system r to the coordinate system of the m^{th} scattered beam r_m, as depicted graphically in Figure A.3. As the transformation is simply a rotation about the y-axis by angle θ_m, the coordinate systems are related via

$$\begin{pmatrix} x_m \\ y_m \\ z_m \end{pmatrix} = \begin{pmatrix} \cos\theta_m & 0 & -\sin\theta_m \\ 0 & 1 & 0 \\ \sin\theta_m & 0 & \cos\theta_m \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix}, \tag{A.13}$$

where

$$\sin\theta_m \equiv \frac{mk}{k_0}. \tag{A.14}$$

Typically, $mk/k_0 \ll 1$ such that

$$\cos\theta_m \approx 1 - \frac{1}{2}\left(\frac{mk}{k_0}\right)^2 \tag{A.15}$$

is a very good approximation. The above coordinate transformation can be written more compactly as

$$r_m = [R(\theta_m)]r, \tag{A.16}$$

where

$$\boldsymbol{R}(\theta) = \begin{pmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{pmatrix} \tag{A.17}$$

is the rotation matrix that rotates the (x, z)-plane about the y-axis by angle θ. Rotation matrices are *orthogonal*, which endows $\boldsymbol{R}(\theta_m)$ with some useful properties [4, Ch. 6]; namely, its inverse is equal to its transpose $[\boldsymbol{R}(\theta_m)]^\mathsf{T}$

$$[\boldsymbol{R}(\theta_m)]^\mathsf{T} = [\boldsymbol{R}(\theta_m)]^{-1} = \boldsymbol{R}(-\theta_m), \tag{A.18}$$

and its determinant is unity

$$\det[\boldsymbol{R}(\theta_m)] = |\boldsymbol{R}(\theta_m)| = 1 \tag{A.19}$$

such that the rotation preserves lengths, i.e. $r_m = r$. It is sufficient to retain terms only to first order in k/k_0 (small scattering angle) and x/z (paraxial limit) for the *amplitude* dependencies of the diffracted field (this is *not*, in general, true for the phase dependencies). Thus, $1/z \approx 1/z_m$ and $x \approx x_m + (mk/k_0)z$ such that the Fraunhofer diffracted field (A.12) can be rewritten as

$$E(\boldsymbol{r}, t) \approx e^{i\overline{\phi}} \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0) \right] e^{-i(\omega_0 + m\omega)t} \\ \times \left[-iE_0 \left(\frac{z_R}{z_m} \right) e^{-(k_0 w_0 \rho_m/2r)^2} e^{ik_0 r} \right], \tag{A.20}$$

where $\rho_m = (x_m^2 + y_m^2)^{1/2}$. Now, the bracketed expression has the form of (A.8) for a far-field Gaussian beam; thus, the diffracted electric field can be more compactly and generally written as

$$E(\boldsymbol{r}, t) \approx e^{i\overline{\phi}} \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0) \right] E_G(\boldsymbol{r}_m) e^{-i(\omega_0 + m\omega)t}. \tag{A.21}$$

Note that (A.21) is valid for $0 \leqslant z < \infty$ rather than only for $z \gg z_R$; that is, computing the far-field diffraction pattern has additionally allowed *inferring* the corresponding near field.

Thus, a sinusoidal phase modulation diffracts an incident Gaussian beam predominantly into downscattered ($m = -1$), unscattered ($m = 0$), and upscattered ($m = 1$) Gaussian beams. The incident beam is coupled into the m^{th} scattered beam with strength $J_m(\widetilde{\phi}_0)$. The

m^{th} scattered beam is Doppler shifted relative to the incident beam by $m\omega$ and propagates at an angle $\theta_m \approx mk/k_0$ relative to the lab-frame optical axis. The adiabatic assumption ($\omega/\omega_0 \ll 1$) is very well-satisfied for a CO_2 probe beam in a typical tokamak plasma (i.e. $\omega/\omega_0 \lesssim 1\,\text{GHz}/28.3\,\text{THz} \sim 10^{-5}$). Thus, the scattering is very nearly elastic, and $|k_{0,m}| = k_0$ is a very good approximation. This constraint of elasticity coupled with knowledge of the scattering angle θ_m allows determination of the scattered wavevector

$$k_{0,m} = (mk)\hat{x} + k_0 \left[1 - \left(\frac{mk}{k_0} \right)^2 \right]^{1/2} \hat{z}. \tag{A.22}$$

Finally, note that the simultaneous presence of both the upscattered and downscattered beams (a key prediction of the Raman-Nath formalism) under typical experimental conditions has been demonstrated empirically [5, Sec. 2.1]. In the above formalism, the simultaneous upscattering and downscattering of the incident probe beam results from the assumption of a sinusoidal phase fluctuation in (A.10); if a complex exponential were assumed instead, one would erroneously deduce that only one such scattering process (either up or down) occurs.

A.4 VALIDITY OF THE RAMAN-NATH FORMALISM

The Raman-Nath formalism employed in Section A.3 is valid as long as the fluctuation wavenumber k is not "too large". This constraint has been rigorously quantified by Bhatia and Noble [6, 7] and is also discussed by Born and Wolf [1, Ch. 12]. Specifically, beam diffraction will be in the Raman-Nath regime when

$$\delta = \frac{\widetilde{n}_e}{\overline{n}_e} \left(\frac{k_0}{k} \right)^2 \gg 1, \tag{A.23}$$

where $\overline{n}_e$ is the bulk plasma density, $\widetilde{n}_e$ is the amplitude of the plasma-density fluctuation, k_0 is the vacuum wavenumber of the probe beam, and k is the wavenumber of the plasma-density fluctuation. Note that the above δ is equivalent to that used by Bhatia and Noble, after

having made the appropriate substitutions for a CO_2 probe beam in a tokamak plasma. Assuming typical tokamak values

$$\frac{\widetilde{n}_e}{\overline{n}_e} \sim 10^{-3},$$
$$k \lesssim 30 \, \mathrm{cm}^{-1}$$

yields $\delta \gtrsim 40$ for a CO_2 probe beam ($k_0 \approx 2\pi \cdot 10^5 \, \mathrm{m}^{-1}$) such that the beam's diffraction is well within the Raman-Nath regime. In the opposite regime ($\delta \ll 1$) the beam can *Bragg scatter* from the phase fluctuation, producing a single, strongly-scattered beam [7] [1, Ch. 12]; such Bragg scattering is the foundation of acousto-optics.

A.5   WAVENUMBER FILTERING OF THE DIFFRACTED FIELD

Examine again the diffracted field (A.21). The spatial dependence of each term in the summation is wholly governed by the factor $E_G(\boldsymbol{r}_m)$, which corresponds to a Gaussian beam that emanates from the beam waist at an angle θ_m relative to the lab-frame z-axis. The coordinate system $\boldsymbol{r}_m$ is defined such that the m^{th} Gaussian beam propagates along the z_m-axis, and it is related to the lab-frame coordinate system $\boldsymbol{r}$ via (A.16). The wavenumber basis $\boldsymbol{k}_m$ that is dual to $\boldsymbol{r}_m$ is similarly related to the lab-frame wavenumber basis $\boldsymbol{k}$ via

$$\boldsymbol{k}_m = [\boldsymbol{R}(\theta_m)]\boldsymbol{k}, \tag{A.24}$$

which naturally results from the geometric constraint that $\boldsymbol{k} \cdot \boldsymbol{r} = \boldsymbol{k}_m \cdot \boldsymbol{r}_m$.

Imagine now that each of the above Gaussian beams is somehow manipulated based upon its Fourier wavenumber content. Assume this manipulation can be described in terms of a transfer function $T(\boldsymbol{k})$, where $\boldsymbol{k}$ is the lab-frame wavenumber basis. A transfer function of this form is appropriate for investigating e.g. diffraction from an aperture or the action of a phase plate. Using the wavenumber basis transformation (A.24), this transfer function can be written in terms of the $\boldsymbol{k}_m$ wavenumber basis as

$$T(\boldsymbol{k}) = T(\boldsymbol{R}_{-m}\boldsymbol{k}_m), \tag{A.25}$$

where the abbreviation $\boldsymbol{R}_m = \boldsymbol{R}(\theta_m)$ has been adopted. The transformed field E_T then has the Fourier representation in the $\boldsymbol{k}_m$ wavenumber basis

$$E_T(\boldsymbol{k}_m) = T(\boldsymbol{R}_{-m}\boldsymbol{k}_m) \cdot E_G(\boldsymbol{k}_m). \tag{A.26}$$

Inverse Fourier transforming $E_T(\boldsymbol{k}_m)$ yields

$$E_T(\boldsymbol{r}_m) = \frac{1}{(2\pi)^3} \int d\boldsymbol{k}_m \left[T(\boldsymbol{R}_{-m}\boldsymbol{k}_m) e^{i\boldsymbol{k}_m \cdot \boldsymbol{r}_m} \right] E_G(\boldsymbol{k}_m). \tag{A.27}$$

Note that each spectral component $E_G(\boldsymbol{k}_m)$ individually satisfies the wave equation. Thus, the above construction of the transformed field $E_T(\boldsymbol{r}_m)$, which consists of a linear combination of the spectral components with arbitrary amplitudes and phases, *also* satisfies the wave equation. Now, explicitly writing $E_G(\boldsymbol{k}_m)$ as the Fourier transform of E_G over a dummy spatial coordinate $\boldsymbol{r}'$, and exchanging the order of integration yields

$$E_T(\boldsymbol{r}_m) = \frac{1}{(2\pi)^3} \int d\boldsymbol{r}' \, E_G(\boldsymbol{r}') \int d\boldsymbol{k}_m \, T(\boldsymbol{R}_{-m}\boldsymbol{k}_m) e^{i\boldsymbol{k}_m \cdot (\boldsymbol{r}_m - \boldsymbol{r}')}. \tag{A.28}$$

Further, for the applications considered here, it is advantageous to change the variables of integration from the $\boldsymbol{k}_m$ wavenumber basis to the lab-frame wavenumber basis $\boldsymbol{k}$. Note that

$$d\boldsymbol{k}_m = \left| \frac{\partial \boldsymbol{k}_m}{\partial \boldsymbol{k}} \right| d\boldsymbol{k} = |\boldsymbol{R}_m| d\boldsymbol{k} = d\boldsymbol{k} \tag{A.29}$$

and that

$$\begin{aligned} \boldsymbol{k}_m \cdot (\boldsymbol{r}_m - \boldsymbol{r}') &= (\boldsymbol{R}_m \boldsymbol{k}) \cdot (\boldsymbol{r}_m - \boldsymbol{r}') \\ &= (\boldsymbol{R}_m \boldsymbol{k})^\mathsf{T} (\boldsymbol{r}_m - \boldsymbol{r}') \\ &= \boldsymbol{k}^\mathsf{T} \boldsymbol{R}_m^\mathsf{T} (\boldsymbol{r}_m - \boldsymbol{r}') \\ &= \boldsymbol{k} \cdot [\boldsymbol{R}_{-m}(\boldsymbol{r}_m - \boldsymbol{r}')] \end{aligned} \tag{A.30}$$

such that the transformed field $E_T(\boldsymbol{r}_m)$ becomes

$$E_T(\boldsymbol{r}_m) = \frac{1}{(2\pi)^3} \int d\boldsymbol{r}' \, E_G(\boldsymbol{r}') \int d\boldsymbol{k} \, T(\boldsymbol{k}) e^{i\boldsymbol{k} \cdot [\boldsymbol{R}_{-m}(\boldsymbol{r}_m - \boldsymbol{r}')]}, \tag{A.31}$$

where, again, the abbreviation $\boldsymbol{R}_{-m} = \boldsymbol{R}(-\theta_m)$ has been adopted and $\boldsymbol{R}(\theta)$ is the rotation matrix given by (A.17).

To make further progress, a particular form of $T(\boldsymbol{k})$ is needed. Assume that wavenumbers are filtered only in the direction of beam scat-

tering i.e. $T(\boldsymbol{k}) = T(k_x)$. (For example, the groove of a phase plate would be aligned with the lab-frame y-axis to effect such filtering). Then

$$\boldsymbol{R}_{-m}(\boldsymbol{r}_m - \boldsymbol{r}') = \begin{pmatrix} (x_m - x')\cos\theta_m + (z_m - z')\sin\theta_m \\ y_m - y' \\ (z_m - z')\cos\theta_m - (x_m - x')\sin\theta_m \end{pmatrix} \quad (A.32)$$

such that (A.31) becomes

$$\begin{aligned}
E_T(\boldsymbol{r}_m) &= \frac{1}{(2\pi)^3} \int d\boldsymbol{r}'\, E_G(\boldsymbol{r}') \int dk_y\, e^{ik_y(y_m - y')} \\
&\quad \times \int dk_z\, e^{ik_z[(z_m - z')\cos\theta_m - (x_m - x')\sin\theta_m]} \\
&\quad \times \int dk_x\, T(k_x) e^{ik_x[(x_m - x')\cos\theta_m + (z_m - z')\sin\theta_m]} \\
&= \frac{1}{2\pi} \int d\boldsymbol{r}'\, E_G(\boldsymbol{r}')\delta(y_m - y') \\
&\quad \times \delta\big((z_m - z')\cos\theta_m - (x_m - x')\sin\theta_m\big) \\
&\quad \times \int dk_x\, T(k_x) e^{ik_x[(x_m - x')\cos\theta_m + (z_m - z')\sin\theta_m]} \\
&= \frac{1}{2\pi} \int dx'\, E_G\big(x', y_m, z_m - (x_m - x')\tan\theta_m\big) \\
&\quad \times \int dk_x\, T(k_x) e^{ik_x(x_m - x')\sec\theta_m}
\end{aligned} \quad (A.33)$$

Contributions to the integral from regions outside of $|x'| \lesssim w(z_m)$ are suppressed by the Gaussian envelope such that

$$\begin{aligned}
w\big(z_m - (x_m - x')\tan\theta_m\big) &\approx w(z_m), \\
R\big(z_m - (x_m - x')\tan\theta_m\big) &\approx R(z_m), \\
\psi\big(z_m - (x_m - x')\tan\theta_m\big) &\approx \psi(z_m)
\end{aligned}$$

are very good approximations. Note that the phase of the Gaussian beam *cannot* be approximated in such a manner; instead, terms up to first order in θ_m must be retained in the phase. After making these approximations, (A.33) reduces to

$$E_T(\boldsymbol{r}_m) \approx E_G(0, y_m, z_m) \cdot \mathcal{E}(\boldsymbol{r}_m, k), \quad (A.34)$$

where

$$\mathcal{E}(\boldsymbol{r}_m, k) = \frac{e^{-imkx_m}}{2\pi}$$
$$\times \int dx' \, \exp\left[\frac{-x'^2}{w(z_m)^2}\right] \exp\left\{i\left[mkx' + \frac{k_0 x'^2}{2R(z_m)}\right]\right\} \quad (A.35)$$
$$\times \int dk_x \, T(k_x) e^{ik_x(x_m - x')}$$

is a complex-valued function that describes the amplitude and phase transformations that result from filtering the scattered radiation by $T(k_x)$. (Note that (A.34) readily reduces to $E_T(\boldsymbol{r}_m) = E_G(\boldsymbol{r}_m)$ when $T(k_x) = 1$, in agreement with expectations). Generalizing (A.21) to allow for such wavenumber-dependent manipulation yields a total diffracted electric field

$$E(\boldsymbol{r}, t) \approx e^{i\overline{\phi}} \sum_{m=-1}^{1} i^m \left[J_m(\widetilde{\phi}_0)\right] E_T(\boldsymbol{r}_m) e^{-i(\omega_0 + m\omega)t}. \quad (A.36)$$

Manipulating the total diffracted electric field in such a manner forms the foundation of phase contrast imaging (PCI), as is discussed in Section 2.5.

BIBLIOGRAPHY

[1] M. Born and E. Wolf. *Principles of Optics*. Cambridge University Press, Cambridge, United Kingdom, 7th (expanded) edition, 1999.

[2] C. V. Raman and N. S. N Nath. The diffraction of light by high frequency sound waves: Part I. *Proc. Indian Acad. Sci.*, A2(4):406–412, 1936.

[3] C. V. Raman and N. S. N Nath. The diffraction of light by high frequency sound waves: Part III. *Proc. Indian Acad. Sci.*, 3(1):75–84, 1936.

[4] J. B. Fraleigh and R. A. Beauregard. *Linear Algebra*. Addison-Wesley Publishing Company, Reading, MA USA, 3rd edition, 1995.

[5] J. R. Dorris. *Phase contrast imaging measurements and modeling of short wavelength turbulence in the DIII-D tokamak*. PhD thesis, MIT, 2010.

[6] A. B. Bhatia and W. J. Noble. Diffraction of light by ultrasonic waves. I. General theory. *Proc. Roy. Soc., A*, 220(1142):356–368, 1953.

[7] A. B. Bhatia and W. J. Noble. Diffraction of light by ultrasonic waves. II. Approximate expressions for the intensities and comparison with experiment. *Proc. Roy. Soc., A*, 220(1142):369–385, 1953.

B

IMAGING SYSTEMS

B.1 GEOMETRIC OPTICS OF IMAGING SYSTEMS

Let the optical axis of an arbitrary optical system lie along the z-axis, and let all optical rays lie in a plane with the optical axis. At a given position z_j, an optical ray is fully described by its transverse distance ρ to the optical axis and its slope $d\rho/dz$. In the paraxial limit $d\rho/dz \approx \theta$, where θ is the angle between the ray and the optical axis. Ray propagation through homogeneous media and refractive interfaces is well-governed by the so-called $ABCD$ ray-matrix formalism [1, Ch. 15]; that is, a ray propagating from point j to point $j + 1$ evolves as

$$\begin{pmatrix} \rho_{j+1} \\ \theta_{j+1} \end{pmatrix} = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} \rho_j \\ \theta_j \end{pmatrix}, \tag{B.1}$$

where the $ABCD$ matrix elements are determined by the optical properties of the media between points j and $j + 1$. Some rudimentary $ABCD$ matrices are given in Table B.1, while more exhaustive lists can be readily found elsewhere [1, Ch. 15] [2].

An imaging system $\mathcal{I}$, by definition, redirects all rays emanating from transverse position $\rho_\mathcal{O}$ in the object plane $S_\mathcal{O}$ to intersect at transverse position

$$\rho_\mathcal{I} = M\rho_\mathcal{O} \tag{B.2}$$

optical element	ABCD matrix
propagation by distance d through medium of constant index of refraction, N	$\begin{pmatrix} 1 & d \\ 0 & 1 \end{pmatrix}$
thin lens with focal length f	$\begin{pmatrix} 1 & 0 \\ -1/f & 1 \end{pmatrix}$

Table B.1: Some useful $ABCD$ ray matrices.

in the image plane $S_{\mathcal{I}}$. Here, M is the *magnification* of the imaging system, and $M < 0$ implies that the image is inverted relative to the object. The imaging system's A, B, and D matrix elements are easily determined by inspection. Recalling the ray-matrix definitions in (B.1), note that $\rho_{\mathcal{I}} = M\rho_{\mathcal{O}} = A\rho_{\mathcal{O}} + B\theta_{\mathcal{O}}$ such that $A = M$ and $B = 0$. Further, assuming the image plane and object plane refractive indices are identical, as is often the case, the determinant of the ray matrix is unity (i.e. $AD - BC = 1$) [3] such that $D = 1/M$. The final matrix element C is determined by the particulars of the imaging system; for propagation through "simple" optical components, such as lenses and homogeneous media, C is constrained to be real. Thus, an imaging system of magnification M is characterized by an ABCD ray matrix of the form

$$\mathcal{I} = \begin{pmatrix} M & 0 \\ C & 1/M \end{pmatrix}. \tag{B.3}$$

The symmetry axis of a Gaussian beam behaves as a ray in the geometric-optics sense [2]. Application of the imaging ABCD ray matrix (B.3) to a beam scattered by θ_m in the object plane (i.e. see (2.42)) indicates that this beam will be rotated by angle θ_m/M relative to the unscattered beam in the image plane, as shown in Figure B.1. Further, the imaging optics do *not* alter the magnitude of the beam's wavevector, i.e. $|\mathcal{I}(k_m)| = k_0$ (this readily follows from the fact that the imaging optics do not alter the energy of the beam's constituent photons). Knowledge of the wavevector's image-plane magnitude and orientation allows determination of the image-plane wavevector as

$$\mathcal{I}(\boldsymbol{k}_{0,m}) = \boldsymbol{k}_{0,m,\mathcal{I}} = (mk_{\mathcal{I}})\,\hat{\boldsymbol{x}} + k_0 \left[1 - \left(\frac{mk_{\mathcal{I}}}{k_0}\right)^2\right]^{1/2} \hat{\boldsymbol{z}}, \tag{B.4}$$

where

$$k_{\mathcal{I}} \equiv \frac{k}{M} \tag{B.5}$$

is the *imaged* wavenumber of the corresponding object-plane phase fluctuation (2.39).

B.2   GAUSSIAN-BEAM TRANSFORMATION IN IMAGING SYSTEMS

In addition to manipulating the ray-like trajectory of a Gaussian beam's symmetry axis, an imaging system also alters other important proper-
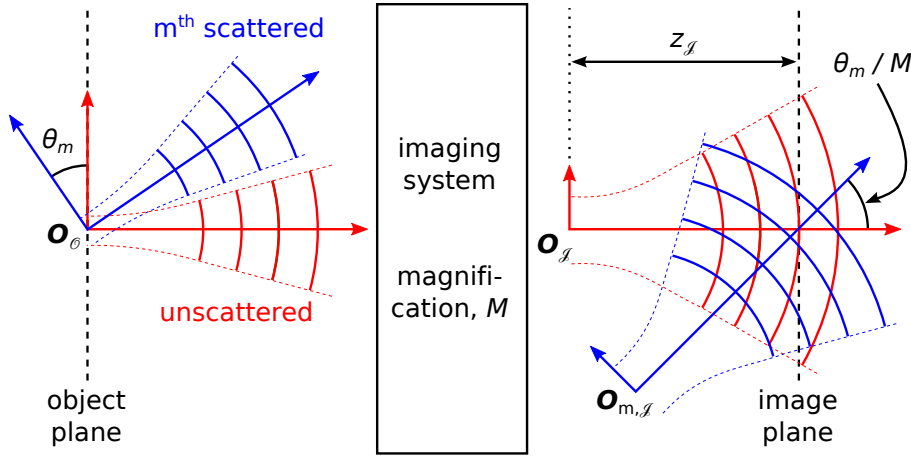
Figure B.1: Beam geometries in an imaging system with magnification M. Beam scattering occurs in the object plane at the probe beam's waist. Thus, the m^{th} scattered beam shares the origin $O_{\mathcal{O}}$ with the unscattered beam but is angularly separated by θ_m. The imaging system redirects all beams emanating from $O_{\mathcal{O}}$ to intersect at angle θ_m/M in the image plane. In general, the image plane does *not* sit at a beam waist such that the post-imaging-system beam waists of the scattered and unscattered beams do not coincide, i.e. $O_{\mathcal{I}} \neq O_{m,\mathcal{I}}$.

ties of an incident Gaussian beam. A Gaussian beam is fully characterized by its in-medium wavelength λ_0/N, its width $w(z_j)$, and its radius of curvature $R(z_j)$ at a single location z_j. These parameters can be conveniently combined to define the so-called complex beam parameter q [1, Sec. 17.1]

$$\frac{1}{q} \equiv \frac{1}{R} - i\left(\frac{\lambda_0}{N\pi w^2}\right). \tag{B.6}$$

Referencing the Gaussian-beam width (2.34) and the Gaussian-beam radius of curvature (2.35), the complex beam parameter can be rewritten as

$$q = z + iz_R, \tag{B.7}$$

where z is the axial distance from the beam waist and z_R is the Rayleigh range (2.37). The Gaussian beam can then be propagated from point j to point $j + 1$ via

$$q_{j+1} = \frac{Aq_j + B}{Cq_j + D}, \tag{B.8}$$

where, amazingly, A, B, C, and D are equal to the corresponding values of the ABCD ray matrix from geometric optics [1, Sec. 20.2] [2].

The beam's transverse and angular displacements relative to the lab-frame optical axis are similarly governed by the so-called "S-parameter transformation" [2]. It is important to note that the complex beam parameter and its evolution are *independent* of the beam's transverse and angular displacements from the lab-frame optical axis (assuming such displacements do not violate the paraxial limit, of course); that is, a scattered beam's width and radius of curvature evolve identically to those of the unscattered beam.

The properties of the image-plane beams are easily determined. Using the ray matrix of an imaging system from (B.3), the image-plane complex beam parameter $q_{\mathfrak{I}}$ is given as

$$q_{\mathfrak{I}} = \frac{M q_{\mathfrak{O}}}{C q_{\mathfrak{O}} + (1/M)},$$ (B.9)

where $q_{\mathfrak{O}}$ is the object-plane complex beam parameter, and M and C are both real. Note that the post-imaging-system beam waists do *not* necessarily sit at the image plane, in which case the beams' native coordinate systems are necessarily *displaced* from each other (i.e. $O_{\mathfrak{I}} \neq O_{m,\mathfrak{I}}$, as indicated in Figure B.1). Examining the image-plane complex beam parameter (B.9) it is easy to see that the beam waists will not sit at the image plane when $|C q_{\mathfrak{O}}| \gg 1/M$ such that $z_{\mathfrak{I}} = \mathrm{Re}(q_{\mathfrak{I}}) \approx M/C \neq 0$.

As the native coordinate systems of the unscattered beam and the m^{th} scattered beam do not align in the image plane, it will be convenient to determine the relevant coordinate transformation. The transformation is derived for the most general case in which the beam waists do not sit at the image plane (i.e. $O_{\mathfrak{I}} \neq O_{m,\mathfrak{I}}$, as indicated in Figure B.1). The coordinate transformation is simply a series of translations and rotations

$$\boldsymbol{r}_{m,\mathfrak{I}} = [\boldsymbol{R}(\theta_m/M)]\,[\boldsymbol{r}_{\mathfrak{I}} - z_{\mathfrak{I}}\hat{\boldsymbol{z}}] + z_{\mathfrak{I}}\hat{\boldsymbol{z}},$$ (B.10)

where

$$\boldsymbol{R}(\theta) = \begin{pmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos\theta \end{pmatrix}$$ (B.11)

is the rotation matrix that rotates the (x, z)-plane about the y-axis by angle θ. Explicitly, the image-plane coordinate transformation (B.10) is

$$
\begin{pmatrix} x_{m,\mathfrak{I}} \\ y_{m,\mathfrak{I}} \\ z_{m,\mathfrak{I}} \end{pmatrix} = \begin{pmatrix} x_{\mathfrak{I}} \cos\left(\frac{\theta_m}{M}\right) \\ y_{\mathfrak{I}} \\ z_{\mathfrak{I}} + x_{\mathfrak{I}} \sin\left(\frac{\theta_m}{M}\right) \end{pmatrix} \approx \begin{pmatrix} x_{\mathfrak{I}} \\ y_{\mathfrak{I}} \\ z_{\mathfrak{I}} + x_{\mathfrak{I}}\left(\frac{\theta_m}{M}\right) \end{pmatrix}, \qquad \text{(B.12)}
$$

where the approximation is valid to first order in θ_m/M.

Imagine now that the detector is located an axial distance $\delta z_{\mathfrak{I}}$ downstream of the image plane (i.e. $z_{\text{det}} = z_{\mathfrak{I}} + \delta z_{\mathfrak{I}}$ such that positive $\delta z_{\mathfrak{I}}$ implies that the detector is downstream of the image plane, and negative $\delta z_{\mathfrak{I}}$ implies that the detector is upstream of the image plane). Coordinate transformation (B.10) then readily generalizes to

$$
\boldsymbol{r}_{m,\text{det}} = [\boldsymbol{R}(\theta_m/M)] \left[\boldsymbol{r}_{\text{det}} - z_{\mathfrak{I}}\hat{\boldsymbol{z}}\right] + z_{\mathfrak{I}}\hat{\boldsymbol{z}}, \qquad \text{(B.13)}
$$

where $\boldsymbol{r}_{\text{det}} = \boldsymbol{r}_{\mathfrak{I}} + (\delta z_{\mathfrak{I}})\hat{\boldsymbol{z}}$ specifies the detector-plane coordinates of the unscattered beam and $\boldsymbol{r}_{m,\text{det}}$ specifies the detector-plane coordinates of the m^{th} scattered beam. Explicitly, the detector-plane coordinate transformation (B.13) is

$$
\begin{pmatrix} x_{m,\text{det}} \\ y_{m,\text{det}} \\ z_{m,\text{det}} \end{pmatrix} = \begin{pmatrix} x_{\text{det}} \cos\left(\frac{\theta_m}{M}\right) - \delta z_{\mathfrak{I}} \sin\left(\frac{\theta_m}{M}\right) \\ y_{\text{det}} \\ z_{\mathfrak{I}} + \delta z_{\mathfrak{I}} \cos\left(\frac{\theta_m}{M}\right) + x_{\mathfrak{I}} \sin\left(\frac{\theta_m}{M}\right) \end{pmatrix}
$$

$$
\approx \boldsymbol{r}_{\text{det}} + \frac{\theta_m}{M} \begin{pmatrix} -\delta z_{\mathfrak{I}} \\ 0 \\ x_{\text{det}} \end{pmatrix} - \frac{1}{2}\left(\frac{\theta_m}{M}\right)^2 \begin{pmatrix} x_{\text{det}} \\ 0 \\ \delta z_{\mathfrak{I}} \end{pmatrix}, \qquad \text{(B.14)}
$$

where the approximation is valid to second order in θ_m/M. As discussed in Section 3.1.5, second-order effects can become significant when the detector is displaced from the image plane (i.e. when $\delta z_{\mathfrak{I}} \neq 0$).

BIBLIOGRAPHY

[1] A. E. Siegman. *Lasers*. University Science Books, Sausalito, CA, USA, 1986.

[2] A. A. Tovar and L. W. Casperson. Generalized beam matrices. IV. Optical system design. *J. Opt. Soc. Am. A*, 14(4):882–894, 1997.

[3] K. Halbach. Matrix representation of Gaussian optics. *Am. J. Phys.*, 32(2):90–108, 1964.

C

SOME IDENTITIES FOR THE PCI WAVENUMBER RESPONSE

C.1 $\mathcal{E}(r_m, k)$ IN THE BEAM'S NEAR FIELD

The effect of wavenumber-dependent manipulation $T(k_x)$ on the m^{th} scattered beam is given by the complex-valued function $\mathcal{E}(r_m, k)$ as defined in (2.45), which is repeated here for completeness

$$
\mathcal{E}(r_m, k) = \frac{e^{-imkx_m}}{2\pi}
$$
$$
\times \int dx' \exp\left[\frac{-x'^2}{w(z_m)^2}\right] \exp\left\{i\left[mkx' + \frac{k_0 x'^2}{2R(z_m)}\right]\right\} \quad \text{(C.1)}
$$
$$
\times \int dk_x\, T(k_x) e^{ik_x(x_m - x')}
$$

The integrals are over the full domain of x' and k_x, but contributions to the integral from regions outside of $|x'| \lesssim w(z)$ are suppressed by the Gaussian envelope such that

$$
\frac{k_0 x'^2}{2R(z_m)} \lesssim \frac{k_0 [w(z)]^2}{2R(z)}, \quad \text{(C.2)}
$$

where the approximations $w(z_m) \approx w(z)$ and $R(z_m) \approx R(z)$ have been used. Now, in the beam's near field ($z \ll z_R$, which is often experimentally relevant), the beam's waist and radius of curvature are

$$
w(z) \approx w_0, \quad \text{(C.3)}
$$
$$
R(z) \approx \frac{z_R^2}{z}, \quad \text{(C.4)}
$$

as can be easily verified by examining the definitions in (2.34) and (2.35), respectively. Thus, (C.2) becomes

$$
\begin{aligned}
\frac{k_0 x'^2}{2R(z_m)} &\lesssim \frac{k_0 [w(z)]^2}{2R(z)} \\
&\approx \frac{k_0 w_0^2}{2(z_R^2/z)} \\
&= \frac{z}{z_R} \\
&\ll 1,
\end{aligned}
\tag{C.5}
$$

where $z/z_R \ll 1$ follows from the near-field assumption. Thus, in the beam's near field, the curvature-induced phase factor is negligible, and $\mathcal{E}(\boldsymbol{r}_m, k)$ reduces to

$$
\begin{aligned}
\mathcal{E}(\boldsymbol{r}_m, k) &\approx \frac{e^{-imk x_m}}{2\pi} \int dx' \, e^{-[x'/w(z_m)]^2} e^{imk x'} \\
&\times \int dk_x \, T(k_x) e^{ik_x(x_m - x')}.
\end{aligned}
\tag{C.6}
$$

This near-field assumption will be implicit in the remainder of the discussion about PCI.

C.2   PCI'S $\mathcal{E}(r_m, k)$

The transfer function of the PCI phase plate can be described as

$$
\begin{aligned}
T(k_x) &= i\sqrt{\eta}\, H(k_g - |k_x|) \\
&+ H(|k_x| - k_g) H(k_D - |k_x|),
\end{aligned}
\tag{C.7}
$$

where $H(x)$ is the Heaviside step function defined as

$$
H(x) = \begin{cases} 0, & x < 0 \\ 1, & x \geqslant 0 \end{cases},
\tag{C.8}
$$

η is the reflectivity of the phase-plate groove, and k_g and k_D are the low-k and high-k cutoffs of the phase plate as defined in (2.86) and (2.89), respectively. Note that the first term on the right-hand side of (C.7) corresponds to reflection from the phase-plate groove, while the

second term corresponds to reflection from the non-grooved portion of the phase plate (i.e. the "face"). Thus, for PCI

$$
\mathcal{E}(r_m, k) = \frac{e^{-imkx_m}}{2\pi} \int dx' \, e^{-[x'/w(z_m)]^2} e^{imkx'}
$$

$$
\times \left\{ \int_{-k_D}^{-k_g} dk_x' \, e^{ik_x'(x_m - x')} \right.
$$

$$
+ i\sqrt{\eta} \int_{-k_g}^{k_g} dk_x' \, e^{ik_x'(x_m - x')}
$$

$$
\left. + \int_{k_g}^{k_D} dk_x' \, e^{ik_x'(x_m - x')} \right\}.
$$

(C.9)

C.3   SOME USEFUL INTEGRALS FOR EVALUATION OF $\mathcal{E}(r_m, k)$

C.3.1   *Finite-domain inverse Fourier transforms of unity*

Note that

$$
\int_{k_1}^{k_2} dk_x e^{ik_x x} = \frac{e^{ik_2 x} - e^{ik_1 x}}{ix}.
$$

(C.10)

Now, if $k_1 = -k_2$, this simplifies to

$$
\int_{-k_2}^{k_2} dk_x e^{ik_x x} = 2k_2 \mathrm{sinc}\left(\frac{k_2 x}{\pi}\right),
$$

(C.11)

where

$$
\mathrm{sinc}(x) = \frac{\sin(\pi x)}{\pi x}
$$

(C.12)

is the normalized sinc function; note that sinc is an *even* function. Finally, note that

$$
\int_{-k_2}^{-k_1} dk_x e^{ik_x x} + \int_{k_1}^{k_2} dk_x e^{ik_x x}
$$

$$
= \int_{-k_2}^{k_2} dk_x e^{ik_x x} - \int_{-k_1}^{k_1} dk_x e^{ik_x x}
$$

$$
= 2k_2 \mathrm{sinc}\left(\frac{k_2 x}{\pi}\right) - 2k_1 \mathrm{sinc}\left(\frac{k_1 x}{\pi}\right).
$$

(C.13)

Using (C.11) and (C.13), it is easy to see that (C.9) becomes

$$
\begin{aligned}
\mathcal{E}(\boldsymbol{r}_{\mathrm{m}},k) = \frac{e^{-imkx_{\mathrm{m}}}}{\pi} & \int dx'\, e^{-[x'/w(z)]^2} e^{imkx'} \\
\times \Bigg\{ & k_{\mathrm{D}}\mathrm{sinc}\left[\frac{k_{\mathrm{D}}}{\pi}(x'-x_{\mathrm{m}})\right] \\
& - k_g\mathrm{sinc}\left[\frac{k_g}{\pi}(x'-x_{\mathrm{m}})\right] \\
& + i\sqrt{\eta}k_g\mathrm{sinc}\left[\frac{k_g}{\pi}(x'-x_{\mathrm{m}})\right]\Bigg\}.
\end{aligned}
\tag{C.14}
$$

C.3.2   *Integral of offset sinc with complex-Gaussian weighting*

Note that (C.14) consists of several integrals of the form

$$
I \equiv \frac{b}{\pi}\int dx\, e^{-ax^2} e^{icx}\mathrm{sinc}\left[\frac{b}{\pi}(x-x_0)\right],
\tag{C.15}
$$

where $a > 0$ and x_0, b, and c are real. While daunting, the integral can be evaluated "analytically" in terms of complex error functions as follows

$$
\begin{aligned}
I &= \frac{b}{\pi} \int dx\, e^{-ax^2} e^{icx} \operatorname{sinc}\left[\frac{b}{\pi}(x - x_0)\right] \\
&= \frac{1}{\pi} \int dx\, e^{-ax^2} e^{icx} \cdot \frac{\sin[b(x - x_0)]}{x - x_0} \\
&= \frac{1}{\pi} \int_0^b d\beta \int dx\, e^{-ax^2} e^{icx} \cos[\beta(x - x_0)] \\
&= \frac{1}{2\pi} \int_0^b d\beta \int dx\, e^{-ax^2} e^{icx} \left[e^{i\beta(x-x_0)} + e^{-i\beta(x-x_0)}\right] \\
&= \frac{1}{2\pi} \int_0^b d\beta \int dx\, e^{-ax^2} e^{i[(\beta+c)x - \beta x_0]} \\
&\quad + \frac{1}{2\pi} \int_0^b d\beta \int dx\, e^{-ax^2} e^{i[-(\beta-c)x + \beta x_0]} \\
&= \frac{1}{2\pi} \int_0^b d\beta \int dx\, e^{-ax^2} e^{i[(\beta+c)x - \beta x_0]} \\
&\quad - \frac{1}{2\pi} \int_0^{-b} d\beta' \int dx\, e^{-ax^2} e^{i[(\beta'+c)x - \beta' x_0]} \\
&= \frac{1}{2\pi} \int_{-b}^b d\beta \int dx\, e^{-ax^2} e^{i[(\beta+c)x - \beta x_0]} \\
&= \frac{1}{2\pi} \int_{-b}^b d\beta\, e^{-(\beta+c)^2/4a} e^{-i\beta x_0} \int dx\, e^{-a[x - i(\beta+c)/2a]^2} \\
&= \frac{1}{2\sqrt{\pi a}} \int_{-b}^b d\beta\, e^{-(\beta+c)^2/4a} e^{-i\beta x_0} \\
&= \frac{1}{2\sqrt{\pi a}} e^{icx_0} \int_{c-b}^{c+b} d\beta'\, e^{-\beta'^2/4a} e^{-i\beta' x_0} \\
&= \frac{1}{2\sqrt{\pi a}} e^{-ax_0^2} e^{icx_0} \int_{c-b}^{c+b} d\beta'\, e^{-(\beta'+i2ax_0)^2/4a} \\
&= \frac{1}{\sqrt{\pi}} e^{-ax_0^2} e^{icx_0} \int_{u(c-b,a,x_0)}^{u(c+b,a,x_0)} du\, e^{-u^2} \\
&= \frac{1}{2} e^{-ax_0^2} e^{icx_0} \left\{ \operatorname{erf}[u(c + b, a, x_0)] \right. \\
&\qquad\qquad \left. - \operatorname{erf}[u(c - b, a, x_0)] \right\},
\end{aligned}
\tag{C.16}
$$

where the error function is defined for complex argument z as

$$
\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt,
\tag{C.17}
$$

and

$$u(\gamma, a, x_0) = \frac{1}{2\sqrt{a}}(\gamma + i2ax_0). \tag{C.18}$$

Now, substituting the appropriate values into (C.16)

$$x_0 \equiv x_m, \qquad a \equiv \frac{1}{w(z)^2}, \qquad b \equiv k_j, \qquad c \equiv mk$$

yields

$$I = \frac{1}{2}e^{-[x_m/w(z_m)]^2}e^{imkx_m}\mathcal{D}(\boldsymbol{r}_m, k, k_j), \tag{C.19}$$

where the difference function $\mathcal{D}$ is defined as

$$\mathcal{D}(\boldsymbol{r}_m, k, k_j) = \mathrm{erf}[u(\boldsymbol{r}_m, k, k_j)] - \mathrm{erf}[u(\boldsymbol{r}_m, k, -k_j)], \tag{C.20}$$

and

$$u(\boldsymbol{r}_m, k, k_j) = \frac{w(z_m)}{2}\left[(mk + k_j) + i\frac{2x_m}{w(z_m)^2}\right]. \tag{C.21}$$

With these definitions (C.14) readily reduces to

$$\begin{aligned}\mathcal{E}(\boldsymbol{r}_m, k) = \frac{1}{2}&e^{-[x_m/w(z_m)]^2} \\ &\times [\mathcal{D}(\boldsymbol{r}_m, k, k_D) + (i\sqrt{\eta} - 1)\mathcal{D}(\boldsymbol{r}_m, k, k_g)].\end{aligned} \tag{C.22}$$

C.4   SYMMETRIES AND DEGENERACIES IN THE IMAGE PLANE

The plasma midplane sits at the object plane of a magnification-M imaging system. The probe beam's waist also nominally sits at this object plane, with a 1/e E radius of $w_{0,\mathcal{O}}$, and the plasma dimensions are far smaller than the corresponding Rayleigh length such that throughout the plasma volume

$$w(z) \approx w_{0,\mathcal{O}}. \tag{C.23}$$

As discussed in Appendix B, the waist of the imaged beam does *not* necessarily occur at the image plane. However, the derivation of (C.6) required assuming that the beam was well within its Rayleigh range (i.e. $|z_{\mathcal{I}}| \ll z_{R,\mathcal{I}}$). Thus, valid application of any of the above derived

results to the image plane requires that the image plane and the waist of the imaged beam approximately overlap such that

$$w(z_\mathrm{J}) \approx w_{0,\mathrm{J}} \approx M w_{0,\mathrm{O}}, \tag{C.24}$$

where $w_{0,\mathrm{J}}$ is the $1/e\,E$ radius of the imaged beam at its waist.

The coordinate system and properties of the m^{th} scattered beam also have certain symmetries in the image plane. In particular, according to the image-plane coordinate transformation in (2.48),

$$\begin{aligned}
x_{-m,\mathrm{J}} &= x_\mathrm{J} \cos\left(\frac{\theta_{-m}}{M}\right) \\
&= x_\mathrm{J} \cos\left(\frac{-\theta_m}{M}\right) \\
&= x_\mathrm{J} \cos\left(\frac{\theta_m}{M}\right) \\
&= x_{m,\mathrm{J}}. \tag{C.25}
\end{aligned}$$

Additionally, again from the image-plane coordinate transformation in (2.48), $z_{-m,\mathrm{J}} = z_\mathrm{J} + x_\mathrm{J} \sin(\theta_{-m}/M)$ such that

$$w(z_{-m,\mathrm{J}}) \approx w(z_\mathrm{J}) \approx w(z_{m,\mathrm{J}}) \tag{C.26}$$

are very good approximations when $\theta_m/M \ll 1$, as is typical. The above symmetries lead to several other symmetries and degeneracies that will be useful for evaluating $\mathcal{E}(r_m, k)$ in the image plane.

C.4.1   *Properties of* u *in the image plane*

The complex-valued function u is defined in (C.21). Referencing symmetry properties (C.25) and (C.26) and cleverly rearranging terms, it readily follows that u is anti-Hermitian with respect to m and $k_{j,\mathfrak{I}}$:

$$
\begin{aligned}
u(\boldsymbol{r}_{-m,\mathfrak{I}}, k_{\mathfrak{I}}, -k_{j,\mathfrak{I}}) &= \frac{w(z_{-m,\mathfrak{I}})}{2}\left[(-mk_{\mathfrak{I}} - k_{j,\mathfrak{I}}) + i\frac{2x_{-m,\mathfrak{I}}}{w(z_{-m,\mathfrak{I}})^2}\right] \\
&\approx \frac{w(z_{m,\mathfrak{I}})}{2}\left[(-mk_{\mathfrak{I}} - k_{j,\mathfrak{I}}) + i\frac{2x_{m,\mathfrak{I}}}{w(z_{m,\mathfrak{I}})^2}\right] \\
&= \frac{w(z_{m,\mathfrak{I}})}{2}\left[-(mk_{\mathfrak{I}} + k_{j,\mathfrak{I}}) + i\frac{2x_{m,\mathfrak{I}}}{w(z_{m,\mathfrak{I}})^2}\right] \\
&= \frac{-w(z_{m,\mathfrak{I}})}{2}\left[(mk_{\mathfrak{I}} + k_{j,\mathfrak{I}}) - i\frac{2x_{m,\mathfrak{I}}}{w(z_{m,\mathfrak{I}})^2}\right] \\
&= \frac{-w(z_{m,\mathfrak{I}})}{2}\left[(mk_{\mathfrak{I}} + k_{j,\mathfrak{I}}) + i\frac{2x_{m,\mathfrak{I}}}{w(z_{m,\mathfrak{I}})^2}\right]^* \\
&= -[u(\boldsymbol{r}_{m,\mathfrak{I}}, k_{\mathfrak{I}}, k_{j,\mathfrak{I}})]^*,
\end{aligned}
\tag{C.27}
$$

where z^* indicates the complex conjugate of z. Further, referencing (C.24), note that

$$
\begin{aligned}
u(\boldsymbol{r}_{m,\mathfrak{I}}, k_{\mathfrak{I}}, k_{j,\mathfrak{I}}) &= \frac{w(z_{\mathfrak{I}})}{2}\left[(mk_{\mathfrak{I}} + k_{j,\mathfrak{I}}) + i\frac{2x_{\mathfrak{I}}}{w(z_{\mathfrak{I}})^2}\right] \\
&\approx \frac{Mw_{0,\mathfrak{O}}}{2}\left[\left(\frac{mk}{M} + \frac{k_j}{M}\right) + i\frac{2Mx_{\mathfrak{O}}}{(Mw_{0,\mathfrak{O}})^2}\right] \\
&= \frac{w_{0,\mathfrak{O}}}{2}\left[(mk + k_j) + i\frac{2x_{\mathfrak{O}}}{(w_{0,\mathfrak{O}})^2}\right] \\
&= u(\boldsymbol{r}_{m,\mathfrak{O}}, k, k_{j,\mathfrak{O}});
\end{aligned}
\tag{C.28}
$$

that is, $u(\boldsymbol{r}_{m,\mathfrak{I}}, k_{\mathfrak{I}}, k_{j,\mathfrak{I}})$ and $u(\boldsymbol{r}_{m,\mathfrak{O}}, k, k_{j,\mathfrak{O}})$ are geometrically *similar*, as is expected in an imaging system.

C.4.2   *Properties of the error function*

The error function has two useful properties that will be exploited shortly. First, the error function is *odd*

$$
\mathrm{erf}(-z) = -\mathrm{erf}(z),
\tag{C.29}
$$

as is easily determined by inspection. Second, the error function *commutes* with complex conjugation

$$\mathrm{erf}(z^*) = [\mathrm{erf}(z)]^*, \tag{C.30}$$

where z^* is the complex conjugate of z.

C.4.3 *Properties of $\mathcal{D}$ in the image plane*

The complex-valued difference function $\mathcal{D}$ is defined in (C.20). Using the anti-Hermitian properties of u from (C.27) and the error-function properties (C.29) and (C.30), it is easy to show that $\mathcal{D}$ is Hermitian with respect to m:

$$\begin{aligned}
\mathcal{D}(\mathbf{r}_{-m,\jmath}, k_\jmath, k_{j,\jmath}) &= \mathrm{erf}[u(\mathbf{r}_{-m,\jmath}, k_\jmath, k_{j,\jmath})] \\
&\quad - \mathrm{erf}[u(\mathbf{r}_{-m,\jmath}, k_\jmath, -k_{j,\jmath})] \\
&= \mathrm{erf}\{-[u(\mathbf{r}_{m,\jmath}, k_\jmath, -k_{j,\jmath})]^*\} \\
&\quad - \mathrm{erf}\{-[u(\mathbf{r}_{m,\jmath}, k_\jmath, k_{j,\jmath})]^*\} \\
&= -\{\mathrm{erf}[u(\mathbf{r}_{m,\jmath}, k_\jmath, -k_{j,\jmath})]\}^* \\
&\quad + \{\mathrm{erf}[u(\mathbf{r}_{m,\jmath}, k_\jmath, k_{j,\jmath})]\}^* \\
&= \{\mathrm{erf}[u(\mathbf{r}_{m,\jmath}, k_\jmath, k_{j,\jmath})] \\
&\quad - \mathrm{erf}[u(\mathbf{r}_{m,\jmath}, k_\jmath, -k_{j,\jmath})]\}^* \\
&= \mathcal{D}^*(\mathbf{r}_{m,\jmath}, k_\jmath, k_{j,\jmath}). \tag{C.31}
\end{aligned}$$

The above symmetry relation also implies a degeneracy when $m = 0$; namely, $\mathcal{D}(\mathbf{r}_{0,\jmath}, k_\jmath, k_{j,\jmath}) = \mathcal{D}^*(\mathbf{r}_{0,\jmath}, k_\jmath, k_{j,\jmath})$, which proves that $\mathcal{D}(\mathbf{r}_{0,\jmath}, k_\jmath, k_{j,\jmath})$ is purely *real*.

C.4.4 *Properties of $\mathcal{E}$ in the image plane*

Eq. (C.31) allows (C.22) to be rewritten as

$$\begin{aligned}
\mathcal{E}(\mathbf{r}_{m,\jmath}, k_\jmath) &= e^{-[x_{m,\jmath}/w(z_{m,\jmath})]^2} \\
&\quad \times [F(\mathbf{r}_{m,\jmath}, k_\jmath) + G(\mathbf{r}_{m,\jmath}, k_\jmath)], \tag{C.32}
\end{aligned}$$

where

$$F(\mathbf{r}_{m,\jmath}, k_\jmath) = \frac{1}{2}\left[\mathcal{D}(\mathbf{r}_{m,\jmath}, k_\jmath, k_{D,\jmath}) - \mathcal{D}(\mathbf{r}_{m,\jmath}, k_\jmath, k_{g,\jmath})\right], \tag{C.33}$$

$$G(\mathbf{r}_{m,\jmath}, k_\jmath) = \frac{i\sqrt{\eta}}{2}\mathcal{D}(\mathbf{r}_{m,\jmath}, k_\jmath, k_{g,\jmath}). \tag{C.34}$$

Here, the notation is mnemonic: the non-grooved portion of the phase plate (i.e. the "face") acts on the m^{th} scattered beam via F, while the phase-plate groove acts on the m^{th} scattered beam via G. Note that F is Hermitian with respect to m

$$F(\boldsymbol{r}_{-m,\jmath}, k_\jmath) = F^*(\boldsymbol{r}_{m,\jmath}, k_\jmath), \qquad (C.35)$$

while G is anti-Hermitian with respect to m

$$G(\boldsymbol{r}_{-m,\jmath}, k_\jmath) = -G^*(\boldsymbol{r}_{m,\jmath}, k_\jmath). \qquad (C.36)$$

Note that the above symmetries also imply a degeneracy when $m = 0$. Specifically, (C.35) states that $F(\boldsymbol{r}_{0,\jmath}, k_\jmath) = F^*(\boldsymbol{r}_{0,\jmath}, k_\jmath)$; that is, $F(\boldsymbol{r}_{0,\jmath}, k_\jmath)$ is purely *real*

$$F(\boldsymbol{r}_{0,\jmath}, k_\jmath) = \mathrm{Re}[F(\boldsymbol{r}_{0,\jmath}, k_\jmath)]. \qquad (C.37)$$

Similarly, (C.36) states that $G(\boldsymbol{r}_{0,\jmath}, k_\jmath) = -G^*(\boldsymbol{r}_{0,\jmath}, k_\jmath)$; that is, $G(\boldsymbol{r}_{0,\jmath}, k_\jmath)$ is purely *imaginary*

$$G(\boldsymbol{r}_{0,\jmath}, k_\jmath) = i \cdot \mathrm{Im}[G(\boldsymbol{r}_{0,\jmath}, k_\jmath)]. \qquad (C.38)$$

Another set of useful degeneracies occurs when the fluctuation wavenumber vanishes (i.e. $k_\jmath = 0$). Note that the $k_\jmath$ dependence of F and G only appears as the *product* $(m \cdot k_\jmath)$. Just as $m = 0$ at finite $k_\jmath$ yields $(m \cdot k_\jmath) = 0$, $k_\jmath = 0$ at finite m also gives $(m \cdot k_\jmath) = 0$; thus,

$$F(\boldsymbol{r}_{m,\jmath}, k_\jmath = 0) = F(\boldsymbol{r}_{0,\jmath}, k_\jmath), \qquad (C.39)$$
$$G(\boldsymbol{r}_{m,\jmath}, k_\jmath = 0) = G(\boldsymbol{r}_{0,\jmath}, k_\jmath), \qquad (C.40)$$

where the approximation $w(z_{m,\jmath}) \approx w(z_{0,\jmath})$ has been used. The PCI amplitude response $A_{pci}(k_\jmath, x_\jmath)$ is given by (2.99); of particular relevance is that $A_{pci}(k_\jmath, x_\jmath) \propto A(k_\jmath, x_\jmath)$, where $A(k_\jmath, x_\jmath) = (A_I^2 + A_Q^2)^{1/2}$ is a real number and A_I and A_Q are given by (2.102) and (2.103), respectively. Note that the notational shorthand $F_m \equiv F(\boldsymbol{r}_{m,\jmath}, k_\jmath)$ and $G_m \equiv G(\boldsymbol{r}_{m,\jmath}, k_\jmath)$ is being used in the expressions for A_I and A_Q. Using (C.39) and (C.40), the expression for A_I in the low-k limit becomes

$$\lim_{k_\jmath \to 0} A_I(k_\jmath, x_\jmath) = \mathrm{Im}(G_0) \left[\lim_{k_\jmath \to 0} \mathrm{Re}(F_1) \right] - \mathrm{Re}(F_0) \left[\lim_{k_\jmath \to 0} \mathrm{Im}(G_1) \right]$$
$$= \mathrm{Im}(G_0)\mathrm{Re}(F_0) - \mathrm{Re}(F_0)\mathrm{Im}(G_0)$$
$$= 0.$$

Similarly, the expression for A_Q in the low-k limit becomes

$$\lim_{k_J \to 0} A_Q(k_J, x_J) = \text{Im}(G_0) \left[\lim_{k_J \to 0} \text{Im}(F_1) \right] + \text{Re}(F_0) \left[\lim_{k_J \to 0} \text{Re}(G_1) \right]$$

$$= \text{Im}(G_0)\text{Im}(F_0) + \text{Re}(F_0)\text{Re}(G_0)$$

$$= \text{Im}(G_0) \cdot 0 + \text{Re}(F_0) \cdot 0$$

$$= 0,$$

where the third line follows from the fact that F_0 is purely real and G_0 is purely imaginary. As A_I and A_Q vanish at $k = 0$, A and A_{pci} must also vanish at $k = 0$. This is in agreement with expectations: as $k \to 0$, the upscattered and downscattered beams fall into the phase-plate groove, reducing the phase contrast, and the response vanishes fully for $k = 0$.

D

A precision oscillator is well-described in the complex plane via

$$A(t) = [A_0 + \varepsilon(t)]\, e^{i[\omega_0 t + \phi(t)]}, \tag{D.1}$$

where A_0 is the nominal amplitude, $\varepsilon(t)$ is the deviation from the nominal amplitude, ω_0 is the nominal angular frequency, and $\phi(t)$ is the phase deviation from the nominal phase $\omega_0 t$ [1]. Within the context of this work, it is sufficient to regard $\varepsilon(t)$ and $\phi(t)$ as zero-mean, stationary, random processes. Note that both $\varepsilon(t)$ and $\phi(t)$ have simple geometric interpretations, as shown in Figure D.1. Typically, $|\varepsilon(t)| \ll A_0$ such that the amplitude deviation can be neglected. However, phase-sensitive instruments, such as interferometers, can be susceptible to small phase deviations. This appendix discusses the IEEE definition of phase noise $\mathcal{L}(f)$ and quantitatively links it to the corresponding noise of a phase-sensitive instrument.
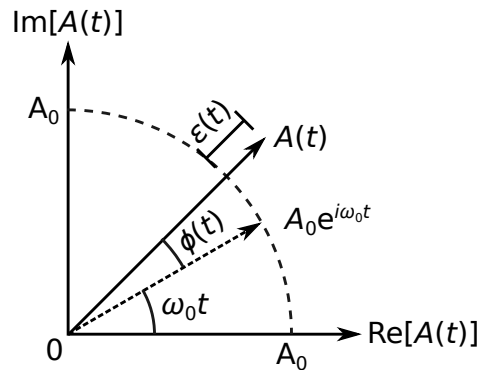
Figure D.1: An oscillator's amplitude deviation $\varepsilon(t)$ and phase deviation $\phi(t)$ can be easily visualized in the complex plane.

D.1 DEFINITIONS

An oscillator's phase noise $\mathcal{L}(f)$ (pronounced as "script-ell of f") is defined as

$$\mathcal{L}(f) = \frac{G_{\phi,\phi}(f)}{2}, \tag{D.2}$$

where $G_{\phi,\phi}(f)$ is the one-sided autospectral density of the phase fluctuations $\phi(t)$ [1]. Note that the one-sided autospectral density $G_{\phi,\phi}(f)$ is related to the two-sided autospectral density $S_{\phi,\phi}(f)$ via

$$G_{\phi,\phi}(f) = \begin{cases} 2S_{\phi,\phi}(f), & f > 0 \\ S_{\phi,\phi}(f), & f = 0 \end{cases} . \tag{D.3}$$

The two-sided autospectral density $S_{\phi,\phi}(f)$ is itself defined as

$$S_{\phi,\phi}(f) = \mathcal{F}\left[R_{\phi,\phi}(\tau)\right](f), \tag{D.4}$$

where $\mathcal{F}$ is the Fourier transform operator and $R_{\phi,\phi}(\tau)$ is the autocorrelation function defined as

$$R_{\phi,\phi}(\tau) = E\left[\phi(t) \cdot \phi(t+\tau)\right]; \tag{D.5}$$

here, E is the expectation-value operator [2, Ch. 5].

D.2   UNITS

Phase noise $\mathcal{L}(f)$ can be expressed in SI units as $\text{rad}^2 \cdot \text{Hz}^{-1}$. However, this is *not* common practice. Instead, it is much more common to express phase noise in units of dBc/Hz, which corresponds to the ratio (expressed in dB) of phase-noise power to carrier-signal power in a one Hz bandwidth. Explicitly,

$$\mathcal{L}(f)\,[\text{dBc/Hz}] = 10\log_{10}\left(\mathcal{L}(f)\,[\text{rad}^2 \cdot \text{Hz}^{-1}]\right). \tag{D.6}$$

While it may initially come as a surprise that phase noise should be expressible in terms of the power in the carrier signal (i.e. the nominal oscillator signal), this choice of units is perfectly natural. To see this, examine Figure D.1: the oscillator's nominal phasor $A_0 e^{i\omega_0 t}$ is displaced in the perpendicular direction by $A(t)\sin\phi(t) \approx A_0\phi(t)$, where the approximation follows for the small amplitude and phase deviations ($|\varepsilon(t)| \ll A_0$, $|\phi(t)| \ll 1$) of practical relevance. Thus, a phase deviation $\phi(t)$ produces a phasor deviation that is proportional to the oscillator's nominal amplitude A_0, making dBc/Hz a natural choice of units for $\mathcal{L}(f)$.

D.3 CONNECTION TO INSTRUMENTAL PHASE NOISE

Numerous phase-sensitive instruments, such as interferometers, demodulators, etc. operate by comparing an oscillator's phase at times t and $t + \tau_j$:

$$\arg[A^*(t) \cdot A(t + \tau_j)] = \omega \tau_j + \delta\phi(t, \tau_j), \tag{D.7}$$

where $\omega \tau_j$ is the expected phase progression and

$$\delta\phi(t, \tau_j) = \phi(t + \tau_j) - \phi(t) \tag{D.8}$$

is the *instrumental* phase noise. Obviously, the instrumental phase noise $\delta\phi(t, \tau_j)$ is very closely related to the oscillator phase deviation $\phi(t)$, and it is the goal of this section to quantitatively relate the autospectral density of the instrumental phase noise $S_{\delta\phi,\delta\phi}(f)$ to the oscillator phase noise $\mathcal{L}(f)$.

D.3.1 *Autocorrelation of instrumental phase noise*

The autocorrelation of the instrumental phase noise $R_{\delta\phi,\delta\phi}(\tau)$ is

$$
\begin{aligned}
R_{\delta\phi,\delta\phi}(\tau) &= E\left[\delta\phi(t, \tau_j) \cdot \delta\phi(t + \tau, \tau_j)\right] \\
&= E\left[\{\phi(t + \tau_j) - \phi(t)\} \cdot \{\phi(t + \tau + \tau_j) - \phi(t + \tau)\}\right] \\
&= E\left[\phi(t + \tau_j) \cdot \phi(t + \tau + \tau_j)\right] - E\left[\phi(t + \tau_j) \cdot \phi(t + \tau)\right] \\
&\quad - E\left[\phi(t) \cdot \phi(t + \tau + \tau_j)\right] + E\left[\phi(t) \cdot \phi(t + \tau)\right] \\
&= 2R_{\phi,\phi}(\tau) - R_{\phi,\phi}(\tau + \tau_j) - R_{\phi,\phi}(\tau - \tau_j), \tag{D.9}
\end{aligned}
$$

where (D.9) follows from the stationary nature of the phase deviation $\phi(t)$ and the invariance of a stationary process's autocorrelation function to translations in time.

D.3.2   *Autospectral density of instrumental phase noise*

The autospectral density of the instrumental phase noise $S_{\delta\phi,\delta\phi}(f)$ is simply the Fourier transform of the corresponding autocorrelation function (D.9). Explicitly,

$$
\begin{aligned}
S_{\delta\phi,\delta\phi}(f) &= \mathcal{F}\left[R_{\delta\phi,\delta\phi}(\tau)\right](f) \\
&= 2\mathcal{F}\left[R_{\phi,\phi}(\tau)\right](f) - \mathcal{F}\left[R_{\phi,\phi}(\tau+\tau_j)\right](f) \\
&\quad - \mathcal{F}\left[R_{\phi,\phi}(\tau-\tau_j)\right](f) \\
&= 2S_{\phi,\phi}(f) - \left(e^{i2\pi f\tau} + e^{-i2\pi f\tau}\right)S_{\phi,\phi}(f) \\
&= 2\left[1 - \cos(2\pi f\tau_j)\right]S_{\phi,\phi}(f) \\
&= 4\sin^2(\pi f\tau_j)S_{\phi,\phi}(f),
\end{aligned}
\tag{D.10}
$$

where the Fourier transform's linear and translational properties have been used. Referencing the definition of oscillator phase noise (D.2), the instrumental phase noise (D.10) can be rewritten (as a one-sided autospectral density) as

$$
G_{\delta\phi,\delta\phi}(f) = 8\sin^2(\pi f\tau_j)\mathcal{L}(f).
\tag{D.11}
$$

BIBLIOGRAPHY

[1] IEEE standard definitions of physical quantities for fundamental frequency and time metrology—random instabilities. *IEEE Std. 1139-2008*, pages c1–35, 2008.

[2] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 4th edition, 2010.

E

SOUND-WAVE CHARACTERIZATION

Sound waves have the well-known dispersion relation

$$\omega = c_s k, \tag{E.1}$$

where ω is the angular frequency of the sound wave, k is the wavenumber of the sound wave, and $c_s = 343\,\mathrm{m \cdot s^{-1}}$ is the sound speed in dry air at sea level and $T = 20\,^\circ\mathrm{C}$. Further, the sound-wave pressure fluctuations produce corresponding fluctuations in the refractive index, making sound waves an ideal tool for probing the wavenumber response of an interferometer. A quantitative comparison between the predicted and observed interferometer responses, however, requires an accurate description of the sound-wave pressure fluctuations.

It is the intent of this appendix to quantitatively characterize the sound-wave pressure fluctuations driven by the speaker used throughout this thesis. Below, Section E.1 describes the speaker and the hardware used to measure the pressure fluctuations of its sound waves. Section E.2 discusses the sound-wave measurements and their implications for the model of the sound-wave pressure fluctuations developed in Section E.3. Then, Section E.4 relates these sound-wave pressure fluctuations to the corresponding refractive-index variations. Finally, Section E.5 computes the phase shift imparted to a CO_2 probe beam by the speaker's sound waves.

E.1 HARDWARE

Section E.1.1 describes the speaker used to drive sound waves through the interferometer probe beam, while Section E.1.2 describes the calibrated microphone used to measure the absolute pressure fluctuations of the sound waves. The speaker and the microphone are mounted in the test stand described in Section E.1.3, allowing simple and accurate measurements of the sound waves.

E.1.1 *Speaker*

Previous sound-wave calibrations were performed with high-frequency speakers suitable for characterizing the medium-k to high-k response

of the PCI. However, to quantify the low-k to medium-k response of the heterodyne interferometer (and simultaneously cross-calibrate the PCI), the speaker should be capable of driving sound waves throughout most of the heterodyne-interferometer wavenumber range. The interferometer high-k cutoff of $5\,\mathrm{cm}^{-1}$ corresponds to a sound wave with frequency $27\,\mathrm{kHz}$. Sound waves of such frequency are typically produced by speakers known as "tweeters".

An economic ($\sim$ \$30) XT25BG60-04 tweeter (Tymphany HK Ltd., Wanchai, Hong Kong; procured through Parts Express) was used throughout this work to quantify the interferometer wavenumber response. The XT25BG60-04 employs a patented dual-concentric 1" diaphragm and a unique waveguide center plug to provide excellent on- and off-axis response; the specified on-axis response is flat from $1.5\,\mathrm{kHz}$ to $20\,\mathrm{kHz}$. The coil impedance is $4\,\Omega$, and the specified RMS power handling is $90\,\mathrm{W}$ (however, operation at significantly lower power levels destroyed the coil of the first XT25BG60-04 procured). Throughout this work, the XT25BG60-04 is driven by a $2\,\mathrm{V}$ peak-to-peak signal.

E.1.2 *Calibrated microphone*

A 378C01 free-field microphone package (PCB Piezotronics, Inc.; Depew, NY, USA) was procured to quantify the on- and off-axis speaker response. The 378C01 consists of a 377C01 microphone and a 426B03 preamplifier. The signal from the microphone package is conditioned with 480C02 battery-powered signal conditioner (also through PCB) prior to signal measurement with an oscilloscope. The microphone has a NIST-traceable calibrated sensitivity of $2.52\,\mathrm{mV}\cdot\mathrm{Pa}^{-1}$ at $251.2\,\mathrm{Hz}$ and $\pm1\,\mathrm{dB}$ sensitivity variation between $20\,\mathrm{Hz}$ and $100\,\mathrm{kHz}$.

E.1.3 *Test stand*

A microphone test stand, machined previously at the MIT PSFC, enables easy and accurate adjustment of the microphone position. Twin aluminum pylons with demarcations in $1\,\mathrm{cm}$ intervals are mounted to a large aluminum base. A small aluminum crossbeam, also with with demarcations in $1\,\mathrm{cm}$ intervals, slides up and down the two pylons and can be easily locked into place at a particular height z above the speaker face. The microphone itself is mounted to a platform that extends $\sim 10\,\mathrm{cm}$ horizontally from the crossbeam. Sound-wave reflections from the platform are minimized with acoustic-damping foam. After centering the speaker below the microphone, the microphone's transverse position can be easily scanned by sliding the mi-
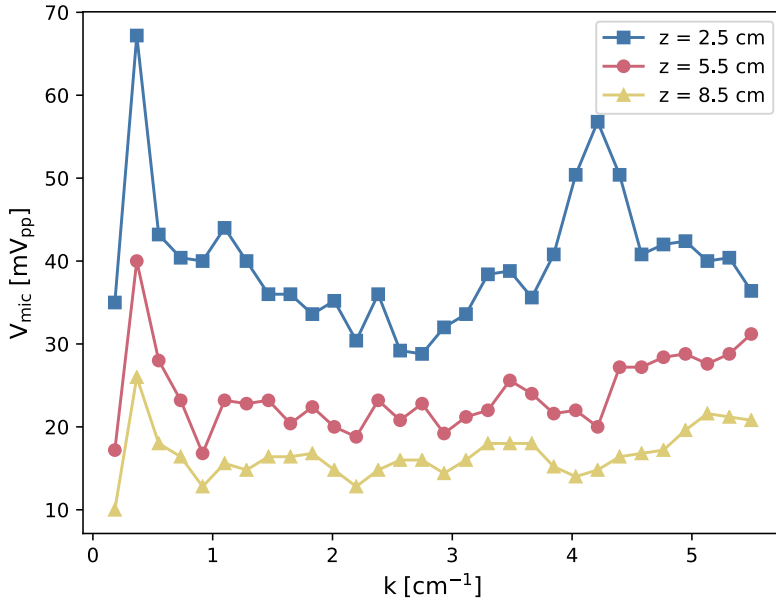
Figure E.1: On-axis amplitude of sound waves as a function of wavenumber k and height z above the speaker face. Note that the amplitude is specified as a *peak-to-peak* value.

crophone platform along the crossbeam. Thus, the test stand allows accurate and independent adjustment of the microphone's height z above the speaker's face and the microphone's radial distance ρ from the speaker's symmetry axis, establishing the (ρ, z) coordinate system for the sound-wave measurements described in Section E.2.

E.2 SOUND-WAVE MEASUREMENTS

To lowest order, the speaker is cylindrically symmetric. Thus, the sound waves are expected to have axial, radial, and frequency dependencies. Sections E.2.1 through E.2.3 summarize these measurements and their implications for the sound-wave model developed in Section E.3. All measurements were made with steady-state speaker drive.

E.2.1 *On-axis amplitude*

After centering the microphone on the speaker's symmetry axis, the on-axis amplitude can be easily characterized by varying both the fre-

quency f of the sound waves and the microphone height z above the speaker face. The frequencies f and heights z are motivated by the parameters of the heterodyne interferometer described in Chapter 4. Specifically, the interferometer spatial bandwidth $|k| \leqslant 5\,\text{cm}^{-1}$ from (4.6) motivates sound-wave measurements at frequencies $f \lesssim 30\,\text{kHz}$ (i.e. $|k| \lesssim 5\,\text{cm}^{-1}$). Further, to produce a robust interference signal during sound-wave calibrations, the speaker is placed very close to the edge of the collimated probe beam, which has $1/e$ E radius $w_0 = 3.4\,\text{cm}$; thus, sound-wave measurements are made at heights spanning the probe-beam profile $z = \{2.5\,\text{cm}, 5.5\,\text{cm}, 8.5\,\text{cm}\}$. The on-axis amplitude of the sound waves as a function of wavenumber k and height z above the speaker face is shown in Figure E.1. As expected, the on-axis amplitude decreases with increasing distance z from the speaker face. Further, the on-axis amplitude has a complicated wavenumber dependence, but it is relatively flat for $1\,\text{cm}^{-1} \lesssim k \lesssim 3.5\,\text{cm}^{-1}$.

E.2.2 *Wavefront phasing*

Characterizing the sound-wave phasing is somewhat more involved than characterizing the on-axis amplitude, as it requires measurements at several radial positions ρ for each frequency f and microphone height z. For this reason, the wavefront-phasing measurements are more coarsely sampled in frequency f than the on-axis amplitude measurements in Section E.2.1. For a given frequency f and height z, the sound-wave phasing is measured by tracking a point of constant phase in the microphone waveform as the radial position ρ is varied; such tracking can be easily accomplished by triggering the oscilloscope with a copy of the waveform that is driving the speaker. To begin the radial scan, the microphone height z is selected, and the microphone is displaced from the speaker's symmetry axis by a few centimeters. Then, in $1\,\text{cm}$ increments, the microphone is moved radially inwards towards the center; upon passing through the center, the radial scan is continued in $1\,\text{cm}$ increments until the sound-wave amplitude becomes negligible. Note that beginning the radial scan with a small displacement from the symmetry axis allows empirical identification of the symmetry-axis location (by e.g. fitting the measured amplitude and/or phasing and identifying the extremum that occurs at the symmetry axis).

At sufficiently large distances, the speaker will behave like a point source, producing sound waves with spherical wavefronts. This point-source approximation is taken as a reasonable starting point for the investigation of the wavefront phasing. If a sound wave is measured
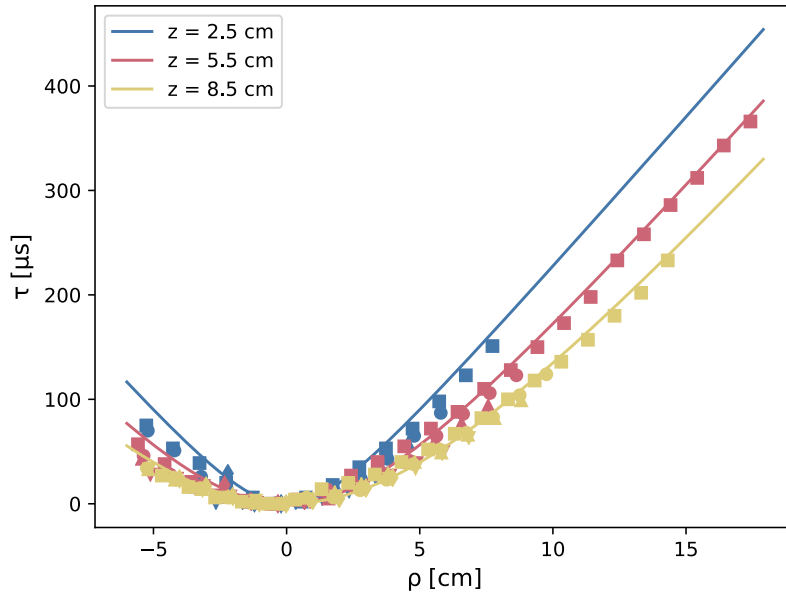
Figure E.2: Wavefront phasing of sound waves. The symbols show the measured time delay τ between the wavefront at height z and radial displacement ρ and the corresponding on-axis wavefront (i.e. same z but $\rho = 0$), with each symbol shape corresponding to particular frequency. The traces correspond to the time delay (E.2) predicted for spherical waves. The close proximity of the measured points to the spherical-wave traces indicates that, to lowest order, the waves are approximately spherical over the spatial domain and frequencies probed.

on axis at height z above the speaker, the corresponding wavefront will subsequently arrive at position $r = (z^2 + \rho^2)$ delayed by a time τ

$$\tau = \frac{r - z}{c_s}, \qquad (E.2)$$

where c_s is the sound speed. Figure E.2 compares the measured time delay to the time delay predicted for spherical waves (E.2) as a function of height z, radial position ρ, and frequency f. Clearly, to lowest order, the waves are approximately spherical over the spatial domain and frequencies probed.

E.2.3 *Spatial envelope*

If the sound-wave amplitude is also measured during the radial scans described in Section E.2.2, the spatial envelope of the sound waves can also be quantified. Figure E.3 displays the spatial envelopes of sound
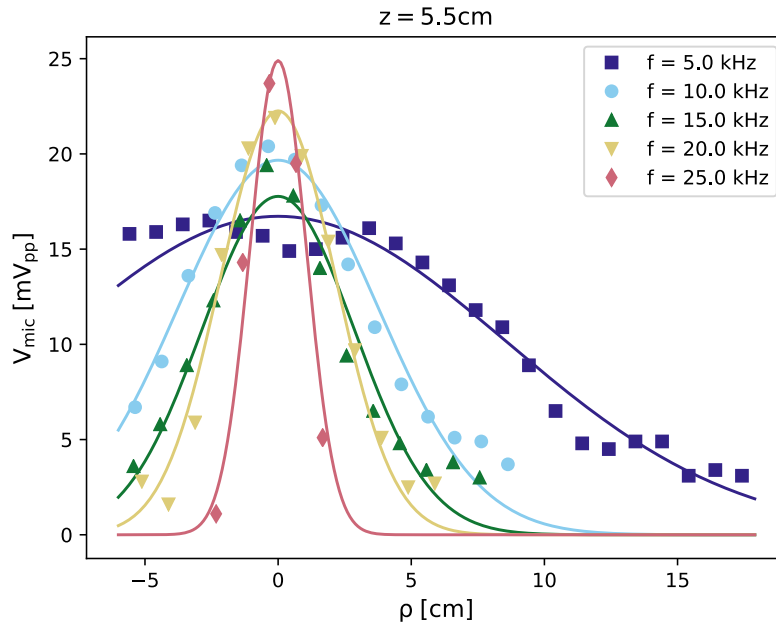
Figure E.3: Sound-wave spatial envelopes for frequencies f at height $z = 5.5$ cm above the face of the speaker. Symbols indicate measurement points, while the traces correspond to Gaussian fits of the form (E.3). Measurements and fits at other heights z exhibit qualitatively similar trends.

waves of various frequencies f at height $z = 5.5$ cm above the face of the speaker. Clearly, the width of the spatial envelope decreases with increasing frequency f. Measurements at other heights z exhibit qualitatively similar behavior.

The narrowing of the spatial envelope with increasing frequency can be quantified by fitting the measurements to an assumed functional form. To lowest order, the spatial envelopes are well approximated by a Gaussian

$$V_{\text{mic}}(\rho) = V_0(z, f) \exp\left[\frac{-\rho^2}{w(z, f)^2}\right], \tag{E.3}$$

where $w(z, f)$ is the $1/e$ radius, which is a function of the height z and the sound-wave frequency f. Gaussian fits to the envelope measurements are also shown in Figure E.3. Deviations from a Gaussian are most apparent at low frequencies; this may be attributable to baffle diffraction across the speaker face but was not further investigated. The approximation of a Gaussian envelope will be sufficiently accurate for the present work. Figure E.4 displays the fitted $1/e$ Gaussian radii w as a function of sound-wave wavenumber k and height
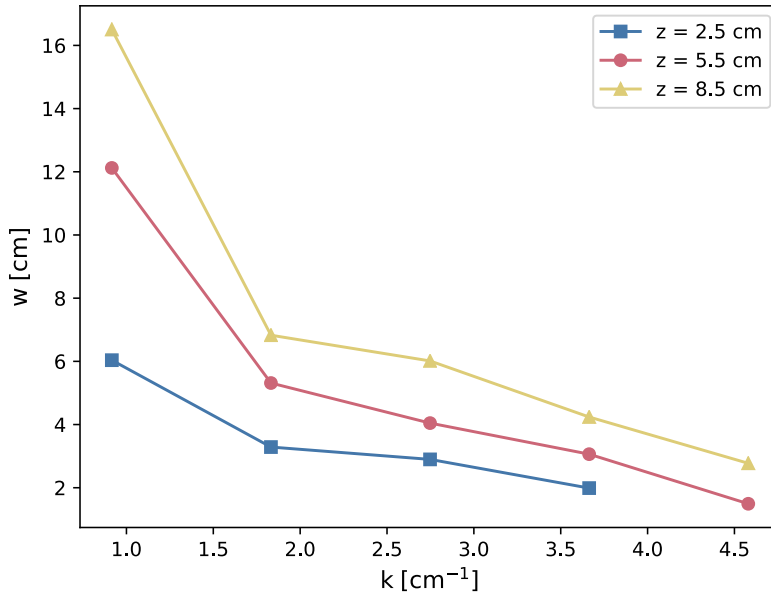
Figure E.4: Fitted $1/e$ Gaussian radii w of sound waves as a function of sound-wave wavenumber k and height z above the face of the speaker.

z above the face of the speaker. As previously and anecdotally noted for $z = 5.5\,\mathrm{cm}$, the width of the spatial envelope w decreases with increasing k for each height z. Further, w increases with increasing z, which results from free-space diffraction of the sound wave.

E.3   SOUND-WAVE MODEL

Quantitatively predicting the response of a heterodyne interferometer to sound-wave pressure fluctuations $\widetilde{p}(r, t)$ requires knowledge of $\widetilde{p}$ at each point along the beam path. Because the measurements described in Section E.2 were made at discrete positions and frequencies, it is necessary to develop an approximate *model* of $\widetilde{p}$. Recall that Section E.2.2 shows that the wavefronts of the sound waves are approximately spherical, while Section E.2.3 shows that the spatial envelope of the sound waves is approximately Gaussian. Thus, the sound-wave pressure fluctuation can be modeled as

$$\widetilde{p}(r, t) = \widetilde{p}_0(z, f) \cdot \exp\left[\frac{-\rho^2}{w(z, f)^2}\right] \cdot \cos(kr - \omega t), \tag{E.4}$$
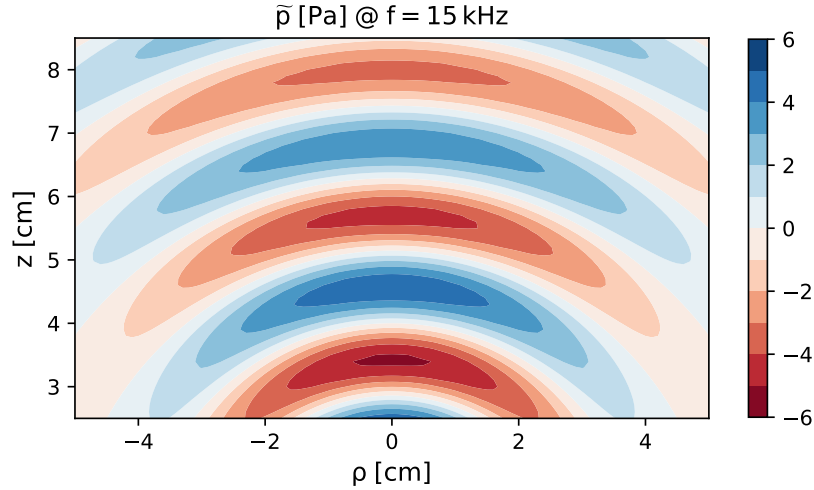
Figure E.5: "Snapshot" (i.e. at a single point in time) of the model's predicted 15 kHz pressure fluctuation $\widetilde{p}$.

where z is the height above the speaker face, $\rho = (x^2 + y^2)$ is the radial distance from the symmetry axis, $\boldsymbol{r} = (x, y, z)$ is the spatial coordinate relative to the center of the speaker face, $r = |\boldsymbol{r}|$ is the corresponding distance from the center of the speaker face, ω is the angular frequency of the sound wave, and k is the wavenumber of the sound wave, as determined from the sound-wave dispersion relation (E.1). Additionally, $\widetilde{p}_0(z, f)$ is the on-axis sound-wave amplitude, and it is a complicated function of z and f that is computed via radial-basis-function (RBF) interpolation [1, Sec. 3.7] [2, 3] of the measured on-axis amplitude shown in Figure E.1 (after conversion from V to Pa using the absolute calibration of the microphone discussed in Section E.1.2). Similarly, $w(z, f)$ is the 1/e Gaussian radius of the sound wave, and it is a complicated function of z and f that is computed via RBF interpolation of the measured 1/e Gaussian radii shown in Figure E.4. As an example, a "snapshot" (i.e. at a single point in time) of the model's predicted 15 kHz pressure fluctuation $\widetilde{p}$ is shown in Figure E.5.

E.4   PERTURBED INDEX OF REFRACTION

Interferometric methods are sensitive to the index of refraction N. Thus, quantitatively predicting the response of a heterodyne interferometer requires computing the perturbed index of refraction $\widetilde{N}$ corresponding to the sound-wave pressure fluctuation $\widetilde{p}$ in (E.4).

While the quantification of air's refractive index in the visible and infrared has been vigorously pursued by the metrology community [4, 5, 6, 7], the lowest-order description will be sufficiently accurate for the present application. Specifically, an electromagnetic wave propagating through air induces a time-varying polarization in the air that alters the propagation of the wave; it is precisely this interaction that the refractive index N quantifies. Clearly, the induced polarization is proportional to the number density n of the air. Thus, air's deviation from the vacuum refractive index of unity is also proportional to n, i.e.

$$N - 1 = \alpha n, \tag{E.5}$$

where α is a complicated function of the atmospheric composition, the vacuum wavelength of the electromagnetic wave, etc. Explicitly writing the refractive index and number density as sums of equilibrium and perturbed components (i.e. $N = \overline{N} + \widetilde{N}$ and $n = \overline{n} + \widetilde{n}$, respectively), one readily sees that the perturbed refractive index can be written as

$$\widetilde{N} = \left(\frac{\overline{N} - 1}{\overline{n}} \right) \widetilde{n}. \tag{E.6}$$

Now, the sound-wave compressions and rarefactions are approximately adiabatic such that the total pressure p and the total number density n are related via $p \propto n^{\gamma}$, where γ is the ratio of specific heats. Thus, the density perturbation $\widetilde{n}$ corresponding to a sound-wave pressure perturbation $|\widetilde{p}| \ll \overline{p}$ (where $\overline{p}$ is the equilibrium pressure) is

$$\widetilde{n} \approx \left(\frac{\overline{n}}{\gamma \overline{p}} \right) \widetilde{p} \tag{E.7}$$

such that the perturbed refractive index (E.6) becomes

$$\widetilde{N} \approx \left(\frac{\overline{N} - 1}{\gamma \overline{p}} \right) \widetilde{p}. \tag{E.8}$$

As DIII-D sits very nearly at sea level, the equilibrium pressure is $\overline{p} = 101\,325$ Pa. Further, for dry air at $T \sim 300$ K, the ratio of specific heats is $\gamma = 1.4$, and the deviation from the vacuum index of refraction for $10.6\,\mu\text{m}$ radiation is $\overline{N} - 1 = 2.7 \times 10^{-4}$ [6, 7, 8]. Thus, the perturbed refractive index (E.8) for a $10.6\,\mu\text{m}$ CO_2 probe beam becomes

$$\widetilde{N} \approx (1.9 \times 10^{-9}) \widetilde{p} \text{ [Pa]}. \tag{E.9}$$
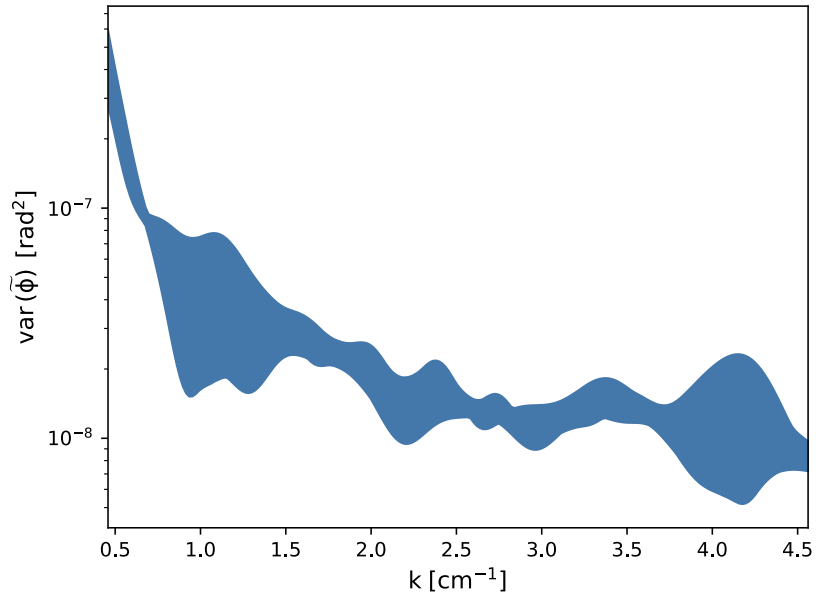
E.5   PHASE SHIFT IMPARTED TO A CO2 PROBE BEAM

A probe beam propagating through a sound-wave pressure perturbation $\widetilde{p}$ will acquire a phase shift

$$\widetilde{\phi} = k_0 \int \widetilde{N} dl, \tag{E.10}$$

where k_0 is the vacuum wavenumber of the probe beam, $\widetilde{N}$ is the perturbed index of refraction (E.8), and the integration is performed along the beam path. This phase shift is, of course, the measurable quantity of a heterodyne interferometer. For a $10.6\,\mu m$ CO_2 probe beam, the phase shift (E.10) becomes

$$\widetilde{\phi}\,[\text{rad}] = (1.1 \times 10^{-5}\,\text{cm}^{-1}) \int (\widetilde{p}\,[\text{Pa}]) dl, \tag{E.11}$$

where (E.9) has been referenced and the differential path length dl must have units of centimeters. The bounds on the variance of the phase shift imparted by the sound-wave pressure fluctuation (E.4) is shown in Figure E.6. To make a quantitative comparison with the corresponding heterodyne-interferometer measurements, the interferometer wavenumber response and noise floor must also be accounted for; such a comparison is performed in Section 4.8.

Figure E.6: The bounds on the variance of the phase shift (E.11) imparted to a $10.6\,\mu m$ CO_2 probe beam by the sound-wave pressure fluctuation (E.4). Because the probe beam has finite cross section and the phase shift varies with the height z above the speaker face, the maximum and minimum values of the variance are indicated by the upper and lower bounds of the shaded region, respectively; the minimum height considered is $z = 2.5\,cm$, and the maximum height considered is $z = 8.5\,cm$. The symmetry axes of the probe beam and the speaker are intersecting and orthogonal. To make a quantitative comparison with the corresponding heterodyne-interferometer measurements, the interferometer wavenumber response and noise floor must also be accounted for; such a comparison is performed in Section 4.8.

BIBLIOGRAPHY

[1] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 3rd edition, 2007.

[2] The Scipy community. *scipy.interpolate.Rbf (v0.19.1)*, June 2017.

[3] How can I perform two-dimensional interpolation using scipy? https://stackoverflow.com/a/37872172/5469497. Accessed: 2017-11-25.

[4] A. T. Young. Refractivity of air. http://aty.sdsu.edu/explain/atmos_refr/air_refr.html. Accessed on 2017-11-26.

[5] J. A. Stone and J. H. Zimmerman. Refractive index of air calculator. http://emtoolbox.nist.gov/Wavelength/Documentation.asp. Accessed on 2017-11-26.

[6] S. Marchetti and R. Simili. Accurate measurement of the refractive index of CO_2, N_2, He, O_2, and air at $10.57\,\mu m$ and T = 23°C. *Infrared Phys. Tech.*, 47(3):263–266, 2006.

[7] R. J. Mathar. Refractive index of humid air in the infrared: model fits. *J. Opt. A: Pure Appl. Opt.*, 9(5):470, 2007.

[8] M. N. Polyanskiy. Refractive index database. https://refractiveindex.info/?shelf=other&book=air&page=Mathar-7.5. Accessed on 2017-11-26.

F

SPECTRAL ESTIMATION

In contrast to deterministic processes, random processes cannot be modeled via an explicit mathematical relationship. Rather, random processes are characterized in terms of probabilities and statistical properties. Any given observation of a random process represents only one of many possible observations; each such observation is referred to as a "sample" or a "realization" of the random process and is denoted as $x_k(t)$. The random process itself consists of the ensemble of all of the potential observations and is denoted as $\{x_k(t)\}$. Random processes can be stationary or nonstationary. The statistical properties of a stationary random process do not vary in time.

The spectral tools discussed below are all developed for analysis of stationary random processes. Section F.1 discusses non-parametric spectral-estimation techniques, which find powerful and versatile application in the analysis of uniformly and richly sampled data. Section F.2 discusses parametric spectral-estimation techniques, which find application in the analysis of data with sharp spectral features, particularly if the data is sparsely sampled. Finally, Section F.3 discusses the hybrid non-parametric-in-time, parametric-in-space technique used to estimate the PCI-measured two-dimensional autospectral density function.

F.1 NON-PARAMETRIC TECHNIQUES

Most of the discussion below is distilled from the seminal work by Bendat and Piersol [1], and inquisitive readers are directed there for a more extensive treatment of the subject.

The windowed, finite Fourier transform $X_k(f, T)$ of a continuous signal $x_k(t)$ sampled for $-T/2 \leqslant t < T/2$ is defined as

$$X_k(f, T) = \int_{-T/2}^{T/2} dt\, [w(t) \cdot x_k(t)]e^{-i\,2\pi ft}, \tag{F.1}$$

where $w(t)$ is an arbitrary windowing function. Typically, the selected windowing function smoothly tapers as $|t| \to T/2$ to minimize side-lobe leakage that results from discontinuities at the start and end of

the sample record. Further, to prevent power loss, the windowing function is also typically normalized such that

$$\frac{1}{T} \int_{-T/2}^{T/2} dt \, [w(t)]^2 = 1. \tag{F.2}$$

The normalized Hanning window is perhaps the most commonly used windowing function, and it is used uniformly throughout this work.

For real-valued, stationary random processes $\{x_k(t)\}$ and $\{y_k(t)\}$, the one-sided *cross-spectral density* function $G_{xy}(f)$ is defined as

$$G_{xy}(f) \equiv \lim_{T \to \infty} \frac{2}{T} E\,[X_k^*(f, T) Y_k(f, T)] \tag{F.3}$$

for $0 < f < \infty$; $G_{xy}(f)$ is not defined for $f < 0$, and it is reduced by a factor of two relative to (F.3) at $f = 0$ (the value of $G_{xy}(0)$ is of little relevance to this work). Note that $E[\cdot]$ is the expectation value operator; this operator averages over all of the realizations in the ensemble, and its application ensures that (F.3) is a statistically consistent definition of the cross-spectral density (that is, ensemble averaging is needed for $G_{xy}(f)$ to approach the true cross-spectral density as $T \to \infty$). If, in addition to being stationary, the random process is also *ergodic*, the ensemble average can be replaced with a time average of $X_k(f, T)$ over successive time slices. If desired, these time slices may partially overlap. Unless otherwise noted, all of the ensemble averages in this work are computed using this assumption of ergodicity, and successive slices are selected to overlap by 50%.

In general $G_{xy}(f)$ is a complex-valued function. This can be made explicit by writing

$$G_{xy}(f) = |G_{xy}(f)| \, e^{i\alpha_{xy}(f)}, \tag{F.4}$$

where $\alpha_{xy}(f)$ is the *cross phase*. For the special case $\{x_k(t)\} = \{y_k(t)\}$, $G_{xx}(f)$ is real-valued (i.e. $G_{xx}(f) = |G_{xx}(f)|$) and is referred to as the one-sided *autospectral density* function.

The degree of correlation between random processes $\{x_k(t)\}$ and $\{y_k(t)\}$ can be easily quantified with the corresponding spectral density functions. In particular, the *magnitude-squared coherence* function $\gamma_{xy}^2(f)$ is defined as

$$\gamma_{xy}^2(f) \equiv \frac{|G_{xy}(f)|^2}{G_{xx}(f) G_{yy}(f)}, \tag{F.5}$$

Spectral estimate	Random error [1]
$G_{xy}(f)$	$\varepsilon[G_{xy}(f)] = \dfrac{1}{|\gamma_{xy}(f)|\sqrt{N_r}}$
$\alpha_{xy}(f)$	$\text{s.d.}[\alpha_{xy}(f)] \approx \dfrac{[1-\gamma_{xy}^2(f)]^{1/2}}{|\gamma_{xy}(f)|\sqrt{2N_r}}$
$\gamma_{xy}^2(f)$	$\varepsilon[\gamma_{xy}^2(f)] = \dfrac{\sqrt{2}[1-\gamma_{xy}^2(f)]}{|\gamma_{xy}(f)|\sqrt{N_r}}$

Table F.1: Random errors in estimates of spectral properties are functions of the number of realizations N_r used in the computation of the ensemble average and the coherence magnitude $|\gamma_{xy}(f)|$. Here, s.d.$[\cdot]$ represents the standard deviation of the estimate, and $\varepsilon[\cdot]$ represents the standard deviation of the estimate *normalized* to the true value of the spectral property.

and it satisfies

$$0 \leqslant \gamma_{xy}^2(f) \leqslant 1 \tag{F.6}$$

for $0 \leqslant f < \infty$. If $\gamma_{xy}^2(f) = 1$, $\{x_k(t)\}$ and $\{y_k(t)\}$ are 100% correlated at frequency f, and if $\gamma_{xy}^2(f) = 0$, $\{x_k(t)\}$ and $\{y_k(t)\}$ are completely uncorrelated at frequency f. Note that the ensemble-averaging operation in (F.3) is paramount to the computation of *informative* values for $\gamma_{xy}^2(f)$; that is, if ensemble averaging is ignored, and only single realizations of the random processes are used, $\gamma_{xy}^2(f) \equiv 1$ for all f, *regardless* of the actual degree of coherence between between $\{x_k(t)\}$ and $\{y_k(t)\}$.

Care should be taken when computing spectral density estimates. Table F.1 summarizes the random errors associated with the estimates of various spectral properties. Note that the number of realizations N_r used in the computation of the ensemble average is a parameter that can be specified at the time of analysis and that increasing N_r reduces the random errors of each spectral estimate. (While increased $\gamma_{xy}^2(f)$ also reduces random errors, $\gamma_{xy}^2(f)$ is an intrinsic property of the data rather than a parameter that can be specified at the time of analysis). Further, in various programming languages, it is not uncommon to "detrend" realizations $x_k(t)$ and $y_k(t)$ by subtracting the signal mean or linear trend prior to application of (F.3). As described in Section 4.6.2, signals are high-pass filtered prior to spectral analysis, and no further detrending is performed.

F.2   PARAMETRIC TECHNIQUES

Parametric techniques attempt to represent a signal with a mathematical model containing a limited number of predefined parameters. For a pedagogical overview of such techniques, the reader is directed to the work of Oppenheim and Schafer [2, Ch. 11] The discussion here is largely synthesized from the discussion by Marple [3].

Many random processes are well approximated by rational transfer-function models. In such models, an input driving sequence n_n and an output (i.e. data) sequence x_n are related via the linear difference equation

$$x_n = \sum_{l=0}^{q} b_l n_{n-l} - \sum_{k=1}^{p} a_k x_{n-k}; \tag{F.7}$$

sequence x_n and n_n each have length N, and successive samples are separated in time by Δt. Such a model is termed an autoregressive-moving average (ARMA) model. Now, if $b_0 = 1$ and the remainder of the b_l are zero, (F.7) reduces to

$$x_n = -\sum_{k=1}^{p} a_k x_{n-k} + n_n \tag{F.8}$$

such that x_n is an autoregression (AR) of order p driven by white noise n_n. The autospectral density $G_{xx}(f)$ of such an AR is

$$G_{xx}(f) = \frac{\sigma^2}{\left| 1 + \sum_{k=1}^{p} a_k \exp(-2\pi k f \Delta t) \right|^2}, \tag{F.9}$$

where σ^2 is the variance of noise term n_n. Because the only frequency dependence in spectral estimate (F.9) appears in the denominator, the AR model is often referred to as an all-pole model. In one dimension, AR spectra are equivalent to spectra computed via the maximum entropy method (MEM). Interestingly, AR spectra do *not* suffer from traditional sidelobes due to windowing.

The AR process is fully characterized by the $(p + 1)$ parameters $(a_1, a_2, \cdots, a_p, \sigma^2)$, and numerous techniques exist for their estimation. The $p + 1$ autocorrelations $\{R_{xx}(0), R_{xx}(\Delta t), \ldots R_{xx}(p \cdot \Delta t)\}$ are related to $(a_1, a_2, \cdots, a_p, \sigma^2)$ via the Yule-Walker normal equations, which the Levinson-Durbin recursion can efficiently solve [2, Sec. 11.6]. However, better AR estimates are often obtained via least-squares linear prediction applied directly to the data x_n. Such techniques may

employ the forward-only linear prediction or some combination of the forward and reverse linear predictions. The covariance method [2, Sec. 11.3.2] is perhaps the most well known forward-only algorithm, while the Burg method is perhaps the most well known forward and reverse algorithm.

Here, the Burg method is briefly reviewed. The algorithm is iterative, so assume that the p^{th} order AR parameters are known. Then, the forward linear prediction is

$$\hat{x}_n = -\sum_{k=1}^{p} a_{p,k} x_{n-k}, \tag{F.10}$$

where $a_{p,k}$ is the k^{th} coefficient of the p-order AR model. The corresponding forward linear-prediction error is

$$\begin{aligned} e_{p,n} &= x_n - \hat{x}_n \\ &= \sum_{k=0}^{p} a_{p,k} x_{n-k}, \qquad p \leqslant n < N, \end{aligned} \tag{F.11}$$

where $a_{p,0} = 1$ by definition and the limits on n are established such that the error is defined only over the available data. Note that $e_{p,n} = n_n$ such that $\sigma^2 = |e_{p,n}|^2$. For a stationary process, the coefficients of the backward linear-prediction error filter are simply conjugated and reversed in time relative to those of the forward linear-prediction error filter such that the backward linear-prediction error is
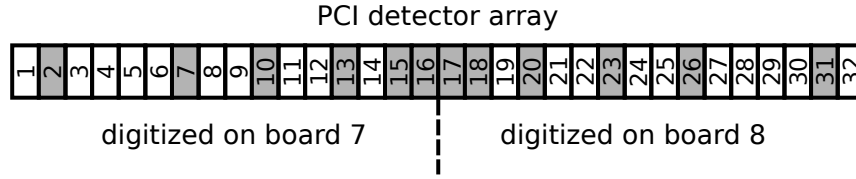
$$b_{p,n} = \sum_{k=0}^{p} a_{p,k}^* x_{n-p+k}, \qquad p \leqslant n < N, \tag{F.12}$$

where z^* indicates the complex conjugate of z. The total forward and backward linear-prediction error is

$$E_p = \sum_{n=p}^{N-1} |e_{p,n}|^2 + \sum_{n=p}^{N-1} |b_{p,n}|^2. \tag{F.13}$$

To estimate the AR parameters, the Burg method minimizes the total error E_p subject to the constraint that the AR parameters satisfy the Levinson-Durbin recursion

$$a_{p,k} = a_{p-1,k} + a_{p,p} a_{p-1,p-k}^*, \qquad 1 \leqslant k \leqslant p. \tag{F.14}$$

PCI detector array

digitized on board 7   digitized on board 8

Figure F.1: Schematic of PCI detector array. Only the gray channels are digitized.

The Levinson-Durbin constraint ensures that the AR filter is stable (i.e. all poles fall within the unit circle) and make E_p a function solely of the unknown coefficient $a_{p,p}$. Setting the derivative of E_p with respect to $a_{p,p}$ to zero then yields

$$a_{i,i} = \frac{-2 \sum\limits_{k=i}^{N-1} b_{i-1,k-1}^* e_{i-1,k}}{\sum\limits_{k=i}^{N-1} \left(|b_{i-1,k-1}|^2 + |e_{i-1,k-1}|^2 \right)}. \tag{F.15}$$

Although the Burg method has satisfactory performance for the present work, it should be noted for completeness that the Burg method suffers from several problems, including spectral line splitting and biases in the frequency estimate. These problems can be corrected via a "least-squares" algorithm (independently developed by Ulrych-Clayton and Nuttall) in which E_p is differentiated with respect to all of the $a_{p,k}$, not just $a_{p,p}$, to obtain a set of normal equations (the Levinson-Durbin constraint is no longer enforced). Application of this modified algorithm is relegated to future work.

F.3   HYBRID TWO-DIMENSIONAL SPECTRAL ESTIMATES

The PCI detector consists of 32 elements arranged in a linear array. However, due to digitization constraints, only a subset of the signals from the linear array are digitized. Further, to achieve reasonable mid-k and high-k response, this subset of channels is non-uniformly spaced, as indicated in Figure F.1. This immediately precludes direct application of the fast Fourier transform (FFT) in the spatial dimension, and more elaborate schemes must be utilized to quantify the spatial spectrum. Below, Section F.3.1 describes the estimation and compensation of the trigger offset between the two PCI digitizer boards. Then, Section F.3.2 discusses estimation of the two-dimensional autocorrelation function, which can be computed even when channels are non-uniformly spaced. Finally, Section F.3.3 discusses calculation of

the two-dimensional autospectral density from the two-dimensional autocorrelation function.

F.3.1 *Estimation & compensation of trigger offset*

Prior to performing any spectral computations, the trigger offset between digitizer board 7 and digitizer board 8 must be estimated and compensated. The general theory is discussed in Appendix G. The trigger offset is estimated by defining

$$\Delta\alpha(f) = \frac{\alpha_{15,16}(f) + \alpha_{17,18}(f)}{2} \tag{F.16}$$

$$\Delta\alpha_{\text{meas}}(f) = \alpha_{16,17}(f), \tag{F.17}$$

and applying (G.10); here $\alpha_{i,j}(f)$ is the cross phase between channels i and j during a stationary portion of the discharge. To minimize random error, at least 1000 realizations are averaged over, and only frequencies with sufficiently high coherence (e.g. $\gamma_{i,j}^2 \geqslant 0.1$) are considered. The offset is rarely larger than one or two timestamps, but even such small offsets can significantly bias spectral estimates. Then, using standard techniques [2, Sec. 4.5], the trigger offset can be compensated easily in post-processing, even if the offset is a non-integer multiple of the sample spacing.

F.3.2 *Two-dimensional autocorrelation function*

The autocorrelation function R_{xx} and the autospectral density function S_{xx} are Fourier transform pairs, i.e.

$$S_{xx}(\xi, f) = \mathcal{F}[R_{xx}(\delta, \tau)](\xi, f), \tag{F.18}$$

$$R_{xx}(\delta, \tau) = \mathcal{F}^{-1}[S_{xx}(\xi, f)](\delta, \tau), \tag{F.19}$$

where $\mathcal{F}$ is the Fourier transform, τ is the temporal lag, δ is the spatial lag, f is the frequency, and ξ is the spatial frequency (related to the wavenumber k via $k = 2\pi\xi$). If the measurements are uniformly sampled in space and time, it is most efficient to estimate the two-dimensional autospectral density S_{xx} using the FFT methods described in Section F.1 and then apply the inverse Fourier transform as in (F.19) to compute the two-dimensional autocorrelation function R_{xx}. However, if the measurements are *not* uniformly spaced, the two

dimensional autocorrelation function can still be estimated via the definition

$$R_{xx}(\delta, \tau) = E[x_k(z, t) \cdot x_k(z + \delta, t + \tau)], \tag{F.20}$$

where $E[\cdot]$ is the expectation-value operator and x_k is the k^{th} realization of the real-valued random process $\{x_k(z, t)\}$. If needed, the autocorrelation can be interpolated onto a uniform grid, and then the autospectral density S_{xx} can be computed by applying the Fourier transform, as in (F.18).

When temporal sampling is uniform but spatial sampling is nonuniform, as in the PCI, it is often convenient to define a "hybrid" two-dimensional autocorrelation function $\widetilde{R}_{xx}$

$$\begin{aligned}
\widetilde{R}_{xx}(\delta, f) &= \mathcal{F}^{-1}[S_{xx}(\xi, f)](\delta) \\
&= \int_{-\infty}^{\infty} d\xi e^{i2\pi\xi\delta} S_{xx}(\xi, f) \\
&= S_{xx}(\delta, f),
\end{aligned} \tag{F.21}$$

where $S_{xx}(\delta, f) = S_{x_i, x_j}(f)$ is the cross-spectral density function between $x_i(t) = x(z, t)$ and $x_j(t) = x(z + \delta, t)$. Note that $S_{x_i, x_j}(f)$ can be efficiently computed via the FFT methods described in Section F.1 such that the "hybrid" autocorrelation function can be estimated as an ensemble average over all of the unique correlation pairs separated by δ

$$\widetilde{R}_{xx}(\delta, f) = \frac{\displaystyle\sum_{i-j=\delta} S_{x_i, x_j}(f)}{\displaystyle\sum_{i-j=\delta} 1}; \tag{F.22}$$

if there are no correlation pairs separated by δ, then $\widetilde{R}_{xx}$ is undefined for this separation. If the variances of x_i and x_j are artificially biased e.g. due to the finite PCI beam width, the variances should be equalized prior to estimating $\widetilde{R}_{xx}$ with (F.22). As discussed in Section F.3.3, the autospectral density can be computed from this $\widetilde{R}_{xx}$ estimate.

Before proceeding, however, it is instructive to consider a few properties of $\widetilde{R}_{xx}$. In general, $\widetilde{R}_{xx}$ is complex-valued. For real-valued x, however, $\widetilde{R}_{xx}$ is Hermitian, i.e. $\widetilde{R}_{xx}(-\delta, -f) = [\widetilde{R}_{xx}(\delta, f)]^*$, where z^* indicates the complex conjugate of z. At any given frequency f, $|\widetilde{R}_{xx}(\delta, f)|$ attains a maximum at $\delta = 0$. The variance of signal $x(z, t)$ is related to $\widetilde{R}_{xx}$ via

$$\text{var}(x) = \int \widetilde{R}_{xx}(0, f) df. \tag{F.23}$$
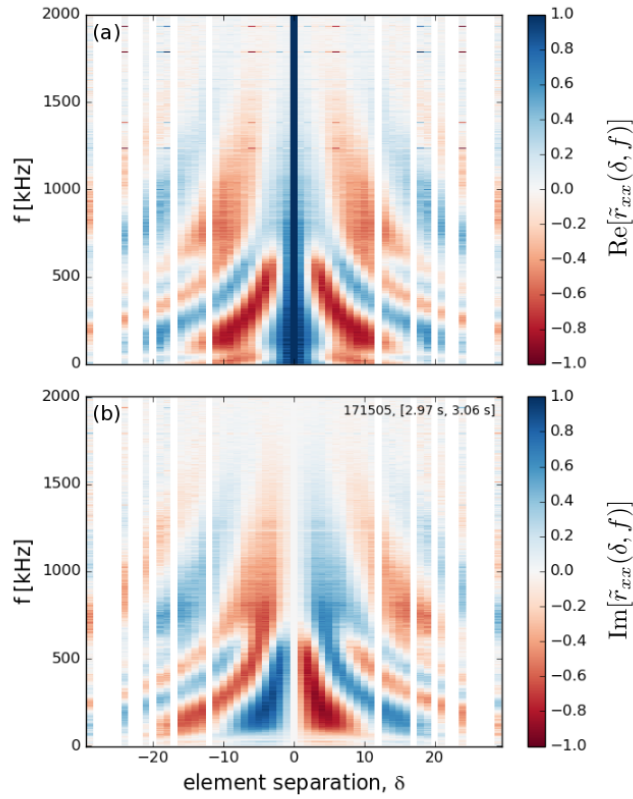
Figure F.2: Example of the (a) real and (b) imaginary components of the PCI-measured normalized, two-dimensional, hybrid autocorrelation function $\tilde{r}_{xx}(\delta, f)$. Note that $\delta = 0$ corresponds to the same detector element, $\delta = 1$ corresponds to adjacent detector elements, etc. The vertical white striations correspond to non-existing element separations (as established by the digitization layout shown in Figure F.1) for which $\tilde{r}_{xx}(\delta, f)$ is not defined. To proceed with the autospectral density estimation in Section F.3.3, the computation must either be restricted to the central, continuous domain, or the autocorrelation function must be interpolated; linear interpolation was found to be sufficient in this work.

Because $|\widetilde{R}_{xx}|$ can vary by several orders of magnitude across the full temporal bandwidth of signal $x(z, t)$, visualizing the spatiotemporal structure of $\widetilde{R}_{xx}$ can be aided by defining the normalized autocorrelation function

$$\tilde{r}_{xx}(\delta, f) = \frac{\widetilde{R}_{xx}(\delta, f)}{\widetilde{R}_{xx}(0, f)}. \tag{F.24}$$

An example of the PCI-measured $\tilde{r}_{xx}(\delta, f)$ is shown in Figure F.2.

The two-dimensional autospectral density function $S_{xx}(k, f)$ can be computed from the hybrid two-dimensional autocorrelation function $\widetilde{R}_{xx}(\delta, f)$. Below and throughout this work, if $\widetilde{R}_{xx}(\delta, f)$ is not defined for a given δ, it is linearly interpolated across the gap. Investigation of more sophisticated interpolation algorithms is relegated to future work.

As suggested by (F.18), computing the spatial Fourier transform of $\widetilde{R}_{xx}(\delta, f)$ produces an estimate of the autospectral density $S_{xx}(k, f)$. Explicitly,

$$
\begin{aligned}
S_{xx}(k, f) &= \mathcal{F}\left[\widetilde{R}_{xx}(\delta, f)\right](k) \\
&= \int_{-\infty}^{\infty} \widetilde{R}_{xx}(\delta, f) e^{-ik\delta} d\delta \\
&\approx \int_{-\Delta/2}^{\Delta/2} \widetilde{R}_{xx}(\delta, f) e^{-ik\delta} d\delta,
\end{aligned}
\tag{F.25}
$$

where $\Delta = 2 \cdot \max(|\delta|)$ is the full span of spatial lags. Of course, for efficiency, the Fourier transform should be computed via the FFT. Further, to minimized sidelobe leakage, $\widetilde{R}_{xx}(\delta, f)$ should be windowed in δ prior to computing the Fourier transform; here, only the Hanning window is considered. Using the $\widetilde{R}_{xx}(\delta, f)$ corresponding to Figure F.2, the resulting Fourier-in-space autospectral density estimate is shown in Figure F.3(a). Clearly, the wavenumber resolution of the Fourier estimate is severely limited by the sparse spatial sampling.

The parametric spectral-estimation techniques discussed in Section F.2 present an alternative method for computing $S_{xx}(k, f)$. For each frequency f in $\widetilde{R}_{xx}(\delta, f)$, a p-order Burg AR can be performed in δ to estimate the autospectral density of the *autocorrelation function* $S_{\widetilde{R}\widetilde{R}}(k, f)$. Now, from the non-parametric methods discussed in Section F.1, the autospectral density of $\widetilde{R}_{xx}$ is defined as

$$
S_{\widetilde{R}\widetilde{R}}(k, f) = \lim_{\Delta \to \infty} \frac{1}{\Delta} E\left[\left|\int_{-\Delta/2}^{\Delta/2} \widetilde{R}_{xx}(\delta, f) e^{-ik\delta} d\delta\right|^2\right]
\tag{F.26}
$$

such that

$$
\int_{-\Delta/2}^{\Delta/2} \widetilde{R}_{xx}(\delta, f) e^{-ik\delta} d\delta \approx \pm \left[\Delta \cdot S_{\widetilde{R}\widetilde{R}}(k, f)\right]^{1/2},
\tag{F.27}
$$
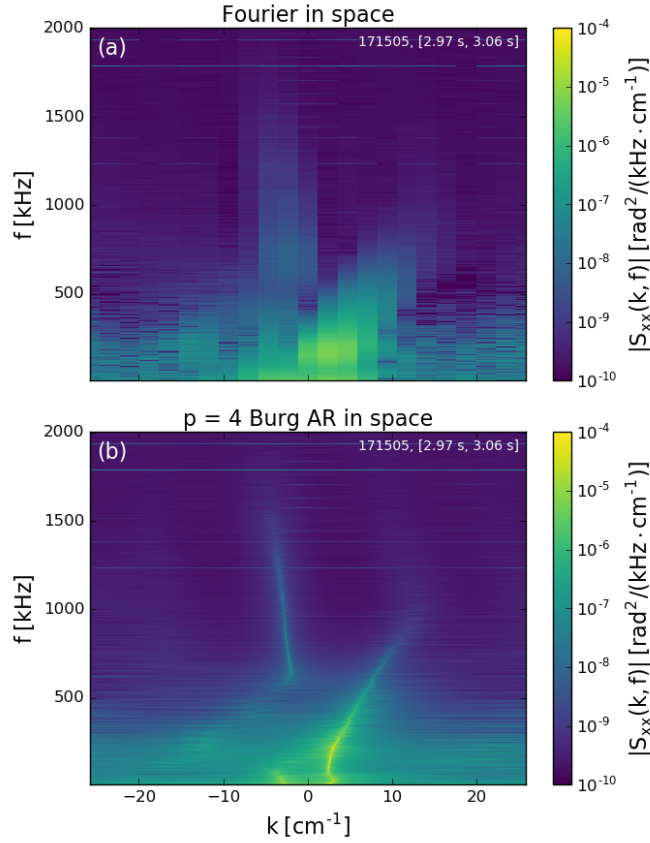
Figure F.3: Two-dimensional autospectral density estimates $S_{xx}(k, f)$. Here, (a) estimates the spatial spectrum using conventional Fourier methods while (b) estimates the spatial spectrum using a $p = 4$ Burg AR evaluated on a uniformly spaced, 1000-element wavenumber grid. It is important to note that the autospectral density estimates in both (a) and (b) correspond to the *same* raw data and the *same* autocorrelation function $\widetilde{R}_{xx}(\delta, f)$ from Figure F.2. However, the Burg AR produces substantially improved wavenumber resolution.

which is real-valued because $S_{\widetilde{R}\widetilde{R}}$ and Δ are real-valued. Substituting (F.27) into (F.25) yields

$$S_{xx}(k, f) \approx \left[\Delta \cdot S_{\widetilde{R}\widetilde{R}}(k, f)\right]^{1/2},\tag{F.28}$$

where the positive root has been selected because $S_{xx}(k, f)$ is positive semidefinite by definition. Power conservation is ensured by normalizing the spectral estimate to the signal variance, i.e.

$$\int S_{xx}(k, f)dkdf = var(x),\tag{F.29}$$

where var(x) is related to $R_{xx}(\delta, f)$ via (F.23). Because the all-pole feature of an AR model is capable of fitting very sharp spectral features and the resulting spectral estimate can be evaluated with an arbitrary resolution, the wavenumber resolution of the Burg AR estimate may be substantially better than that of the corresponding Fourier estimate, as shown in Figure F.3(b).

BIBLIOGRAPHY

[1] J. S. Bendat and A. G. Piersol. *Random Data: Analysis and Measurement Procedures*. John Wiley & Sons, Inc., Hoboken, New Jersey, USA, 4th edition, 2010.

[2] A. V. Oppenheim and R. W. Schafer. *Discrete-Time Signal Processing*. Pearson, Upper Saddle River, NJ, USA, 3rd edition, 2010.

[3] S. L. Marple. A tutorial overview of modern spectral estimation. In *International Conference on Acoustics, Speech, and Signal Processing*, pages 2152–2157 vol.4, May 1989.

SYNCHRONIZATION OF DIGITAL RECORDS

Digital signal processing is often foundational to signal analysis. Of course, application of such techniques requires converting an analog signal to a digital record. Efficient conversion requires both quantization of the signal magnitude and temporal sampling [1]. When examining the phasing between multiple digital records, the synchronization of this temporal sampling is of paramount importance.

This appendix discusses post-processing synchronization of digital records. Below, Section G.1 defines temporal sampling and mentions caveats regarding nominal and actual sampling parameters. Section G.2 then discusses various digitization schemes, highlighting which schemes allow synchronization. Finally, Section G.3 details the synchronization of phase-locked digital records.

G.1 TEMPORAL SAMPLING

Typically, temporal sampling of signal $x_j(t)$ occurs at a fixed sampling rate F_j such that successive points in the digital record are separated in time by $1/F_j$. Digitization begins at the trigger time $t_j[0]$ such that the m^{th} digitized point is sampled at time

$$t_j[m] = t_j[0] + \frac{m}{F_j}. \tag{G.1}$$

Ideally, the *realized* sampling rate F_j and trigger time $t_j[0]$ are equal to their *nominal* values F_j^{nom} and $t_j^{nom}[0]$, respectively. However, short-term jitter, long-term drifts, and constant offsets often plague real-world digitization such that $F_j \neq F_j^{nom}$, $t_j[0] \neq t_j^{nom}[0]$, and

$$t_j[m] \neq t_j^{nom}[0] + \frac{m}{F_j^{nom}}; \tag{G.2}$$

that is, the actual sample times of the digital record differ from their nominal values. In a properly operating digitizer, these discrepancies are typically small, and an autospectral-density estimate (for example) of $x_j(t)$ from its digital record will be negligibly compromised. When estimating the *phasing* between $x_j(t)$ and $x_k(t)$ for $j \neq k$, how-

ever, identifying and correcting such timebase discrepancies becomes paramount in importance.

G.2   WHICH DIGITAL RECORDS CAN BE SYNCHRONIZED?

The digitization scheme determines whether or not digital records $\{x_j[m]\}$ and $\{x_k[m]\}$ can be synchronized. The cleanest, simplest, and most problem-free scheme is to digitize $x_j(t)$ and $x_k(t)$ on the *same* system such that the actual sample rates and trigger times of both digital records are identical (i.e. $F_j = F_k$ and $t_j[0] = t_k[0]$, respectively). However, such a scheme is not always feasible. Further, note that multiple digitizer boards operating in a master-slave configuration can still suffer from trigger-time offsets, despite nominally being part of the same digitization system. The next-best scheme is to use phase-locked digitizers such that

$$\frac{F_j}{F_k} = \frac{F_j^{nom}}{F_k^{nom}} = \text{constant (for \textit{phase-locked} digitizers)}, \tag{G.3}$$

regardless of any short-term jitter or long-term drift in the digitizer clocks. As shown in Section G.3, the actual sampling times of phase-locked digital records differ (at most) by a constant "trigger offset", which can be compensated easily. Finally, the least-desirable scheme is to use free-running digitizers such that $F_j/F_k \neq F_j^{nom}/F_k^{nom}$; it may be impossible to synchronize records from free-running digitizers. While the below discussion considers synchronization via post-processing, it should be noted for completeness that hardware solutions for synchronization also exist [2].

G.3   SYNCHRONIZATION OF PHASE-LOCKED DIGITAL RECORDS

This section details the synchronization of phase-locked digital records. Specifically, Section G.3.1 defines the "trigger offset" between phase-locked digital records, and Section G.3.2 discusses the phase bias produced by a finite trigger offset. Section G.3.3 describes methods for estimating the trigger offset. Then, using standard techniques [3, Sec. 4.5], the trigger offset can be compensated easily in post-processing, even if the offset is a non-integer multiple of the sample spacing.

Phase-locked digitizers may suffer from a deleterious "trigger offset". To see this, consider two digitizers j and k. Assume that the digitizers have different nominal trigger times $t_j^{nom}[0] \neq t_k^{nom}[0]$ but the same nominal sampling rate $F_j^{nom} = F_k^{nom}$. (If the digitizers have different nominal sampling rates, however, the records from one of the digitizers can be digitally resampled [3, Sec. 4.6] with the sampling rate of the other digitizer, and then the presentation below proceeds unchanged). Because the nominal trigger times of digitizers j and k differ, their m^{th} nominal timestamps also differ, i.e. $t_j^{nom}[m] \neq t_k^{nom}[m]$. Instead, $t_j^{nom}[m] = t_k^{nom}[n]$, where

$$n = m + F_j^{nom}\left(t_j^{nom}[0] - t_k^{nom}[0]\right), \tag{G.4}$$

and the equality of the nominal sampling rates has been utilized. Now, for digitizer j define $\delta t_j = t_j[0] - t_j^{nom}[0]$ to be the difference between the actual and nominal trigger times, $\delta F_j = F_j - F_j^{nom}$ to be the difference between the actual and nominal sampling rates, and $\overline{\delta F_j} = \delta F_j/F_j^{nom}$ to be the normalized difference between the actual and nominal sampling rates ($|\overline{\delta F_j}| \ll 1$). Similar definitions apply for digitizer k. Because the digitizers satisfy the phase-locked constraint (G.3),

$$\overline{\delta F_j} = \overline{\delta F_k}. \tag{G.5}$$

(Note that equality (G.5) holds even if the nominal sampling rates are different). Then, to first order in $\overline{\delta F_j}$, the actual sampling times $t_j[m]$ are related to the nominal sampling times $t_j^{nom}[m]$ via

$$t_j[m] \approx t_j^{nom}[m] + \delta t_j - \frac{m \cdot \overline{\delta F_j}}{F_j^{nom}}. \tag{G.6}$$

Thus, trigger-time discrepancy δt_j produces a constant offset between the actual and nominal sampling times of digitizer j, while sampling-rate discrepancy δF_j produces a linear ramp between the actual and nominal sampling times of digitizer j.

Now, in some situations it is of the utmost importance that the actual sampling times of digitizers j and k align. (Inferring the spatial structure of a signal from its phasing between two spatially separated sensors is one such example). The *nominal* sampling times are aligned
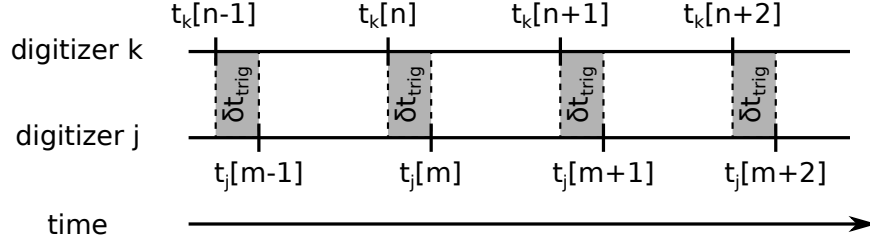
Figure G.1: The "trigger offset" δt_{trig} between phase-locked digitizers. Here, the n^{th} nominal sampling time of digitizer k equals the m^{th} nominal sampling time of digitizer k (i.e. $t_j^{nom}[m] = t_k^{nom}[n]$), where n and m are related via (G.4). However, a finite trigger offset (G.7) produces a discrepancy between the *actual* sampling times of the two digitizers (i.e. $t_j[m] \neq t_k[n]$), which can bias phase measurements, as discussed in Section G.3.2.

via (G.4), but any remaining discrepancy between the *actual* sampling times is

$$\begin{aligned} \delta t_{trig} &= t_j[m] - t_k[n] \\ &= \left(\delta t_j - \delta t_k\right) + \overline{\delta F_j}\left(t_j^{nom}[0] - t_k^{nom}[0]\right); \end{aligned} \tag{G.7}$$

this situation is shown schematically in Figure G.1. Here, the first term on the right-hand side of (G.7) corresponds to the difference between trigger-time discrepancies of each digitizer, while the second term on the right-hand side corresponds to the relative sampling-rate discrepancy weighted by the difference in nominal trigger times. As both effects are related to triggering, this timestamp discrepancy is referred to as the "trigger offset" δt_{trig}. Note that δt_{trig} is a single, constant value for any given pair of phase-locked digital records.

G.3.2 *Effect of the trigger offset*

The trigger offset (G.7) biases the measured phasing between the digital records. To see this, let $x_j(t)$ be a coherent mode of angular frequency ω such that the corresponding digital record is

$$\begin{aligned} x_j[m] &= x_j(t_j[m]) \\ &\propto X_j(\omega)e^{i\omega t_j[m]} \\ &= |X_j(\omega)|e^{i\{\alpha_j(\omega)+\omega t_j[m]\}}, \end{aligned} \tag{G.8}$$

where $|X_j(\omega)|$ is the Fourier amplitude and $\alpha_j(\omega)$ is the Fourier phase. (Here, the Fourier-transform kernel is $\propto e^{-i\omega t}$ and the inverse Fourier-transform kernel is $\propto e^{i\omega t}$, in accord with the NumPy implementa-

tion of the fast Fourier transform (FFT) [4]; if the opposite convention is used for FFT computations, then the substitution $\delta t_{\text{trig}} \to -\delta t_{\text{trig}}$ should be made in equations (G.9), (G.10), and (G.11)). Then, after aligning *nominal* sampling times via (G.4), the *measured* phase difference $\Delta\alpha_{\text{meas}}$ between digital records $\{x_j[m]\}$ and $\{x_k[n]\}$ is

$$
\begin{aligned}
\Delta\alpha_{\text{meas}} &= \arg\left(x_k^*[n] \cdot x_j[m]\right) \\
&= \left[\alpha_j(\omega) - \alpha_k(\omega)\right] + \omega\left(t_j[m] - t_k[n]\right) \\
&= \Delta\alpha(\omega) + \left(\omega \cdot \delta t_{\text{trig}}\right),
\end{aligned} \tag{G.9}
$$

where $\Delta\alpha(\omega) = \alpha_j(\omega) - \alpha_k(\omega)$ is the true phase difference and δt_{trig} is defined in (G.7). Thus, non-zero trigger offset δt_{trig} biases the measured phase difference $\Delta\alpha_{\text{meas}}$ away from the true phase difference $\Delta\alpha$. The above argument readily extends to broadband signals.

G.3.3 *Estimating the trigger offset*

Clearly, a finite trigger offset is undesirable. In some situations, it is possible to estimate the trigger offset. Then, using standard techniques [3, Sec. 4.5], the trigger offset can be compensated easily in post-processing, even if the offset is a non-integer multiple of the sample spacing.

If the true phase difference $\Delta\alpha(\omega)$ is known *a priori* (from e.g. another measurement), solving for δt_{trig} in (G.9) yields an estimated trigger offset

$$
\delta t_{\text{trig}} = \frac{\Delta\alpha_{\text{meas}}(\omega) - \Delta\alpha(\omega)}{\omega}. \tag{G.10}
$$

Although *a priori* knowledge of $\Delta\alpha$ may make (G.10) seem rather academic, it does find real-world application. For example, imagine the signals from a regularly spaced array of channels are digitized across multiple digitizer boards. The intra-board trigger offsets are negligible such that the true phase difference $\Delta\alpha$ can be accurately estimated from adjacent channels digitized on the same board; comparing this estimate of $\Delta\alpha$ to the measured phase difference $\Delta\alpha_{\text{meas}}$ between adjacent channels digitized on different boards via (G.10) then yields an estimate of the trigger offset between the boards. This methodology is used to estimate the trigger offset between the two boards of the DIII-D PCI digitizer.

In addition to requiring *a priori* knowledge of the true phase difference $\Delta\alpha$, trigger-offset estimate (G.10) also suffers from aliasing. That is, $\Delta\alpha_{\text{meas}}$ is only measured modulo 2π such that (G.10) specifies an

infinite set of potential trigger offsets, with adjacent values spaced by $2\pi/\omega$. This is particularly troublesome for "large" trigger offsets.

Under certain circumstances, the trigger offset can be estimated in an alternative, alias-free manner. For example, consider a coherent mode with time-dependent angular frequency $\omega(t)$. If the angular frequency ramps linearly in time (i.e. $\dot{\omega} = d\omega/dt = $ constant) and the true phase difference $\Delta\alpha$ does *not* vary in time, taking the time derivative of (G.9) and solving for δt_{trig} yields

$$\delta t_{trig} = \frac{1}{\dot{\omega}} \frac{d\left[\Delta\alpha_{meas}(\omega)\right]}{dt} = \frac{d\left[\Delta\alpha_{meas}(\omega)\right]}{d\omega}. \tag{G.11}$$

Because (G.11) depends on the derivative of $\Delta\alpha_{meas}$, it is an alias-free estimate of the trigger offset. Further, (G.11) does *not* require *a priori* knowledge of the true phase difference $\Delta\alpha$ (other than requiring that it be constant in time). This methodology is used to estimate the trigger offset between DIII-D's two toroidally separated interferometers.

BIBLIOGRAPHY

[1] W. R. Bennett. Spectra of quantized signals. *Bell Sys. Tech. J.*, 27(3):446–472, July 1948.

[2] J. Stillerman, W. Burke and B. Labombard. An optical timing verification system for Alcator C-Mod. *Fusion Engr. Design*, 85(3):367 – 369, 2010.

[3] A. V. Oppenheim and R. W. Schafer. *Discrete-Time Signal Processing*. Pearson, Upper Saddle River, NJ, USA, 3rd edition, 2010.

[4] The Scipy community. *numpy.fft (v1.13.0)*, June 2017.