Ameliorating Systematic Errors in Full-Field AMCW Lidar

A thesis submitted in fulfilment
of the requirements for the degree of

Doctor of Philosophy

in Electronic Engineering

by

John Peter Godbaz

THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2012

Abstract

This thesis presents an analysis of systematic error in full-field amplitude modulated continuous wave range-imaging systems. The primary focus is on the mixed pixel/multipath interference problem, with digressions into defocus restoration, irregular phase sampling and the systematic phase perturbations introduced by random noise. As an integral part of the thesis, a detailed model of signal formation is developed, that models noise statistics not included in previously reported models.

Prior work on the mixed pixel/multipath interference problem has been limited to detection and removal of perturbed measurements or partial amelioration using spatial information, such as knowledge of the spatially variant scattering point spread function, or raytracing using an assumption of Lambertian reflection. Furthermore, prior art has only used AMCW range measurements at a single modulation frequency. In contrast, in this thesis, by taking multiple measurements at different modulation frequencies with known ratio-of-integers frequency relationships, a range of new closed-form and lookup table based inversion and bounding methods are explored. These methods include: sparse spike train deconvolution based multiple return separation, a closed-form inverse using attenuation ratios and a normalisation based lookup table method that uses a new property we term the characteristic measurement. Other approaches include a Cauchy distribution based model for backscattering sources which are range-diffuse, like fog or hair. Novel bounding methods are developed using the characteristic measurement and attenuation ratios on relative intensity, relative phase and phase perturbutation.

A detailed noise and performance analysis is performed of the characteristic measurement lookup table method and the bounding methods using simulated data. Experiments are performed using the University of Waikato Heterodyne range-imager, the Canesta XZ-422 and the Mesa Imaging Swissranger 4000 in order to demonstrate the performance of the lookup table method. The lookup table method is found to provide an order of magnitude improvement in ranging accuracy, albeit at the expense of ranging precision.

Acknowledgements

Firstly, a big thank-you to my supervisors Drs. Michael Cree and Adrian Dorrington for their support and guidance. This document has been a long and painful time in development and would not have been possible without their moral, analytical and financial support.

I am deeply indebted to all those who have participated in the Range-Imaging research group over the past seven years: in particular to Dr. Andrew Payne for all his technical assistance and to Lee Streeter, who's grasp of mathematics greatly exceeds my own.

I would also like to acknowledege the support of a Teriary Education Commission Top Achiever Doctoral Scholarship over the first three years of my study; this was a great enabler and it is a huge shame that the programme has been cancelled.

Lastly, thank-you to my parents Peter and Miranda, and to my sister Susannah for putting up with all my comings-and-goings at odd hours over the past four or so years.

Contents

List of Figures

List of Tables

List of Symbols

$f[x]$	A function of the discretely valued parameter x (eg. $x \in \mathbb{Z}$)
$f(x)$	A function of the continuously valued parameter x (eg. $x \in \mathbb{R}$)
$\forall_a b$	Universal instantiation

Unary Operations

a^{T}	Transpose of a
a^*	Complex conjugate or Hermetian transpose of a
$\ln a$	Natural logarithm of a
$\Re(a)$	Real part of a
$\Im(a)$	Imaginary part of a
$\arg(a)$	Complex argument of a (phase angle)
$\lvert a \rvert$	Absolute value of a
$\lfloor a \rfloor$	Floor of a
$\lceil a \rceil$	Ceiling of a
$\lVert a \rVert_p$	$\mathcal{L}_p$ norm of a

Binary Operations

$a * b$	Convolution of a by b
$a \otimes b$	Hadamard product of a with b (elementwise multiplication)
$a \oslash b$	Elementwise division of a by b
$a \sim b$	a is sampled from probability distribution b
$\mathbb{E}[a\lvert b]$	The expected value of a given b
$\mathcal{L}(a\lvert b)$	The likelihood of a given b
$a \wedge b$	Logical and
$a \vee b$	Logical or
$a \Rightarrow b$	Implies operator
$a \Leftrightarrow b$	If and only if
$a \in b$	a is a member of set b
$a \backslash b$	The set a with the elements of set b removed
$a \approx b$	a is approximately equal to b

Subscripts and Superscripts

$\hat{a}$	Normalised version of measurement or component return a
$\tilde{a}$	Measurement a, subject to systematic errors
$\square_{sv}$	Indicates that operation $\square$ is spatially variant

Distributions

$\mathcal{B}(a)$	Bernoulli distribution, parameterised by a
$\mathcal{N}(\mu, \Sigma)$	Gaussian distribution with mean μ and variance Σ
$\mathcal{N}_{\mathbb{C}}(\mu, \sigma^2)$	Circularly symmetric complex Gaussian distribution
$\mathcal{P}(\lambda)$	Poisson distribution for mean λ
$\mathcal{U}(a, b)$	Uniform distribution between a and b

Variables and Functions

a_0	Absolute amplitude of primary component return (usually brightest)
a_1	Absolute amplitude of second component return
a_i	Absolute amplitude of ith component return
$a_{\mathrm{p}}(b_h)$	Expected amplitude perturbation from circularly symmetric perturbation for specified relative intensity
$a_{\mathrm{q}}(\sigma_\xi)$	Expected amplitude perturbation from a circularly symmetric noise distribution with Rayleigh density
a_ρ	Truncated triangle parameter (amplitude)
b	Relative amplitude of secondary component return
b_h	Relative amplitude of aliasing harmoninc or perturbing component
Δb_h	Change in relative amplitude of perturbing component
b_ρ	Truncated triangle parameter (peak intensity)
$b_{\min}$	Minimum bound on the relative phase of the second component return
$\mathrm{cov}(x, y, \varkappa)$	Covariance of x and y over regions of an image masked by $\varkappa$ implemented in the Fourier domain
c	Relative amplitude of secondary component return (as a function parameter)
c_ρ	Truncated triangle parameter (constant offset)
c_ϵ	Naïve phase error improvement by taking the minimum of multiple samples from a distribution
D	Aperture diameter
d	Radial distance
d_i	Radial distance to the ith component return
d_c	Demodulation contrast
d_ϕ	Range as a parameter
$E(\phi, k)$	Incomplete second order elliptic integral
$E(k)$	Complete second order elliptic integral
e	Euler's number, base for natural logarithm
$F[x, y](u)$	Continuous Fourier transform of signal return model versus range
$F_\nu[u]$	Fourier transform of base-ratio model for correlation waveform
$f[x, y](d)$	Signal return model (discrete pixels, but continuous over range)
$f_{\mathcal{N}}(x\|\mu, \sigma^2)$	PDF for normal distribution
$f_{\mathrm{p}}[i]$	Truncated triangle model for correlation waveform (hetrodyne)
$f_{\mathrm{Q}}(x)$	Asymmetric function that demonstrates systematic perturbations from the impact of random noise
f_ξ	The signal return model
$f_{\mathcal{L}}(u\|0, \lambda_{\mathcal{C}})$	PDF of a Laplace distribution as a function of the spread/width parameter of the corresponding Fourier transformed Cauchy distribution

$f_\nu[i]$	Linear interpolation (base-ratio) model for the correlation waveform
h_i	Duty cycle of illumination waveform
h_s	Duty cycle of sensor waveform
H_ψ	Matrix transform corresponding to cyclic convolution by the reference waveform ψ
H_ν	Matrix transformation corresponding to the linear basis for the base-ratio correlation waveform model given a coarse phase step
$\mathcal{I}_n$	$n \times n$ identity matrix
I_i	Imaginary part of ith measurement
i	Counter (typically used over each component return)
J	Jacobian matrix
j	Imaginary unit, $\sqrt{-1}$
K_l	Set of aliasing correlation waveform harmonics for measurement l
k	Counter (typically used over correlation waveform harmonics)
L	2D Laplacian filter
L_θ	Set of stationary points of the modulus of the characteristic measurement with respect to relative phase θ
l	Counter (typically used over each measurement)
$l_k(y_0, y_1)$	Phase calculation function from location in PCA decomposed form
M	Number of instances
m	Number of instances
N	Number of instances (typically used for number of measurements)
n	Number of instances (typically used for number of component returns)
$P_{\mathrm{q}}(\Delta\theta)$	Reference waveform sampling function
$p(w)$	LUT projection function
p_r	Probability of a spike per sample in a sparse spike train
p_s	TT. Param
Q_{m}	Expected minimum absolute value of m samples from a Gaussian distribution with unit variance
$Q_{\mathrm{r}}(\Delta\theta)$	Calibrated measurement response as a function of TOF induced phase shift
R_i	Real part of ith measurement
$R(\Xi)$	Function that inverts the multiple return problem to determine H
$\hat{R}(\chi)$	Normalised inverse of multiple return problem
R_{X}	Rotation and scaling matrix corresponding to the calibration for the phase and amplitude of the negative fundamental frequency of the reference waveform
R_ξ	Rotation and scaling matrix that rotates a range measurement to $(1, 0)^T$
r_0	Relative frequency of measurement ξ_0
r_1	Relative frequency of measurement ξ_1
r_l	Relative frequency of measurement ξ_l
r_h	Truncated triangle waveform shape parameter
$\mathrm{rect}(x)$	Rect function
$S(u)$	Fourier transform of the continuous correlation waveform
$S[u]$	Fourier transform of the discrete correlation waveform

$\tilde{S}[u]$	Fourier transform of the discrete correlation waveform subject to aliasing
$s(\Delta\phi)$	Continuous correlation waveform measured at phase step $\Delta\phi$
$s[i]$	Sample i of the correlation waveform
$\tilde{s}[i]$	Sample i of the correlation waveform subject to aliasing and systematic error
$\mathrm{sinc}(x)$	Normalised sinc function
$T_n(x)$	The nth order Chebyshev polynomial of the first kind
$t_\mathrm{P}(\xi)$	Amplitude and phase tuple as a function of a complex domain range measurement
$U_n(x)$	The nth order Chebyshev polynomial of the second kind
u	Spatial frequency (as a function parameter)
v	Reference spatial frequency (ξ_l is a sample of spatial frequency $r_l v$), also used for vertical spatial frequency in cases of 2D Fourier transform
w	Utility variable (local context)
x	Horizontal pixel location or a utility variable (local context)
x_i	ith eigenvector of a matrix
x_0	Eigenvector corresponding to direction of maximum variance
x_1	Eigenvector perpendicular to x_0
y	Vertical pixel location
y_0	Distance along the direction of the eigenvector x_0
y_1	Distance along the direction of the eigenvector x_1
z_Q	Intermediate variable for calculating stationary points of Characteristic Measurement
α	Utility variable
$\alpha_{k,l}$	Aliasing coefficient for harmonic k in measurement l
α_DN	Photons per DN (or was it the other way around?)
β	Utility variable
β_B	Relative gain mismatch between channels A and B on sensor
β_l	Statistical bound for LF deconvolution algorithm in standard deviation units
γ	Regularisation constant
γ_B	Relative offset for inverted modulation waveform for B channel on sensor (vs $S(0)$)
Γ	The set of all phase steps in the context of a reference waveform sampling function
$\delta(x)$	Dirac delta function
ϵ_A	Error vector for κ_A
ϵ_B	Error vector for κ_B
ϵ_M	Noise in a complex domain range measurement
ϵ_X	Overall error in a measurement
ϵ_r	Read noise in a measurement
ϵ_s	Photon shot noise in a measurement
ϵ_u	Uncertainty bounds in context of LF deconvolution method
H	Set of all component returns
$\hat{\mathrm{H}}$	Set of normalised component returns

η_0	Primary component return (usually brightest)
η_1	Second component return
η_i	ith component return
$\hat{\eta}_0$	Normalised primary component return (usually brightest)
$\hat{\eta}_i$	ith normalised component return
θ	Relative phase of secondary component return (vs. primary). Also, phase of correlation waveform.
$\Delta\theta$	Relative phase induced by TOF to object
θ_f	Phase perturbation of the primary component return by the secondary component return
θ_0	Absolute phase of the primary component return
θ_1	Absolute phase of the second component return
θ_i	Absolute phase of the ith component return
θ_h	Integration variable and for relative phase of aliasing harmonic
θ_ξ	Phase of a complex domain range measurement
θ_min	Minimum bound on the relative phase of the second component return
θ_high	Relative phase corresponding to maximal modulus of the characteristic measurement
θ_low	Relative phase corresponding to minimum modulus of the characteristic measurement
θ_b	Relative phase corresponding to the bounding case for the characteristic measurement
$\hat{\kappa}$	Total integrated intensity normalised by η_0
κ_A	Vector of sampled correlation waveform samples subject to noise (A channel)
κ_B	Vector of sampled correlation waveform samples subject to noise (B channel)
$\Lambda_\mathrm{f}(c, \phi_\mathrm{f})$	Relative perturbation of the primary component return by the second for specific relative phase and amplitude values
$\Lambda_\chi(c, \phi_\mathrm{f})$	The characteristic measurement for specific relative phase and amplitude values
λ	AMCW modulation wavelength
λ_i	ith eigenvalue of a matrix
λ_ξ	Fine (subsample) location for linear interpolation (base-ratio) model of correlation waveform
$\lambda_\mathcal{C}$	Width parameter of Cauchy distribution
μ_l	Measurement normalisation coefficient at spatial frequency $r_l v$
μ_n	Measurement normalisation coefficient at spatial frequency v
μ_0	Measurement normalisation coefficient at spatial frequency $r_0 v$
ν	Phasor describing a diffuse-range return in polar form, encoding the width and phase of the component return
ν_i	ith diffuse-range polar phasor, encoding phase and range-width
ν_ρ	Vector of base-ratio parameters (basis vector weightings for decomposition)
ν_ξ	Coarse sample location (discrete samples) for linear interpolation (base-ratio) model of correlation waveform

$\nu_{\mathcal{L}}$	Width parameter for a Cauchy Cartesian diffuse-range return model
$\nu_{\mathcal{N}}$	Width parameter for a Gaussian Cartesian diffuse-range return model
Ξ	Set of all complex domain range measurements
$\hat{\Xi}$	Set of normalised measurements
ξ_0	First normalised range measurement (usually lowest frequency)
ξ_1	Second complex domain range measurement
ξ_l	lth complex domain range measurement
$\hat{\xi}_0$	First normalised range measurement (usually lowest frequency)
$\hat{\xi}_l$	lth normalised range measurement
π	Ratio of circle circumference to diameter in Euclidian space
π_{B}	Absolute offset of inverted B channel sensor modulation waveform
π_α	PSF blur model – scaling parameter
π_β	PSF blur model – distance from front principal plane to the in-focus point along the optical axis
π_{R}	Real part of a measurement at the zeroth spatial frequency
π_{I}	Imaginary part of a measurement at the zeroth spatial frequency
ρ_{h}	Waveform shape modelling – period held high
ρ_{l}	Waveform shape modelling – period held low
ρ_{s}	Waveform shape modelling – transition period
ρ_{U}	Relative attenuation constant for uniform diffuse-range component return model versus a Dirac model
ρ_X	Calibration coefficient for the phase and amplitude of the negative fundamental frequency of the reference waveform
ρ_I	Offset due to ambient light
σ_p	PSF blur radius
σ_r^2	Variance of read noise for a measurement
σ_α	PSF blur radius for image α
σ_β	PSF blur radius for image β
$\sigma^2_{\Re\Re}$	Variance of the real part of an unrotated complex domain range measurement
$\sigma^2_{\Re\Im}$	Covariance of the real part of an unrotated complex domain range measurement with the imaginary
$\sigma^2_{\Im\Re}$	Covariance of the real part of an unrotated complex domain range measurement with the imaginary
$\sigma^2_{\Im\Im}$	Variance of the imaginary part of an unrotated complex domain range measurement
σ^2_ξ	Variance of a complex domain range measurement in the circularly symmetric case
Σ_ρ	Unrotated covariance of a complex domain range measurement
Σ_θ	Covariance of a range measurement in polar form
Σ_ν	Covariance of the measured data for least squarse base-ratio model fitting to the correlation waveform
τ_0	Attenuation ratio at relative frequency r_0
τ_1	Attenuation ratio at relative frequency r_1
τ_l	Attenuation ratio at relative frequency r_l

$\upsilon_\lambda[i]$	Model of linear interpolation (base-ratio model) as a function of fine step size (λ_ξ)
$\Upsilon_\lambda[u]$	Fourier transform of model of linear interpolation as a function of fine step size (λ_ξ)
ϕ_d	Phase step
$\Delta\phi_\mathrm{d}$	Phase step size
ϕ_h	Heterodyning phase domain (phase domain to integrate over at each step)
ϕ_f	Relative phase of a component return (as a function parameter)
χ	Characteristic measurement
$\Psi(u)$	Fourier transform of reference waveform
$\Psi_\mathrm{i}(u)$	Fourier transform of illumination waveform
$\Psi_\mathrm{s}(u)$	Fourier transform of sensor response waveform
$\psi(\phi)$	Reference waveform
$\psi_\mathrm{i}(\phi)$	Illumination waveform
$\psi_\mathrm{s}(\phi)$	Sensor response waveform
ω	Pixel weightings
ω_λ	Phase step perturbation for unequal phase step simulation

Chapter 1

Introduction

If ranging is the science of measuring distances, then range-imaging is the science of systematically sampling the range to objects over a particular field of view. Instead of a single range measurement, full-field image ranging devices produce two dimensional matricies of range data, which can be utilised for such varied applications as machine vision, process-line quality control and computer gaming. For many years range-imaging techniques have been applied to problems such as 3D modelling for computer games, manufacturing quality control and robotic vision, however applications such as computer gaming and human-computer interfaces have been limited by the inability to produce sequences of high resolution range-images at adequate frame rates in real-time. Techniques have either been capable of real-time application, but not reliable or accurate enough – such as stereo vision and depth-from-focus – or have been slow but accurate, like pattern projection, triangulation line scanners and lidar point scanning systems.

Recently, the development of full-field systems has accelerated as systems have shed the necessitity for complicated optical mixers, such as image intensifiers, in favour of CMOS sensors specifically designed for the purpose. As Microsoft's Kinect camera has demonstrated, there is a huge untapped market for full-field ranging devices. With the potential for mass-manufacture to reduce costs, it is reasonable to expect range-imaging to become a prominant part of many every-day consumer devices. One of the biggest advantages of full-field systems is the lack of moving parts, whereas point scanners require the manufacture of expensive precision machinery. However, unless the multipath interference induced by the transition to full-field systems is ameliorated, point scanners will continue to command much of the ranging market.

With the increasing utlisation of full-field range-imagers the importance of removing systematic errors from these systems has grown simultaneously. While many

full-field amplitude modulated continuous wave ranging systems, such as the SwissRanger 4000, advertise centimetre level precision the accuracy of these systems is commonly an order of magnitude or more worse, predominantly due to the mixed pixel effect or multipath interference. In order to extend the use of these systems to as wide a range of real-world problems as possible, it is necessary to understand and address these errors.

Previous research into mixed pixels and multipath interference in full-field AMCW range-imagers has been relatively limited. While the formation of mixed pixels is relatively simple, the vast majority of prior art has been restricted to the detection and removal of perturbed range estimates. For multipath interference generated by scattering within the optics of full-field systems, there has also been the limited application of interative restoration techniques, modelling the scattering as a spatially variant point spread function (PSF) and applying deconvolution to partially ameliorate the erroneous range estimates.

This thesis focusses upon modelling and removal sources of systematic error via sagacious use of post-processing. While primarily focussing on the mixed pixel and multipath interference problems, we also digress onto the systematic impacts of noise, uneven phase sampling and defocus. Ultimately, the most important research is the development of new methods for the determination of the range and amplitude characteristics of component returns within mixed pixels; we believe these to be significant contributions to the literature. Whereas prior art has attempted to ameliorate the mixed pixel/multipath interference problem using measurements taken at a single modulation frequency, our research overcomes the limitations of previous methods, amoung other techniques, by taking multiple measurements at different modulation frequencies.

1.1 Overview of the Thesis

In Chapter 2 we introduce different types of lidar systems and their properties. We review prior art, using it to explain the different types of systematic errors characteristic of AMCW lidar systems; this leads to the explaination of the mixed pixel interference problem, which is the focus of this document. The chapter demonstrates how previous research has not fully addressed the fundamental nature of the mixed pixel problem. A description of the lidar systems utilised for experiments in this thesis is given, aiming to give the reader an enhanced understanding of the specifics involved in the implementation of the hardware.

In Chapter 3 a fundamental model for AMCW range measurement formation

is developed. While this builds on material already in the literature, much of the analysis is new. Topics include correlation waveform formation and spatiotemporal variance, signal return models, aliasing, random error and the impacts of uneven phase sampling.

The remainder of the main body of the thesis is devoted to developing methods to separate out the component returns within mixed pixels. Chapter 4 develops methods for fitting the shape of the correlation waveform; relying on harmonic content to separate out the component returns. We develop three separate methods: the first using an off-the-shelf sparse spike train deconvolution method, another using optimisation to fit piecewise functions to the waveform and lastly a simple approach using a Moore-Penrose pseudo-inverse to determine phase.

Chapter 5 develops two different analytically invertable systems of equations for modelling the multiple return problem. The first is based around the concept of range-diffuse sources, such as fog or hair, and allows the determination of the range, amplitude and spread of up to two components within each pixel from range measurements at four different frequencies. The second approach uses a concept called attenuation ratios – which are related to the cancellation properties of component returns at a particular measurement frequency. By calculating the attenuation ratio at two different measurement frequencies, it is possible to determine the amplitude and range of up to two components within each pixel. However, in practical implementation determination of the attenuation ratio requires two measurements: one of the total integrated intensity of the component returns and one at the specified modulation frequency. As a result, this method is over-determined. Determining the range and amplitude of two component returns within a mixed pixel requires two measurements of different, non-zero spatial frequencies in the exactly-determined case (this is demonstrated using optimisation in Chapter 6). In fact, neither method is exactly-determined and we have been unable to find a closed-form inverse for any exactly-determined case. While it may be possible, it is a *non-trivial* exercise.

In Chapter 5 we additionally develop the concepts of range measurement normalisation and component return normalisation. These allow the modelling of the relationships between range measurements and component returns, free of any scaling or range translation effects using a quantity we call the characteristic measurement.

Chapter 6 introduces the characteristic measurement, which is a quantity that describes the relationship between two complex domain range measurements taken at different modulation frequencies. The chapter further develop the mathematical properties of the characteristic measurement, and relates them to the relative amplitude and phase of the component returns. Using these properties a set of bounding

methods are derived using a mixed theoretical and empirical approach.

Chapter 7 focusses on evaluation of the bounding and separation methods presented in Chapter 6. After introducing the simulatory assumptions, the chapter systematically evaluates the impact of relative phase, amplitude, SNR, frequency ratio and many other parameters on restoration performance. The chapter also includes an analysis of characteristic measurement and attenuation ratio based bounding methods and a comparison between the characteristic measurement lookup table method from Chapter 6 and the attenuation ratio polynomial method from Chapter 5.

Chapter 8 gives a brief overview of future options for ongoing research.

In the appendicies we cover additional material that was not included in the main body of the thesis for the sake of cohesiveness. Appendicies A and B develop new approaches to defocus restoration unique to full-field AMCW lidar systems and appendix C develops a scene texture/patterning approach to determination of range (range/phase from texture) and scattered light estimation. These three appendicies are copies of published papers Godbaz *et al.* (2010), Godbaz *et al.* (2011) and Godbaz *et al.* (2009b) respectively, with minor alterations.

1.2 Contributions of the Thesis

This thesis makes a number of original contributions, including:

- Discovery of the range/amplitude normalisation that defines the characteristic measurement, the latter of which is a also new concept.
- Development of the characterisic measurement lookup table inversion method; an exactly-determined inversion method. This includes an extremely detailed performance analysis and demonstration on real data.
- Development of a significant number of charactertic measurement and attenuation ratio based bounds on parameters such as relative amplitude, relative phase and the phase perturbation of the primary component return.
- Development of a Cauchy distribution based model for diffuse-range component returns, which allows determination of amplitude, phase and spread in specific over-determined cases using simple closed-form equations.
- Development of a closed-form attenuation ratio inversion method. In practice, this method is over-determined since attenuation ratios are currently unable to be measured directly.

- Development of correlation waveform shape based methods for separation of component returns and improving ranging precision. This includes development of a general equation for truncated-triangle waveforms that generalises the standard triangle wave assumption to any modulation duty cycle and quantification of the impact of duty cycle on phase precision.
- Discovery of the perturbing nature of the second harmonic of the correlation waveform, which results in the noise characteristics of complex domain range measurements being non-circularly symmetric due to shot noise statistics. Previous analyses have ignored the influence of harmonics in the correlation waveform; albeit only the second harmonic results in any net perturbation.
- Discovery that a non-circularly symmetric noise distribution results in a systematic phase perturbation, particularly if using the naïve approach of averaging phase/range rather than in the complex domain.
- A brief analysis of the systematic perturbations introduced by irregular phase steps.
- Development of a simple approach for blindly determining the focal parameters of a full-field lidar system – in other words, the distance at which objects are in-focus and a scaling coefficient. The focal parameters jointly describe the characteristics of the range variance of the point spread function of a simple optical system under geometric optics.
- Development of a proof-of-concept restoration method for defocussed range images. This approach utilises the range data intrinsic to the range image to determine the scale of the point spread function across the image and iteratively improve the quality of the range image. Whereas coded-apertures have previously been applied to standard intensity imaging, they have never been applied to range imaging; thus allowing us to make the spatially variant deconvolution problem more well-posed.
- Development of a simple method for determining true range from a range image using scene texture/patterning, irrespective of scattered light.

1.3 Publications Arising from this Thesis

In addition to the conference papers listed here, additional material by the author includes a Journal paper entitled ‘AMCW Lidar Mixed Pixel/Multipath Interference Bounding and Separation Using the Characteristic Measurement’ currently in

preparation and a full US patent (No. 627543001800 – 'Apparatus and Method for Measuring the Distance and/or Intensity Characteristics of Objects').

Journal Paper

- Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Understanding and ameliorating non-linear phase and amplitude responses in AMCW lidar. *Remote Sensing*, **4**, pp. 21–42 (2012b)

Conference Papers

As primary author:

- Godbaz, J. P., M. J. Cree, A. A. Dorrington, and R. Kunnemeyer. Defocus restoration for a full-field heterodyne ranger via multiple return separation. In: *Image and Vision Computing New Zealand 2007 (IVCNZ'07)*, pp. 52–57. University of Waikato, Hamilton, New Zealand (2007)

- Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Mixed pixel return separation for a full-field ranger. In: *Image and Vision Computing New Zealand 2008 (IVCNZ'08)*, pp. 1–6 (2008)

- Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Multiple return separation for a full-field ranger via continuous waveform modelling. In: *SPIE 7251 - Image Processing: Machine Vision Applications II*, volume 7251, p. 72510T. SPIE, San Jose, CA, USA (2009a)

- Godbaz, J. P., M. J. Cree, A. A. Dorrington, and A. D. Payne. A fast maximum likelihood method for improving amcw lidar precision using waveform shape. In: *IEEE Sensors 2009*, pp. 735–738 (2009c)

- Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Undue influence: Mitigating range-intensity coupling in amcw 'flash' lidar using scene texture. In: *Image and Vision Computing New Zealand 2009 (IVCNZ09)*, pp. 304–309 (2009b)

- Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Blind deconvolution of depth-of-field limited full-field lidar data by determination of focal parameters. In: *SPIE 7533 - Computational Imaging VIII*, volume 7533, p. 75330B. SPIE, San Jose, California, USA (2010)

- Godbaz, J., M. Cree, and A. Dorrington. Extending amcw lidar depth-of-field using a coded aperture. In: *ACCV 2010*, edited by R. Kimmel, R. Klette, and

A. Sugimoto, volume 6495 of *Lecture Notes in Computer Science*, pp. 397–409. Springer Berlin / Heidelberg (2011)

- Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Closed-form inverses for the mixed pixel/multipath interference problem in AMCW lidar. In: *Computational Imaging X, Proceedings of the SPIE, Vol. 8296* (2012a)

As primary originator of algorithm/methodology, but not primary author of paper:

- Dorrington, A. A., J. P. Godbaz, M. J. Cree, A. D. Payne, and L. V. Streeter. Separating true range measurements from multi-path and scattering interference in commercial range cameras. In: *SPIE 7864 - Three-Dimensional Imaging, Interaction, and Measurement*, edited by J. A. Beraldin, G. S. Cheok, M. B. McCarthy, U. Neuschaefer-Rube, A. M. Baskurt, I. E. McDowall, and M. Dolinsky, p. 786404. San Francisco Airport, California, USA (2011)

As a contributing author:

- Dorrington, A. A., M. J. Cree, D. A. Carnegie, A. D. Payne, R. M. Conroy, J. P. Godbaz, and A. P. P. Jongenelen. Video-rate or high-precision: a flexible range imaging camera. In: *SPIE 6813 - Image Processing: Machine Vision Applications*, volume 6813, p. 681307. SPIE, San Jose, CA, USA (2008)

- Larkins, R. L., M. J. Cree, A. A. Dorrington, and J. P. Godbaz. Surface projection for mixed pixel restoration. In: *Image and Vision Computing New Zealand 2009 (IVCNZ09)*, pp. 431–436 (2009)

- Cree, M. J., J. P. Godbaz, R. L. Larkins, W. H. Round, L. Streeter, A. A. Dorrington, R. Knnemeyer, A. D. Payne, and D. Worsley. Computer vision and image processing at the university of waikato. In: *Image and Vision Computing New Zealand (IVCNZ'10)*. Queenstown (2010)

Chapter 2

Background and Ranging System Overview

This chapter gives a brief overview of the background necessary to understand the body of the thesis. The first section of this chapter introduces full-field amplitude modulated continuous wave range imaging and explain how it works, including comparisons to related modulation techniques such as frequency modulated continuous wave and range-gating. The mixed pixel/multipath interference problem is explained as well as published mitigation techniques. The second section of the chapter is devoted to a literature review of lidar systems and studies on systematic error sources. The remainder of the chapter discusses the specific systems used for experimentation in this thesis.

2.1 Ranging Techniques, Mixed Pixels and Multipath Interference

There are many different types of ranging techniques, they can be categorised into passive and active techniques. Passive techniques typically rely on triangulation or analysis of light conditions and include depth-from-defocus (Pentland, 1987), stereo vision (Dhond and Aggarwal, 1989) and shape-from-shading (Zhang *et al.*, 1999). Active techniques rely on the active illumination of a scene, these techniques include pattern projection (Salvi *et al.*, 2004), active depth-from-defocus (Watanabe *et al.*, 1996) and time-of-flight Kolb *et al.* (2009). In this thesis we are interested in active time-of-flight range imaging; in particular, full-field amplitude modulated continuous wave (AMCW) lidar. AMCW lidar has many advantages over other active approaches including being faster than pattern projection while being more

accurate than active-focus methods.

There have been many reviews of the state-of-the-art in active range-finding. Kolb *et al.* (2009) discusses full-field time-of-flight (TOF) sensors in more of an applications context; in particular, range-RGB image fusion, real-time image segmentation using range information and determination of the plenoptic function, describing the light field, using range data. Other relevant reviews include: Foix *et al.* (2011), MacKinnon *et al.* (2008), Stefan May and Surmann (2007), Blais (2003) and the now extremely out-of-date Besl (1988).

2.1.1 Lidar Modulation Techniques

All lidar systems operate using the TOF principle. Because light travels at a finite, fixed speed, it is possible to determine to the range to an object using an active illumination source and by measuring the range-dependent propagation delay induced in the illumination signal as it travel to objects in the scene.

Fig. 2.1 is a basic overview of the time-of-flight principle. An illumination source sends out pulses of light, in order for the light to be measured by a sensor co-located with the illumination source it must travel twice the distance to the target object, in this case the stick figure. If a shrub is placed half way between the target and the illumination source, then the light from the shrub is reflected back to the camera and arrives before light from the stick figure. An example range-image is shown in Fig. 2.2.

While the concept of TOF is simple, the implementation is more subtle. One way to model the light measured by the camera is using a backscattering function. For a single pixel, the backscattering function models the intensity of light measured by the camera as a function of range; the ultimate aim of a lidar system is to reconstruct this function for each pixel. Fig. 2.1 shows the formation of the backscattered illumination waveform. Fig. 2.1a is an example backscattering function; in this case it is a scaled and translated Dirac delta function, which is the simplest possible model for the backscattering function. In order to measure the time delay in an illumination signal it is necessary to modulate the illumination. Fig. 2.1b is an example rectangular illumination modulation signal. The light which is backscattered by the scene and measured at the camera is given by the convolution of the backscattering function with the illumination modulation. In the case of Fig. 2.1c, where there is only a single backscattering source, the backscattering intensity and range are easy to extract. However, from a signal processing point-of-view, only spatial frequencies present in the illumination modulation waveform are present in the backscattered illumination waveform. This means that even if the backscattered illumination could

be perfectly recorded by the camera, only a limited amount of information is available about the backscattering function – and this depends on the frequency content of the illumination waveform. Fig. 2.4 gives an example of three different types of modulation; the first is AMCW or amplitude modulated continuous wave, shown in Fig. 2.4a; the second is FMCW or frequency modulated continuous wave, shown in Fig. 2.4c; and the final is range-gating or pulsed modulation, shown in Fig. 2.4e.

AMCW works by producing a continuous sinusoid modulation signal, rather than a pulse. The brightness of the illumination changes over time without changing the frequency of the modulation at any point, hence amplitude modulated. As the range to objects in the scene changes, the TOF changes, which results in the phase of the sinusoid changing.

In order to measure the range to an object, AMCW measures the phase of the sinusoidal modulation of the backscattered illumination[1]. One side-effect is that after 2π radians the range measurements repeat, giving a fixed ambiguity interval. Measurement of the waveform phase and amplitude is typically achieved by correlating the returned illumination signal with a reference sinusoid.

While lidar modulation techniques are exactly the same as those utilised in radar, AMCW is typically regarded as a 'toy system' in radar; one of the primary reasons being the issue of multipath interference. Multipath interference and mixed pixels are caused when a single range measurement contains light from more than one backscattering source; in an AMCW system this results in erroneous range and amplitude measurements. Why this occurs is explained in Section 2.2.1. As a result, little has been published on its intricacies outside the lidar world. Particularly for full-field systems, AMCW offers advantages over FMCW and range-gating systems because it requires much lower data rates, which is particularly important when trying to produce full-field range images at a high frame rate.

The most common implementation of AMCW range-imaging indirectly measures the illumination waveform by correlating it with another waveform using an optical mixer or modulated sensor at 90 degree phase steps. Because both the illumination and sensor modulation waveforms are sinusoidal, determining the range and brightness of an object becomes equivalent to determining the amplitude and phase of a sinusoidal signal. Fig. 2.5 shows this simple case, where four measurements, $I_0, I_1, I_2, I_3 \in \mathbb{R}$, have been taken at 90 degree phase steps. While there are many possible approaches to determing the phase and amplitude of the waveform from these measurements, the simplest approach is to Fourier transform the samples and find the phase and modulus of the Fourier bin corresponding to the sinusoidal signal.

[1] In reality, the waveform is not a sinusoid, but a sinusoid is conceptually simpler.

(a) Step One: Modulated LEDs or lasers illuminate the scene.

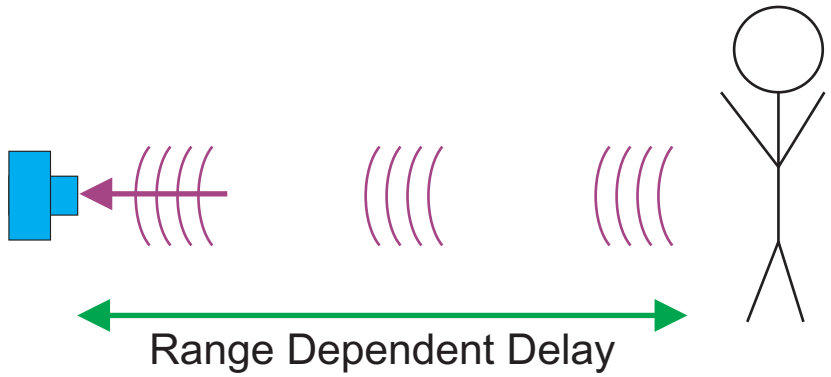

(b) Step Two: A range-camera measures the range induced delay in the modulated illumination signal, allowing estimation of range.

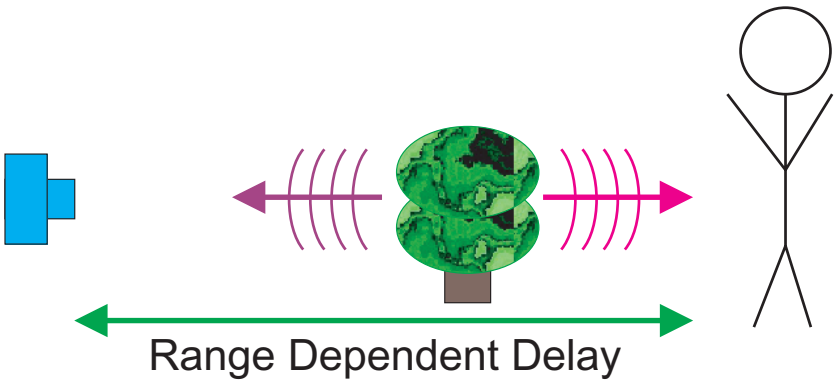

(c) Example: Two Objects. While light has already been reflected by the tree, the illumination signal has yet to reach the person.

Figure 2.1: Range measurement by determination of the time-of-flight of an active illumination signal.

(a) Photograph (b) 3D visualisation of the range-image

Figure 2.2: An example range-image (Copyright Andrew Payne).

By convention, we take the phase and amplitude of the negative fundamental frequency. Negative, in this case, because it means that the phase angle is proportional to range, rather than negatively proportional.

The illumination waveform, after having been sampled by correlation with the sensor modulation waveform, is known as the correlation waveform. We denote the actual correlation waveform, subject to systematic errors, as $\tilde{s}$. If the correlation waveform is sampled four times per cycle, then the range is given by

$$r \propto \arg\left(\tilde{S}\left(-\frac{1}{4}\right)\right) \tag{2.1}$$

$$\propto \operatorname{atan2}(I_3 - I_1, I_2 - I_0), \tag{2.2}$$

where $\tilde{S}(-\frac{1}{4})$ is the Fourier transform bin corresponding to the negative fundamental frequency of the correlation waveform. Equivalently, the amplitude is given by the modulus of the same Fourier bin:

$$a \propto \left|\tilde{S}\left(-\frac{1}{4}\right)\right| \tag{2.3}$$

$$\propto \sqrt{(I_3 - I_1)^2 + (I_2 - I_0)^2}. \tag{2.4}$$

This is the standard measurement approach for AMCW lidar: AMCW measurement formation is discussed in greater detail in Chapter 3.

FMCW operates by modulating the illumination signal with a chirp, which in

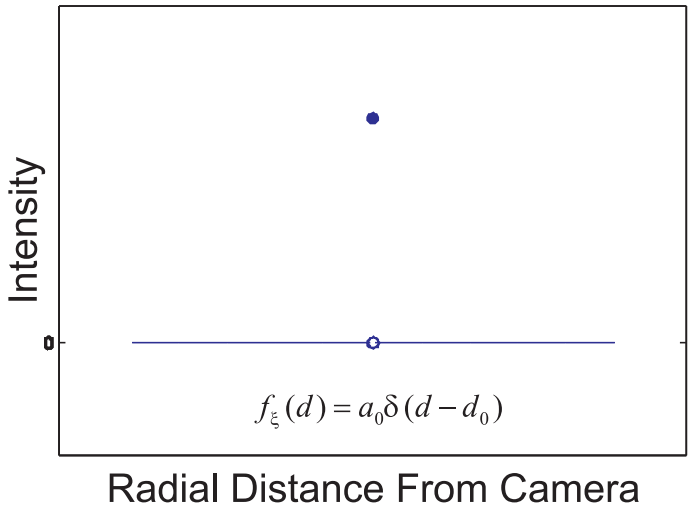

(a) The backscattering function represents the backscattered intensity as a function of range. It is typically modelled as a Dirac delta over range.

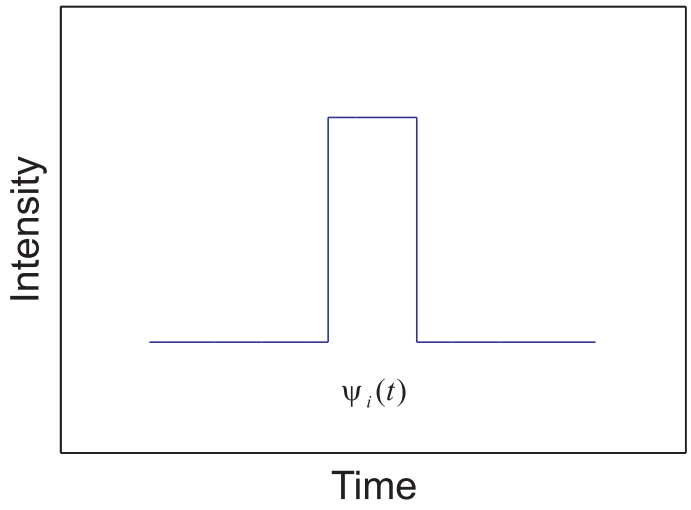

(b) The illumination modulation waveform.

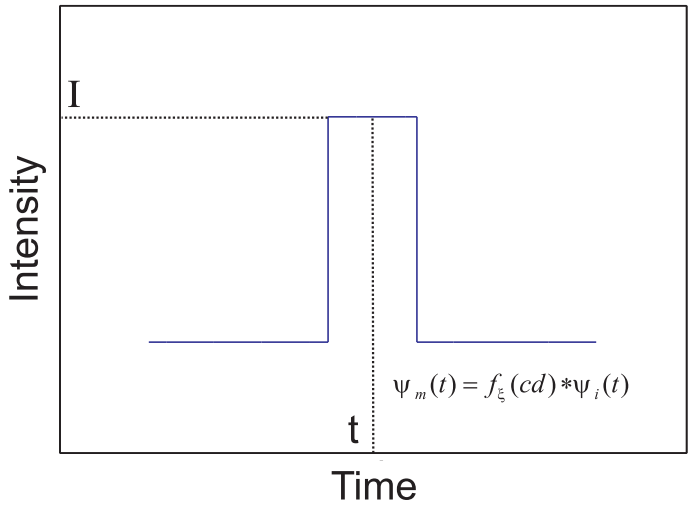

(c) The backscattered illumination waveform is the convolution of the backscattering function with the illumination modulation waveform, from whence the intensity of (I), and time delay to (t) the backscattering source can be determined

Figure 2.3: Formation of the backscattered illumination waveform.

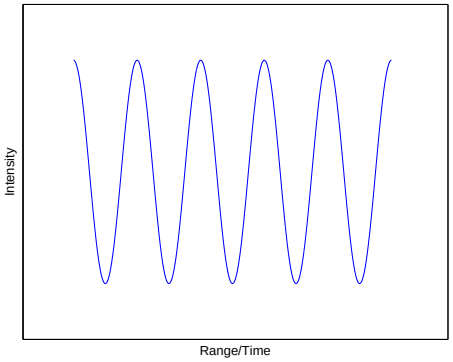

(a) AMCW Illumination Waveform

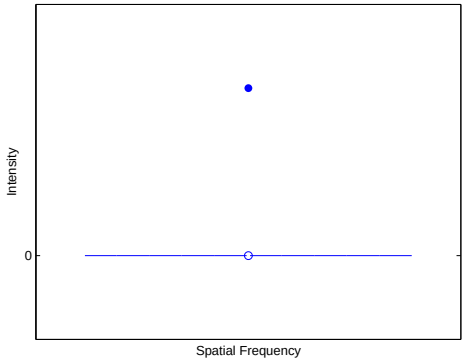

(b) AMCW Spatial Frequencies

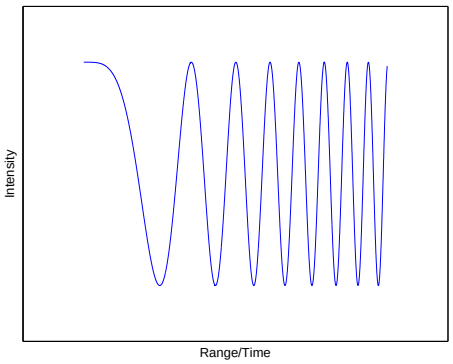

(c) FMCW Illumination Waveform

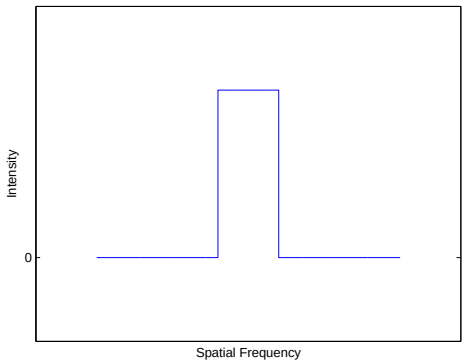

(d) FMCW Spatial Frequencies

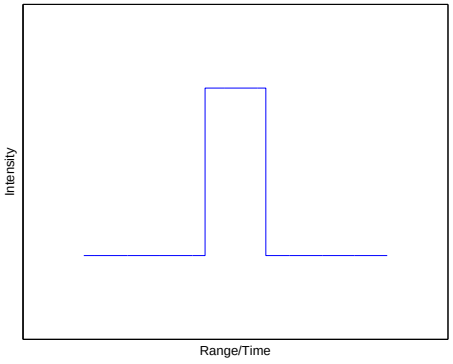

(e) Range-gating Illumination Waveform

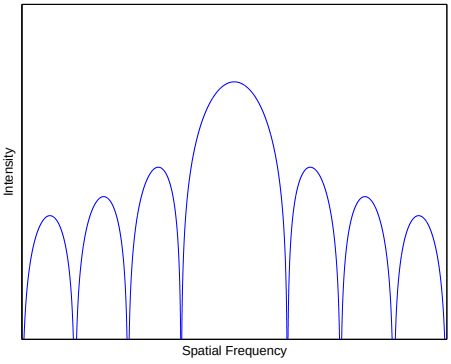

(f) Range-gating Spatial Frequencies

Figure 2.4: A comparison of the illumination waveforms and the spatial frequencies they implicitly sample for different lidar modulation techniques.

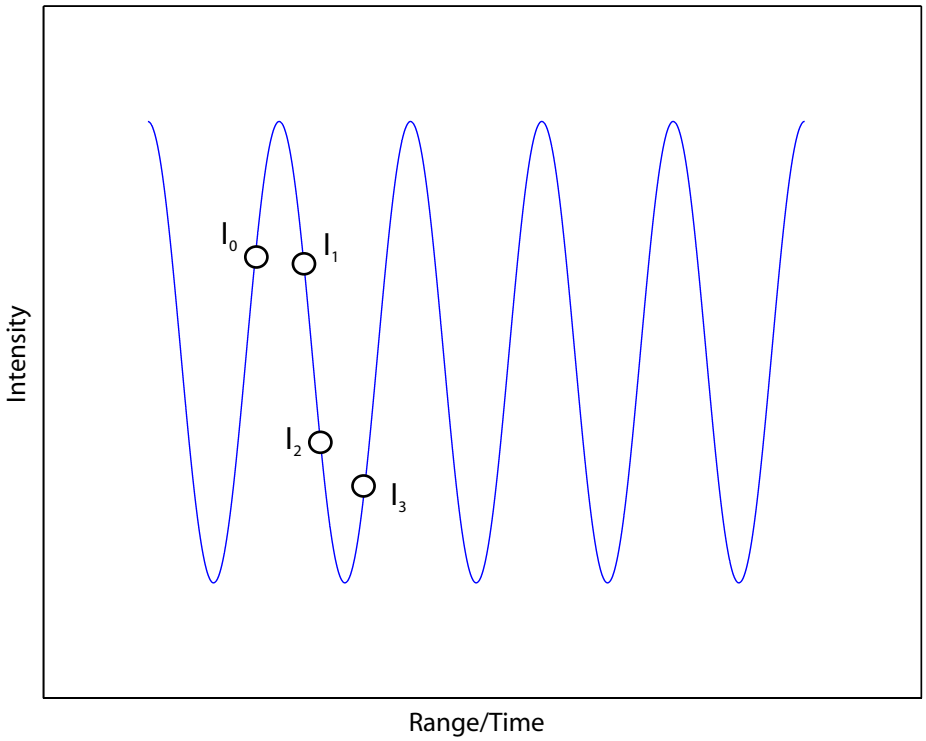

Figure 2.5: Sampling a sinusoid in order to determine amplitude and phase, thus the brightness and range corresponding to a backscattering source/target.

linear frequency modulated systems is a sinusoid where the frequency is a linear function of time (Jankiraman, 2007). There is much more inherent information in an FMCW lidar measurement: this can be demonstrated by comparing the spatial frequency content which the FMCW lidar illumination waveform implicitly samples of the scene, as shown in Fig. 2.4d. Whereas AMCW modulation is equivalent to sampling a single discrete spatial frequency of the scene, FMCW samples across a range of frequencies – as a consquence FMCW techniques require much higher data rates, nevertheless they have the ability to distinguish between multiple sources, thus do not suffer from multipath interference in the same way that AMCW does.

There are a number of different possible approaches to actually sensing the FMCW chirp, depending on the precise nature of the modulation; one common approach is to mix a measured illumination signal with a reference chirp. Another technique that is effectively equivalent to FMCW is to take an AMCW system and sweep across a frequency range, taking measurements at closely spaced regular intervals – also known as stepped AMCW. Some of the inverse methods presented in Chapters 5 and 6 could be considered to be a type of stepped AMCW, where the

number of measurements has been reduced to the absolute minimum mathematically required for a fully-determined solution.

The final modulation technique is range-gating, in which the illumination is modulated using a rectangular waveform with an extremely short duty cycle (Schilling *et al.*, 2002). Conceptually, this is the simplest approach to TOF: a pulse of light is sent out, and the time for the pulse to return to the sensor is measured. Like FMCW, this technique handles multipath, but requires much higher data rates.

2.2 Mixed Pixels and Multipath Interference

2.2.1 The Origins of Mixed Pixels

Mixed pixels occur in AMCW systems, but not in FMCW systems or range-gating systems, because AMCW systems only sample a single discrete spatial frequency (see Fig. 2.4). A simpler way to visualise this is to consider the signal intensity of the illumination as it reaches the range-sensor.

Fig. 2.6a is a plot of the signal returns within a single pixel; this corresponds to the backscattered illumination intensity as a function of range from the camera. We refer to this as the signal return model. The illumination signal intensity at the sensor is given by the convolution of the signal return model by the illumination modulation waveform. Fig. 2.6b shows the illumination signal in the case of a range-gating system. Because the illumination waveform is a short rectangular pulse, each backscattering source within the signal return model is clearly visible; as a result it is a relatively trivial matter to determine the range and amplitude of the backscattering sources. On the other hand, Fig. 2.6c shows the illumination signal intensity at the sensor in the case of AMCW modulation. The first component within the signal backscatters a sinusoidal signal, as does the second component, but with a different intensity and a range-induced phase shift. The illumination signal intensity at the sensor is given by the sum of the two sinusoids; because they are both at exactly the same frequency, the result is also a sinusoid. Fundamentally, there is no signal processing method capable of reconstructing the phase and amplitude of the two contributing sources from this illumination signal without extra information; it is an underconstrained problem with no inverse. This is the AMCW mixed pixel/multipath inteference problem.

An example of what the mixed pixel problem does to the measured phase of the AMCW modulated illumination in given in Fig. 2.10. This is a slice through a simulated full-field range-image before and after the application of a spatially variant blur, which simulated limited depth-of-field. Around the edges of objects,

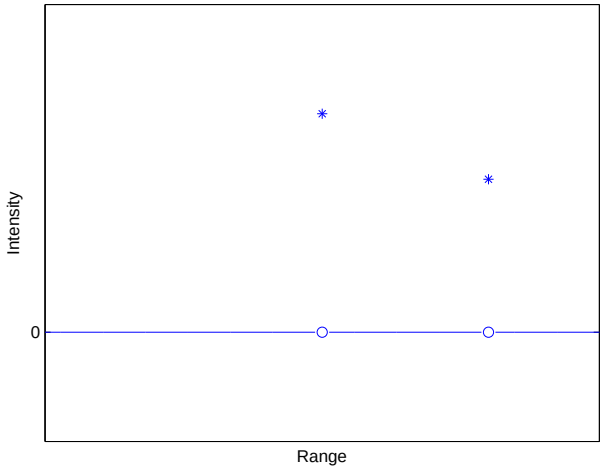

(a) Component Returns (Backscattering Function)

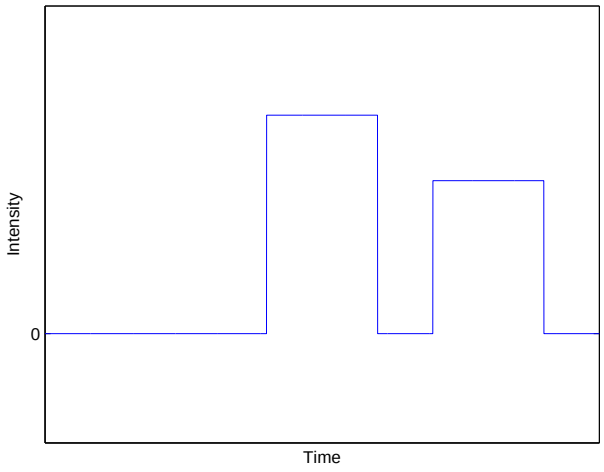

(b) Measured Illumination (Range-gating)

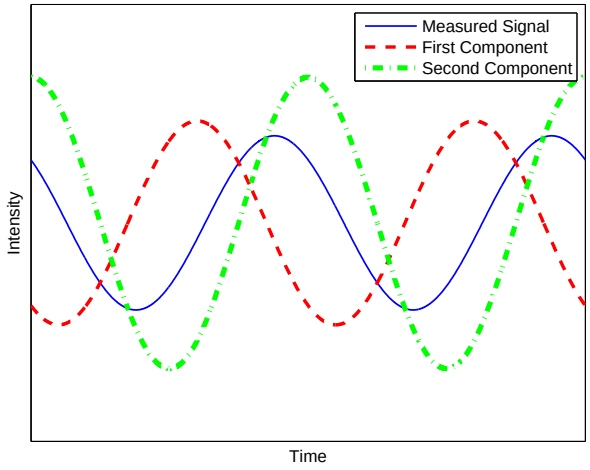

(c) Measured Illumination (AMCW)

Figure 2.6: The backscattered illumination waveforms for different modulation techniques in the two component case.

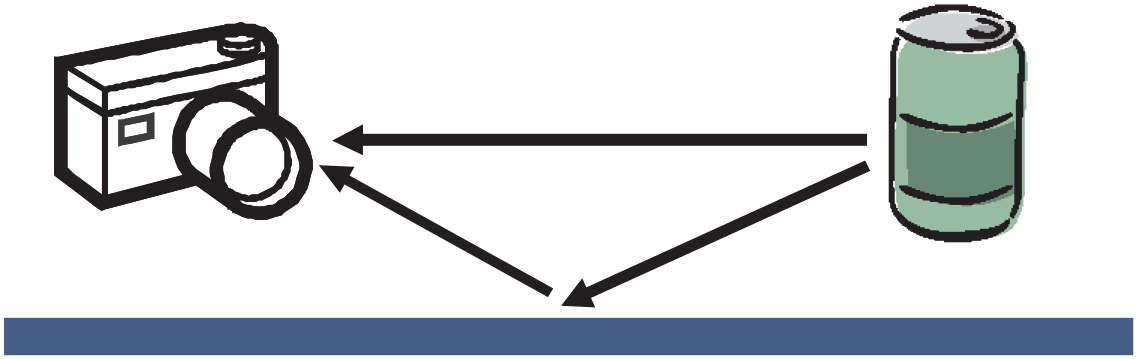

(a) Intra-scene multipath interference occurs due to reflections within the scene, resulting in mixing of light from multiple different sources

(b) Photo of scene subject to intra-scene reflections

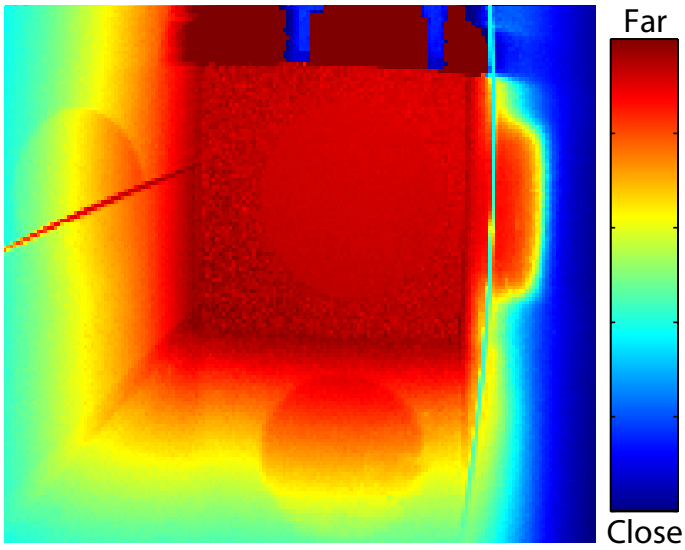

(c) Range data from scene; the reflection of the circle can be clearly seen in the floor and walls

Figure 2.7: The formation of intra-scene multipath interference.

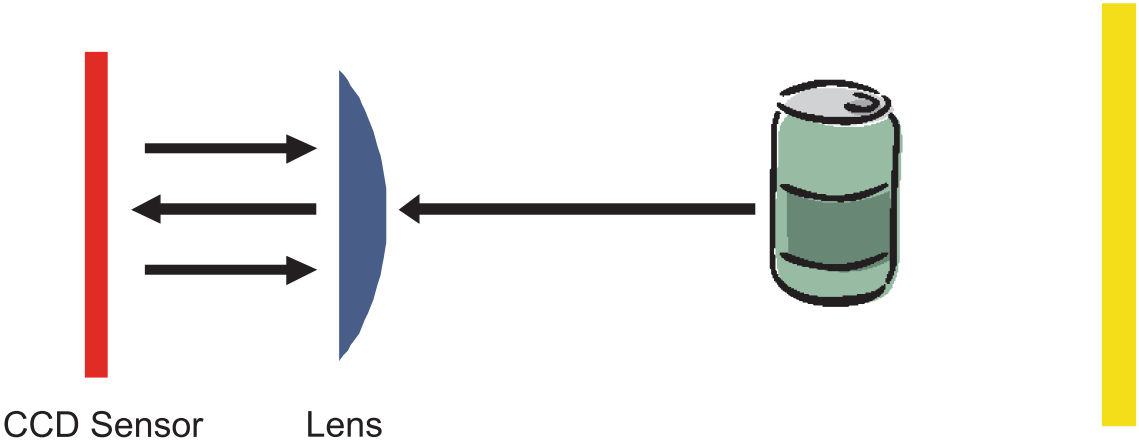

Figure 2.8: Formation process of intra-camera scattering induced multipath interference; light from the foreground target is scattered onto the pixels sensing the background object due to reflections off the sensor and lens assembly.

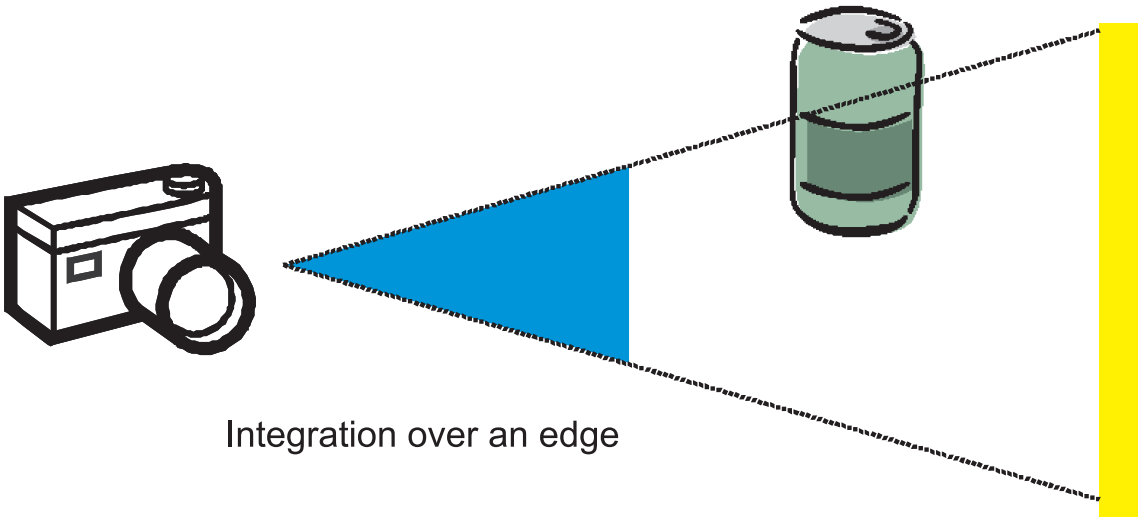

Figure 2.9: Formation process of mixed pixels; any pixel integrating over an edge mixes light from more than one object.

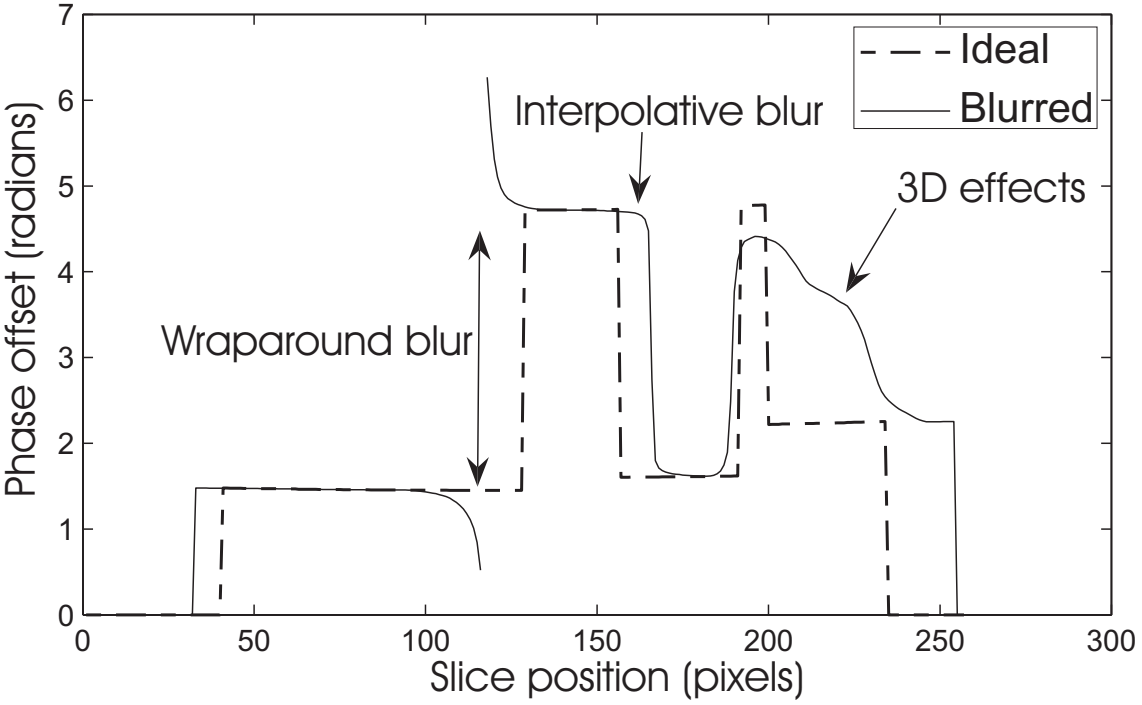

Figure 2.10: A simulated example of mixed pixels caused by defocus in an AMCW lidar system.

pixels contain range data from more than one source, resulting in either interpolative or wraparound blur, depending on whether the returns are within π radians of each other.

The mixed pixel problem is not specific to range imaging. It may occur in any imaging science, where a pixel or integration/sampling region contains data from more than one discrete source. For example, Chang *et al.* (1998) discusses the problem of decomposing mixed pixels in multispectral/hyperspectral images as linear combinations of discrete signatures. One possible application being determination of precise land use statistics from LandSat images, despite significant quantites of mixed pixels. The difference between the LandSat decomposition problem and the AMCW decomposition problem is that each component return within a mixed AMCW measurement is composed of continuously variable amplitude and range values. While it is possible to model AMCW measurements in terms of linear combinations of discrete states (as demonstrated in Chapter 4), there are other more direct/efficient approaches.

Multipath interference is a more general term than the mixed pixel problem and refers to the broader case including intrascene reflections. For example, when imaging the corner of a room it is common for light to be reflected several times, resulting in blurry range-data. Multipath interference is also caused by reflections within the optics of full-field range-imagers, resulting in crosstalk between measurements.

Kweon *et al.* (1991) deals with the experimental characterisation of a point scanning AMCW system and makes the first known (passing) reference to the AMCW lidar mixed pixel problem. A closely coupled paper from the same research group that references the same problem is Hebert and Krotkov (1992). Although point scanners do not integrate over large spatial regions in the same manner that full-field systems do, any point scanning system illuminates a particular solid angle at any time; this solid angle may integrate over multiple objects at non-identical ranges in the same manner as a pixel in a full-field system. Both papers discuss other issues common to full-field systems such as temporal drift.

Adams (1993) develops a method for detecting the range discontinuities that cause mixed pixels in scanning systems. Adams and Probert (1996) embelishes this method and also includes a discussion of measurement noise. They also suggests median filtering as a simple mitigation method, although it can 'fail catastrophically'.[2]

There are also a small number of more recent papers on mixed pixels in point

[2]There are some sources referenced in this paper, in particular Adam's PhD Thesis (dated 1992) and other papers by Hebert (dated 1991), which have been difficult to source. We consider the effort required to procure them to be disproportionate to their probable value.

scanning lidar systems. Tuley *et al.* (2004, 2005) emperically characterise the perturbations induced by mixed pixels and evaluate several filters for mixed pixel removal. Tang *et al.* (2007) also evaluate some simple discontinuity detection algorithms for mixed pixel removal. Ultimately, while mixed pixel detection and removal algorithms might be adequate for point scanning systems, where only a limited number of points are perturbed, because almost all measurements in a full-field system are perturbed to some extent, mixed pixel/multipath interference detection and removal in the same vein might entail throwing away every range-image in its entirity. As a result, it is necessary to develop an approach that restores perturbed points, rather than merely detecting them. Tang *et al.* (2009) takes a more constructive approach in processing point scanner data, by attempting to correct for the damage to measurements of height, width and other properties by estimating the edge loss based on geometric principles.

2.2.2 Mixed Pixels and Multipath Interference in Full-field Systems

There have been a number of papers published on the multipath interference problem in full-field AMCW lidar: many of these papers have been written without knowledge of the previous literature on mixed pixels in point scanning AMCW systems and without awareness of previous research. As a result, some of these papers have inadequately characterised the nature of the systematic perturbations induced by multiple backscattering sources.

Falie and Buzuloiu (2007) gives some initial work on noise, and recognition of a systematic error at low intensities caused by other objects in the scene. Falie (2008) identifies one source of interference being due to internal reflections within the camera and posits the use of structured light in order to determine the true range, albeit, they did not implement such a system. This is similar in concept to our method given in appendix C, using patterning/texture already present in the scene to determine true range. Falie and Buzuloiu (2008a) and Falie (2009) used calibration squares with known reflectivities on objects in the scene to remove the impact of any relatively homogeneous scattering.

Mure-Dubois and Hügli (2007a,b) modelled the scattering process as a 2D Gaussian PSF. Utilising the separable nature of the Gaussian, an inverse filter was developed to partially mitigate the impact of scattered light for a SwissRanger SR3100. Because the inverse filter was implemented in the Fourier domain, it was suitable for real-time implementation, with moderate success.

Karel *et al.* (2010) gave a limited emperical characterisation of impact of intra-camera scattering, but did not develop any scattering model or restoration method.

Guomundsson *et al.* (2007) identified the multipath problem in the case of looking at the corners of a room, but did not address restoration or recovery of correct range information. Oprisescu *et al.* (2007) identified an intensity related error and attempted a fixed calibration using the erroneous assumption that the calibration is scene independent.

Kavli *et al.* (2008) developed a restoration method for removing the impact of scattering on an image using the assumption of a spatially variant scattering function and a simple iterative method. The proposed method was demonstrated on a SwissRanger SR3000, sampling the scattering PSF using retroreflective dots. This paper demonstrated that the SwissRanger SR3000 scattering PSF was highly anisoplantatic.

Jamtsho and Lichti (2010) emperically characterises intra-lens scattering and then applies two different compensation methods using SwissRanger imagers (SR3000 and SR4000). Unfortunately, the paper gives a very incomplete description of how the compensation is achieved. Despite claiming an 80% improvement in range error in a test of the first method, because of ambiguity in the paper it is difficult to ascertain as to whether the fundamental assumptions behind the restoration are correct. In particular, the emphasis on the analysis of amplitude and phase perturbations raises questions about whether the compensation was implemented in the complex domain, as would be required for a physically correct compensation. The authors also note that the SR4000 produces less scattering distortion than the SR3000; this could potentially be due to the utilisation of better anti-reflection coatings in the optics.

Fuchs (2010) develops the most advanced model for multipath effects yet; the author models multipath reflections within the FOV by assuming that each surface in the FOV is a Lambertian reflector and scatters light onto all the other objects in the scene. Using the perturbed range and amplitude measurements it is possible to estimate the multipath perturbation for each pixel and subtract it from the range-image. There are several significant flaws with this approach, including the use of perturbed range data, which may result in erroneous estimates of the sum reflected light. Even using an iterative approach, where the corrected range-data is used to reestimate the scattered light, convergence is not necessarily guaranteed, especially for highly perturbed points. Other issues include specular reflection, reflections and scattered light from objects outside the camera FOV and execution time (approximately ten minutes per image for their particular implementation).

There is no discussion of occlusion in the paper, so accurate restorations are most probably limited to simple scene configurations like the internal corners of rooms; a more advanced approach would probably require ray tracing. Despite these issues, it is an interesting and original approach.

Given that each AMCW range measurement is equivalent to sampling a particular spatial frequency of the signal returns, by taking large sequences of measurements at different modulation frequencies it is possible to recover a model for the signal returns. Simpson *et al.* (2005) achived this by taking a sequence of 20 measurements between 10MHz and 200MHz in an harmonic sequence; using an inverse Fourier transform it was possible to recover an extremely low resolution model. However, this stepped AMCW technique is highly limited due to the extremely large number of measurements and the modulation frequency bandwidth required. However, the approach suggests that taking multiple measurements at different modulation frequencies may be key to separating out different backscattering sources.

2.3 Systematic Errors and Calibration

There are many sources of systematic errors in full-field AMCW lidar systems and calibration is inextricably a part of reducing these errors. There are many papers in the literature on correcting systematic errors, often mixing discussions of perspective deprojection, aliasing, bias and flat-field correction. Other related issues include motion correction and phase unwrapping, all of which utilise post-processing involving significant computational complexity.

While many papers discuss aliasing calibration there appears to have been limited exploration of the mathematical properties of aliasing. The literature focusses almost entirely on emperical calibartion and removal of the perturbation, rather than understanding its impact fully. This appears to be a flaw common to a number of papers in the literature, which identify apparent systematic errors and attempt to calibrate them out without fully examining the possible causes of the perturbation.

2.3.1 General Errors Sources and Calibration

Sample formation and the basic characteristics of correlation waveform aliasing in homodyne modulated AMCW systems is addressed by Lange (2000) and Luan (2001) developing basic equations for each. In order to determine the phase of the correlation waveform, both rely on the equations derived by Xu *et al.* (1995) for the phase and amplitude of harmonics in the stationary Gaussian noise case.

The integration time dependent error mentioned by Foix *et al.* (2011) has also been analysed by Kahlmann *et al.* (2006) and Lindner and Kolb (2007). Kahlmann *et al.* found that the linearity error changes with integration time. While neither paper posits a credible cause, one possibility is that the sensor response waveform changes with different integration times, with a possible causal mechanism being changes in sensor temperature. Linder and Kolb also found a related coupling of amplitude with distance errors, however it is unclear from the paper whether there may have been confounding multipath interference.

Karel (2008) developed a calibration method for radial distortion using a sequence of range measurements of a known scene using a hand-held camera. One particularly interesting factor noted was a temporal periodicity in range and amplitude measurements; although this was a very small variation, it could potentially be related to the manner in which the modulation signals are generated in the Swiss-Ranger SR-3000.

Lindner *et al.* (2010) calibrates for radial distortion, perspective projection, aliasing non-linearity and reflectivity induced errors. The paper develops a model based on a fourth order model of radial distortion and B-Spline based correction for aliasing range non-linearity. This paper builds off Lindner and Kolb (2006), which implemented a version of the aliasing range non-linearity correction and Lindner and Kolb (2007), who developed the concept of 'intensity-related distance error'. One issue with the paper is that the evidence of 'intensity-related distance error' is entirely emperical. Although the existence of changes in range-response due to changes in integration period are plausible, neither Lindner and Kolb (2007) nor Lindner *et al.* (2010) posit mechanisms for the generation of systematic errors as a function of reflectivity – for example, due to scattered light. Because the journal article does not specifically acknowledge aliasing induced intensity scaling, it also possible that the authors may be conflating measured range as a function of reflectivity with measured range and intensity as a function of range (essentially, mistaking correlation for causation). Another paper by the same author, Lindner *et al.* (2008), discusses the same B-Spline method but also develops a range measurement method specific to perfectly triangular correlation waveforms.

Lottner *et al.* (2008) discussing the difficult problem of developing a calibration method to handle the use of multiple illumination sources for a single scene. Steiger *et al.* (2008) discusses calibration, including addressing calibration for temperature changes, which in general is rarely addressed by papers on calibration. Kang *et al.* (2011) discusses a range calculation method specific to 50% duty cycle square wave illumination modulation waveform and sensor response waveform.

Lichti (2008) developed a method for calibration of a 3D range camera, including aliasing induced non-linearity, lens distortions and delays across the sensor using a sequence of over thirty images of a specially designed reference scene.

Rapp *et al.* (2008) presented a basic analysis of error in AMCW lidar systems, including a simple emperical analysis of precision and accuracy for a PMD sensor. The paper also discusses how differential measurement systems partially cancel out bias and gain variations between channels within a pixel. Fuchs and Hirzinger (2008) present another description of systematic errors in AMCW range imaging, specifically mentioning aliasing effects on amplitude and range and signal latency across the sensor, in application to range measurements involving an industrial robot. Keller and Kolb (2009) and Schmidt and Jähne (2009) both developed low-level models of the behaviour of an AMCW lidar sensor, based around the PMD design.

Hegde and Ye (2008), while not developing a noise model, developed a method for filtering range-image surface normals to remove random errors using singular value decomposition. This is in the context of edge detection and segmentation, where the high level of noise in estimates of surface normals can result in erroneous edge detection. An alternative approach to filtering the surface normals would be to detect the presence of mixed pixels at the edges of objects. We propose reliable methods of achiving this in Chapter 6. Detecting mixed pixels could conceivably significantly reduce the impact of random noise while only requiring single measurements at each of two different measurement frequencies.

Hussmann and Edeler (2009) simplified the standard four phase step measurement algorithm. In its simplest form, this involves taking measurements at only zero and ninety degree phase offsets, therefore the phase measurement problem reduces to $\phi = \text{atan2}(I_1, I_0)$. While the paper demonstrated the algorithm using real data from a PMD camera, the paper was rather unconvincing as to whether the resultant data was accurate, as it did not discuss possible systematic errors due to pixel bias and other inhomogeneities. While the use of only two phase steps may halve the time required to produce a range-image, the systematic error is likely to be difficult to calibrate out, as bias is typically a function of sensor temperature. A subsequent paper, Schmidt *et al.* (2011), developed a method for dynamic determination of bias and gain correction coefficients. This more recent research potentially enables approaches such as Hussman's to be implemented without bias and gain variation induced systematic errors.

MacKinnon *et al.* (2008) investigated how to measure the quality of a range image, including the impact of factors such as spatial resolution, Gaussian noise, non-Lambertian reflection, ambient light and volumetric scattering (diffuse-range

sources). Foix *et al.* (2011) provide an up-to-date review of full-field AMCW systems, including noise sources. Foix notates correlation waveform aliasing as 'wiggling' or 'circular error', an unexplained integration time related systematic error, saturation, temperature drift and multipath interference. In particular, Foix *et al.* mentioned multipath errors induced by concavities, such as corners, which are more difficult to model due to the returns within a pixel being range-diffuse.

2.3.2 Motion

Because most commercial full-field systems use four sequential captures at ninety degree phase offsets, motion in one of the frames invalidates the relationship of the frames before and after the motion (or during, for a continuous movement). While we do not directly address the problem of motion in this thesis, it is an important source of error in range measurements and combined with data rate limitations it is a major constraining factor on the availability of data to solve the mixed pixel separation problem. If it were possible to capture all the phase step measurements simultaneously, possibly by either temporal multiplexing, then the lateral motion problem reduces to the mixed pixel problem. Thus, given potential hardware development, solving the mixed pixel problem may be equivalent to solving the non-axial motion problem.

Lottner *et al.* (2007) investigated motion artefacts in range images; he noted several viable approaches to the detection of motion artefacts including comparison to previous images in the case of a sequence and the utilisation of edge detection algorithms. The paper also identified occlusion of the light source as a major contributing factor to motion artefacts.

Lindner and Kolb (2009) developed an optical-flow based correction method for both axial and lateral motion in a scene. In general most AMCW range-imagers use differential measurements, measuring $A - B$, where A and B correspond to two different charge collection regions within a pixel, with the modulation waveform being inverted for the second. The most fundamental operating assumption in optical flow is conservation of intensity; in general for a range image, $A - B$ is not conserved. This is especially true given a sequence of measurements at different phase steps. In order to address this, Lindner instead implemented an optic flow algorithm that utilised common mode measurements, $A + B$ – which correspond to the total integrated intensity plus offsets from ambient light and pixel bias. Common mode intensity is largely conserved, even between measurements at different phase steps. Once flow had been determined, Lindner implemented an unexplicated resampling step to correct the range measurements. The axial correction was implemented as

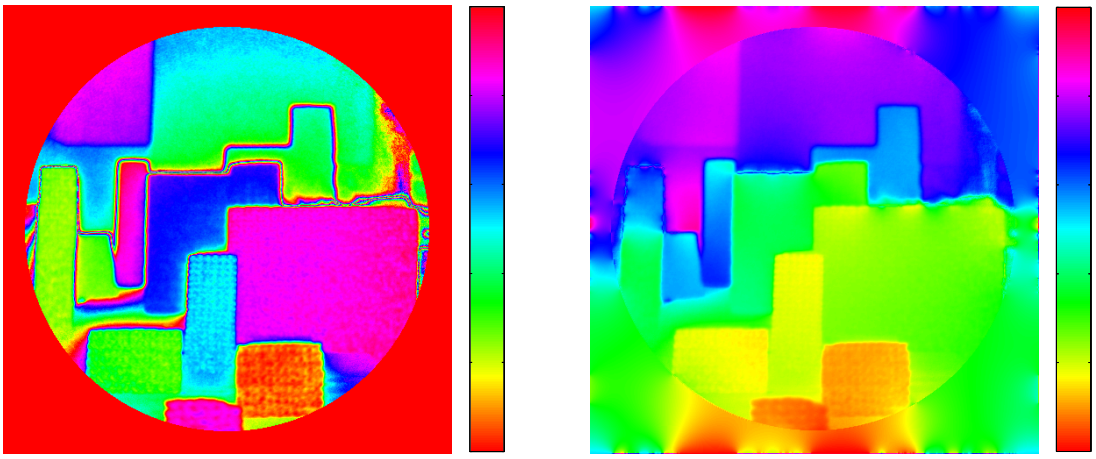

(a) Wrapped (b) Unwrapped

Figure 2.11: Example phase unwrapping of a range-image using our own implementation of unweighted phase unwrapping in the Fourier domain using Poisson's equation.

a separate second processing step.

One approach to determination of pixel subject to motion artefacts which does not appear to have been discussed in the literature is use of the second harmonic. Most AMCW systems use four phase steps; from a Fourier perspective this gives four pieces of information: the mean intensity, a phasor encoding range and amplitude and a scalar coefficient representing the amplitude of the second harmonic. From an information content perspective, for a single isolated range measurement, the only evidence of whether the measurement is subject to motion artefacts or not is encoded in the zeroth and second harmonics of the sampled correlation waveform. In practice, the ratio of the intensity of the second harmonic to the modulus of the first harmonic (amplitude of the range measurement) is a highly useful indicator of lateral motion.

2.3.3 Phase Wrapping

Phase unwrapping is sometimes referred to as range dealiasing or described as a range extension method. It is common in both one and two dimensional cases, the latter particularly useful in synthetic aperture radar. Whether it should be regarded as a systematic error is debatable.

There is a significant body of work on the use of spatial constraints to produce phase unwrapped images. These methods typically operate by calculating the spatial derivatives of phase and then re-integrating them without any wrapping at 2π. Often scenes cannot be trivially integrated due to residues; that is, infinite loops in the

spatial constraints akin to those found when taking the contour integral of the complex logarithm around zero. Dealing with these residues introduces a significant amount of additional complexity. For example, Goldstein *et al.* (1988) present an algorithm that detects residues and determines branch cuts, so as to attempt to minimise the size of the discontinuous regions. Other methods include framing the problem using Poisson's eqution, such as in Ghiglia and Romero (1994). It is possible to solve the problem using inversion in the Fourier domain, which is essentially a deconvolution operation. An example phase unwrapping using our own implementation is given in Fig. 2.11.

There are substantial issues with spatial approaches, including such fundamental problems as the ambiguous nature of the unwrapping; it may be possible to say that one object in in front of another, but it is comparatively easy to underestimate the phase shift between two objects by a multiple of 2π. Given the comparative ease of taking images at different modulation frequencies, multiple frequency measurements have become the favoured approach in full-field lidar. A common technique for this is the synthetic wavelength technique, described by Dändliker *et al.* (1998), often utilised in interferometry. For measurements at wavelengths λ_0 and λ_1, a measurement can be synthesised as if directly made at a wavelength of $|\lambda_0 - \lambda_1|$.

Falie and Buzuloiu (2008b) give a very brief discussion of phase unwrapping of images using dual measurement frequencies of 19 MHz and 21 MHz. The method solves a system of pseudo-Diophantine equations in order to determine the unwrapped phase[3]. Although the authors do not specify the exact solution method, it is probably brute-force enumerative. Jongenelen *et al.* (2010) discuss ambiguity interval extension by taking two measurements simultaneously, a similar approach is discussed by Payne *et al.* (2010).

Bamji (2010) is a Canesta patent for a lossless method of phase unwrapping using two synthetic wavelengths; while some of the claims might be considered prior art, the main concept is interesting, if simple. Whereas the standard synthetic wavelength technique uses a single synthetic wavelength at a longer wavelength than either of the physical measurements, this technique also uses a synthetic higher frequency, which results in a greater phase precision. The final range measurement is produced by phase unwrapping the high frequency synthetic phase using the information contained in the lower frequency synthetic phase. However, the method is certainly not lossless, as claimed in the patent title.[4] Consider the phases of the two measurements $\theta_0 \sim \mathcal{N}(x, \sigma_0^2)$ and $\theta_1 \sim \mathcal{N}(y, \sigma_1^2)$ and the phases of the synthetic

[3]Diophantine equations are polynomials with integer variables

[4]Patent Title: Method and System for Lossless Dealiasing in Time-of-Flight (TOF) Systems.

measurements $\phi_0 = \theta_0 - \theta_1$ and $\phi_1 = \theta_0 + \theta_1$, then the covariance of the tuple (ϕ_0, ϕ_1) is

$$\Sigma = \sigma_0^2 \begin{pmatrix} 1 \\ -1 \end{pmatrix} \cdot \begin{pmatrix} 1 & -1 \end{pmatrix} + \sigma_1^2 \begin{pmatrix} 1 \\ 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 1 \end{pmatrix}. \tag{2.5}$$

For equal phase precision in the physical measurements, $\sigma_0^2 = \sigma_1^2$, the covariance matrix is diagonal. This means that simply unwrapping the higher frequency synthetic phase (ϕ_1) using the lower frequency synthetic phase (ϕ_0) is not noise optimal; there is additional remaining statistically independent information available. A noise optimal measurement[5] can be achieved by taking a noise weighted average of the range estimates produced from each of the unwrapped synthetic phases. Admittedly, the weighting is significantly asymmetric as the range estimate from the lower synthetic phase is much more noisy, but it is not a completely insignificant improvement.

Frey *et al.* (2001) and Droeschel *et al.* (2010) use a more complicated approach to phase unwrapping: loopy belief propagation on what is essentially a Markov Random Field, although they do not explicitly describe it as such. The MRF is composed of an undirected graph of jump and curl nodes, that model the irrotational and rotational components of the data respectively. Constraining curl to zero removes residues and allows determination of a plausible model for the unwrapped phase.

2.4 Experimental Full-field Lidar Systems

While there are a wide variety of full-field lidar systems, modern commercial systems have converged on integrated CMOS sensors as the most cost effective approach. Other systems typically use separate modulated optical mixers such as Pockels cells (Xu *et al.*, 1995), or image intensifiers (Christie *et al.*, 1995), followed by a standard imaging system, like a CCD camera. In this section we introduce and discuss the range-imaging systems utilised experimentally in this thesis.

The first system is the University of Waikato Range-Imager, which is a custom system based around an image-intensifier optically coupled to a charge coupled device (CCD) camera. The other two systems are pre-existing commercial systems: the Mesa Imaging SwissRanger SR4000, which is available off-the-shelf, and the second the Canesta XZ-422 Demonstrator Camera, which is not. Whereas the UoW system uses visible laser illumination, the other systems use LED illumination systems, operating in the near infrared above 800nm. The systems have varying resolutions:

[5]That is: noise optimal if one makes the assumption that the error in the measured phase can be described by a Gaussian. This assumption is discussed in Chapter 3.

the UoW system has a resolution of 512 × 512 pixels, the SwissRanger 4000 has a resolution of 175 × 144 and the Canesta system has a resolution of 160 × 120.

Apart from Mesa Imaging and Canesta, Photonic Mixer Devices (PMD or PMDTechnologies) are another company that produce full-field AMCW lidar systems. They currently offer the highest resolution off-the-shelf imaging device at 204 × 204 pixels (PhotonICs 41X-S). Several papers have been published on their chip design including papers by Xu *et al.* (1998) and Schwarte *et al.* (2004). PMD not only sell end-to-end systems, but also individual sensor chips/bring-out boards; this allows research groups to develop their own driving circuitry for the purposes of testing more advanced modulation techniques, for example, by our own research group (Jongenelen *et al.*, 2009; Whyte *et al.*, 2010; Conroy *et al.*, 2011). Whyte *et al.* (2010) developed a pseudo-random encoding technique to ameliorate inference between multiple simultaneously operating cameras, and Conroy *et al.* (2011) developed a new resonance based power saving modulation technique.

Other research using PMD sensors includes papers by Plaue (2006), who analysed the impact of different modulation signals on range linearity, also noting the existance of a slight range-intensity coupling effect. Radmer *et al.* (2008) produced an additional analysis of offsets induced by amplitude and changes in the integration period. Use of the term PMD is sometimes confused due to general usage of the term in a general sense for modulated CMOS sensors as well as systems made specifically by PMDTechnologies GmbH.

2.4.1 University of Waikato Range Imager (Heterodyne Image Intensifier System)

The University of Waikato Range-Imager is a custom system based around a modulated image-intensifier. A charge coupled device (CCD) camera is optically coupled to the rear of the image intensifier and scenes are illuminated using divergant laser illumination in the visible domain. While I did not contribute to the hardware, I developed the software system required for controlling and recording measurements prior to my PhD candidacy.[6] Payne (2008) gives a detailed account of the hardware development of the range-imager system. Other published work includes (Dorrington *et al.*, 2007; Payne *et al.*, 2005; Cree *et al.*, 2006; Payne *et al.*, 2006a,b, 2008a). This system is the primary experimental device utilised herein.

Fig. 2.12a shows the entire range-imager system in operation during testing, in this particular case using only a single illumination source. The lasers are mounted

[6]Including associated visualisation tools, the overall software system was approxmately 30,000 lines of code.

on the PCB in the front left, and the illumination is transmitted via optical fibres to divergant lenses mounted around the primary optics. Immediately behind the primary optics is the image intensifier, which has a silver electronics box mounted directly below. Image intensifiers are high voltage devices, in this particular case a power supply capable of generating 5000 volts is required, although it is only necessary to modulate an 80 volt signal in order to achieve shuttering of the intensifier. The necessary power supplies are contained in the low profile boxes at the rear left of the assembly. At the very back is the CCD camera, a Pantera TF-1M60, which is mounted with heatsinks in order to achieve best possible noise performance. The CCD camera is optically coupled to the back of the image intensifier; this configuration is shown in Fig. 2.13. Fig. 2.12b indicates the typical coaxial illumination method. While each laser diode is rated to 80mW, the total illumination intensity is typically less than 100mW.

Fig. 2.14 is an overview of the most important modulation and control signal flows. The direct digital synthesiser (DDS) is a system for generating analogue clock signals; in particular, a DDS is an extremely useful approach to generating high frequency signals that need to be synchronised. At every clock cycle an accumulator is incremented by a fixed 'frequency tuning word,' which determines the frequency of the output. An analogue output signal is generated by using the accumulator value as the index to a lookup table modelling a sinusoid. The lookup table values specify the input value to an analogue-to-digital converter, that generates the output signal. For low frequency signals this enables extremely accurate synthesis of a pure sinusoid, however given that the clock frequency is 400MHz and the modulation signals generated are commonly around 80MHz, there is significant harmonic content generated. These harmonics are ameliorated by passing the signal through a low pass filter.

In the UoW range-imager, the DDS generates three different synchronised signals: one triggers CCD frame acquisition, one is used for illumination modulation and the other is used to modulate the image intensifier. The frame acquisition and image intensifier modulation signals are further processed by a field programmable gate array (FPGA), which gives extra flexibility. In order to vary the duty cycle of the (theoretically) rectangular illumination modulation, the filtered modulation signal is passed through a comparator, where the signal intensity is compared to a variable threshold voltage. The amplified modulation is then used to drive the laser diodes. The time of flight of the modulated illumination results in a phase shift in the illumination, which is measured by the remainder of the range-imager system.

The primary optics focus an image of the scene onto the photocathode of the

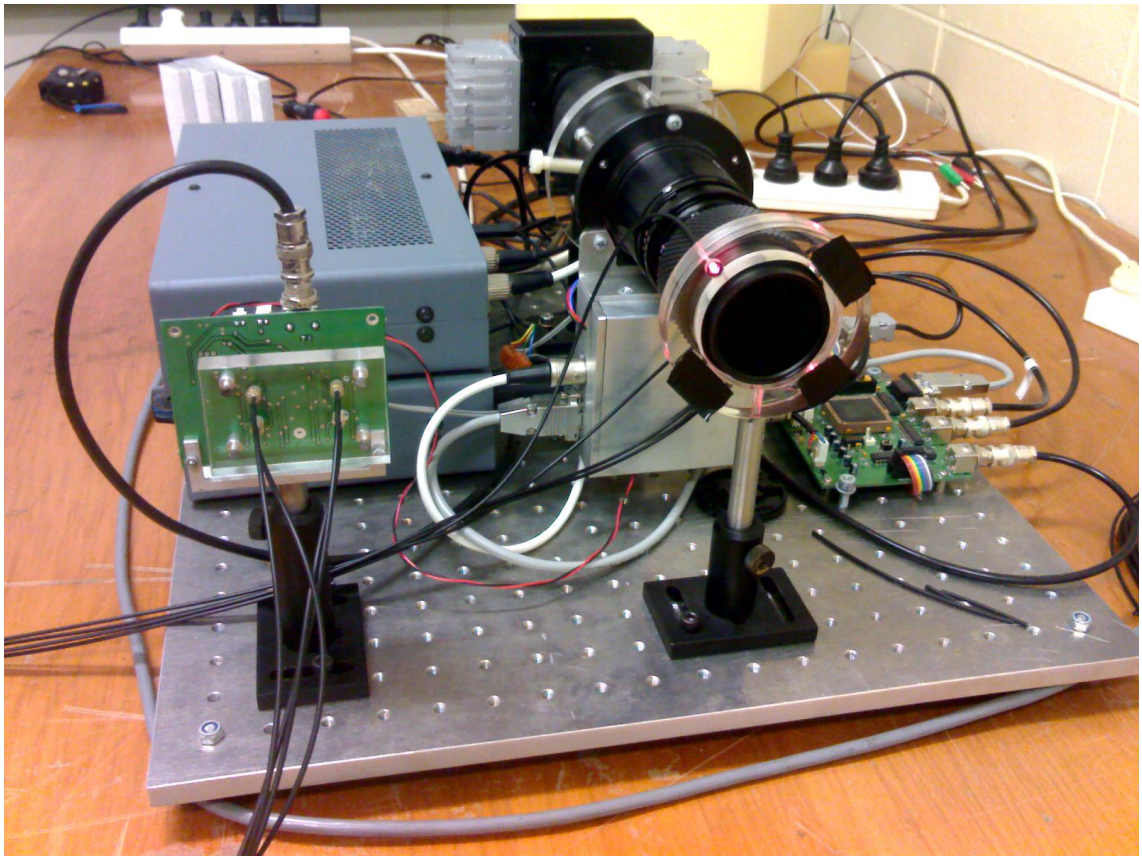

(a) Entire System

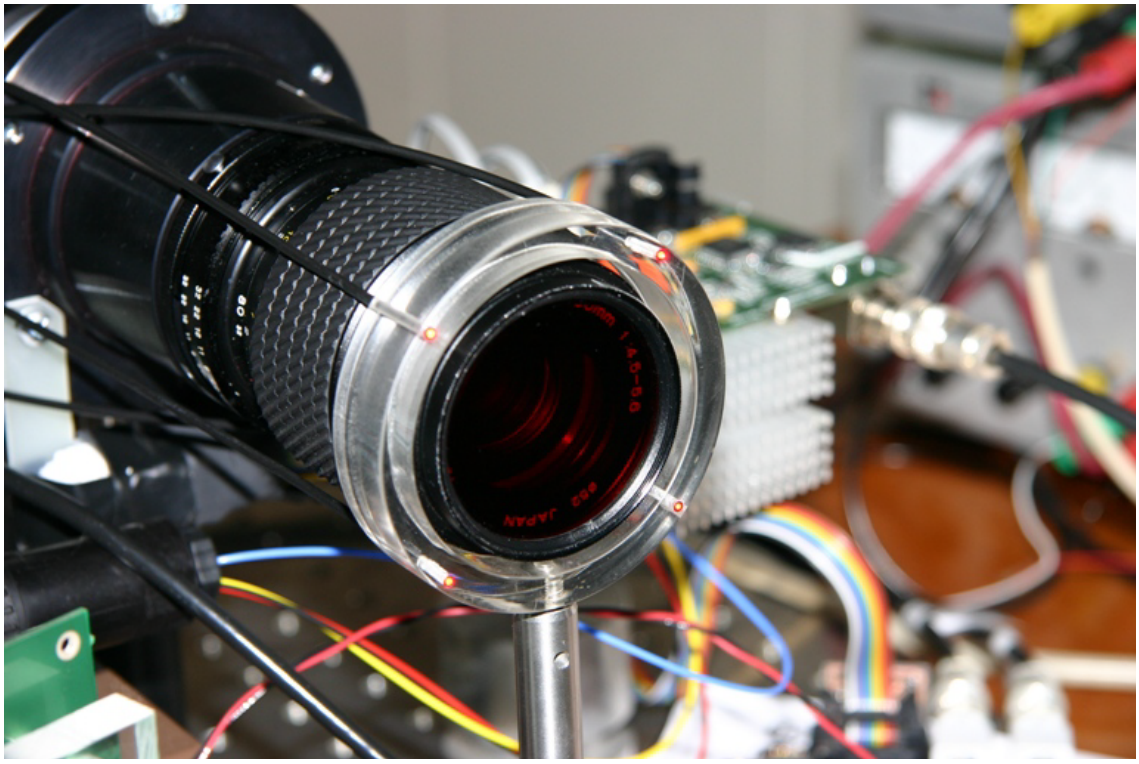

(b) Typical Coaxial Illumination Configuration*

Figure 2.12: The UoW range-imager system. * Copyright Chronoptics Group.

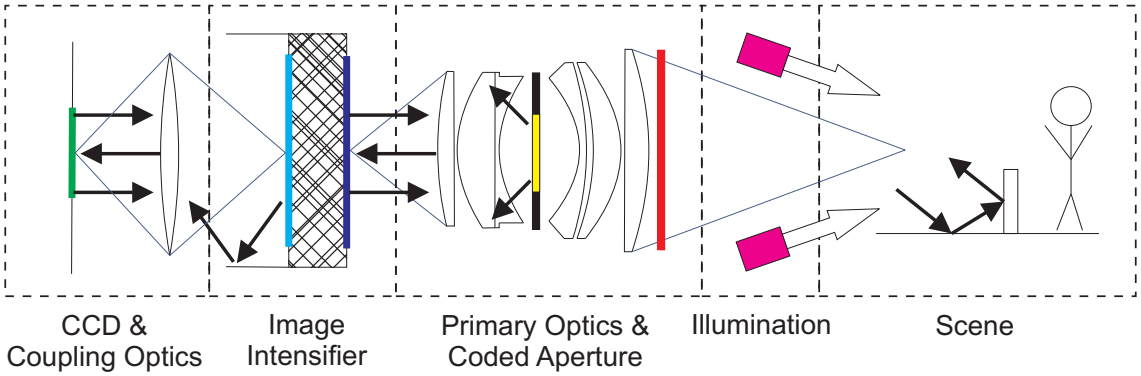

Figure 2.13: The optical configuration of the UoW range-imager when equipped with a coded aperture. Key: modulated lasers (magenta), narrowband filter (red), coded aperture (yellow), image intensifier photocathode (blue), phosphor screen (cyan), CCD image sensor (green). Black arrows represent sources of multipath.

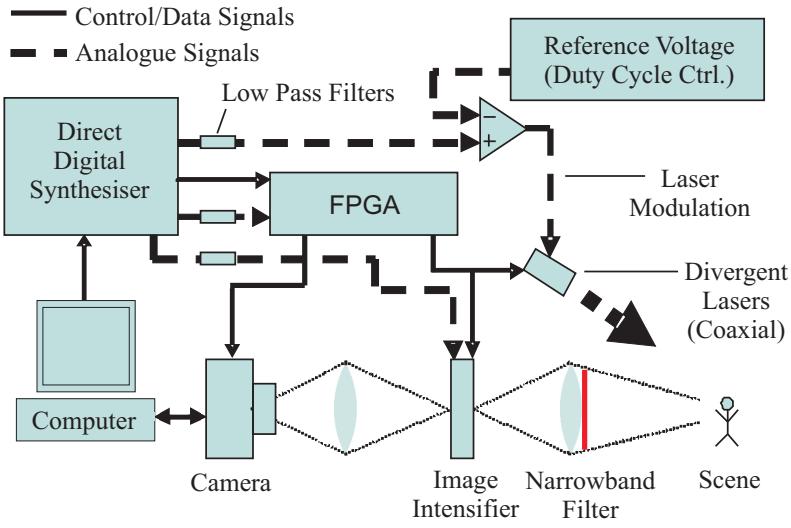

Figure 2.14: Overview of the UoW range-imager modulation and control signal flow.

image intensifier. An image intensifier is a device that is used to electro-optically amplify light signals; the most common application is night-vision goggles, which operate by amplifying visible and near-infrared light at night and in other low-light conditions. Photons hitting the photocathode liberate electrons via the photoelectric effect, these photons are directed using electric fields towards the microchannel plate, which as as an electron multiplier. By varying the electric potential across the microchannel plate, the image intensifier can be modulated to act as a shutter, rather than merely a fixed gain amplifier. A phosphor screen produces an image which is coupled to the CCD using an additional lens. While many intensified CCDs use fibre-optic coupling of the CCD to the intensifier, this particular system does not; as a result there are some minor issues with light scattering post-intensifier. A much larger concern is the extraordinarily high reflectivity of the photocathode;

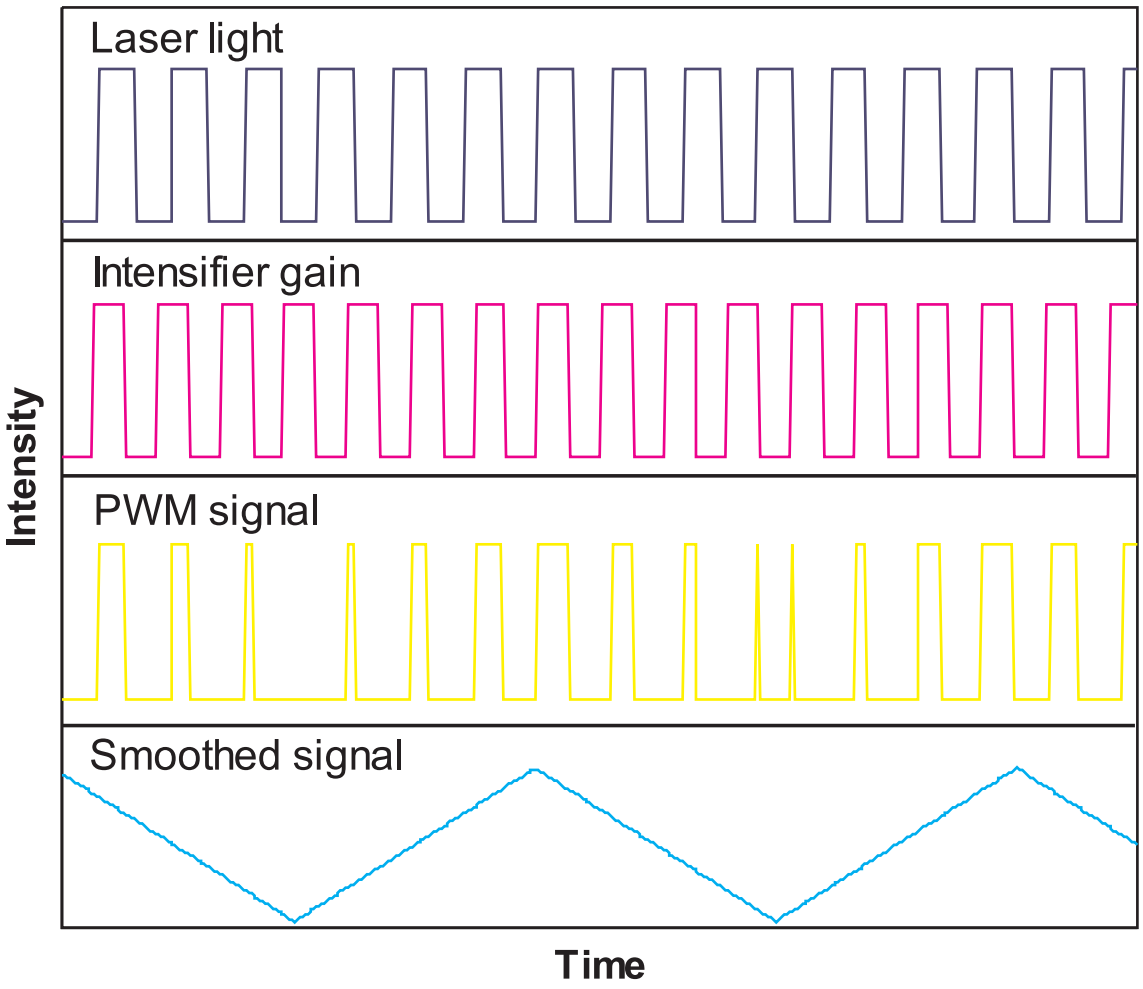

Figure 2.15: The generation of the beat waveform in a heterodyning AMCW system.

despite antireflection coatings on the primary optics, a significant amount of light is scattered. Appendix B discusses this issue in greater depth.

While the system is capable of either homodyne or heterodyne operation, we most commonly operate it in the latter. Homodyning involves modulating the image intensifier and illumination at the same modulation frequency, whereas heterodyning involves modulating them at different frequencies. Fig. 2.15 illustrates how heterodyning works. Two modulation signals are produced at near identical frequencies: for example, 80MHz and 80.0000001MHz. At the image intensifier, the laser illumination, incorperating a range dependent phase shift, is electro-optically mixed with the image intensifier modulation; this results in a pulse-width modulated signal, where the pulse width oscillates as the laser and intensifier modulation signals move in and out of phase. The pulse width modulated signal is low-pass filtered by the response of the phosphor screen, which results in a smoothed signal. For rectangular modulation of both the image intensifier and illumination, the resulting waveform is

a truncated triangle waveform. This 'beat waveform' has a frequency equivalent to the different between the two modulation frequencies; for the above example, that is 0.1Hz. Most importantly, the phase of the waveform is proportional to the phase offset of the illumination waveform relative to that of the image intensifier modulation; therefore, given an absolute reference one can determine the range to the imaged object. The 'beat waveform' shape is the correlation of the illumination modulation waveform with the image intensifier shutter response, hence is often referred to as the 'correlation waveform;' albeit there is some ambiguity of the use of the term due to the use in both homodyning and heterodyning contexts. The difference being that in the heterodyning context, each sample of the waveform is integrated over a range of phase shifts, whereas in the homodyne case each sample corresponds to a discrete phase shift; this subtly impacts on the shape of the correlation waveform, in particular the extent to which the system is impacted by aliasing due to frequencies above the Nyquist frequency.

In our particular system, the waveform is sampled temporally by the CCD camera at a fixed sample rate which is constrained to have some ratio of integers relationship to the beat frequency. In the simplest case, the complex domain range measurement for each pixel corresponds to the negative fundamental frequency of a Fourier transform of the correlation waveform versus time. This is developed in greater detail in Chapter 3; for now it is sufficient to know that it is possible to measure amplitude and range from the Fourier transform of the signal. Unlike more advanced CMOS sensors, which cancel out the effects of ambient light, the image intensifier system accumulates a significant offset from ambient light, which decreases SNR; this is partially mitigated using a narrow-band filter, but is difficult to resolve completely.

Images taken by the range-imager are calibrated for bias and flat-fielded. Fig. 2.16a shows a bias image for the CCD sensor. At certain modulation frequencies inteference is visible (Fig. 2.16b) due to coupling of the modulation with the readout clock on the camera; this can be avoided by choosing alternative modulation frequencies. One of the big issues with the use of an image intensifier is burn-in; that is, over time the intensity response of the system changes. This is illustrated in Fig. 2.16c, which is an image of a defocussed, uniformly illuminated planar object, utilised as a flat-field image. The circle corresponds to the optical output of the image intensifier, dark and light regions within this correspond to burn-in patterns. The white band at the top of the image is the result of an experiment where a board was imaged over a long period of time; the region at the top, above where the board was placed, has not been subject to the same level of degradation as the lower regions. There

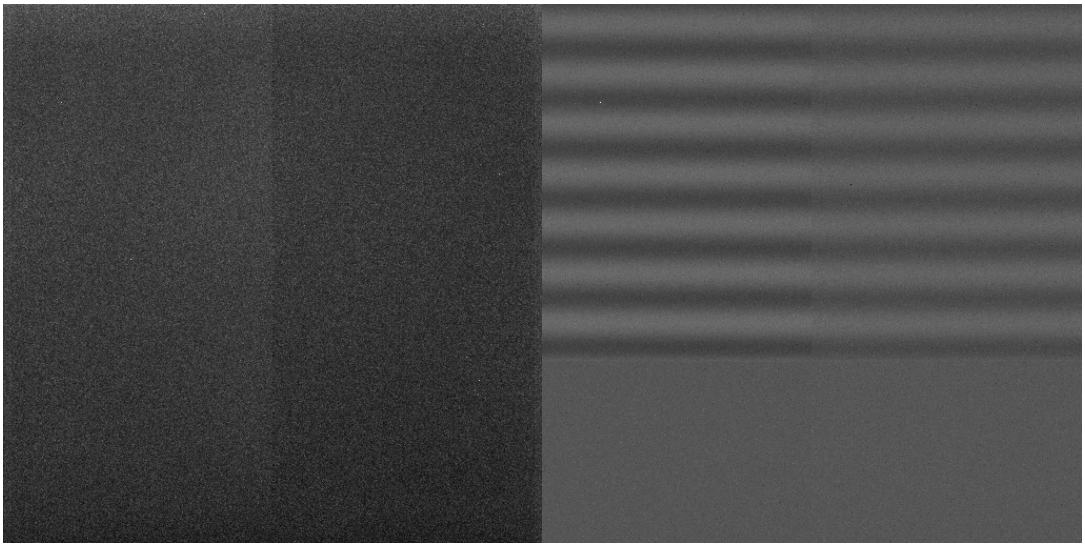

(a) Bias (b) Interference Pattern

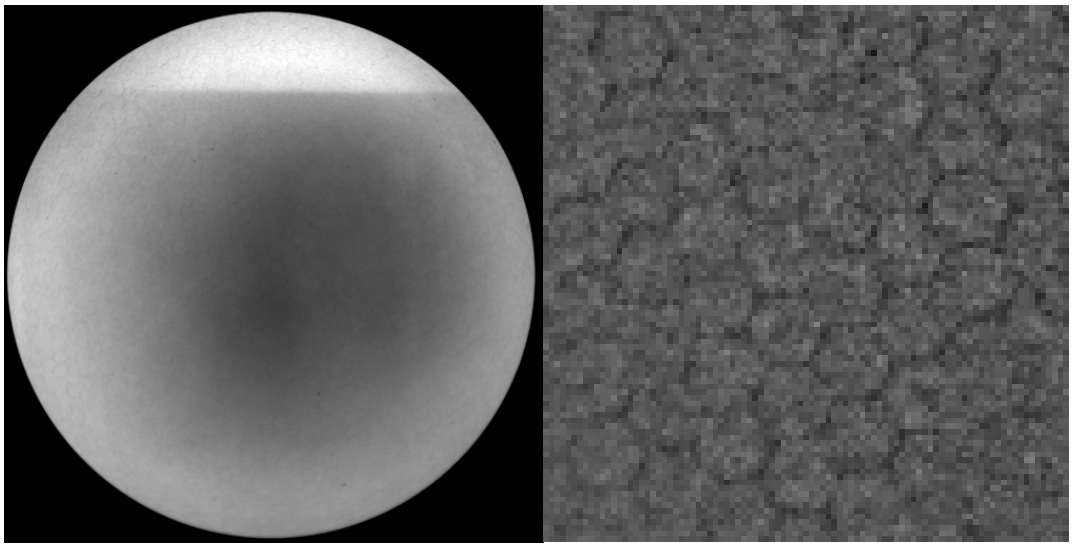

(c) Flat-field (d) MCP Patterning

Figure 2.16: Image intensifier system calibration images and behaviour.

is also a dark region in the centre, which is most probably due to a combination of increased illumination in the centre of the image and a greater concentration of objects placed close to the camera, rather than any particular singular experiment. The phenomenon of burn-in is a known characteristic of image intensifier systems and is merely one reason why manufacturers have transitioned to CMOS sensors. The final image is a contrast-enhanced subregion of the flat-field image. The hexagonal patterning is due to the method of construction of the MCP in the image intensifier. Bundles of glass rods are assembled into hexagonal assemblies and drawn to lengthen and shrink them. These bundles are stacked to form the MCP;

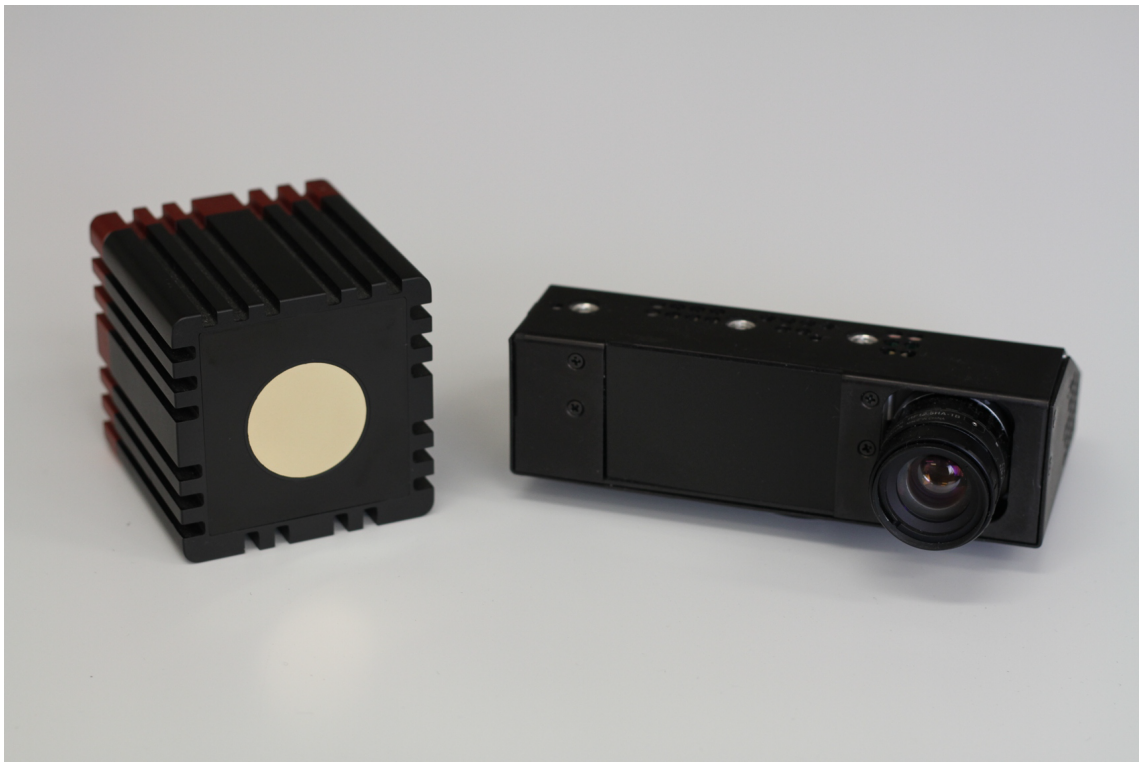

Figure 2.17: The commercial ranging systems: on the left the SwissRanger SR4000 and on the right, the Canesta XZ-422 demonstrator. (Copyright Andrew Payne and Robert Larkins)

there are slight irregularities and dark regions around the boundaries of the bundles, which results in these artefacts.

One of the most frustrating aspects of dealing with this system is the poor coupling between the CCD and the image intensifier; bumping the table can knock out the relative alignment, which requires another flat-field image to be taken. In retrospect, this is one of the biggest faults with the paper presented in appendix A; some of the images were calibrated with flat-fields that had been accidentally invalidated, making defocus restoration a foolhardy exercise. Another issue with this system is buggy firmware in the camera, that results in low intensity regions of images being truncated at certain frame rates and resolution settings and occasional dropped frames.

2.4.2 Canesta XZ-422 Demonstrator (Integrated CMOS Sensor)

The Canesta XZ-422 is a four-sample homodyne system based around a modulated CMOS sensor and infrared LED illumination. Although holding over 37 patents on

range-imaging technologies, there are few academic papers on Canesta cameras.[7] Hsu *et al.* (2006) and Gokturk *et al.* (2004) focus on automotive applications; the latter, in particular, discussing effects such as correlation waveform aliasing, motion artefacts and the noise impact of ambient light on range data. Our own research group has published several papers utilising the XZ422 demonstrator camera; Payne *et al.* (2011) concentrated on optimising the duty cycle of modulation waveforms to reduce systematic error and Dorrington *et al.* (2011) demonstrate a simple, optimisation based version of the 2 : 1 frequency ratio mixed pixel separation algorithm presented in Chapter 6.

The XZ-422 typically uses differential measurements to remove ambient light from the raw data. A picture of the system is given in Fig. 2.17, which shows that the LED illumination ensemble is offset to one side of the imaging optics. Unfortunately this results in shadowing, where the illumination is occluded by objects in the scene. This can potentially result in erroneous range measurements.

The XZ-422 can be operated over a wider range of customisable frequencies, but the maximum utilised for this thesis was 36MHz. This particular frequency was chosen because it enabled measurements to be made at 18, 24 and 36 MHz without changing the internal oscillator frequency of the camera; only the digital divisor settings. This gave frequency ratios of 2 : 1, 3 : 2 and 4 : 3 for the purposes of our testing; with a focus on 2 : 1 in this thesis. In practice, the greatest limiting factor for modulation frequency is the response of the LED illumination system, as LEDS become more difficult to modulate at frequencies above 50 MHz.

2.4.3 SwissRanger 4000 (Integrated CMOS Sensor)

The SwissRanger 4000 is a four-sample homodyne system that is broadly similar to the Canesta XZ-422 except that it uses coaxial LED illumination and an optical feedback system. The optical feedback redirects some of the LED illumination directly back to a region of the sensor, allowing the removal of drift in phase and amplitude.

There is a large corpus of publications regarding the various SwissRanger systems. SwissRanger systems are one of the most common full-field lidar systems used for academic research – primarily because they can be bought off-the-shelf[8]. Kahlmann and Ingensand (2005) discussed calibration, including the distribution of

[7] This is now a moot point given that Canesta have been bought out by Microsoft; any new devices will be branded as Microsoft devices. A similar fate befell 3DV, who's system was primarily aimed at the movie and gaming markets; Iddan and Yahav (2001) and Yahav *et al.* (2007).

[8] Many other manufacturers, such as Canesta, are primarily focussed on the OEM mass-market and providing the necessary customer support for low volume orders is not seen as profitable.

the illumination intensity, photogrammatic calibration, range non-linearity and the impact of temperature. This was subsequently discussed again by Kahlmann *et al.* (2006) and Kahlmann and Ingensand (2008). In particular, Dorrington *et al.* (2010) found the SR4000 built-in photogrammetric calibration was erroneous, whereas the built-in XZ422 Canesta camera calibration was of comparatively high quality.[9] Robbins *et al.* (2008) analysed the accuracy of SwissRanger SR3100 measurements, although not specifically identifying the mixed pixel/multipath effect, noting the impact of bright objects in the scene on range-intensity coupling. In addition to these, there are several papers on pixel design (Büttgen *et al.*, 2006; Büttgen *et al.*, 2006; Büttgen and Seitz, 2008) and overviews of the camera systems Oggier *et al.* (2004, 2005).

The SwissRanger LEDs appear quite bright, despite the extremely low sensitivity of the human retina at this particular wavelength; this is somewhat disturbing, although we have verified that the device is eye-safe. The SR4000 was originally capable of 60MHz, but a firmware 'upgrade' restricted the usable frequencies to a maximum of 31MHz. The first time the system was operated at this particular setting resulted in one quarter of the LEDs becoming inoperable; admittedly, the setting was marked in the documentation as 'This frequency is for internal testing and should not be used.' As a result, it is not particularly surprising that this setting was disabled, albeit extremely frustrating in that it was seriously detrimental to the performance of the lookup table based mixed pixel separation algorithm of Chapter 6.

[9]It was later determined that there was a minor coding error in the libraries distributed with the SwissRanger.

Chapter 3

Modelling Measurement Formation

In this chapter we address the mechanisms behind the formation of the correlation waveform, how it is sampled and the impact of noise and systematic errors on measured phase and amplitude.

The first section describes how the correlation waveform is formed in the case of a perfect point-return. Whereas other works have modelled the correlation waveform as a triangle wave, we create an extended model for a truncated triangle waveform that can be generalised to rectangular illumination modulation and sensor response waveforms of any duty cycle. Examples are provided of temporal and spatial variability in phase and amplitude across sensors, in addition to the determination of the impulse response describing the spatial sensor responsivity variation.

The concept of a set of signal returns within each pixel is developed, in summation forming a signal return model, which is convolved by a reference waveform in order to form the correlation waveform. We describe how the correlation waveform is sampled in both the homodyne and heterodyne cases and how harmonics in the signal alias onto the fundamental.

An analysis of the primary noise sources in AMCW lidar is given; as a result of these analyses we discover that the presence of a second harmonic in the correlation waveform can result in systematic phase and amplitude perturbations, even if the second harmonic is cancelled out by a four phase sample differential measurement system. Finally we analyse the impact of irregular phase steps on phase and amplitude response, showing that this makes the sampling function more broadband, thus more highly susceptible to perturbations from harmonics in the correlation waveform.

3.1 Elementary Measurement Model

The model presented here uses our own notation, rather than copying that of any other author. Whereas previous work, such as Xu *et al.* (1995), Lange (2000) and Luan (2001), has described the correlation waveform in terms of the sum of translated and scaled cosine functions, here we describe the correlation waveform in terms of a Fourier transform, including negative frequencies.

3.1.1 An Ideal Point Measurement

We model the range measurement of a point (infinitesimal angular diameter) at a radial distance d from the camera with an intensity of a. The point is illuminated using an amplitude modulated light source, the illumination waveform notated as $\psi_i(\theta)$. The illumination is reflected back to the sensor by the point, resulting in a phase shift being introduced into the illumination waveform by total TOF, given by

$$\Delta\theta = 4\pi\frac{d}{\lambda}, \tag{3.1}$$

where λ is the illumination modulation wavelength. The illumination at the sensor is given by

$$s_i(\theta) = a\psi_i(\theta + \Delta\theta), \tag{3.2}$$

where θ is the phase of the illumination modulation waveform and $\Delta\theta$ is the phase shift introduced by TOF. AMCW range imaging sensors operate by correlating this returned illumination with another signal (also occasionally called gain mixing), the sensor shutter waveform – denoted by $\psi_s(\phi)$. The mixed signal is given by

$$s_c(\theta, \phi) = a\psi_i(\theta + \Delta\theta)\psi_s(\phi), \tag{3.3}$$

where ϕ is the phase of the sensor modulation waveform. By varying the relationship between θ and ϕ and integrating over time it is possible to efficiently measure $\Delta\theta$. In the homodyne case, for a modulation frequency of f_m and a phase difference between the modulation signals of $\phi_d = \phi - \theta$, a measurement at a particular phase step ϕ_d is notated as

$$\begin{aligned} s(\phi_d) &= \frac{1}{2\pi}\int_{\theta=0}^{2\pi} s_c(\theta, \theta + \phi_d)d\theta & (3.4)\\ &= a\psi(\phi_d + \Delta\theta) & (3.5) \end{aligned}$$

where $\psi(\,)$ is referred to as the reference waveform and $s(\,)$ as the correlation waveform. The distinction between the two is not made in other papers due to the assumption that there is only a single component return.

The Fourier transform of the correlation waveform is notated as

$$S(u) = a\mathrm{e}^{-8\pi^2 jud/\lambda}\Psi(u), \tag{3.6}$$

where $\Psi(u)$ is the Fourier transform of the reference waveform and $\mathrm{e}^{-8\pi^2 jud/\lambda}$ is the phase shift introduced into each spatial frequency by the TOF. A complex domain range measurement, ξ, is typically produced from the measured Fourier transform by

$$\xi = \frac{S(-\frac{1}{2\pi})}{\Psi(-\frac{1}{2\pi})}, \tag{3.7}$$

which calibrates for the relative amplitude and phase of the negative fundamental frequency in the sampled reference waveform. By convention we utilise the negative fundamental frequency in order to ensure that the phase shift is a positive linear function of distance. In this ideal case, the original intensity and distance is recovered from a complex domain range measurement by

$$a = |\xi| \tag{3.8}$$

$$d = \frac{\lambda}{4\pi}\arg(\xi) = \Delta\theta, \tag{3.9}$$

which is the relationship described in the introductory section for the four sample case (Eqns. 2.2 and 2.4).

3.1.2 The Formation of the Reference Waveform

The reference waveform, ϕ, is formed as the correlation of the illumination and sensor shutter waveforms. Many systems use approximately rectangular illumination and sensor modulation. Whereas Kang *et al.* (2011) developed a triangular model specific to 50% duty cycle modulation signal, here we develop a new model that generalises to modulation signals of any duty cycle.

Taking $\theta \in [-\pi, \pi)$, the modulation waveforms are idealised as

$$\psi_i(\theta) = \mathrm{rect}\left(\frac{\theta}{2\pi h_i}\right) \tag{3.10}$$

$$\psi_s(\theta) = \mathrm{rect}\left(\frac{\theta}{2\pi h_s}\right), \tag{3.11}$$

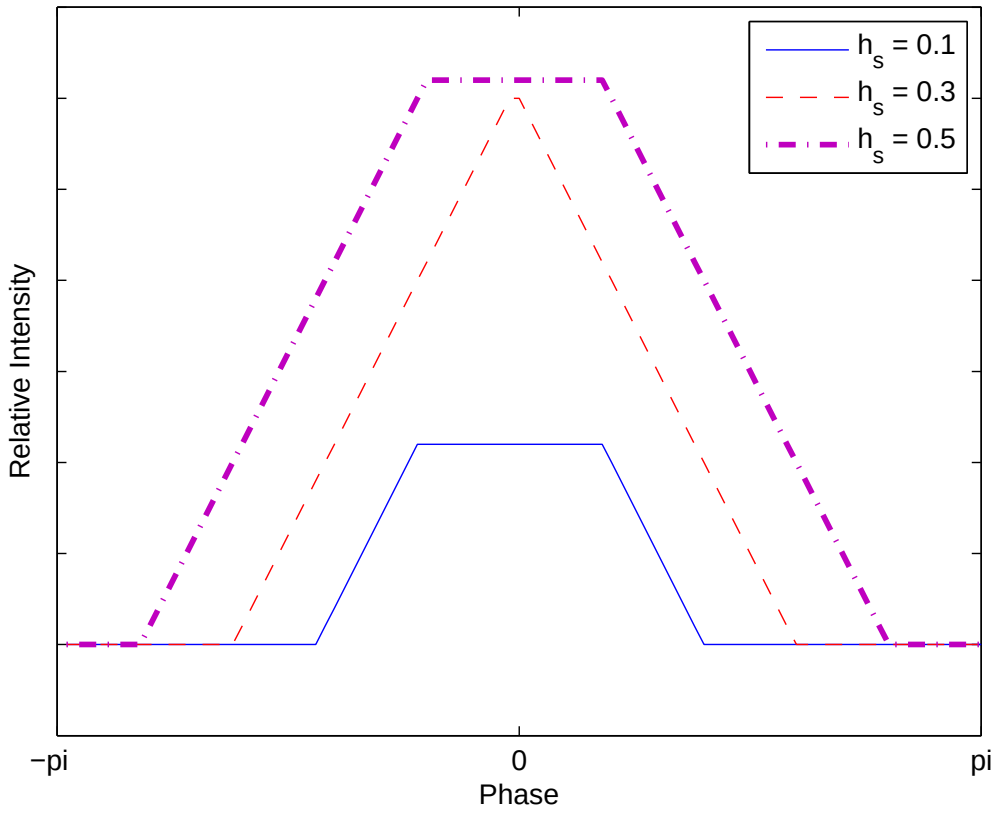

Figure 3.1: The reference waveform as a function of h_s, given $h_i = 0.3$. This figure details how the waveform shape changes with duty cycle under the assumption of rectangular illumination and sensor shutter waveforms.

where $h_i, h_s \in [0, 1]$ are the duty cycles of the illumination and sensor modulation respectively. The reference waveform is then,

$$\psi(\phi_d) = \frac{1}{2\pi} \int_{\theta=0}^{2\pi} \text{rect}\left(\frac{\theta}{2\pi h_i}\right) \text{rect}\left(\frac{\theta + \phi_d}{2\pi h_s}\right) d\theta \tag{3.12}$$

$$\Psi(u) = 2\pi \text{sinc}(2\pi h_i u) \text{sinc}(2\pi h_s u) \tag{3.13}$$

Varying the duty cycle gives a number of different truncated triangle reference waveform shapes as shown in Fig. 3.1. At $h_i = h_s = 0.5$, the waveform degenerates to triangular. Kang *et al.* (2011) developed a simple piecewise linear approach in the case that produces aliasing free range measurements in the perfect triangle waveform case; no literature appears to address the general truncated-triangle form. One possible reason for this is that it has a far more complicated form and is difficult to fit. We now develop this general form, which we apply to range measurement in Section 4.2.

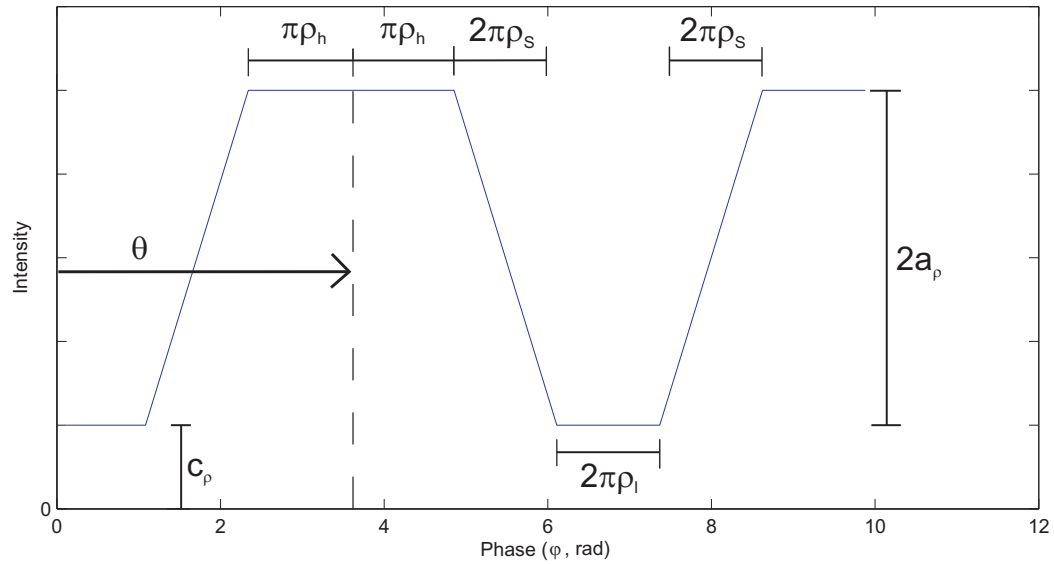

Figure 3.2: The parameters of a truncated triangle waveform

3.1.3 A Truncated Triangle Reference Waveform

A truncated triangle waveform, ignoring phase and intensity, can be uniquely described by the tuple $(\rho_h,\ \rho_l,\ \frac{c_\rho}{a_\rho})$ which may be calculated directly from the duty cycles of the laser and intensifier waveforms or fit to empirical data. The values ρ_h and ρ_l represent the proportion of time spent at peak and trough intensity respectively and $\frac{c_\rho}{a_\rho}$ represents the relative size of any constant offset to the amplitude of the waveform (see Fig. 3.2). While the tuple (h_i, h_s) is a complete representation in the case of waveform formation from correlation of two variable duty cycle rect functions, the addition of an extra parameter allows modelling of any additional constant offsets in $\psi_i(\theta)$ and $\psi_s(\theta)$.

Using knowledge of the duty cycle of the illumination and sensor shutter modulation, one can now calculate the parameters that describe the shape of the beat waveform, ultimately modelling the impulse response with a continuous piecewise function. Firstly calculating the waveform peak intensity, b_ρ, and the period of peak intensity, ρ_h. The waveform peak intensity is the maximal value of $\psi(\theta)$ for a particular pixel over time, where $b_\rho = 2a_\rho + c_\rho$.

$$b_\rho = \min(h_l, h_r) \tag{3.14}$$

$$\rho_h = |h_l - h_r| \tag{3.15}$$

The gradient of the waveform is quantised, with only $\Delta\psi(\theta) \in \{0, -\frac{1}{2\pi}, \frac{1}{2\pi}\}$ permitted. This implies that the time taken to descend from an intensity of b to zero is $2\pi b_\rho$. While the waveform does not necessarily reach zero, knowledge of the phase

shift required in order to achieve this allows us to determine if there is any constant offset c_ρ. We define r_h as the combined phase shift required for full ascent from zero intensity to a, peak intensity, and descent to zero again. This defines a special case, because if $r_h > 2\pi$ then the impulse response never takes a value of zero.

$$r_h = 2\pi\rho_h + 4\pi b_\rho \tag{3.16}$$

The proportional size of the flat lower region of the truncated triangle wave, if any, is given by

$$p_l = \begin{cases} 0, & r_h > 2\pi \\ 1 - \frac{r_h}{2\pi}, & r_h \leq 2\pi \end{cases} \tag{3.17}$$

If the value of r_h is too high, there is a constant offset c_ρ to the waveform

$$c_\rho = \begin{cases} \frac{r_h}{4\pi} - 1, & r_h > 2\pi \\ 0, & r_h \leq 2\pi. \end{cases} \tag{3.18}$$

With knowledge of ρ_l, the proportion of sloped ascent or descent region, ρ_s, can be determined by

$$\rho_s = \frac{1 - \rho_h - \rho_l}{2}, \tag{3.19}$$

and the amplitude a as

$$a_\rho = \frac{b_\rho - c_\rho}{2}. \tag{3.20}$$

Which enables calculation of the last of the defining parameters, $\frac{c}{a}$, which provides the relative size of any constant offset.

For $\theta \in [0, 2\pi)$, the reference waveform may be rewritten as

$$\psi(\phi_d) = \frac{1}{2\pi}\int_{\theta=0}^{2\pi} \text{rect}\left(\frac{\theta}{2\pi h_i}\right) \text{rect}\left(\frac{\theta + \phi_d}{2\pi h_s}\right) d\theta \tag{3.21}$$

$$= \frac{c_\rho}{b_\rho} + \begin{cases} 1, & \theta < \pi\rho_h \\ 1 - \frac{\theta - \pi\rho_h}{\rho_s}, & \pi\rho_h \leq \theta < \pi(\rho_h + 2\rho_s) \\ 0, & \pi(\rho_h + 2\rho_s) \leq \theta < \pi(2 - 2\rho_s - \rho_h) \\ \frac{\theta - \pi(2 - 2\rho_s - \rho_h)}{\rho_s}, & \pi(2 - 2\rho_s - \rho_h) \leq \theta < \pi(2 - \rho_h) \\ 1, & \pi(2 - \rho_h) \leq \theta < 2\pi, \end{cases} \tag{3.22}$$

which gives a complete parametric model of the reference waveform shape in the case of rectangular modulation signals.

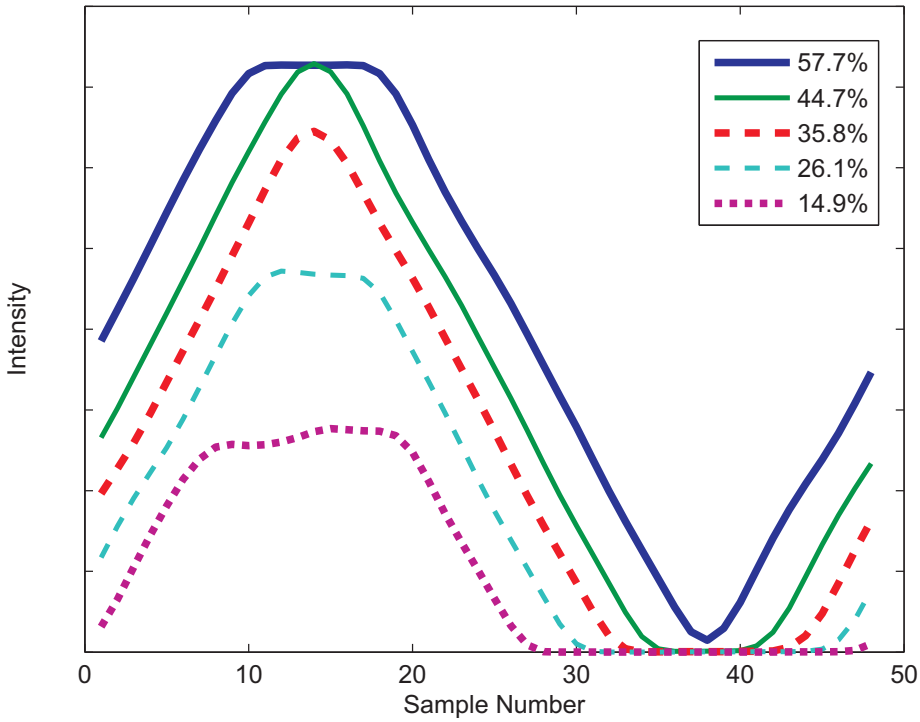

Figure 3.3: Measured reference waveform shape versus laser duty cycle for the UoW heterodyne range-imager.

3.1.4 The Measured Reference Waveform

For the University of Waikato Image Intensifier system, the measured reference waveform as the laser illumination duty cycle is varied are shown in Fig. 3.3. Apart from very low laser modulation duty cycles, the waveform is well approximated by a truncated triangle shape. This allows the use of Eqn. 3.22 as a model of the reference waveform, in order to determine the phase and amplitude of multiple component returns within each pixel (see Chapter 4).

Using a photodiode it is possible to trivially sample the actual laser illumination waveform; however, sampling the image intensifier shutter response is a serious challenge. In a previous paper Payne *et al.* (2008a) measured the image intensifier shutter response using an FPGA synchronised high speed pulsed laser; this required complex additional hardware. However, using our knowledge of the reference waveform formation process the shutter response waveform can be determined via deconvolution using a known illumination waveform and a measurement of the reference waveform. The Richardson-Lucy deconvolution algorithm (Richardson, 1972; Lucy, 1974) was used, which is a maximum-likelihood method designed for the presence of Poisson distributed noise. The resultant deconvolved shutter response is given in

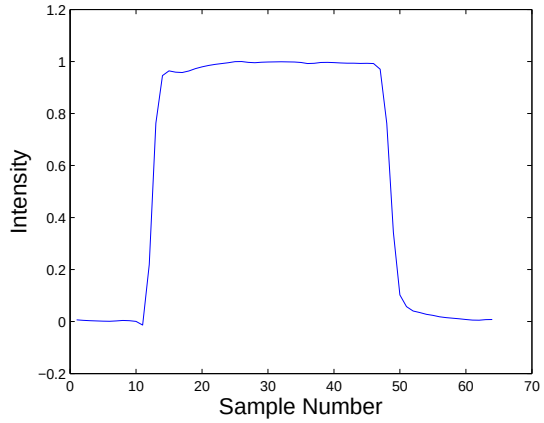

(a) Laser Waveform

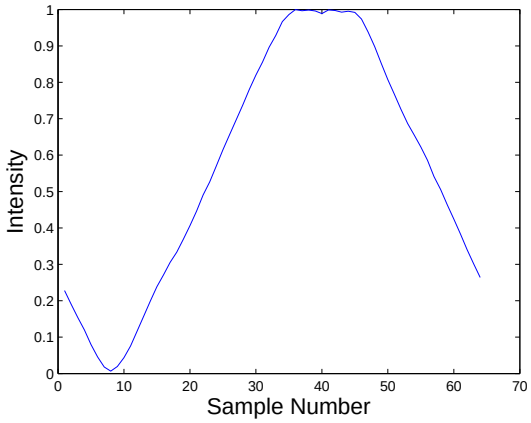

(b) Reference Waveform

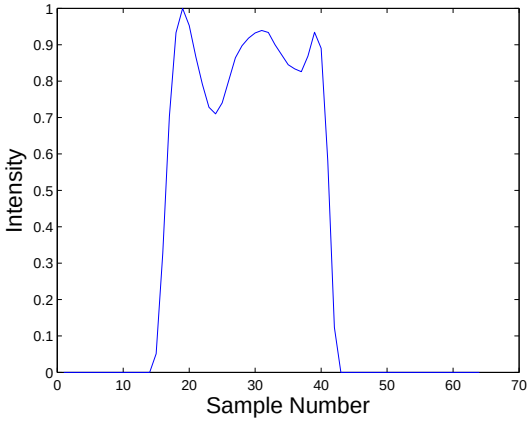

(c) Intensifier Shutter Waveform

Figure 3.4: Laser illumination waveform, recorded reference waveform and the intensifier shutter response, the latter recovered using 30 iterations of the Richardson-Lucy deconvolution algorithm. It shows image intensifier ringing consistant with Payne *et al.* (2008a).

Fig. 3.4 – this response is consistent with the results given by Payne *et al.* (2008a).

These in turn are formed by the convolution of the illumination and sensor modulation waveforms with the illumination and sensor responses. In practice, these underlying modulation waveforms are rectangular – although there are some advantages to sinusoidal modulation, it is more challenging to implement than rectangular modulation.

3.1.5 Spatial Variance and Specifics of Formation

An ideal abstraction of the sensor might assume that the reference waveform is unvarying; however, in reality the waveform shape changes both over time and

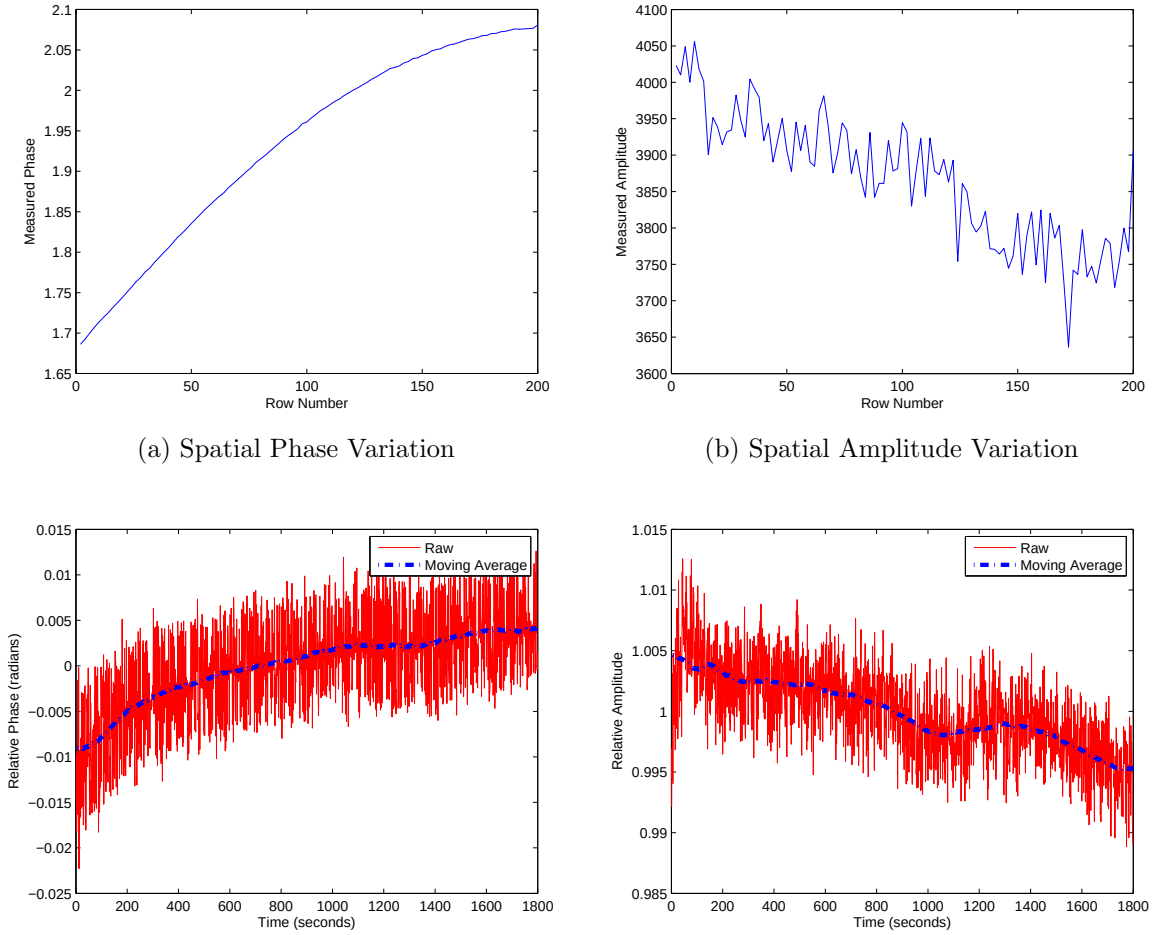

(a) Spatial Phase Variation

(b) Spatial Amplitude Variation

(c) Temporal Phase Variation

(d) Temporal Amplitude Variation

Figure 3.5: Spatiotemporal variation in the phase and amplitude of the negative fundamental frequency of the correlation waveform for a CMOS sensor. Raw data provided by Andrew Payne, with additional processing and analysis by the author.

spatially. Typically a modulated CMOS sensor is driven from one side of the chip, as the signal travel down sensor columns it becomes phase shifted and attenuated. The most appropriate electrical model is that of a transmission line. Each pixel essentially acts as an RLC network and depending on the sensor design and operating frequency, there is the potential for reflections and even standing waves. In contrast to modulated CMOS sensors, image intensifiers are typically driven from the outside, resulting in a characteristic 'irising effect' (Payne *et al.*, 2008a).

The presence of a phase shift across CMOS sensors has been acknowledged in other work: Fuchs and Hirzinger (2008) models the time delay across an IFM O3D100 ToF sensor as a linear function of row and column number. Fig. 3.5a shows the phase shift of the negative fundamental frequency of the sensor response wave-

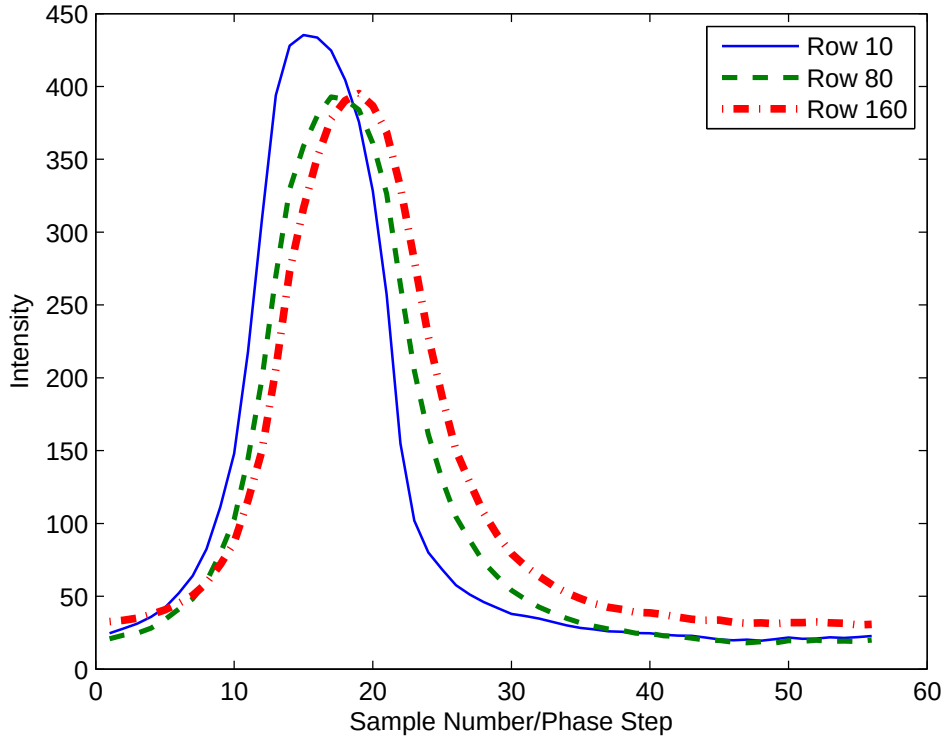

(a) Spatial Variation in Reference Waveform Shape

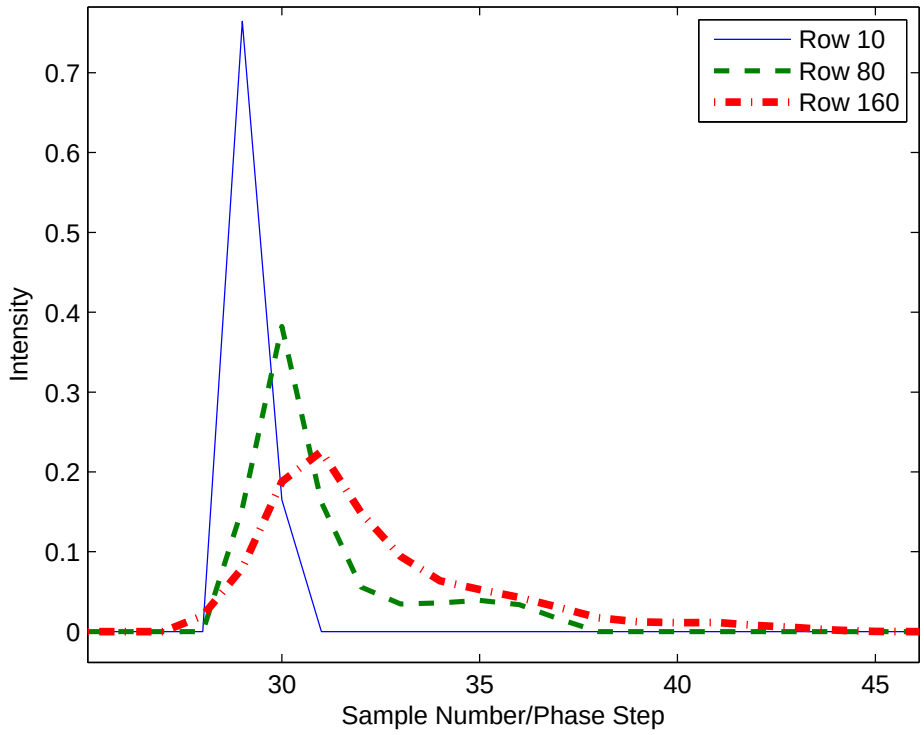

(b) Estimated Impulse Response Versus Row 2

Figure 3.6: Analysis of the spatial variation in reference waveform shape and the impulse response relating the waveform at the top of the sensor to that further down. Raw data provided by Andrew Payne, with additional processing and analysis by the author.

form down a sensor column in a modulated CMOS sensor[1]; it is clearly not a linear function of row number for this particular sensor. Fig. 3.5b shows the change in amplitude, an effect that does not appear to have been modelled in previous work. The data appears much more noisy due to the lack of any flat-field calibration; were the data not averaged across each row, the spatial amplitude variation due to transmission line effects would be far less significant than the pixel gain inhomogeneities.

Not only is the phase and amplitude of the negative fundamental frequency of the sensor response waveform perturbed, but the entire waveform is filtered and phase shifted. Fig. 3.6a shows the correlation waveform shape averaged across different rows in the sensor; the higher the row number, the more attenuated and delayed the waveform is. By deconvolving the waveforms for rows 10, 80 and 160 by the waveform for row 2 using the Richardson-Lucy algorithm, the impulse responses given in Fig. 3.6b were determined. The impulse responses become progressively more delayed and long tailed as the row number increases. Because the waveform shape changes across the sensor, the aliasing calibration ought to as well. However, no previous research appears to have addressed this aspect of calibration.

Another important effect is temporal variation in the amplitude and phase of the sensor response; this is most probably temperature related. Depending upon the exact method of signal generation, it is quite plausible for the entire reference waveform to change shape as the temperature changes. This has the potential to make aliasing calibration substantially more difficult, unless using approaches such as harmonic cancellation (see Payne *et al.* (2008b)). While optical feedback, like that used by the SwissRanger 4000, can calibrate for changes in amplitude and phase, it cannot easily calibrate for changes in harmonic content. Fig. 3.5c and 3.5d show phase and amplitude measurements over a period of approximately half an hour. It is also known that integration time can impact on linearity (Kahlmann *et al.*, 2006), although the precise mechanism is as yet unexplained.

3.2 Modelling Multiple Component Returns

In Section 3.1 we modelled the correlation waveform in the case of an ideal point-return. In this section we develop the idea of a continuous function called a signal return model, which describes the sources of the backscattered illumination recorded by each pixel. Previous work using AMCW lidar has either modelled the correlation waveform as a single translated and scaled copy of the reference waveform (Xu

[1]The raw measurements were provided to the author by Andrew Payne; the dataset was further analysed in order to produce Figs. 3.5 and 3.6. The author offers gratuitous thanks.

et al., 1995; Lange, 2000; Luan, 2001) or roughly estimated the impact of mixed pixels/multipath interference on the negative fundamental frequency using spatial information (Mure-Dubois and Hügli, 2007a; Kavli *et al.*, 2008; Falie, 2009; Fuchs, 2010).

While each AMCW range measurement provides only limited information about the signal return model, an ensemble of measurements has the potential to enable the determination of the phase and amplitude of the component returns making up the signal return model. We briefly identify different approaches to modelling the signal returns, including both point-return and diffuse-range models. These models are ultimately applied in Chapters 4 through 7.

3.2.1 Forming the Correlation Waveform

Section 3.1 described the point-return case. When an individual pixel contains a single backscattering source at a discrete range, the correlation waveform is a scaled and translated reference waveform. If the backscattering sources within the pixel are modelled as a function of phase, in this case the signal return model is given by

$$f_\xi(d_\phi) = a\delta(d_\phi - d), \tag{3.23}$$

where f_ξ is the signal return model for the pixel. The correlation waveform, $s(\phi_\mathrm{d})$, is formed as a convolution of the reference waveform with the signal return model, viz

$$s(\phi_\mathrm{d}) = f_\xi\left(\phi_\mathrm{d}\frac{\lambda}{4\pi}\right) * \psi(\phi_\mathrm{d}), \tag{3.24}$$

where λ is the modulation wavelength.

Ignoring intracamera scattering, intrascene reflections and semi-transparent objects, the light backscattered by a scene can be modelled by two discontinuous functions $d_\xi(x, y)$, which corresponds to the radial distance to the backscattering source at subpixel location $(x, y) \in \mathbb{R}^2$, and $I_\xi(x, y)$, which corresponds to the backscattered intensity. The continuous signal return model is then given by

$$f_\xi(x, y, d) = I_\xi(x, y)\delta(d_\xi(x, y) - d). \tag{3.25}$$

However, cameras do not measure a continuous 2D image, instead each pixel integrates over a known solid angle. The resultant sampled signal return model is

notated as

$$f_\xi[x,y](d) = \int_{x-0.5}^{x+0.5} \int_{y-0.5}^{y+0.5} f_\xi(\alpha,\beta,d)\, d\alpha\, d\beta. \tag{3.26}$$

It is this 2D dimensional integration operation that results in the phenomenon of mixed pixels, as each pixel can integrate over multiple unrelated surfaces. Also, because surfaces are not necessarily orthogonal to the optical axis, component returns have the quality of range-width in addition to range and amplitude. We call these diffuse-range returns; they can occur due to sloped objects or volumetric scattering sources such as fog or hair. They are addressed in Chapter 5.

Given that in the ideal case a calibrated complex domain range measurement is equivalent to sampling the spatial frequency of the signal return model corresponding to the fundamental frequency of the reference waveform, an ideal complex domain range measurement is equivalent to sampling the Fourier transform of the signal return model for each pixel. Via Eqn. 3.7,

$$\xi = \frac{S(-\frac{1}{2\pi})}{\Psi(-\frac{1}{2\pi})} \tag{3.27}$$

$$= F_\xi\left(\frac{2}{\lambda}\right). \tag{3.28}$$

Thus, given a large sequence of measurements, it is possible to reconstruct the signal return model by explicitly sampling each of the spatial frequencies. Stepped AMCW lidar has previous been combined with an inverse Fourier transform in order to achieve this; Simpson *et al.* (2005) uses an harmonic sequence of 20 measurements. Unfortunately, the number of measurements required to achieve a high quality model is extraordinarily high, and the modulation frequency bandwidth required for high resolution is extremely difficult to meet.

Another plausible approach is to utilise the harmonic content in the correlation waveform. Assuming that the reference waveform is sufficiently broadband, it is possible to express the signal return model as a deconvolution problem. This is the approach we demonstrate in Chapter 4. However, AMCW systems are not designed for harmonic content. Often they are designed to minimise harmonic content because harmonics typically lead to aliasing induced phase non-linearity. Deconvolution is sometimes used in full-waveform lidar systems in order to help isolate and determine component returns (Stilla and Jutzi, 2009). Measurements made by FMCW and range-gating systems can also be modelled in a similar manner as the convolution of an impulse response with a signal return model.

3.2.2 Modelling the Signal Returns

In order to recover the signal return model without requiring large sequences of measurements, it is necessary to make some simplifying assumptions about the nature of the signal return model. Rather than considering it as a continuous (or discrete) function where the intensity of the backscattering has an unrestricted domain, it is possible to limit the dimensionality of the problem by considering the signal return model as the sum of translated and scaled parametised distributions.

Fitting a model to the measured waveform is a common problem in full-waveform lidar. Models such as sums of scaled Gaussian functions (Hofton *et al.*, 2000), and scaled Generalised Gaussian functions (Chauve *et al.*, 2007), are fit to the captured data in order to determine the range, amplitude and width of the component returns. The simplest possible model, albeit naïve, is a point-return model using scaled and translated Dirac delta functions (also referred to as a sparse spike train). Apart from some analyses in Chapter 5, for the purpose of further discussions we make the assumption that the signal return model for a single pixel can be approximated by

$$f_{\xi}(d_{\phi}) = \sum_{i=0}^{n-1} a_i \delta(d_{\phi} - d_i), \tag{3.29}$$

where n is the number of returns, a_i is the amplitude of the ith return and d_i is the distance to the ith return.

3.3 Sampling of the Correlation Waveform

In the previous two sections we addressed the formation of the ideal continuous correlation waveform. However, practical AMCW range-imagers sample the correlation waveform using techniques such as homodyning and heterodyning. In this section we address these methods of sampling the correlation waveform and the impact they have on the measured data, including effects such as aliasing and the implementation of ambient light cancellation.

3.3.1 Aliasing and Modulation Techniques

In Eqn. 3.7 we showed how a complex domain range measurement is related to the Fourier transform of the correlation waveform. In practice, the Fourier transform cannot be measured directly, and is instead calculated from samples of the correlation waveform. There are two primary approaches to the sampling problem, depending on whether each measurement is made at a fixed relative phase, ϕ_d, or

integrated over a range of relative phases. The former corresponds to the simple homodyne modulation technique, and the latter corresponds to a wider group of techniques including heterodyning and harmonic cancellation (Payne *et al.*, 2008b).

Measurements subject to aliasing and other systematic errors are henceforth notated using a tilde. For example, $\tilde{s}$ is the correlation waveform after being sampled and subject to aliasing, whereas s is the ideal unaliased correlation waveform.

In the simple homodyne case, the illumination and sensor modulation frequencies are identical. Each measurement, taken at a fixed relative phase step, is given by

$$\tilde{s}[i] = s(\Delta\phi_d i), \tag{3.30}$$

where ϕ_d is the phase step size. From sampling theory, any harmonics above the Nyquist frequency alias onto those below. For a sampling frequency f_s, harmonics of frequency f are inseparable from those of frequency $|f + nf_s|$ for $n \in \mathbb{Z}$. The Fourier transform of a correlation waveform subject to aliasing, $\tilde{s}[i]$, can be notated as

$$\tilde{S}[u] = \frac{\Delta\phi_d}{2\pi} \sum_{k=-\infty}^{\infty} S\left(\frac{u}{\phi_d} + kf_s\right) = \frac{\Delta\phi_d}{2\pi} \sum_{k=-\infty}^{\infty} S\left(\frac{u+k}{\phi_d}\right). \tag{3.31}$$

The conventional approach is to extract a complex domain range measurement from the negative fundamental harmonic, as performed by Lange (2000). In the aliased case the measured complex domain range measurement is erroneous and can be written in terms of the ideal, unperturbed measurement as

$$\tilde{\xi} = \sum_{k=-\infty}^{\infty} \alpha_{nk+1} \frac{\xi^{nk+1}}{|\xi|^{nk}} \tag{3.32}$$

where α_{nk+1} is the aliasing coefficient calculated by

$$\alpha_x = \frac{\Psi(\frac{-x}{n\phi_d})}{\Psi(-\frac{1}{2\pi})}. \tag{3.33}$$

The problem with Eqn. 3.32 is that it is extremely complicated and non-linear. Since most practical systems use four samples of the correlation waveform at 90 degree phase steps and contain significant harmonic content, the raw range data produced by these imagers has a significantly non-linear range and amplitude response. Typically for a near 50% duty cycle, all the odd harmonics of the correlation waveform alias onto the fundamental. Given that a trivial inverse does not appear to exist, most practical implementations rely on lookup tables or other simple mathematical approximations such as B-Splines (Lindner *et al.*, 2010), to correct measured values.

Examples and discussion of the impact of aliasing on phase linearity are given in Section 3.5.2 (see Figs. 3.9 and 3.10).

One approach to partially mitigate the impact of aliasing is to use heterodyning. Heterodying works by running the illumination and sensor modulation waveforms at fractionally different modulation frequencies. As a result, each sample of the correlation waveform is integrated over a range of relative phases. In this case the measured aliased waveform is given by

$$\tilde{s}[i] = \int_{-\phi_h/2}^{\phi_h/2} s(\Delta\phi_d i + \theta_h)\mathrm{d}\theta_h, \tag{3.34}$$

where $0 \leq \phi_h \leq \Delta\phi_d$ is the range of phases over which the integration is performed. The case when $\phi_h = 0$ is the degenerate homodyne case, whereas $\phi_h = \Delta\phi_d$ requires that the modulation signals be paused or reset during the sensor readout period. Heterodyning has two major impacts on the waveform: it immediately attenuates many of the higher harmonics, and it changes the nature of motion during the measurement period – pure radial motion results in a change in the frequency of the correlation waveform.

The Fourier transform of the correlation waveform under heterodyning becomes

$$\tilde{S}[u] = \frac{\Delta\phi_d}{2\pi} \sum_{k=-\infty}^{\infty} \mathrm{sinc}\left(\frac{\phi_h(u+k)}{\phi_d}\right) S\left(\frac{u+k}{\phi_d}\right) \tag{3.35}$$

3.3.2 Measurement Techniques and Ambient Illumination Rejection

As illustrated in Fig. 2.13 for the UoW range-imager, typical range-imagers have narrow band filters on their primary optics in order to heavily attenuate the contribution of ambient light to measurements. For systems using infrared illumination, this is assisted by the reduced intensity of ambient infrared contributions versus visible. However, a small amount of ambient light is still detected despite attempts at amelioration. Let ρ_I be the constant offset due to ambient light in the correlation waveform for a single measurement.

The measurement technique utilised by the University of Waikato Full-Field Heterodyne Range-Imager is sequential sampling of the correlation waveform using a single channel. For each pixel this can be considered to produce a raw measurement vector, κ_A, given by

$$\kappa_\mathrm{A} = (\tilde{s}[0], \tilde{s}[1], ...\tilde{s}[n-1])^T + \epsilon_\mathrm{A} + \rho_\mathrm{I}, \tag{3.36}$$

where $\epsilon_{\mathrm{A}} \in \mathbb{R}^n$ is a vector representing the impact of noise. An aliasing perturbed complex domain range measurement may be calculated using the discrete Fourier transform bin corresponding to the negative fundamental frequency, namely the scaled inner product given by

$$\tilde{\xi} = \rho_{\mathrm{x}}(\mathrm{e}^{-2\pi j 0/n}, \mathrm{e}^{-2\pi j 1/n}, ...\mathrm{e}^{-2\pi j (n-1)/n}) \cdot \kappa_{\mathrm{A}}, \tag{3.37}$$

where ρ_{x} is the scaling coefficient given by

$$\rho_{\mathrm{x}} = \frac{\Delta\phi_d}{2\pi\Psi(-\frac{1}{2\pi})}. \tag{3.38}$$

The mean correlation waveform intensity is determined by averaging the samples, in the single component return case giving

$$\bar{\kappa}_{\mathrm{A}} = \frac{1}{n}\sum_{i=0}^{n-1}\kappa_a[i] = \frac{\Psi(0)}{|\Psi(-\frac{1}{2\pi})|}|\xi| + \rho_{\mathrm{I}}. \tag{3.39}$$

While the mean correlation waveform intensity is useful for noise modelling, as the variance of photon shot noise is a function of non-differential intensity, it is not typically considered a useful value by the average user of a range imaging system. In particular, if a range-imager is being used under high ambient light conditions such as outside in full sun-light, it is possible for the majority of the dynamic range of the measurements to be used up by the ambient light offset. In order to achieve sufficiently high measurement precision it is necessary to have extremely high bit count Analogue-To-Digital Converters (ADC), resulting in increased complexity and bitrate. A preferable alternative is to supress background light using differential measurements.

3.3.3 Differential Measurements

Whereas the Waikato heterodyne range-imager takes a single measurement at a single time, most integrated CMOS sensors take two measurements simultaneously for each pixel. These two measurements are generally referred to as the 'A' and 'B' channels. One of the biggest issues with active illumination based sensors is that an increase in ambient light contributes to measured values; as the ambient light intensity increases, the usable dynamic range of the measurements decreases. In order to remove the influence of ambient light, albeit not the corresponding increase in photon shot noise, most CMOS sensors modulate the A and B channels 180 degrees out of phase or modulate B with the inverse of the modulation signal for

A. Often, a differential sensor will subtract the value of channel B from channel A in the analogue domain, and only read out the combined value; however, depending on the sensor it may be also be possible to read A and B individually or A plus B. One approach to ambient light suppression is given by Schmidt and Jähne (2009), where whenever one of the channels approaches saturation, a constant charge is subtracted from both of the channel A and B quantum wells; since only the difference is measured, the change in constant offset can be ignored.

We now assume that channel A is modulated by $\psi_s(\theta)$ and channel B is modulated by $\pi_\mathrm{B} - \psi_s(\theta)$ for some offset π_B such that $\pi_\mathrm{B} - \psi_s(\theta) \geq 0$ for all values of θ. From a digital perspective, the B channel sensor modulation waveform is an inverted form of the A channel waveform – except that in practice there may be non-linear responses which result in other shape differences. Assuming κ_A to be the measurement vector for the A channel, the B channel vector is given by

$$\kappa_\mathrm{B} = \gamma_\mathrm{B} S(0) - (\tilde{s}[0], \tilde{s}[1], ...\tilde{s}[n-1])^T + \epsilon_\mathrm{B} + \rho_\mathrm{I}, \tag{3.40}$$

where the coefficient γ_B is given by

$$\gamma_\mathrm{B} = \frac{\Delta\phi_\mathrm{d}}{2\pi} \cdot \frac{\pi_\mathrm{B} n - \Psi(0)}{\Psi(0)}. \tag{3.41}$$

Since the subtraction is generally computed in the analogue on-chip, in an ideal case the combined measurement vector is given by

$$\kappa_\mathrm{A} - \kappa_\mathrm{B} = 2(\tilde{s}[0], \tilde{s}[1], ...\tilde{s}[n-1])^T - \gamma_\mathrm{B} + \epsilon_\mathrm{A} - \epsilon_\mathrm{B}. \tag{3.42}$$

The unknown constant offset caused by ambient light has been cancelled out and replaced with another offset $\gamma_\mathrm{B} S(0)$; ideally this offset is minimised through judicious choice of π_B. A complex domain range measurement can now be calculated using Eqn. 3.37, replacing κ_A by $\kappa_\mathrm{A} - \kappa_\mathrm{B}$.

There are practical issues with the implementation of the A and B channels. Depending on the precise hardware implementation, it is possible to have spatial coupling, where the A and B channels may have varying sensitivities to light incident on different regions within the pixel due to the extent of fringing fields. This has the potential to cause artefacts around the edges of objects or on patterned surfaces. Another concern is gain mismatch between the A and B channels; by using four phase steps this mismatch is automatically eliminated. If we assume that the relative gain

of the B channel is $\beta_{\mathrm{B}} \approx 1$, we can rewrite the differential measurement as

$$\kappa_{\mathrm{A}} - \beta_{\mathrm{B}}\kappa_{\mathrm{B}} = (1+\beta_{\mathrm{B}})(\tilde{s}[0], \tilde{s}[1], ...\tilde{s}[n-1])^T - \beta_{\mathrm{B}}\gamma_{\mathrm{B}}S(0) + \epsilon_{\mathrm{A}} - \beta_{\mathrm{B}}\epsilon_{\mathrm{B}} + (1-\beta_{\mathrm{B}})\rho_{\mathrm{I}}. \quad (3.43)$$

This latter equation shows that the ambient light has no-longer been completely cancelled out – but the remaining levels are likely to have much less of an impact on dynamic range. For cameras such as the Swissranger 4000, which provide no facility for making common mode (sum of channels A and B) or non-differential measurements, this may provide a mechanism for determining the mean correlation waveform intensity for noise analysis purposes. However, the uncertainty in this value is likely to be very high. In general, most papers ignore the possibility of gain mismatch.

While some cameras such as the Canesta XZ-422 provide the ability to use measurements derived from two differential phase steps with a 90 degree phase shift, it is not a reliable technique under most circumstances. Even ignoring the need for per-pixel bias calibrations, unless all measurements are taken in a completely darkened room there are likely to be systematic errors due to the presence of uncancelled ambient light in the measurements. While Hussmann and Edeler (2009) demonstrated this using a PMD system, without calibration, high quality ranging results may require a bias/gain model parametised by sensor temperature. The complex domain range measurement, ignoring noise and fixed offsets, can be notated as

$$\tilde{\xi} = \rho_{\mathrm{x}}(\tilde{s}[0] - \beta_{\mathrm{B}}\tilde{s}[2] + (1-\beta_{\mathrm{B}})\rho_{\mathrm{I}} + j(\tilde{s}[1] - \beta_{\mathrm{B}}\tilde{s}[3] + (1-\beta_{\mathrm{B}})\rho_{\mathrm{I}})). \quad (3.44)$$

3.4 Random Noise Sources

In this section we discuss different noise sources, and the impact they have on range measurements, in particular photon shot noise and read noise. We do not discuss noise due to jitter/oscillator frequency stability, although it can be a significant source of error in some systems. Additionally, we have ignored more complicated design specific error sources such as noise due to ADC instabilities/non-linearities, CMOS pixel design issues, row/column issues etc. Ideally, these would be largely resolved at a hardware level so that they are not present in the output data.

3.4.1 Photon Shot Noise

The arrival and detection of photons by a CCD or CMOS sensor is a discrete stochastic process which can be modelled by a Poisson distribution. The Poisson distribution is a parametric distribution that is described by a single parameter λ, corresponding to the expected number of detections over a fixed period of time. The probability density function over k, the number of detections, is given by

$$f(k|\lambda) = \frac{\lambda \mathrm{e}^{-\lambda}}{k!}. \tag{3.45}$$

As λ increases above 20, a Poisson distribution approaches a Gaussian distribution, viz

$$\lambda > 20 \Rightarrow \mathcal{P}(\lambda) \approx \mathcal{N}(\lambda, \lambda), \tag{3.46}$$

for a Poisson distribution $\mathcal{P}(\lambda)$ and Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$. Below approximately 20 detections per sample, the increased kurtosis of a Poisson distribution versus a Gaussian makes the approximation relatively poor.

For a sufficiently high intensity measurement, x, the variance, σ_{x}, is given by the relation

$$(\alpha_{DN}\sigma_{\mathrm{x}})^2 = \alpha_{DN}x, \tag{3.47}$$

where α_{DN} is the number of photons per DN (Digital Number: the quantised value output by the ADC.)

3.4.2 Read Noise

Read noise is generally modelled as a zero centred Gaussian distribution, but comes from a number of different sources including dark current, reset noise, amplifier noise and quantisation error. The particular applicable noise sources differ depending on

the design of the sensor, however the combined impact may still be modelled in the same manner. There is potential for any fixed calibrations to be invalidated by changes in sensor temperature. Read noise is primarily a consideration at low light levels as photon noise tends to dominate at higher light levels.

3.4.3 Noise in the Measurement Domain

The contribution of random noise, ϵ_{x}, to a measurement $\tilde{\xi}$ can be modelled as the sum of component noise sources, namely

$$\epsilon_{\mathrm{x}} = \epsilon_{\mathrm{r}} + \epsilon_{\mathrm{s}}, \tag{3.48}$$

where $\epsilon_{\mathrm{r}} \sim \mathcal{N}(0, \sigma_{\mathrm{r}}^2)$ is the contribution from read noise and $\epsilon_{\mathrm{s}} \sim \mathcal{N}(0, x/\alpha_{DN})$ is the contribution from shot noise. Assuming that the noise sources are uncorrelated, then $\epsilon_{\mathrm{x}} \sim \mathcal{N}(0, \sigma_{\mathrm{r}}^2 + x/\alpha_{DN})$.

For a measurement vector κ_{A}, the noise component is given by

$$\epsilon_{\mathrm{A}} \sim \mathcal{N}\left(0, \sigma_{\mathrm{r}}^2 + \frac{1}{\alpha_{DN}}\left(\rho_{\mathrm{I}} + \begin{bmatrix} \tilde{s}[0] & 0 & \cdots & 0 \\ 0 & \tilde{s}[1] & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \tilde{s}[n-1] \end{bmatrix}\right)\right). \tag{3.49}$$

One way to measure α_{DN} and σ_{r} is to produce a Photon Transfer Curve (Janesick, 1997). A PTC is a plot of the amount of noise in a measurement versus the mean intensity of that measurement. In addition to allowing determination of these parameters, a PTC can show the effects of uncalibrated fixed pattern noise, non-linearity and other sensor problems. Fig. 3.7 shows the results of a PTC for the image intensifier system when calibrated for bias and flat-fielded. A sudden change in slope in Fig. 3.7a may possibly be due to non-linear characteristics of the image intensifier above a certain image intensity. Fig. 3.7b shows a fit to the lower region of the curve, giving an estimate of approximately 3 DN per photon.

3.4.4 Noise in the Complex Domain

There have been a number of low level analyses of the impact of noise on range measurements, such as those of Xu *et al.* (1995) and Frank *et al.* (2009). However, these analyses have made the assumption that the noise distribution is isotropic across all the phase samples, which is not necessarily true. While photon shot noise results in a circularly symmetric complex Gaussian distribution for range measurement error

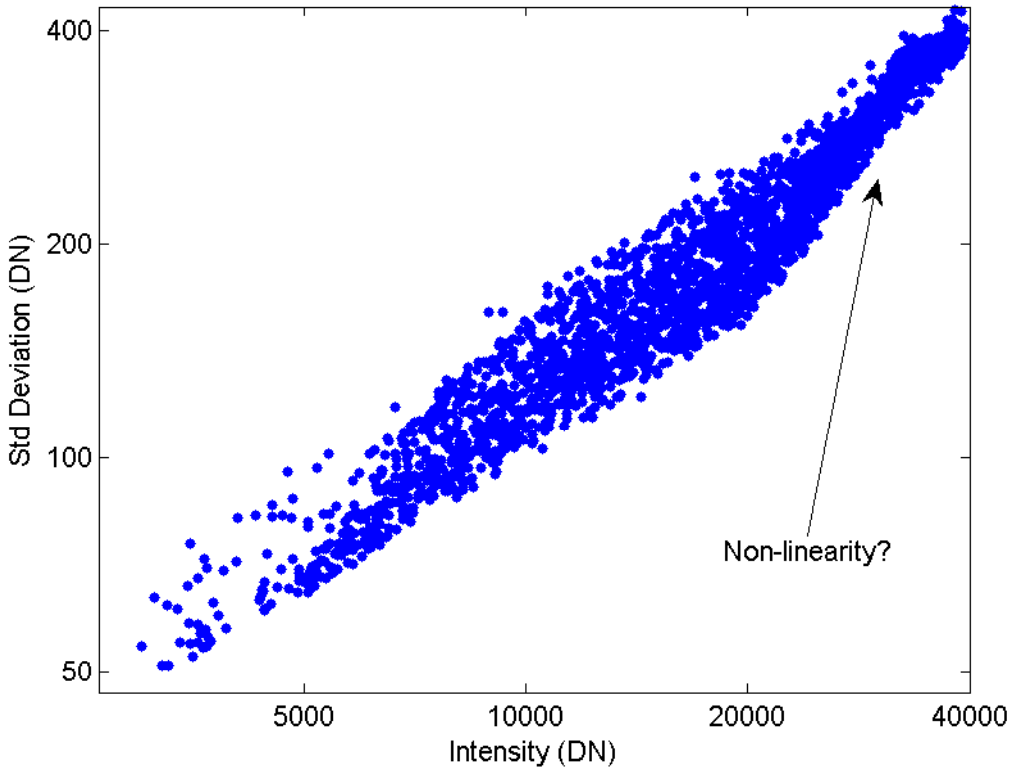

(a) Entire PTC Curve

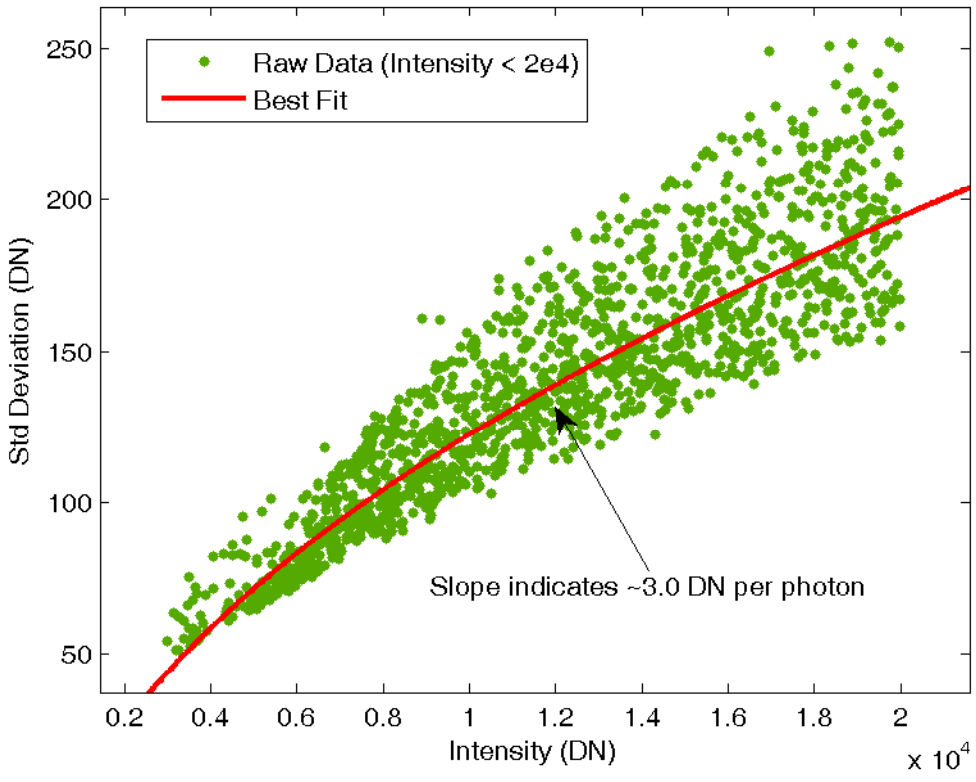

(b) Best Fit Model

Figure 3.7: Photon transfer curves for the image intensifier system.

in the case of a perfect sinusoid correlation waveform, the presence of harmonics changes this relationship. In this section we develop a new model that takes into account photon shot noise and the impact of correlation waveform harmonics.

Using Eqn. 3.37 one can propagate the impact of noise on κ_{A} in the non-differential measurement case through to the resultant complex domain range measurement $\tilde{\xi}$, which is sampled from a Gaussian distribution such that $(\Re(\tilde{\xi}), \Im(\tilde{\xi}))^T \sim \mathcal{N}(0, \Sigma_\xi)$.

The covariance matrix, Σ_ξ, is given by

$$\begin{aligned}
\Sigma_\xi &= \begin{pmatrix} \mathrm{var}(\Re(\tilde{\xi})) & \mathrm{var}(\Re(\tilde{\xi}), \Im(\tilde{\xi})) \\ \mathrm{cov}(\Im(\tilde{\xi}), \Re(\tilde{\xi})) & \mathrm{var}(\Re(\tilde{\xi})) \end{pmatrix} && (3.50)\\
&= R_{\mathrm{x}}^{-1} \left(\sigma_{\mathrm{r}}^2 \mathbf{I}_2 + \frac{1}{\alpha_{DN}} \sum_{i=0}^{n-1} \begin{pmatrix} \cos(\frac{2\pi i}{n}) \\ \sin(\frac{2\pi i}{n}) \end{pmatrix} \cdot \begin{pmatrix} \cos(\frac{2\pi i}{n}) \\ \sin(\frac{2\pi i}{n}) \end{pmatrix}^T (\rho_{\mathrm{I}} + \tilde{s}[i]) \right) R_{\mathrm{x}} && (3.51)\\
&= R_{\mathrm{x}}^{-1} \left(\sigma_{\mathrm{r}}^2 \mathbf{I}_2 + \frac{1}{\alpha_{DN}} \Sigma_\rho \right) R_{\mathrm{x}}, && (3.52)
\end{aligned}$$

where R_{x} is a rotation and scaling matrix corresponding to complex multiplication by ρ_{x} and is given by

$$R_{\mathrm{x}} = |\rho_{\mathrm{x}}| \begin{pmatrix} \cos(\arg(\rho_{\mathrm{x}})) & \sin(\arg(\rho_{\mathrm{x}})) \\ -\sin(\arg(\rho_{\mathrm{x}})) & \cos(\arg(\rho_{\mathrm{x}})) \end{pmatrix} \tag{3.53}$$

and Σ_ρ is the unrotated, unscaled covariance matrix representing the contribution from shot noise given by

$$\Sigma_\rho = \begin{pmatrix} \sigma_{\Re\Re}^2 & \sigma_{\Re\Im}^2 \\ \sigma_{\Im\Re}^2 & \sigma_{\Im\Im}^2 \end{pmatrix}. \tag{3.54}$$

The diagonal elements of Σ_ρ can be rewritten using the trigonometric identities

$$\begin{aligned}
\cos(\theta)\cos(\phi) &= \frac{\cos(\theta-\phi) + \cos(\theta+\phi)}{2} && (3.55)\\
\sin(\theta)\sin(\phi) &= \frac{\cos(\theta-\phi) - \cos(\theta+\phi)}{2} && (3.56)
\end{aligned}$$

giving

$$\begin{aligned}
\sigma_{\Re\Re}^2 &= \sum_{i=0}^{n-1} \cos^2\left(\frac{2\pi i}{n}\right) (\rho_{\mathrm{I}} + \tilde{s}[i]) && (3.57)\\
&= \sum_{i=0}^{n-1} \frac{1 + \cos(\frac{4\pi i}{n})}{2} (\rho_{\mathrm{I}} + \tilde{s}[i]) && (3.58)
\end{aligned}$$

and

$$\sigma^2_{\Im\Im} = \sum_{i=0}^{n-1} \sin^2\left(\frac{2\pi i}{n}\right)(\rho_{\mathrm{I}} + \tilde{s}[i]) \tag{3.59}$$

$$= \sum_{i=0}^{n-1} \frac{1-\cos(\frac{4\pi i}{n})}{2}(\rho_{\mathrm{I}} + \tilde{s}[i]). \tag{3.60}$$

In the same manner the remaining elements can be rewritten using the trigonometric identity

$$\cos(\theta)\sin(\phi) = \frac{\sin(\theta+\phi) - \sin(\theta-\phi)}{2} \tag{3.61}$$

to give

$$\sigma^2_{\Re\Im} = \sigma^2_{\Im\Re} = \sum_{i=0}^{n-1} \cos\left(\frac{2\pi i}{n}\right)\sin\left(\frac{2\pi i}{n}\right)(\rho_{\mathrm{I}} + \tilde{s}[i]) \tag{3.62}$$

$$= \sum_{i=0}^{n-1} \frac{\sin(\frac{4\pi i}{n}) - 0}{2}(\rho_{\mathrm{I}} + \tilde{s}[i]). \tag{3.63}$$

Using these results, Σ_ρ can be rewritten in terms of the Fourier transform of the ideal, non-noise perturbed, sampled correlation waveform $\tilde{S}[u]$ as

$$\Sigma_\rho = \frac{1}{2}\begin{pmatrix} \rho_{\mathrm{I}} + \tilde{S}[0] + \Re(\tilde{S}[2]) & \Im(\tilde{S}[2]) \\ \Im(\tilde{S}[2]) & \rho_{\mathrm{I}} + \tilde{S}[0] - \Re(\tilde{S}[2]) \end{pmatrix}. \tag{3.64}$$

Since $\tilde{s}$ or ρ_{I} cannot be measured directly, a good approximation of Σ_ρ is given by using the κ_{A} in lieu, which is the noise perturbed version of $\tilde{s} + \rho_{\mathrm{I}}$.

These results indicate that if the second harmonic in the sampled correlation waveform has a non-zero amplitude, then the noise in the complex domain range measurement is not circularly symmetric. This includes cases where there is no natural second harmonic, only another frequency aliasing onto it. However, extending Eqn. 3.64 the differential measurement case, all the terms involving the second harmonic become weighted by $1 - \beta_{\mathrm{B}}$, which significantly reduces the asymmetry.

If the second harmonic is of relatively low intensity it is reasonable to model the noise as a circularly symmetric complex Gaussian distribution. This can be written in the non-differential case as

$$\begin{pmatrix} \Re(\tilde{\xi}) \\ \Im(\tilde{\xi}) \end{pmatrix} \sim \mathcal{N}\left(0, \sigma^2_{\tilde{\xi}}\mathbf{I}_2\right) \tag{3.65}$$

or

$$\tilde{\xi} \sim \mathcal{N}_{\mathbb{C}}\left(0, \sigma_{\xi}^2\right), \tag{3.66}$$

where

$$\sigma_{\xi}^2 = |\rho_{\mathrm{x}}| \left(\sigma_{\mathrm{r}}^2 + \frac{\rho_{\mathrm{I}} + \tilde{S}(0)}{2\alpha_{DN}} \right). \tag{3.67}$$

Although this is not always a correct assumption, it is the assumption made by previous authors such as Frank *et al.* (2009).

3.4.5 Noise in the Polar Domain

The conversion from a complex domain range measurement to polar form is a non-linear transformation given by

$$(a, \theta_{\xi}) = t_{\mathrm{p}}(\tilde{\xi}) = (|\tilde{\xi}|, \arg(\tilde{\xi})). \tag{3.68}$$

Frank *et al.* (2009) modelled the transformation from raw four phase step measurements to polar form as a linear transformation. This is a first order approximation in which error in the phase angle and error in the amplitude are approximated as orthogonal. We extend this to a more general form, by operating on the complex domain measurement covariance matrix. The Jacobian of $t_{\mathrm{p}}(\tilde{\xi})$ is

$$J = \begin{pmatrix} \frac{\partial a}{\partial \Re(\tilde{\xi})}, \frac{\partial a}{\partial \Im(\tilde{\xi})} \\ \frac{\partial \theta_{\xi}}{\partial \Re(\tilde{\xi})}, \frac{\partial \theta_{\xi}}{\partial \Im(\tilde{\xi})} \end{pmatrix} \tag{3.69}$$

$$= \begin{pmatrix} \cos(\theta_{\xi}) & \sin(\theta_{\xi}) \\ -\frac{1}{a}\sin(\theta_{\xi}) & \frac{1}{a}\cos(\theta_{\xi}) \end{pmatrix}. \tag{3.70}$$

Given a complex domain measurement covariance matrix Σ_{ξ}, the polar form covariance matrix is

$$\Sigma_{\theta} = J \Sigma_{\xi} J^T. \tag{3.71}$$

While in practice this is generally a reasonable approximation, in reality neither of the elements in the tuple (a, θ_{ξ}) are Gaussian distributed. If the complex Gaussian distribution is circularly symmetric then noise in the amplitude estimate is sampled from a Rayleigh distribution with spread parameter σ_{ξ}^2. Noise in the phase estimate is sampled from a projected complex Gaussian distribution. For large SNRs this noise is approximately Gaussian, however as SNR decreases a Gaussian noise distribution becomes a progressively worse noise model. An illustration is given in Fig. 3.8.

The impact of these noise distributions on optimal approaches to averaging data and the introduction of systematic errors is discussed in Section 3.5.3.

While it has yet to be investigated, an interesting extension to this noise model would be to attempt to analyse what effect jitter/frequency stability in concert with the shape of the correlation waveform has on the noise distribution. In the case of jitter, the noise in each phase sample is proportional to the derivative of the correlation waveform at that point. This is likely to be significantly more difficult to analyse mathematically than shot noise statistics, although still relatively easy to simulate computationally.

3.5 Systematic Errors

There are many sources of systematic error in a range image. Here we discuss issues such as aliasing of the correlation waveform harmonics, systematic perturbations from non-circularly symmetric random noise and uneven phase sampling. The perturbations due to mixed pixels and multipath interference are discussed in detail in Chapter 6, while defocus is discussed in appendices A and B.

There are also a number of other different systematic errors, such as saturation, sensor non-linearities, systematic errors due to jitter and occlusion, which we do not detail with in detail. Occlusion, in particular, may be a significant source of error in systems where the illumination is non-coaxial, such as the Canesta XZ-422. As a result, regions in the shadow of other objects from the illumination source can still be imaged by the sensor. Often valid measurements are returned by the camera, due to intracamera light scattering.

3.5.1 Modelling Perturbation of the Primary Component Return

The perturbation of the primary component return in either the two component return case, or the single alising harmonic case can be modelled by a useful function $\Lambda_{\mathrm{f}}(b, \theta)$:

$$\Lambda_{\mathrm{f}}(b, \theta) = 1 + b\mathrm{e}^{j\theta}. \tag{3.72}$$

This function represents the relative perturbation in phase and amplitude of a return by a second component with relative amplitude, b, and phase, θ. This function is used below in Section 3.5.2 to analyse the properties of aliasing and in Chapters 5 and 6 to analyse the perturbation of a primary component return by a secondary component return.

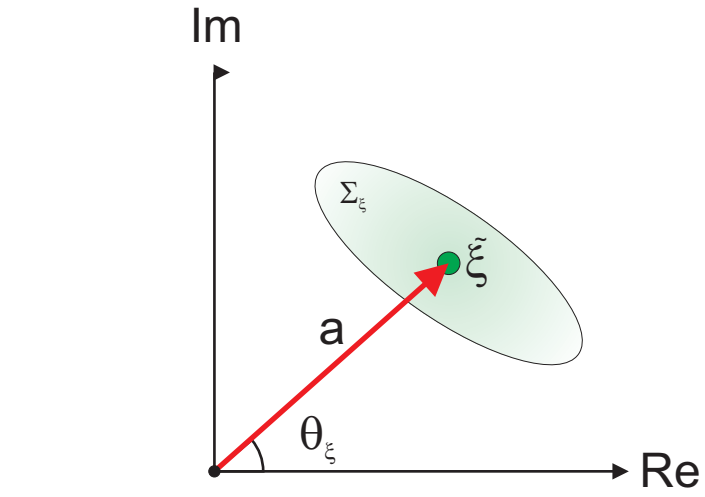

(a) Phasor Perturbation

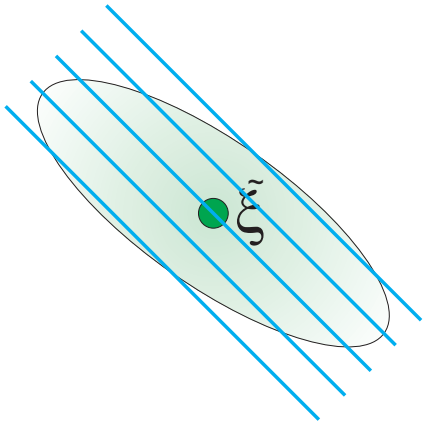

(b) Amplitude (High SNR)

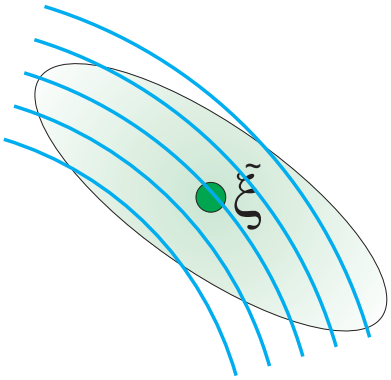

(c) Amplitude (Low SNR)

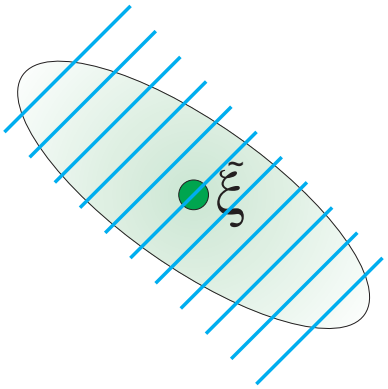

(d) Phase (High SNR)

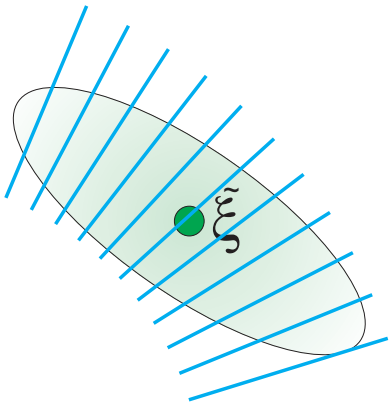

(e) Phase (Low SNR)

Figure 3.8: Illustration of systematic errors due to non-orthogonality of amplitude and phase. Fig. a shows the uncertainty in a measurement, $\tilde{\xi}$, with non-circularly symmetric covariance Σ_ξ. As the SNR decreases and the region of uncertainty grows, lines of constant amplitude become more curved (Figs. b and c) – this also occurs with lines of constant phase (Figs. d and e). At high SNRs, amplitude and phase are near-orthogonal making a first order approximation appropriate. As SNR decreases, errors in amplitude and phase become coupled and a Gaussian approximation can lead to greater errors.

A shorthand is sometimes utilised for the phase perturbation of the brightest component return by the second brightest. The parameter θ_{f} at some relative frequency r_l, is defined as

$$\theta_{\mathrm{f}} = \arg(\Lambda_{\mathrm{f}}(b, r_l\theta)) = \tan^{-1}(b\sin(r_l\theta), 1 + b\cos(r_l\theta)). \tag{3.73}$$

This can be considered to be the amount by which the phase of a first component return, 1, is perturbed by a second component return, $b\mathrm{e}^{j\theta}$. In general, unless otherwise specified, θ_{f} is notated at a relative frequency of one.

3.5.2 Aliasing

In Section 3.3.1 it was described how aliasing impacts the formation of the sampled correlation waveform $\tilde{s}$. An example of the impact of aliasing perturbations on measured amplitude and phase is presented in Figs. 3.9 and 3.10. This section elaborates on some of the specific properties of aliasing on complex domain range measurements. While many papers have discussed aliasing and aliasing calibration methods, very few have presented any other analysis. An exception is Payne *et al.* (2011), which analyses the impact of aliasing on demodulation contrast, including determination of the mean, maximum and minimum demodulation contrast as a function of illumination modulation duty cycle. We now develop some general properties of aliasing.

Recall Eqn. 3.32

$$\tilde{\xi} = \sum_{k=-\infty}^{\infty} \alpha_{nk+1} \frac{\xi^{nk+1}}{|\xi|^{nk}}, \tag{3.74}$$

which models $\tilde{\xi}$ as a function of ξ. It is demonstrable that for some values of α_k there are multiple possible solutions to the inverse. For a harmonic at a relative frequency $m \in \mathbb{Z}$, the aliased measurement is given by

$$\begin{aligned} \tilde{\xi} &= \xi + \alpha_m \frac{\xi^m}{|\xi|^{m-1}} && (3.75)\\ &= \xi\Lambda_{\mathrm{f}}(b_h, \theta_h) && (3.76) \end{aligned}$$

for the relative intensity $b_h = |\alpha_m|$ of the aliasing coefficient α_m and the relative phase $\theta_h = (m-1)\arg(\xi) + \arg(\alpha_m)$. The phase angle of the aliased measurement is

$$\arg(\tilde{\xi}) = \arg(\xi) + \arg(\Lambda_{\mathrm{f}}(b, (m-1)\arg(\xi) + \arg(\alpha_m))). \tag{3.77}$$

If $\frac{d\arg(\tilde{\xi})}{d\arg(\xi)} \leq 0$ for any value of $\arg(\xi)$, then more than one value of $\arg(\xi)$ maps onto

the same value of $\arg(\tilde{\xi})$ giving multiple solutions for any aliasing correction method. If $\theta_\xi = \arg(\xi)$, the condition can be rewritten as

$$0 \geq \frac{d\tilde{\xi}}{d\theta_\xi} = 1 + \frac{d\Lambda_\mathrm{f}(b, \theta_\xi)}{d\theta_\xi}. \tag{3.78}$$

The second derivative is required to find the minimal value of $\frac{d\Lambda_\mathrm{f}(b,\theta_\xi)}{d\theta_\xi}$, given by

$$\frac{d\Lambda_\mathrm{f}(b, \theta_\xi)}{d\theta_\xi} = \frac{b(\cos(\theta_\xi) + b)}{1 + 2b\cos(\theta_\xi) + b^2} \tag{3.79}$$

$$\frac{d\Lambda_\mathrm{f}(b, \theta_\xi)^2}{d^2\theta_\xi} = \frac{b(b-1)(b+1)\sin(\theta_\xi)}{(1 + 2b\cos(\theta_\xi) + b^2)^2} \tag{3.80}$$

Thus the zeros of the second derivative occur at $\theta_\xi \in \{0, \pi\}$. If $k > 0$, then the minimum value of the first derivative occurs at $\theta_\xi = \pi$, at which

$$\frac{d\Lambda_\mathrm{f}(b, \theta_\xi)}{d\theta_\xi} = \frac{b}{b-1}. \tag{3.81}$$

This gives

$$1 \leq mb, \tag{3.82}$$

as the conditions for more than one solution. If $k < 0$, then the minimum value of the first derivative occurs at $\theta_\xi = 0$, giving

$$1 \leq -mb. \tag{3.83}$$

The generalised condition for multiple solutions can be given as

$$1 \leq |m|b. \tag{3.84}$$

Thus for the third harmonic $b \geq 1/3$ will result in multiple solutions, or for the fourth $b \geq 1/4$. In practice this is an uncommon occurrence, nevertheless theoretically possible. For more complicated cases involving the interaction of more than one aliasing harmonic it is necessary to consider the relative phases of the harmonics in order to produce a tight bound. For a precise bound, inspection is a good option, alternatively an approximate bound can be calculated below which a single solution is guaranteed and above which multiple solutions could plausibly occur. This bound can be found by considering the worst case scenario in which all the harmonics align,

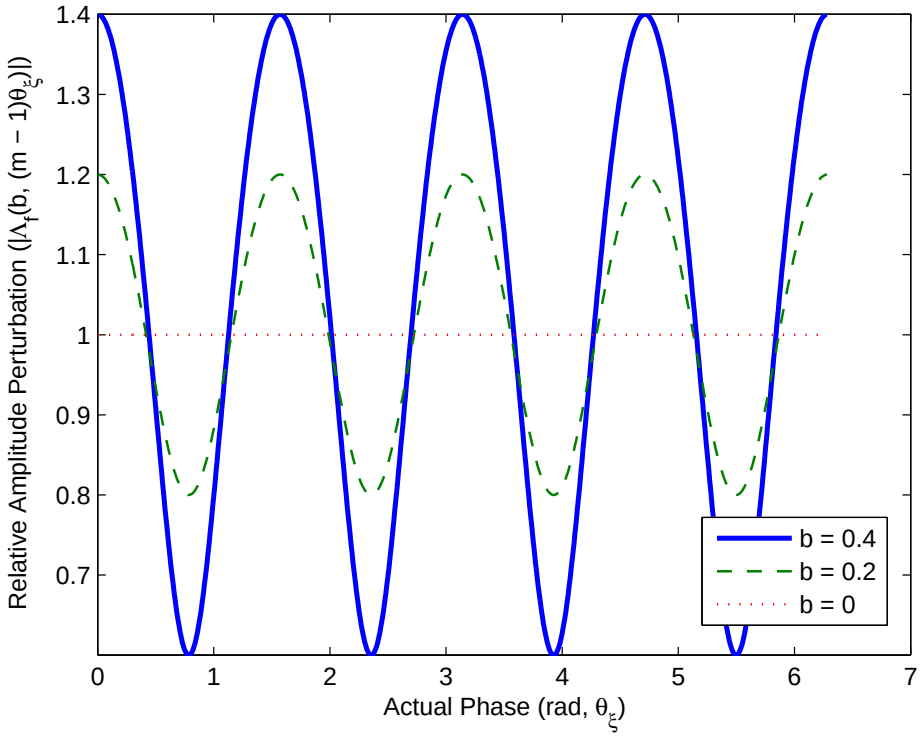

Figure 3.9: The impact of aliasing on amplitude measurements. As the relative amplitude of the aliasing harmonic (b) increases, the perturbation increases. Simulated for the third harmonic aliasing onto the negative fundamental.

giving multiple possible solutions under the conditions

$$0 \geq 1 + \sum_{k=-\infty}^{\infty} (k-1)\frac{|\alpha_k|}{|\alpha_k| - 1}. \tag{3.85}$$

We now analyse some of the more subtle systematic impacts of aliasing. Given the assumption of a uniform prior over $\theta_h \in [\theta_{\min}, \theta_{\max}]$, the expected complex domain range measurement over the specified domain is given by

$$\mathbb{E}\left[\Lambda(b_h, \theta_h)|\theta_h \sim \mathcal{U}(\theta_{\min}, \theta_{\max})\right] = \int_{\theta_{\min}}^{\theta_{\max}} \frac{\Lambda_{\mathrm{f}}(b_h, \theta_h)}{\theta_{\max} - \theta_{\min}} d\theta_h = \left[\frac{\theta_h - jbe^{j\theta_h}}{\theta_{\max} - \theta_{\min}}\right]_{\theta_{\min}}^{\theta_{\max}} \tag{3.86}$$

Integrating over 2π gives

$$\mathbb{E}\left[\Lambda(b_h, \theta_h)|\theta_h \sim \mathcal{U}(-\pi, \pi)\right] = \int_{0}^{2\pi} \frac{\Lambda_{\mathrm{f}}(b_h, \theta_h)}{2\pi} d\theta_h = 1. \tag{3.87}$$

In other words, on average in the complex domain circularly symmetric perturba-

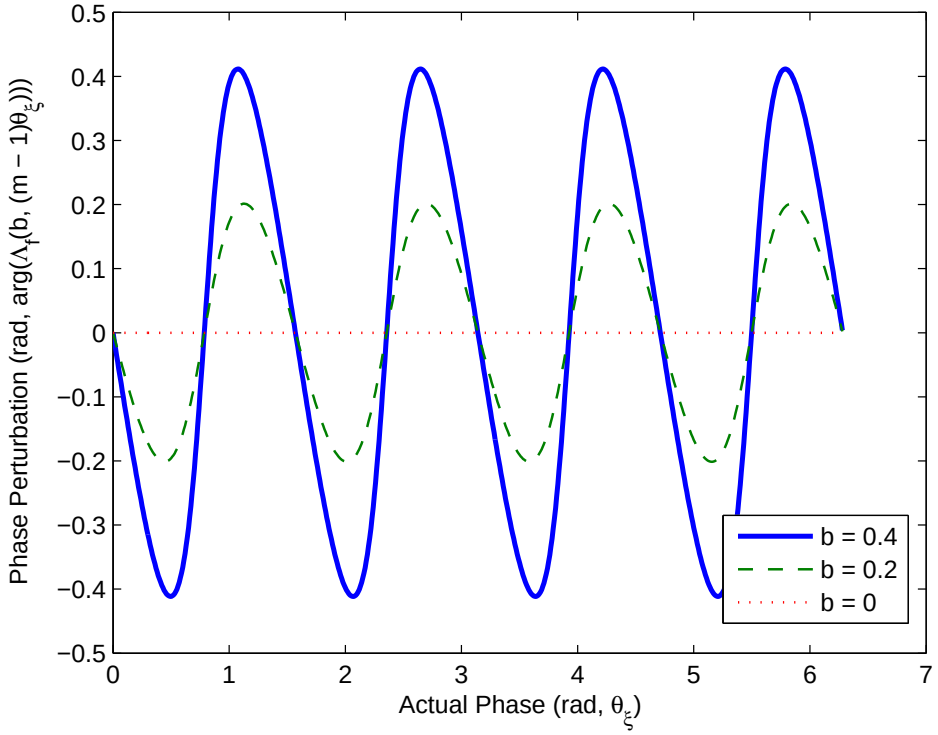

(a) Phase Perturbation

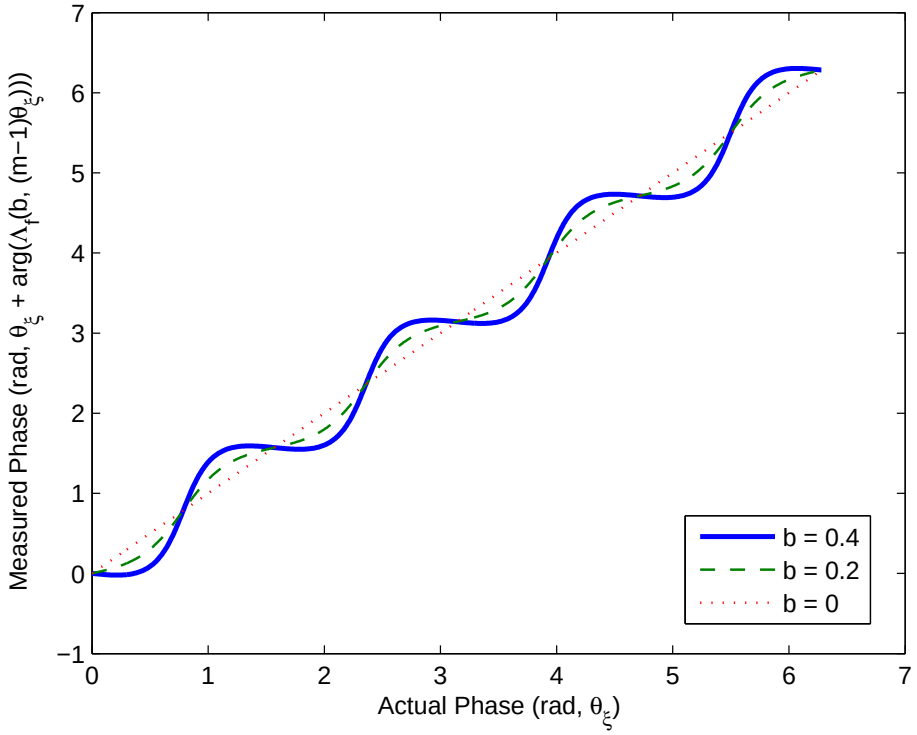

(b) Actual Measured Phase

Figure 3.10: The impact of aliasing on phase measurements. As the relative amplitude of the aliasing harmonic (b) increases, the perturbation increases. Simulated for the third harmonic aliasing onto the negative fundamental.

tions cancel out. This also holds for the expected phase using symmetry, viz

$$\begin{aligned}
\mathbb{E}[\arg(\Lambda_{\mathrm{f}}(b_h,\theta_h))|\theta_h \sim \mathcal{U}(-\pi,\pi)] &= \int_{-\pi}^{\pi} \frac{\arg(\Lambda_{\mathrm{f}}(b_h,\theta_h))}{2\pi} d\theta_h \\
&= \int_{-\pi}^{0} \frac{\arg(\Lambda_{\mathrm{f}}(b_h,\theta_h))}{2\pi} d\theta_h \\
&\quad + \int_{0}^{\pi} \frac{\arg(\Lambda_{\mathrm{f}}(b_h,\theta_h))}{2\pi} d\theta_h \\
&= \int_{-\pi}^{0} \frac{\arg(\Lambda_{\mathrm{f}}(b_h,\theta_h))}{2\pi} d\theta_h \\
&\quad - \int_{-\pi}^{0} \frac{\arg(\Lambda_{\mathrm{f}}(b_h,\theta_h))}{2\pi} d\theta_h \\
&= 0. \qquad (3.88)
\end{aligned}$$

However, this is *not* true for amplitude:

$$\begin{aligned}
\mathbb{E}[|\Lambda(b_h,\theta_h)|\,|\theta_h \sim \mathcal{U}(-\pi,\pi)] &= \int_{-\pi}^{\pi} \frac{|\Lambda_{\mathrm{f}}(b_h,\theta_h)|}{2\pi}\, d\theta_h \\
&= \int_{-\pi}^{\pi} \frac{\sqrt{1+b_h^2+2b_h\cos(\theta_h)}}{2\pi}\, d\theta_h \\
&= \frac{\sqrt{1+b^2}}{2\pi} \int_{-\pi}^{\pi} \sqrt{1+\beta\cos(\theta_h)}\, d\theta_h \\
&= \frac{4\sqrt{(1+b^2)(\beta+1)}}{2\pi} E\left(\frac{2\beta}{\beta+1}\right) \\
&= a_{\mathrm{p}}(b_h), \qquad (3.89)
\end{aligned}$$

where $E(k)$ is the complete elliptic integral of the second kind, $\beta = \frac{2b_h}{1+b_h^2}$ and we have defined a function $a_{\mathrm{p}}(b_h)$ which returns the mean amplitude perturbation for a relative intensity b. The consequences of this revelation include that any circularly symmetric noise distribution results in a systematic increase in the measured amplitude.

The relationship between mean amplitude and b_h can be viewed in two very different ways; either through the impact on demodulation contrast or on raw amplitude. Payne *et al.* (2011) discussed this systematic amplitude bias through its impact on demodulation contrast, which is the ratio between the amplitude of the correlation waveform frequency being measured and the mean offset due to modulated light (ignoring any ambient contribution). If there are only two harmonics, demodulation contrast, d_c can be calculated by

$$d_c = \frac{a_{\mathrm{p}}(b_h)}{1+b_h}. \qquad (3.90)$$

As demodulation contrast decreases, there is an increase in shot noise relative to the amplitude of the frequency being measured, so SNR decreases. This is one reason why all averaging of range measurements needs to be done in the complex domain, not of individual phase or amplitude measurements.

3.5.3 Systematic Errors Due to Random Noise

The previous section discussed how aliasing could induce a systematic error in the expected amplitude. We now use these results to analyse the impact of Gaussian noise on the expected amplitude and phase.

Given a measurement, $\tilde{\xi}$, and complex Gaussian distributed noise, $\epsilon_\xi \sim \mathcal{N}_\mathbb{C}(0, \Sigma_\xi)$, the expected amplitude of $\tilde{\xi} + \epsilon_\xi$ is

$$\mathbb{E}[|\tilde{\xi} + \epsilon_\xi| \,|\epsilon_\xi \sim \mathcal{N}_\mathbb{C}(0, \Sigma_\xi)] = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} |\xi + R + jI| \, f_\mathcal{N}\left(\binom{R}{I} | \binom{0}{0}, \Sigma_\xi\right) \, dI \, dR, \quad (3.91)$$

where $f_\mathcal{N}(x|\mu, \Sigma)$ is the probability density function for the Gaussian distribution parametised by mean μ and covariance Σ. Given that this equation is *somewhat* difficult to integrate analytically, we must either use a numerical approximation or make generalisations about the behaviour of the equation. In the circularly symmetric case using Eqn. 3.89, we can model the amplitude perturbation induced by the noise distribution as

$$\begin{aligned} \mathbb{E}\left[\frac{|\tilde{\xi} + \epsilon_\xi|}{|\xi|}\middle| \epsilon_\xi \sim \mathcal{N}_\mathbb{C}(0, \Sigma_\xi)\right] &= \int_0^\infty a_\mathrm{p}\left(\frac{x}{|\xi|}\right) f_\mathcal{R}(x|\sigma_\xi) dx & (3.92) \\ &= a_\mathrm{q}(\sigma_\xi^2) & (3.93) \end{aligned}$$

where $f_\mathcal{R}(x, \sigma^2)$ is the probability density function for a Rayleigh distribution and we have defined a new function $a_\mathrm{q}(\sigma_\xi^2)$.

Given that $a_\mathrm{p}(0) = 1$ and since Fig. 3.11 illustrates that

$$a_\mathrm{p}(b_h + \Delta b_h) > a_\mathrm{p}(b_h) \quad (3.94)$$

for $\Delta b_h \in [0, 1]$ (i.e. that $a_\mathrm{p}(b_h)$ is a strictly increasing function), it is possible to make some new observations about averaging before and after the transformation into polar form. Averaging $M > 1$ measurements of $\tilde{\xi}$, each subject to a peturbation from noise, $\epsilon_\xi[i] \sim \mathcal{N}_\mathbb{C}(0, \Sigma_\xi)$, gives

$$\frac{1}{M}\sum_{i=0}^{M-1} \xi + \epsilon_\xi[i] = \xi + \epsilon_\mathrm{M}, \quad (3.95)$$

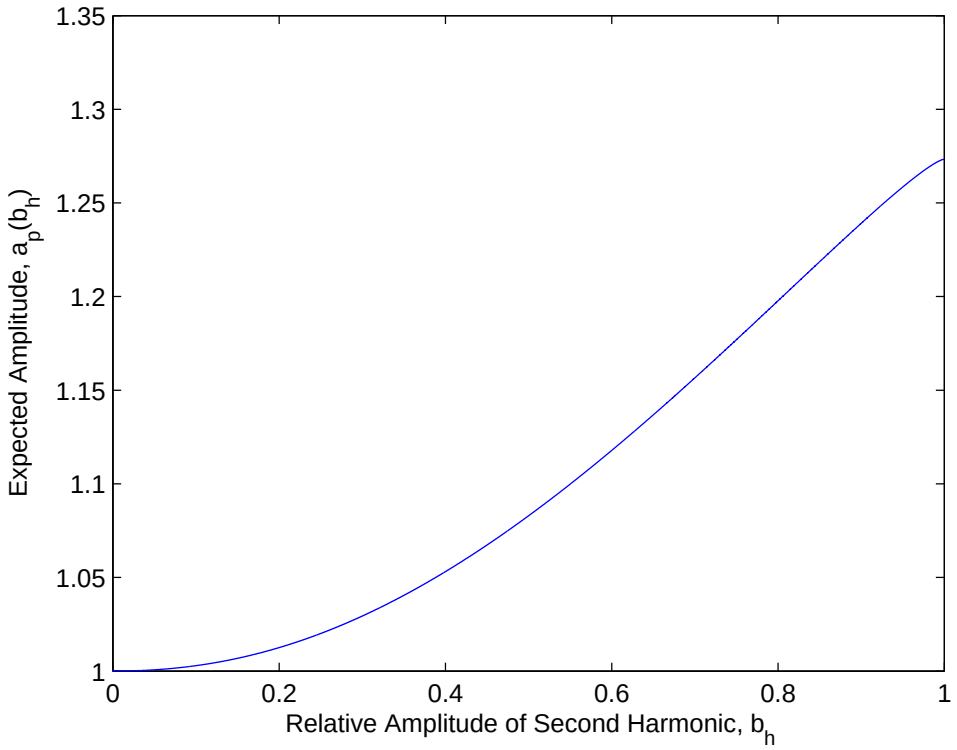

Figure 3.11: The expected measured amplitude in the case of a circularly symmetric perturbation of relative amplitude, b_h, given a uniform prior over relative phase, θ_h.

where $\epsilon_{\mathrm{M}} \sim \mathcal{N}_{\mathbb{C}}(0, \Sigma_\xi^2/M)$ is the averaged noise. The expected amplitude when calculating amplitude pre- and post- averaging is given by

$$\mathbb{E}\left[\frac{1}{M}\sum_{i=0}^{M-1}|\xi + \epsilon_\xi[i]|\,\middle|\,\epsilon_\xi[i] \sim \mathcal{N}_{\mathbb{C}}(0,\Sigma_\xi)\right] = |\xi| a_{\mathrm{q}}(\sigma_\xi^2) \tag{3.96}$$

$$\mathbb{E}\left[\left|\frac{1}{M}\sum_{i=0}^{M-1}\xi + \epsilon_\xi[i]\right|\,\middle|\,\epsilon_\xi[i] \sim \mathcal{N}_{\mathbb{C}}(0,\Sigma_\xi)\right] = |\xi| a_{\mathrm{q}}(\sigma_\xi^2/M). \tag{3.97}$$

From eqn. 3.94,

$$a_{\mathrm{q}}(\sigma_\xi^2) > a_{\mathrm{q}}(\sigma_\xi^2/M), \tag{3.98}$$

hence, averaging before calculation of amplitude results in a reduction in the systematic amplitude perturbation. Although not proven here, the same perturbation and averaging relationships hold for non-circularly symmetric covariance matrices.

While Eqn. 3.88 showed that the expected phase perturbation for circularly symmetric noise sources is zero, this does not hold for non-circularly symmetric noise sources.

In order to simplify the problem, let R_ξ be a rotation and scaling matrix such

that $R_\xi \tilde{\xi} = (1,0)^T$. Thus

$$R_\xi = \frac{1}{|\tilde{\xi}|}\begin{pmatrix}\cos(\arg(\tilde{\xi})) & -\sin(\arg(\tilde{\xi}))\\ \sin(\arg(\tilde{\xi})) & \cos(\arg(\tilde{\xi}))\end{pmatrix}. \tag{3.99}$$

We now perform an eigen decomposition of $R_\xi \Sigma_\xi R_\xi^{-1}$, ie. find the eigenvalues λ_0 and λ_1 for eigenvectors x_0 and x_1, such that $\lambda_0 > \lambda_1$:

$$(R_\xi \Sigma_\xi R_\xi^{-1} - \lambda_i \mathbf{I}_2)x_i = 0. \tag{3.100}$$

Since Σ_ξ is symmetric, x_0 and x_1 are orthonormal vectors with x_0 pointing in the direction of maximum variance. This is essentially performing principal components analysis (Bishop, 2006) on the noise.

There are three distinct special cases for the analysis of phase perturbation. In the first two cases, $x_0 = (\pm 1, 0)^T$ and $x_0 = (0, \pm 1)^T$, there is no systematic phase perturbation because the axis of maximum variance is either aligned with, or perpendicular to, the direction of $\tilde{\xi}$. We now analyse the most interesting case, namely when the axis of maximum variance is offset.

The expected phase perturbation is given by

$$\begin{aligned}\mathbb{E}\left[\arg\left(\frac{\xi+\epsilon_\xi}{\xi}\right)\middle|\epsilon_\xi[i] \sim \mathcal{N}_\mathbb{C}(0,\Sigma_\xi)\right] &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\arg(1+(1,j)\cdot(x_0y_0+x_1y_1))\\ &\qquad f_\mathcal{N}\left(x_0y_0+x_1y_1\middle|\begin{pmatrix}0\\0\end{pmatrix},\Sigma_\xi\right)dy_0\,dy_1\\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\arg(1+(1,j)\cdot(x_0y_0+x_1y_1))\\ &\qquad f_\mathcal{N}(y_0,0,\lambda_0)f_\mathcal{N}(y_1|0,\lambda_1)dy_0dy_1\\ &= \int_{-\infty}^{\infty} l_\mathrm{j}(y_1)f_\mathcal{N}(y_1|0,\lambda_1)\,dy_1 \end{aligned}\tag{3.101}$$

where the bivariate Gaussian distribution has been decomposed into the product of two univariate Gaussian distributions and a new function $l_\mathrm{j}(y_1)$ has been defined to model the integration along the axis of maximum variance. This is given by

$$\begin{aligned} l_\mathrm{j}(y_1) &= \int_{-\infty}^{\infty}\arg(1+(1,j)\cdot(x_0y_0+x_1y_1))f_\mathcal{N}(y_0|0,\lambda_0)dy_0 &(3.102)\\ &= \int_{-\infty}^{\infty} l_\mathrm{k}(y_0,y_1)f_\mathcal{N}(y_0|0,\lambda_0)dy_0 &(3.103)\end{aligned}$$

where the phase calculation function $l_\mathrm{k}(y_0, y_1)$ is defined by

$$l_\mathrm{k}(y_0,y_1) = \arg(1+(1,j)\cdot(x_0y_0+x_1y_1)), \tag{3.104}$$

which is the phase angle of a line in 2D space. We now analyse the stationary points in order to determine the point at which maximum phase perturbation occurs, which can be analysed in order to determine the direction of any systematic perturbation. The first derivative of $l_{\mathrm{k}}(y_0, y_1)$ relative to y_0 is given by

$$\frac{dl_{\mathrm{k}}(y_0, y_1)}{dy_0} = \frac{x_0^T \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} c_k}{c_k^T c_k + 2y_0 x_0^T c_k + y_0^2 x_0^T x_0} \tag{3.105}$$

where c_k is given by

$$c_k = \begin{pmatrix} 1 \\ 0 \end{pmatrix} + y_1 x_1. \tag{3.106}$$

The stationary points of $\frac{dl_{\mathrm{k}}(y_0,y_1)}{dy_0}$ can be found by setting the second derivative to 0. The second derivative is

$$\frac{dl_{\mathrm{k}}^2(y_0, y_1)}{d^2 y_0} = \frac{x_0^T \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} c_k (x_0^T c_k + y_0 x_0^T x_0)}{(c_k^T c_k + 2y_0 x_0^T c_k + y_0^2 x_0^T x_0)^2}. \tag{3.107}$$

Setting the derivative to zero gives

$$\frac{dl_{\mathrm{k}}^2(y_0, y_1)}{d^2 y_0} = 0 \Rightarrow y_0 = -\frac{x_0^T c_k}{x_0^T x_0} = y_{\mathrm{m}}, \tag{3.108}$$

where y_{m} is the value of y_0 at the stationary point. The stationary point can be shown to be a maximum by taking the third derivative. Since $x_0 \cdot x_1 = 0$, y_1 has no influence on the value of y_{m}; it solely depends on the value of x_0. We now reason about $l_{\mathrm{k}}(y0, y1)$ using y_{m}.

The point $c_k(1, j) \cdot (y_{\mathrm{m}} x_0)$ corresponds to the point on the line which is at minimal distance from the origin, thus

$$(c_k + y_0 x_0)^T \cdot x_0 = 0. \tag{3.109}$$

From simple geometry, the second derivative, $|\frac{dl_{\mathrm{k}}^2(y_0,y_1)}{dy_0^2}|$, is an odd function around the point y_{m}, ie.

$$\frac{dl_{\mathrm{k}}^2(z - y_{\mathrm{m}}, y_1)}{d^2 z} = -\frac{dl_{\mathrm{k}}^2(y_{\mathrm{m}} - z_1 y_1)}{d^2 z} \tag{3.110}$$

for a free variable $z \in \mathbb{R}$.

Given that there is only one stationary point, $|\Delta l_{\mathrm{k}}(y_0, y_1)|$ decreases monotonically either side of the point y_{m}. This implies that

$$\mathrm{sgn}(l_{\mathrm{k}}(y_0, y_1)) = \mathrm{sgn}(y_{\mathrm{m}}). \tag{3.111}$$

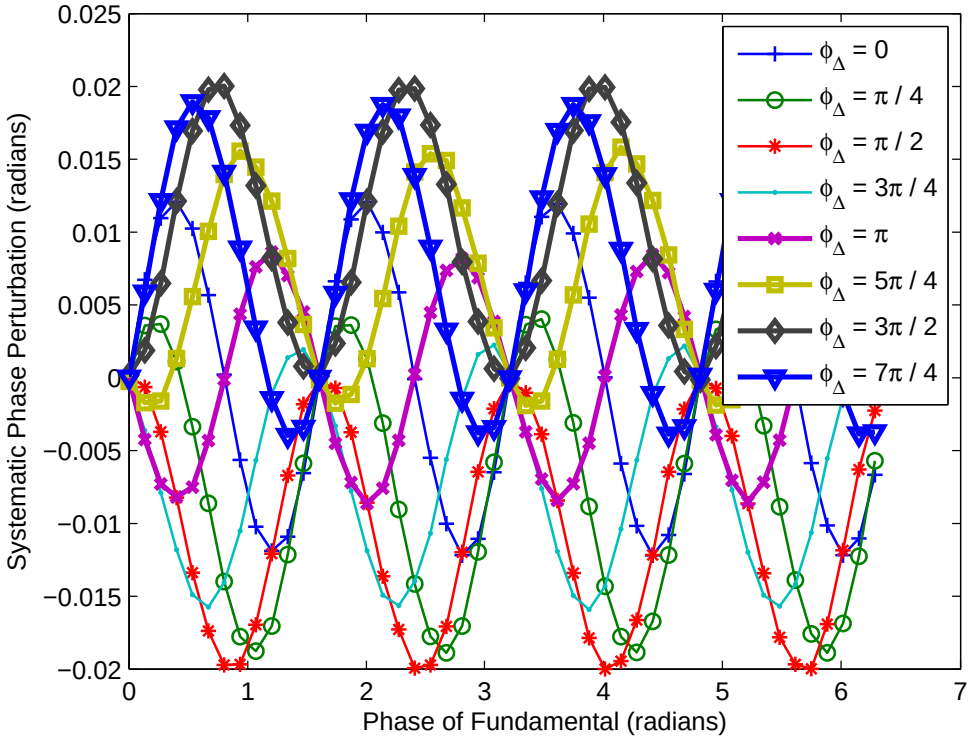

(a) By Phase of Fundamental

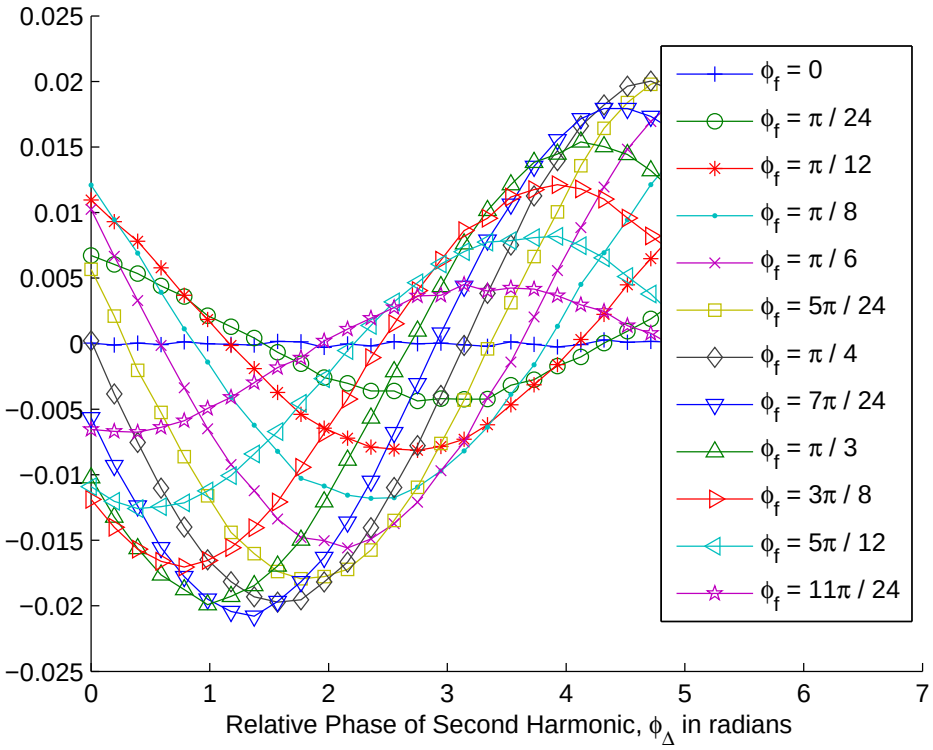

(b) By Relative Phase of Second Harmonic

Figure 3.12: Simulation of the systematic perturbation introduced by non-circularly symmetric error in the complex domain range measurement using four phase steps.

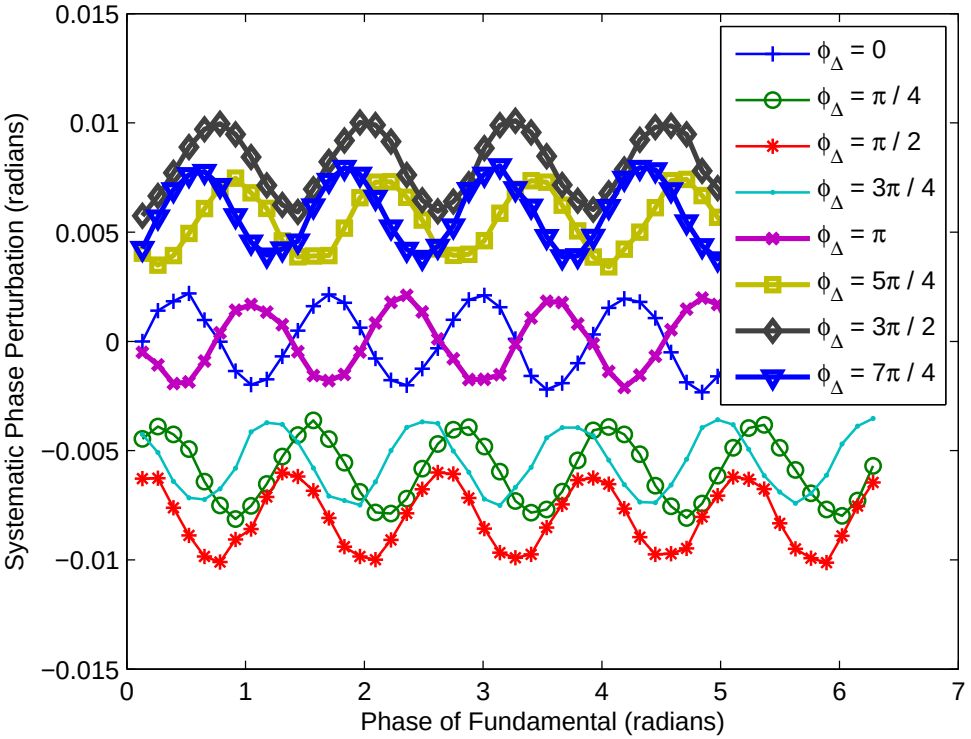

(a) By Phase of Fundamental

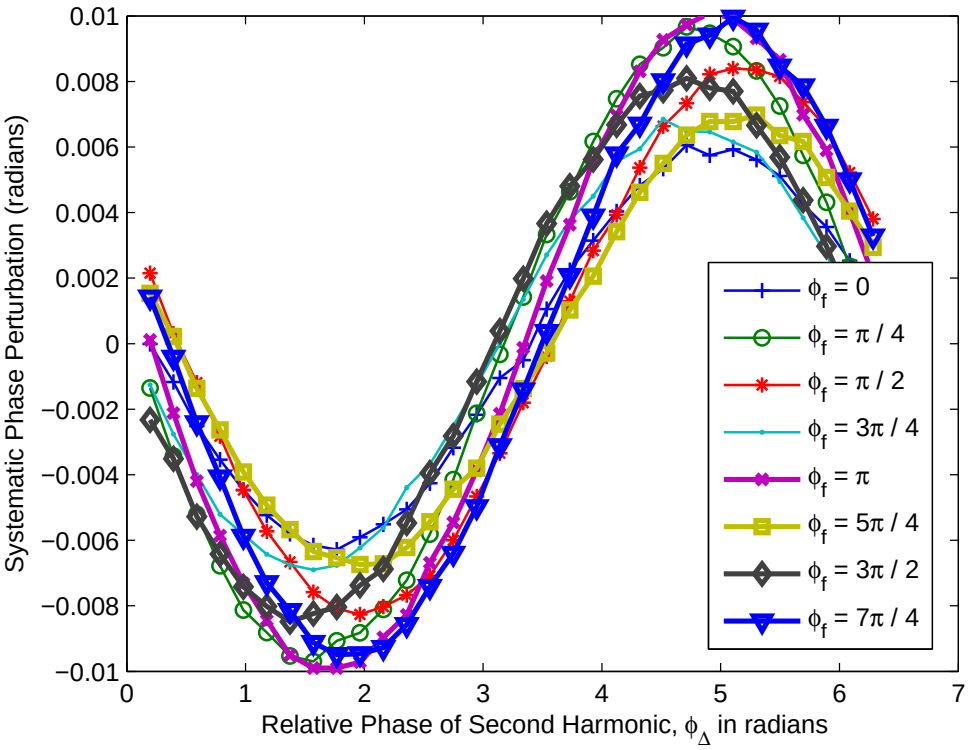

(b) By Relative Phase of Second Harmonic

Figure 3.13: Simulation of the systematic perturbation introduced by non-circularly symmetric error in the complex domain range measurement using five phase steps.

Propagating this through Eqn. 3.103 gives

$$\mathrm{sgn}(l_{\mathrm{j}}(y_1)) = \mathrm{sgn}(y_{\mathrm{m}}). \tag{3.112}$$

Since $\mathrm{sgn}(y_{\mathrm{m}}) = 0$ only occurs if the axis of maximum variance is aligned with, or perpendicular to $\tilde{\xi}$, we conclude that if the axis of maximum variance is offset, the expected phase is perturbed in the direction given by

$$-\mathrm{sgn}\left(x_0^T \cdot \begin{pmatrix}1\\0\end{pmatrix}\right). \tag{3.113}$$

Figs. 3.12 and 3.13 give the systematic phase perturbation for a simulation of complex domain range measurements subject to a non-circularly symmetric noise distribution. The simulation assumes that the correlation waveform is given by

$$\tilde{s}[i] = \alpha\left(1 + \cos\left(\phi_{\mathrm{f}} + \frac{2\pi}{n}i\right)\right) + \beta\left(1 + \cos\left(2\left(\phi_{\mathrm{f}} + \frac{2\pi}{n}i\right) + \phi_{\Delta}\right)\right), \tag{3.114}$$

where $\alpha = \beta = 100$ photons, and the only noise source is assumed to be photon shot noise. The choice of such a high relative amplitude for the second harmonic somewhat exaggerates the results, but is useful for demonstrative purposes. Fig. 3.13 is indicative of the general behaviour at six or more phase steps, whereas Fig. 3.12 is unique due to the second harmonic being defined solely by its amplitude, with no intrinsic phase information. For greater than four phase steps, there is typically a relatively large fixed offset in the phase measurements, which is a function of the relative phase of the second harmonic (ϕ_{Δ}) with a relatively small oscillation around this value as a function of the phase of the fundamental (θ_{f}).

As with aliasing, it appears that a fixed lookup table model ought to be able to remove the systematic perturbations, however this would require the inclusion of intensity as parameter in addition to phase.

3.5.4 Uneven Phase Sampling

In this section we develop a simple model for the impact of uneven phase sampling and present the results of a simulation demonstrating the effects. Previous models for lidar measurement formation have assumed that the phase samples are equally spaced, which is not always a realistic assumption due to electronic limitations. One likely reason why it has not been discussed before is that it manifests in exactly the same manner as aliasing, as a phase linearity error; as a result it can also be compensated for in exactly the same manner as aliasing, through a lookup table or

B-Spline model. Given a known reference waveform, ψ, and a sampling function, P_{q}, we can model the calibrated measurement response as

$$Q_{\mathrm{r}}(\Delta\theta) = \frac{\tilde{\xi}(\Delta\theta)}{\alpha \mathrm{e}^{j(\Delta\theta)}} \tag{3.115}$$

$$= \rho_{\mathrm{x}} \mathrm{e}^{-j\Delta\theta}(\psi * P_{\mathrm{q}})(\Delta\theta), \tag{3.116}$$

where $\tilde{\xi}(\Delta\theta)$ is the value of $\tilde{\xi}$ expressed as a function of the phase shift induced by the backscattering source in the illumination/correlation waveform, $\Delta\theta$, where ρ_{x} calibrates for the phase and amplitude of the fundamental of the correlation waveform. In the ideal continuous case the sampling function corresponds to a continuous Fourier transform of the negative fundamental frequency, giving

$$P_{\mathrm{q}}(\Delta\theta) = \mathrm{e}^{j\Delta\theta}, \tag{3.117}$$

which implies that

$$Q_{\mathrm{r}}(\Delta\theta) = 1. \tag{3.118}$$

However, the general case is not necessarily so ideal. Whereas thus far we have made the assumption that the phase steps are equally spaced, that is not necessarily so. One potential cause in practical implementations is coupling between the power supplies of the illumination and sensor modulation; even the slightest bleed from one signal into the other can result in a subtle phase shift that results in a systematic phase and amplitude error. Making the assumption that the processing mistakenly considers the the phase steps to be evenly spaced, we model the sampling function as

$$P_{\mathrm{q}}(\Delta\theta) = \sum_{i=0}^{M-1} \mathrm{e}^{2\pi j\theta/M} \delta(\Delta\theta - \gamma_i), \tag{3.119}$$

where γ_i is actual phase of the ith phase step and we define $\Gamma = \{\gamma_0, \gamma_1, ...\gamma_{M-1}\}$ as the set of all phase steps.

In order to illustrate the impact of irregular phase steps, we implemented a discrete simulation of Eqns. 3.116 and 3.119 using an 840 element vector to represent the sampling function. This allows phase shifts in each sample to be represented in a quantised, Kronecker delta form. While the exact nature of any perturbation in the phase steps is situation dependant, it is necessary to assume a specific model for the purposes of simulation; this does not necessarily reflect the perturbations that

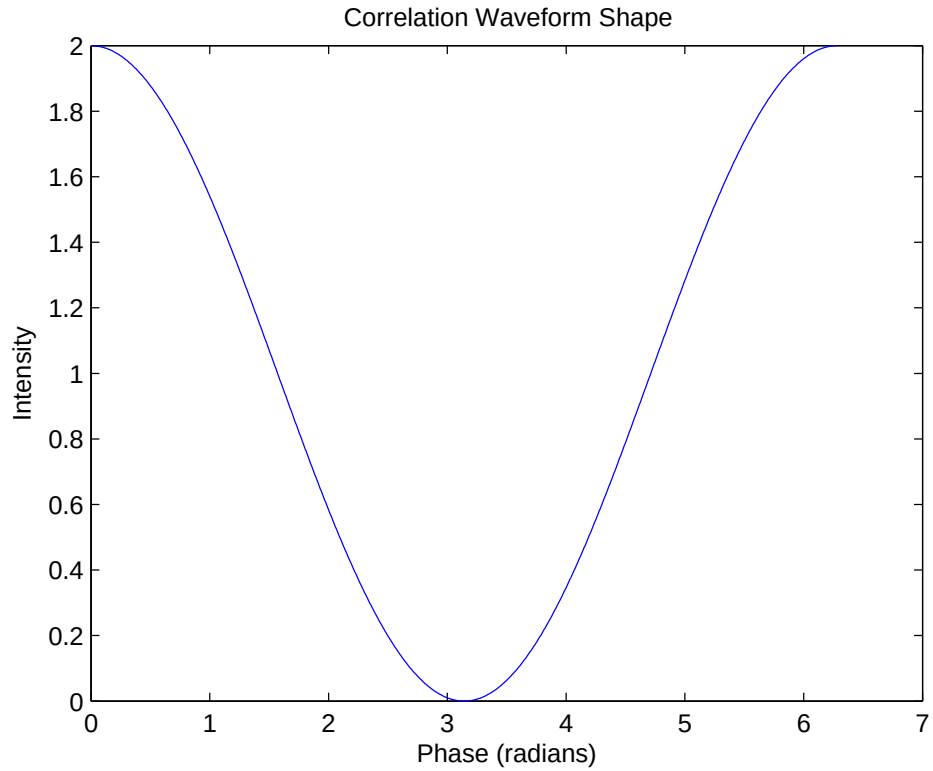

(a) Sinusoidal

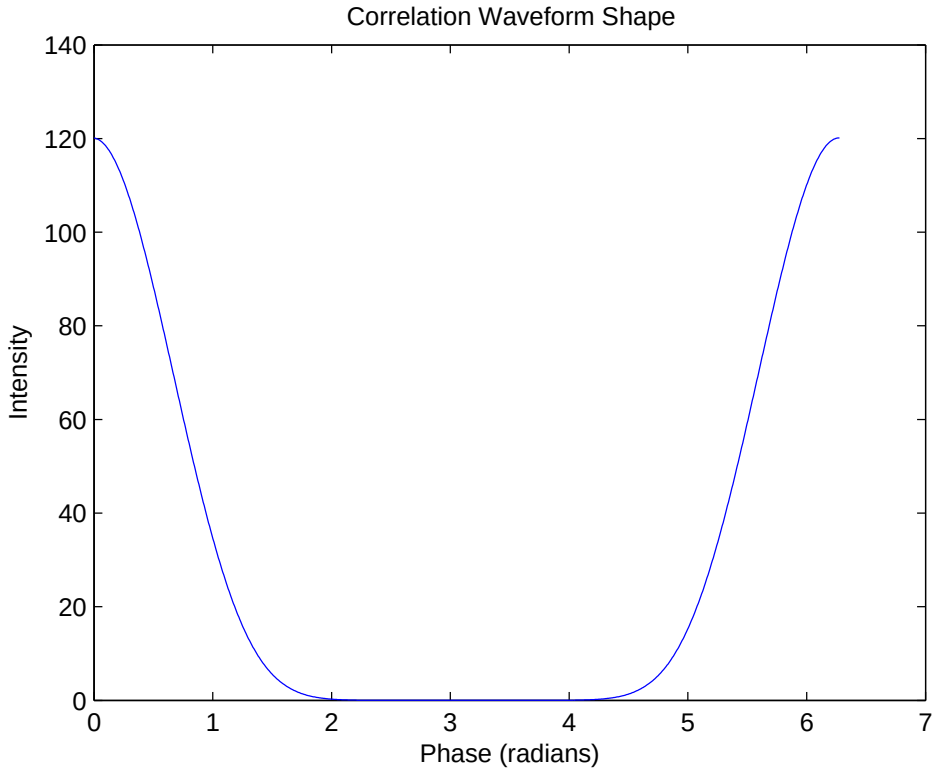

(b) Non-sinusoidal

Figure 3.14: Correlation waveform shapes for the uneven phase sampling experiment.

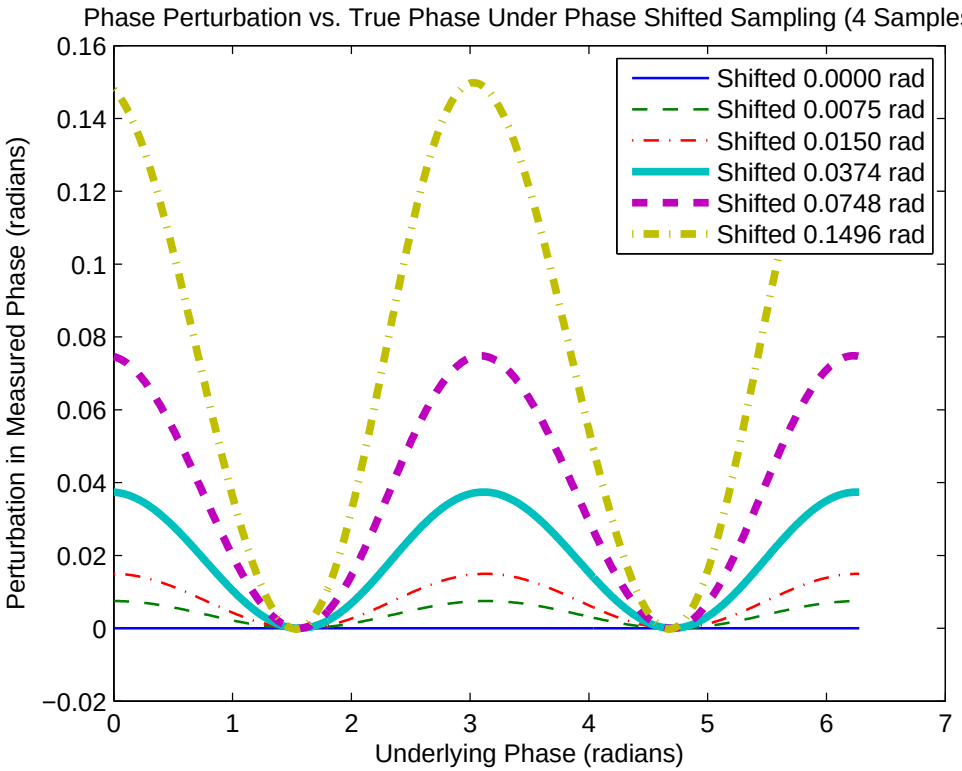

(a) Four Sample Phase Response

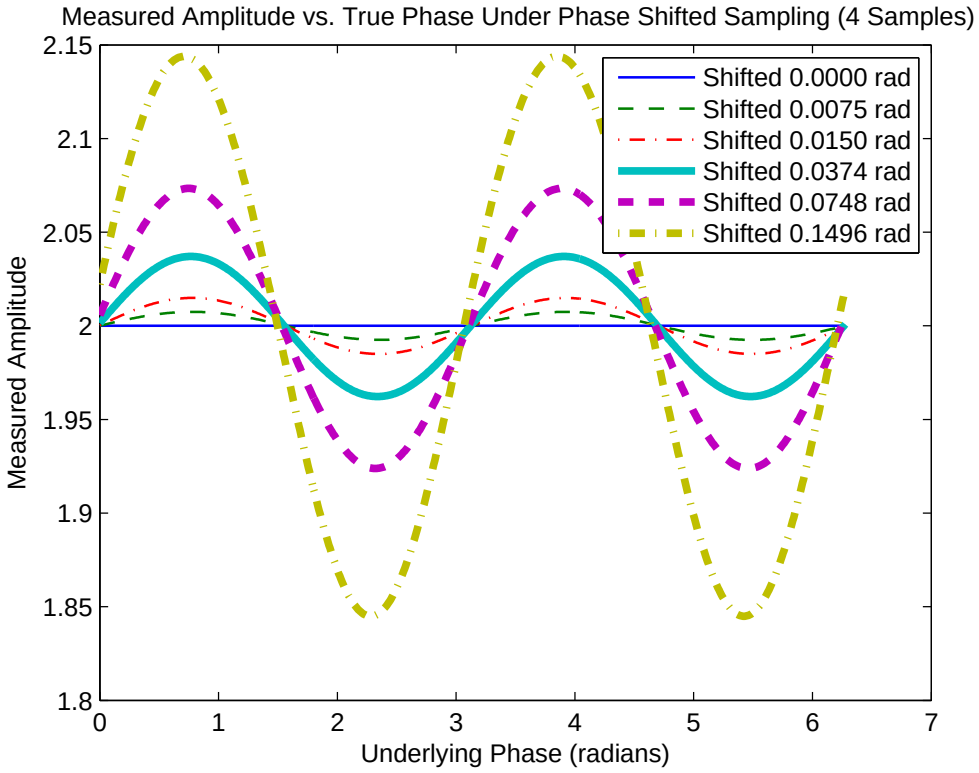

(b) Four Sample Amplitude Response

Figure 3.15: Uneven phase sampling experiment: phase and amplitude response for a sinusoidal correlation waveform as a function of the the phase shift.

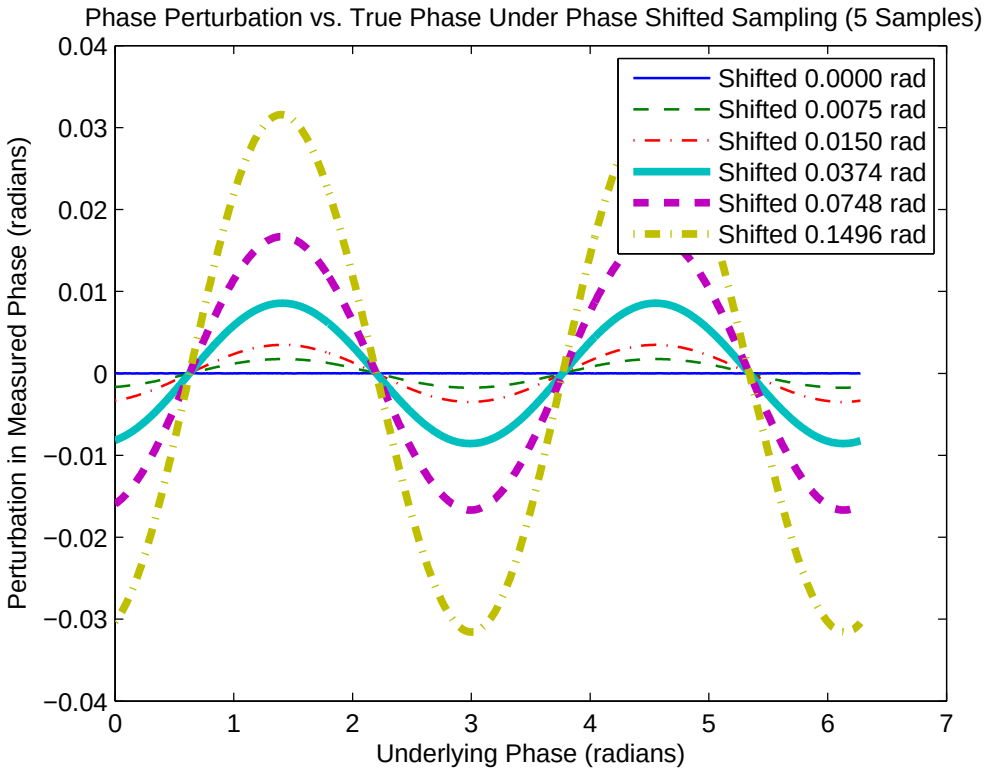

(a) Five Sample Phase Response

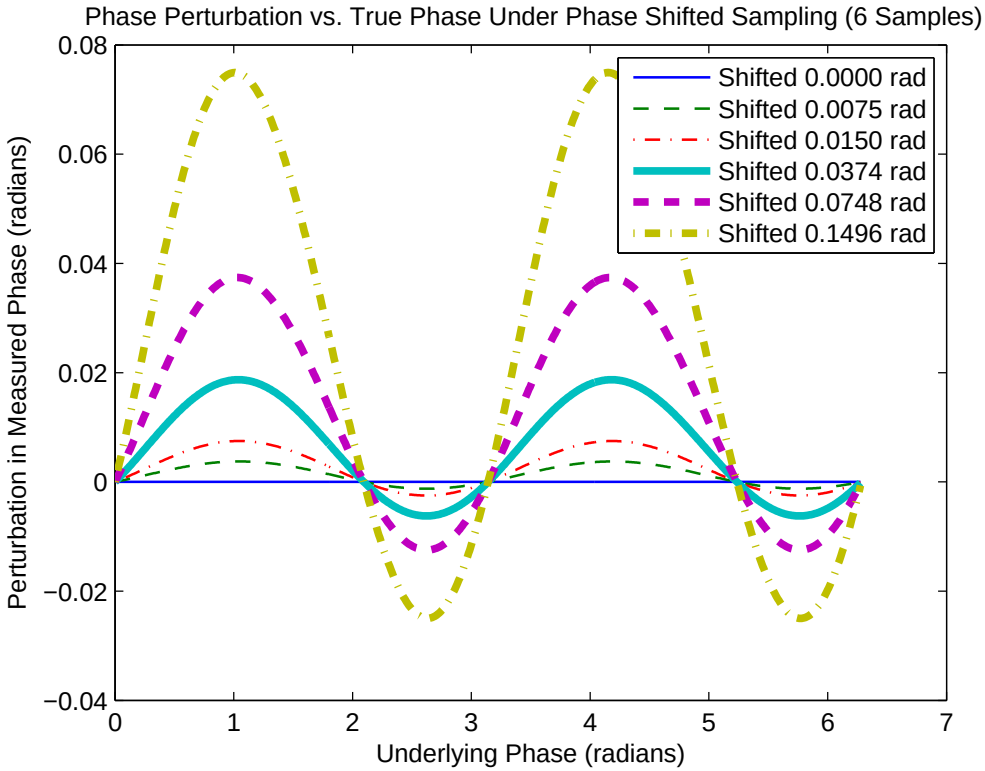

(b) Six Sample Phase Response

Figure 3.16: Uneven phase sampling experiment: phase response for a sinusoidal correlation waveform as a function of the the phase shift.

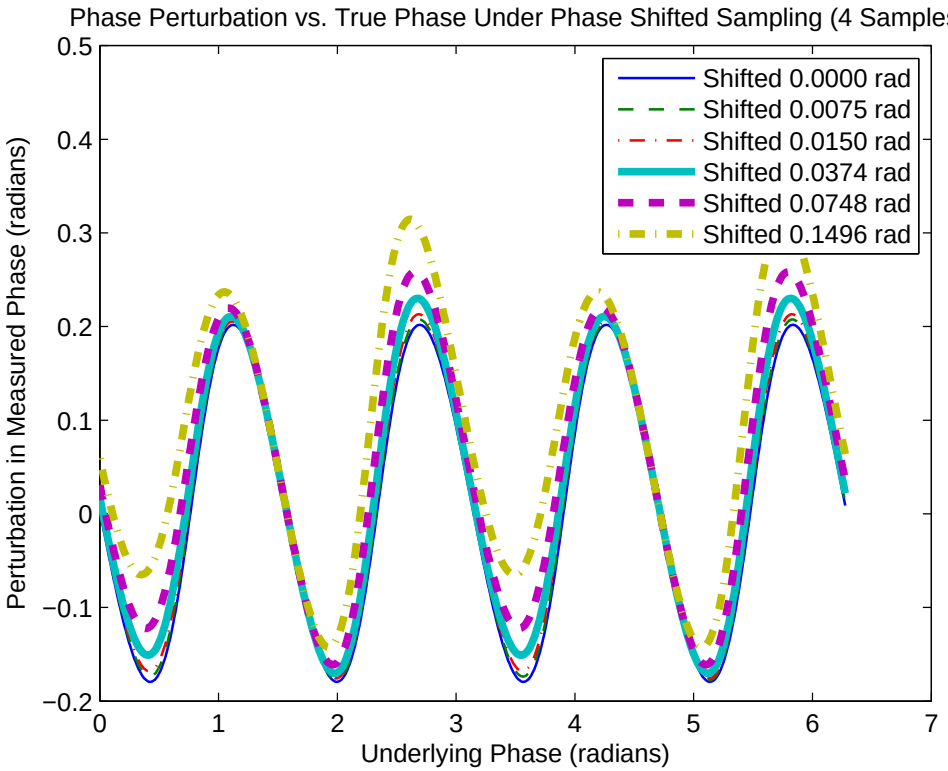

(a) Four Sample Phase Response

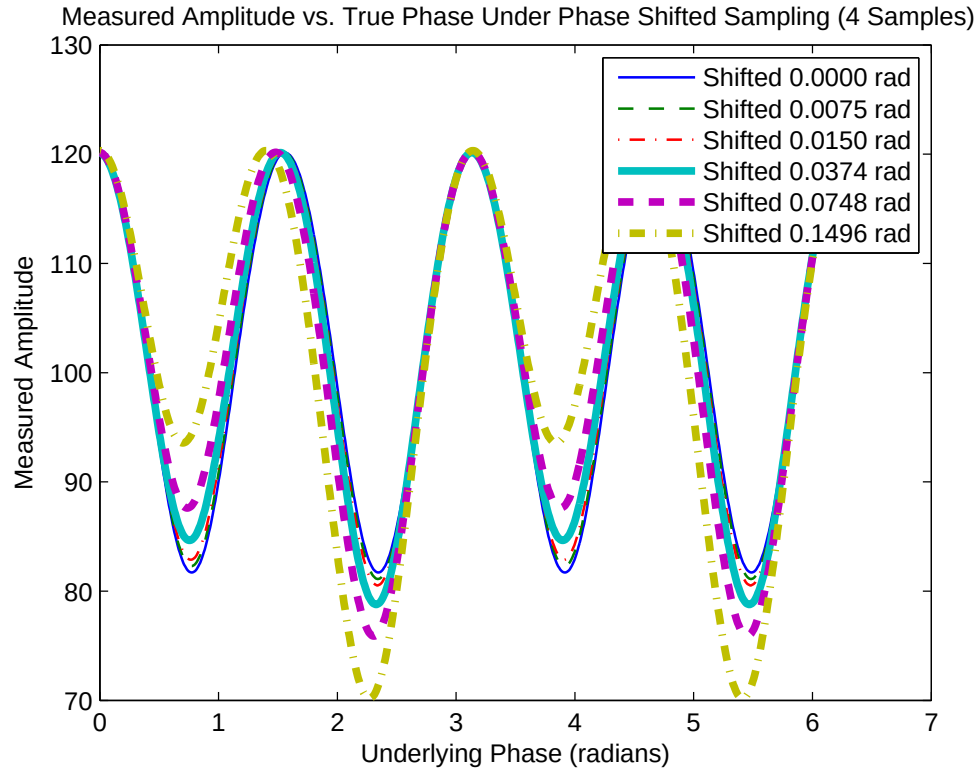

(b) Four Sample Amplitude Response

Figure 3.17: Uneven phase sampling experiment: phase and amplitude response for a non-sinusoidal correlation waveform as a function of the the phase shift.

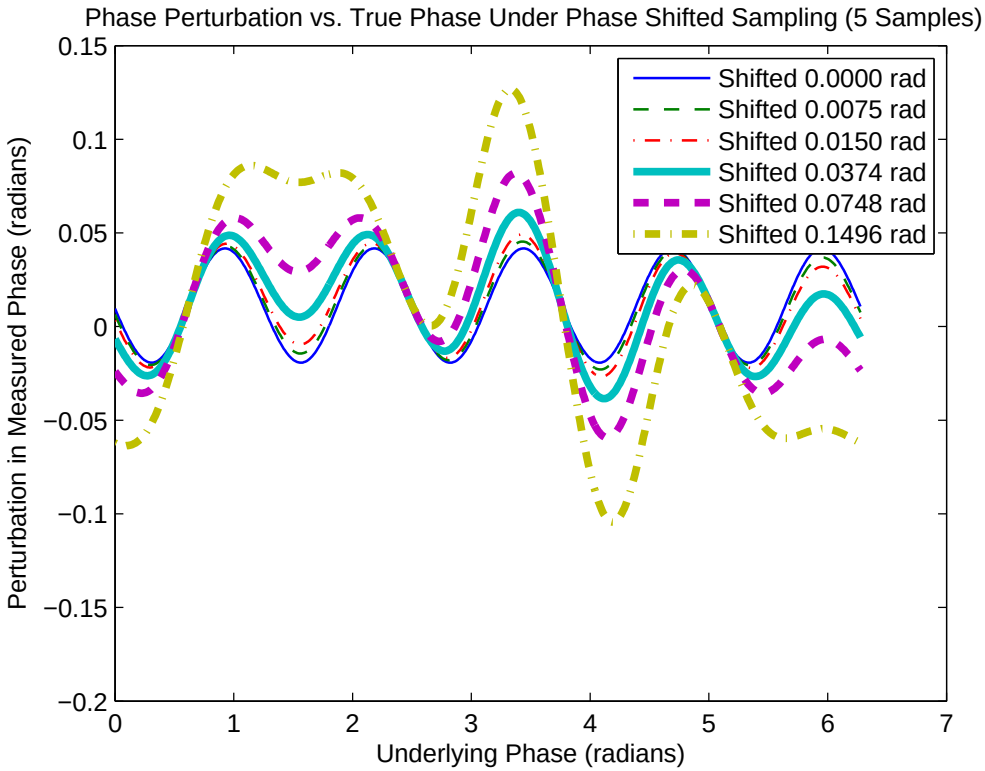

(a) Five Sample Phase Response

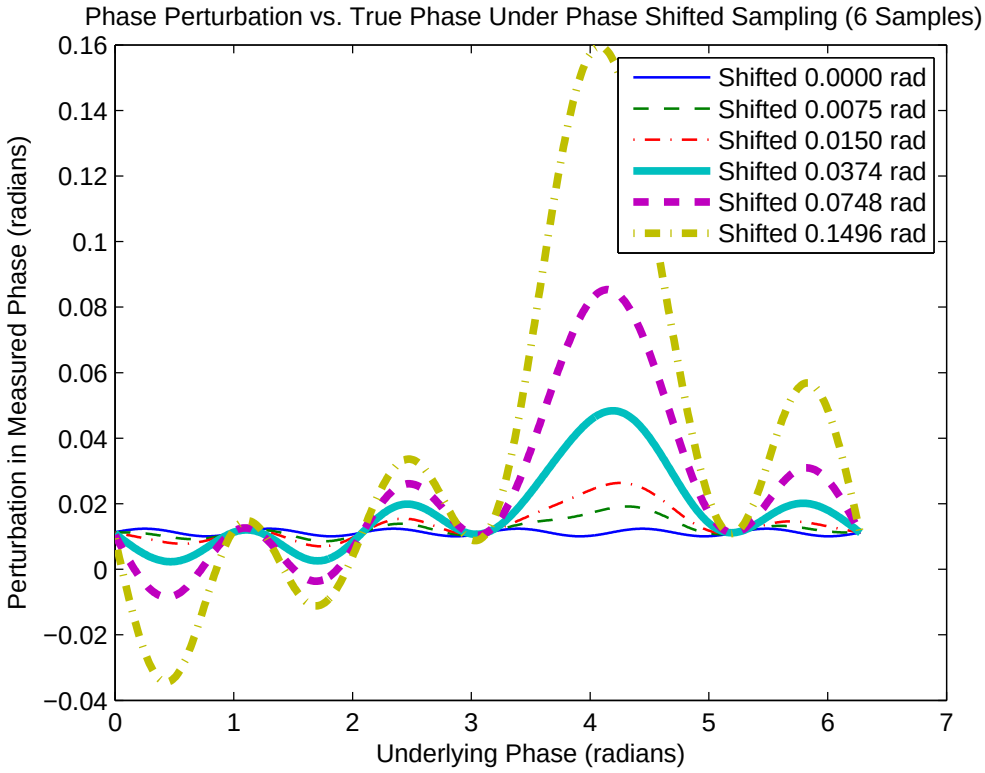

(b) Six Sample Phase Response

Figure 3.18: Uneven phase sampling experiment: phase response for a non-sinusoidal correlation waveform as a function of the the phase shift.

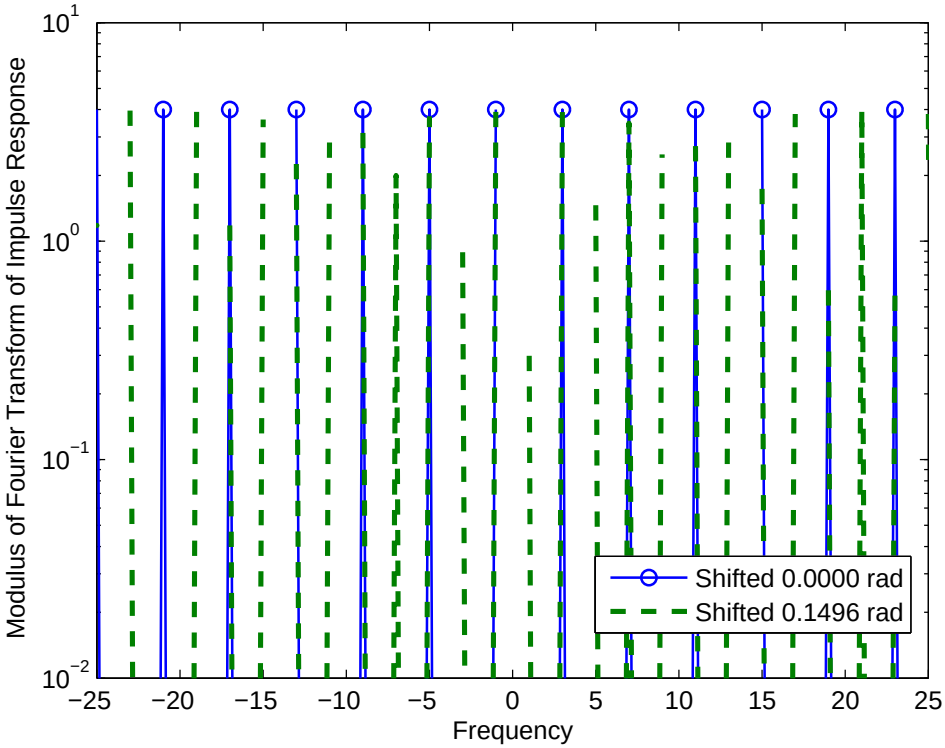

(a) Four Samples

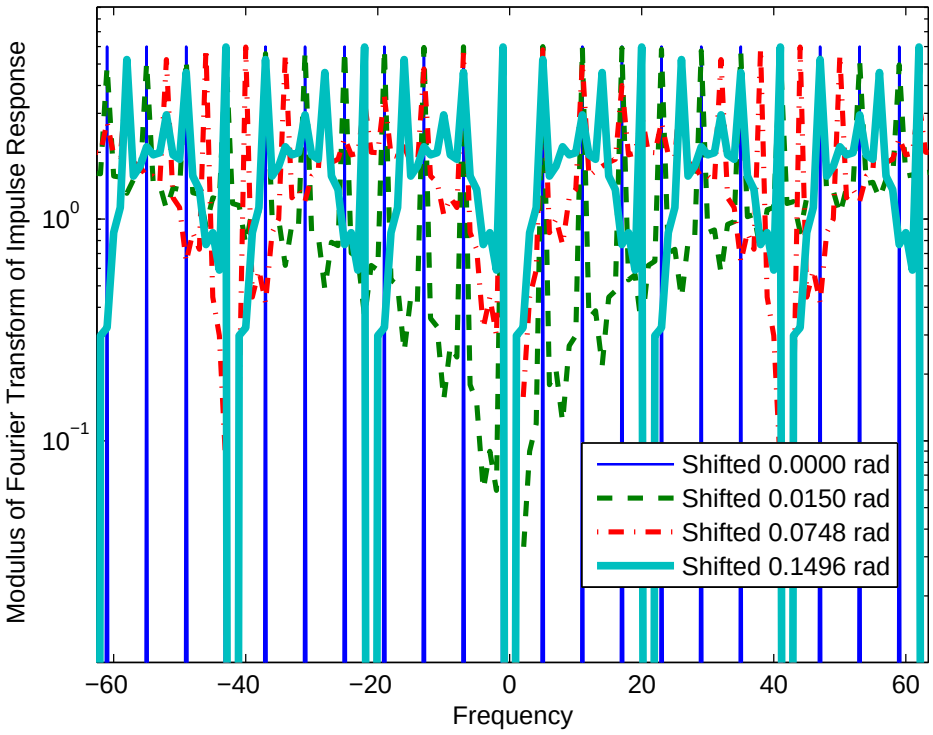

(b) Six Samples

Figure 3.19: The frequency response of the sampling function impulse response versus phase step shift.

occur in an actual system. The purpose is merely to illustrate the behaviour in a 'commonplace' circumstance. In the four sample case, $M = 4$, we assumed

$$\Gamma = \left\{0, \frac{\pi}{2} + \omega_\gamma, \pi, \frac{3\pi}{2} + \omega_\gamma\right\}, \tag{3.120}$$

where ω_γ is the amount of phase perturbation. In the five sample case, $M = 5$, we assumed

$$\Gamma = \left\{0, \frac{2\pi}{5} - \omega_\gamma, \frac{4\pi}{5}, \frac{6\pi}{5} + \omega_\gamma, \frac{8\pi}{5}\right\}. \tag{3.121}$$

In the six sample case, $M = 6$, we assumed

$$\Gamma = \left\{0, \frac{\pi}{3} + \omega_\gamma, \frac{2\pi}{3}, \pi + \omega_\gamma, \frac{4\pi}{3}, \frac{5\pi}{3} - \omega_\gamma\right\}. \tag{3.122}$$

It is interesting to note that the behaviour in latter case if $\gamma_5 = \frac{5\pi}{3} + \omega_\gamma$ is completely different, so the behaviour is very specific to the nature of the perturbation.

Fig. 3.14 gives the two correlation waveform shapes utilised for the uneven phase sampling experiment: one sinusoidal and one non-sinusoidal. Figs. 3.15 and 3.16 plot the amplitude and phase response, $|\arg(P_\mathrm{q})|$ and $\arg(P_\mathrm{q})$, for a pure sinusoidal correlation waveform, while Figs. 3.17 and 3.18 plot the response for a waveform with significant harmonic content. The first case is not subject to any aliasing, so as ω_γ increases, the systematic phase and amplitude error increases from a value of zero. Despite the completely different nature of the phase step perturbations in the four, five and six sample cases, the phase and amplitude error always remain cyclic with a frequency of two (although, in the six sample case the phase error is quite asymmetric.) It is interesting to note that in practical four sample systems, there is no plausible mechanism for aliasing to produce a cyclic error with a frequency of two, therefore this could be a plausible indicator of irregular sampling.

Figs. 3.17 and 3.18 display significant aliasing in all cases except for at six phase steps. Irregular phase steps have much less of an impact on the four step case relative to the existing aliasing error than in the five and six step cases, as a result it would be easily for this source of systematic error to be missed in an aliasing calibration. Whereas the sinusoidal waveform resulted in a pure cyclic error, the presence of significant harmonics produces what might be termed in technical signal processing jargon as 'a mess'. In order to understand what is occurring, it is necessary to analyse the frequency content of the sampling function. Fig. 3.19 compares the frequency content of the sampling function for different amounts of perturbation.

For $\omega_\gamma = 0$, the frequency content is zero except for at the correlation waveform frequencies which aliasing onto the fundamental. As the phase step error increases, the amplitude at those aliasing frequencies decreases, but the frequency content become more broadband; as a result, frequencies in the correlation waveform which would normally have no impact on phase and amplitude measurements result in systematic perturbations. In the four phase step case, the third harmonic aliasing onto the negative fundamental results in a cyclic error with a frequency of four; exactly the same relationship holds for other frequencies in the correlation waveform, explaining Figs. 3.15 and 3.16. In the pure sinusoid case, due to the broadband sampling function the positive fundamental is mixed with the negative fundamental, resulting in a cyclic error with a frequency of two.

Chapter 4

Correlation Waveform Shape Fitting Approaches to Mixed Pixel Separation

In Chapter 3 we introduced the concept of the correlation waveform as the convolution of a signal return model, f_ξ, with the reference waveform, ψ; it follows logically, that in order to determine the signal return model for a pixel one should formulate the problem as a deconvolution operation. This chapter develops this concept, attempting to utilise the harmonic content of the correlation waveform in order to determine the amplitude and phase of the component returns within each pixel.

Analyses in the literature using the harmonic content of the correlation waveform are extremely limited. While Xu *et al.* (1995) discusses determination of the phase and amplitude of the harmonics in the correlation waveform, it does not actually give any sort of meaningful equation to determine the phase from the combination of harmonics, only from each harmonic individually. Both Luan (2001) and Frank *et al.* (2009) reference Xu's work, restating many of the equations. One practical solution involves approximating the the distribution of phase error as normal and then performing an SNR weighted average over the unwrapped phase values (this method is detailed in Section 7.2.6). It does not appear to have been explicitly posited in the literature, although it is an obvious and elementary approach. In addition, Frank *et al.* (2009) mentions the use of psuedoinverses for phase and amplitude determination, although only in the pure sinusoid case, which simplifies to measuring the negative fundamental frequency bin of the discrete Fourier transform – the standard method for phase and amplitude determination.

A significant amount of research has been performed into fitting models in full-waveform lidar: for example, Hofton *et al.* (2000) and Chauve *et al.* (2007) use

numerical methods to fit Gaussian and generalised Gaussian models respectively to the measured waveform. Stilla and Jutzi (2009) discusses a variety of different techniques, including: additional models for the waveform pulse, such as rectangular functions; analysis of component returns, parametising them by time, width and amplitude; component return detection methods, including peak detection, leading edge detection, constant fraction detection, centre of gravity detection; and the application of deconvolution methods. The main problem with the most common approaches to fitting distributions to recorded waveforms is that they necessitate iterative solutions, which are highly computationally intensive when performed across a 2D matrix of measurements. Also, the majority of the models rely on the shape being a nicely-defined continuous function, like a Gaussian distribution; the correlation waveform is better modelled by a piecewise model, such as Eqn. 3.22.

Another common signal processing approach utilised for the processing of FMCW range-data is matched filter processing (Soumekh, 1999; Jankiraman, 2007). However, matched filter processing does not reconstruct missing spatial frequencies in the same manner that well-designed deconvolution methods do.

In this chapter we develop three difference general approaches: the first uses a sparse spike train deconvolution technique and extracts component returns from the deconvolved data; the second approach uses one of two waveform models, either a piecewise truncated-triangle model or a piecewise linear interpolation model; the final approach expands on the linear interpolation model to determine phase, amplitude and the offset due to ambient light in the single component case. While the methods presented in this chapter are limited in their application due to noise constraints and long integration times, they illustrate some of the issues involved in component return separation and lead into the more useful methods in the following chapters.

4.1 Sparse Spike Train Deconvolution

4.1.1 Deconvolution Approaches

The primary limitation that a deconvolution approach faces is the limited harmonic content in the correlation waveform. AMCW systems are not designed for correlation waveform harmonic content; in the limit, designing a system for maximal correlation waveform harmonic content produces a range-gating system, the design of which is ultimately outside the scope of this thesis. In order to determine the location and amplitude of component returns within a pixel, it is necessary to implicitly synthesise the missing, unmeasured spatial frequencies; this is a difficult problem.

The approach we take is to model the component returns as a sparse spike train (SST), that is, the sum of a number of translated and scaled Dirac delta functions. Since volumetric scattering and the resultant diffuse-range component returns are less common, this is not an unreasonable assumption.

SST deconvolution methods have been widely applied in geophysics and have similarity to 2D deconvolution methods utilised in radio astronomy, such as Högbom cleaning (Högbom, 1974). There are a number of common approaches, including $\mathcal{L}_1$ norm regularisation and Bernoulli-Gaussian distribution based search methods. The former involves minimising a cost function given by

$$\epsilon(f_\xi) \quad = \quad ||H_\psi f_\xi - s||_2^2 + \lambda ||f_\xi||_1, \tag{4.1}$$

where f_ξ is the signal return model, H_ψ is a Toeplitz matrix modelling convolution by the reference waveform ψ, s is the correlation waveform and λ is a regularisation parameter. Similar problem formulations are used in basis persuit (Donoho *et al.*, 2003) and Lasso regression (Tibshirani, 1994). From a Bayesian viewpoint, this is solving a maximum a posteriori problem assuming a Gaussian error distribution and a Laplace prior distribution for the signal return model intensities. The relationship between the standard deviation of the noise and the scale parameter of the Laplacian distribution is modelled by the regularisation parameter. In many applications positivity is constrained, hence the prior distribution might alternatively be modelled as exponential.

In the most general case, regularisation using any $\mathcal{L}_p$ norm can be considered to assume an exponential power distribution as prior, sometimes referred to as the generalised error distribution (Tadikamalla, 1980). While it is more difficult to determine the global minimum in this case, it allows more flexibility in designing the prior to suit the application domain; depending on the dataset, it can also be possible to dynamically estimate the hyperparameter p from the data. Regularisation using the $\mathcal{L}_{0.8}$ norm of the spatial derivatives in an image has been used to constrain the 2D deconvolution problem so as to retain object edges while minimising noise (Levin *et al.*, 2007).

Fig. 4.1 demonstrates our own implementation of $\mathcal{L}_1$ norm regularisation, solved using gradient descent implemented using Landweber iterations (Landweber, 1951), a method commonly used for solving deconvolution problems. Fig. 4.1a gives the impulse response and 4.1b gives a demonstration problem. In this case, the original sparse spikes were clearly separated and remain so, even after convolution and the addition of significant noise. Because the convolved peaks are so clearly separated, one trivial approach to determination of the location of each component would be to

set a threshold, presumably determined using knowledge of the noise distribution, and then find the centroid of each peak which projects above this threshold. However, this trivial approach breaks down as the peaks become more closely entangled; this is where deconvolution methods become far more valuable.

Fig. 4.1c shows an example situation before and after processing by our implementation of $\mathcal{L}_1$ norm regularisation. The recovered spikes are broader than the idealised source data, but have retained the same location. Fig. 4.1d shows a case where the spikes are so close together that the deconvolution has not separated the component returns. While ad-hoc alteration of the regularisation parameters and number of iterations might improve the results, there is always a point beyond which two spikes can no-longer be clearly separated for a given SNR. Figs. 4.1e and 4.1f demonstrate the impact of SNR on the ability to separate two component returns. In the former case the component returns can be recovered, but in the latter case where the noise standard deviation has been doubled, attempting to deconvolve the data results in a large number of subsidiary spikes. Given no other information, it would be impossible to know which spikes are correct; given the SNR, the deconvolution algorithm has overfitted the data.

One particularly popular model for a sparse spike train is that of a Bernoulli-Gaussian distribution. There are a number of approaches utilising this, including Kormylo and Mendel (1982) and Kaaresen (1997). The unnormalised posterior probability for the amplitude and spike locations given a measurement vector, s, modelled as a Bernoulli-Gaussian distribution, subject to Gaussian noise, is

$$p(a, x|s) \propto p(s|a, x)p(a|x)p(x), \tag{4.2}$$

where $a \in \mathbb{R}^n$ is the set of amplitudes for each spike and x is a vector of boolean values, $x_i \in \{0, 1\}$, describing whether sample i corresponds to a spike or not. The probability of the existence of an individual spike, $p(x_i)$, is given by a Bernoulli distribution

$$x_i \sim \mathcal{B}(p_\mathrm{r}), \tag{4.3}$$

where p_r is the mean probability of the existance of spike. The probability of a known spike to have a particular amplitude, $p(a_i|x_i)$, is given by a Gaussian distribution such that

$$(x_i = 0 \Rightarrow a_i = 0) \wedge (x_i = 1 \Rightarrow a_i \sim \mathcal{N}(0, \sigma_a^2)), \tag{4.4}$$

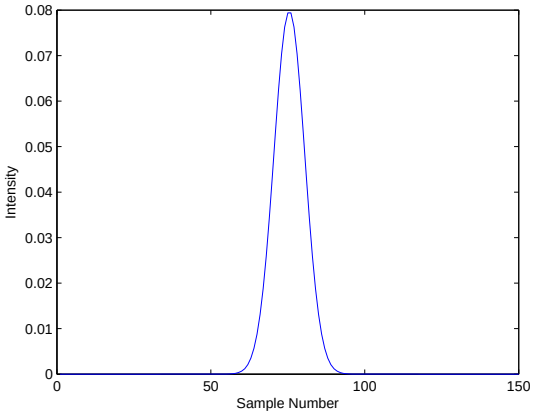

(a) Impulse Response

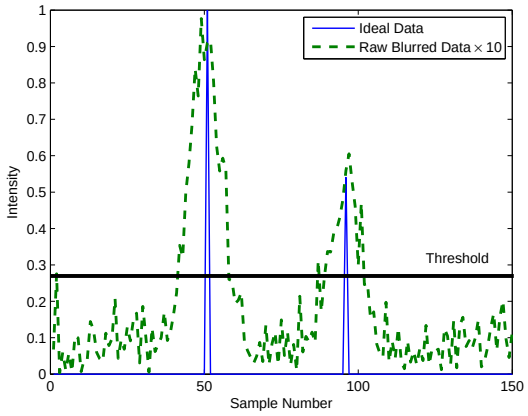

(b) Clear Separation (45 Sample Distance)

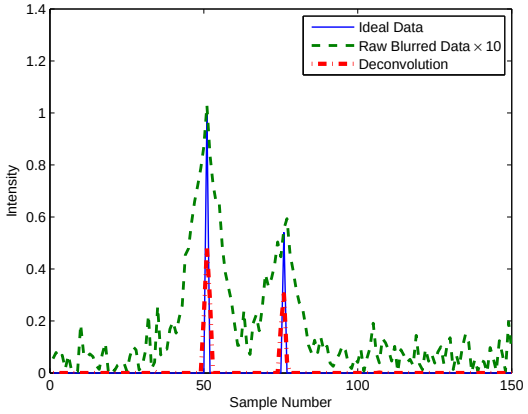

(c) Separated (25 Sample Distance)

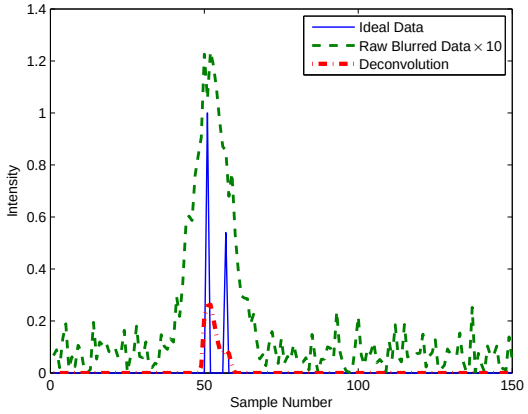

(d) Inseparable (6 Sample Distance)

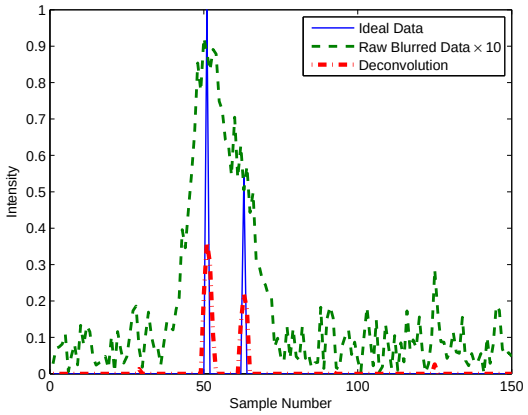

(e) Separated (12 Sample Distance)

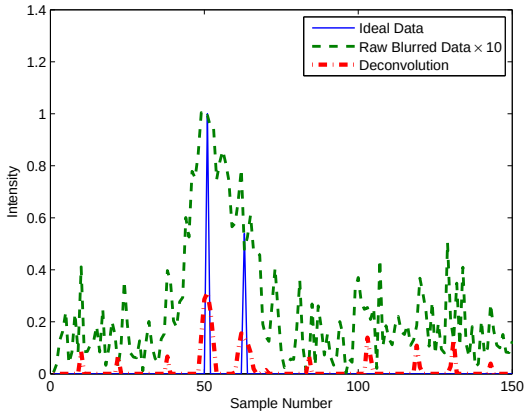

(f) Overfitted/Too Noisy (12 Sample, 2× Noise SD)

Figure 4.1: Demonstration of $\mathcal{L}_1$ norm regularisation. The second component is 54% the intensity of the primary.

where σ_a^2 is the variance of the spike amplitude distribution. Finally, the probability of the data given the component returns, $p(s|a, x)$, is defined by the noise distribution

$$s \sim \mathcal{N}(H_\psi a, \sigma_s^2 \mathcal{I}_n), \tag{4.5}$$

where σ_s^2 is the noise variance. For any given x, determination of a is a trivial least squares problem that can be solved using a Moore-Penrose pseudoinverse. The simplest possible approach is to brute force search through every possible permutation of x and find the value that has the maximum unnormalised posterior probability; however, this is only achievable when the total number of spikes is limited, or the number of elements in the representation is small. Instead, most practical implementations use a tree search, where each step corresponds to the addition, removal or transposition of a spike. For example, Kaaresen (1997) uses a windowing technique where the algorithm operates on a limited subsection of the data at a time, limiting the extent to which data has to be recomputed at each step. Even so, because of the highly non-linear nature of the problem, it is extremely difficult to ensure that both the global maximum is found and that the execution time is feasible.

There are several issues with the most common approaches to SST deconvolution; one of these is the assumption that a spike corresponds to a Kronecker delta. Because we are modelling samples of the convolution of a Dirac delta with an impulse response, it is more appropriate to assume that a single component return will be spread across two adjacent samples; this is discussed in Section 4.1.3. Another issue is the suitability of the prior distributions; while it is nice to be able to model the probability of a sample being a spike using a single parameter, the multiple return problem is difficult to parametise in this manner. It is difficult to even define precisely how to determine the number of returns within a signal and, in reality, the probability of a sample containing a component return is a function of the existence of other component returns. In this thesis we avoid the issue by primarily addressing the question of whether there are one or two components within each pixel; oftentimes we assume that there are two, even when there is no evidence of that (this is the assumption used for the lookup table mixed pixel recovery algorithm discussed in Chapter 6).

4.1.2 The Levy-Fullagar Algorithm

Another, more obscure, approach is the Levy-Fullagar algorithm (Levy and Fullagar, 1981); this utilises noise based constraints in the Fourier domain and finds the solution within these constraints which minimises the $\mathcal{L}_1$ norm. Whereas the other

methods we have discussed require either the assumption of a regularisation parameter, the behaviour of this algorithm is defined by statistical bounds in the Fourier domain.

The constraints are of the form

$$\Re\left(\frac{\tilde{S}[u]}{\Psi[u]}\right) - \epsilon_u \leq \Re(F_\xi[u]) \leq \Re\left(\frac{\tilde{S}[u]}{\Psi[u]}\right) + \epsilon_u \tag{4.6}$$

$$\Im\left(\frac{\tilde{S}[u]}{\Psi[u]}\right) - \epsilon_u \leq \Im(F_\xi[u]) \leq \Im\left(\frac{\tilde{S}[u]}{\Psi[u]}\right) + \epsilon_u, \tag{4.7}$$

where $\tilde{S}$ is the fourier transform of the raw data, Ψ is the Fourier transform of the reference waveform and ϵ_u is an error bound on spatial frequency u, given by

$$\epsilon_u = \frac{\beta_1 \sigma_1 \sqrt{n}}{\sqrt{2}|\Psi[u]|}, \tag{4.8}$$

where β_1 defines Gaussian confidence limits. These error bounds incorporate the implicit assumption that the noise is white and Gaussian. Using linear programming and assuming a positivity constraint, the Levy-Fullagar algorithm finds the conforming value of f_ξ that minimises the $\mathcal{L}_1$ norm. An implementation of the algorithm can use any subset of the measured spatial frequencies, depending on whether they contain useful data and the requisite execution speed. Any linear programming method is acceptable and the results are always deterministic. An extremely useful property of the Levy-Fullagar approach is that the zeroth spatial frequency (mean offset from zero) can be trivially ignored by not including it as a constraint. This is highly important when sampling the correlation waveform using a non-differential system, where there is a highly significant offset in the waveform liable to cause other algorithms to fail.

The choice of algorithm to utilise is fairly arbitrary, although some may give better results than others depending on the input data. For the implementation of the first mixed pixel separation approach we have utilised the LF algorithm because it is relatively easy to implement, does not require a computationally complex tree search and does not involve the explicit assumption of a prior distribution over the component returns and determination of the associated regularisation parameter.

4.1.3 Modelling Continuous-Valued Phases

An estimate of the signal return model, f_ξ, is broadly equivalent to the measurements captured directly by a range-gating system. However, whereas AMCW systems

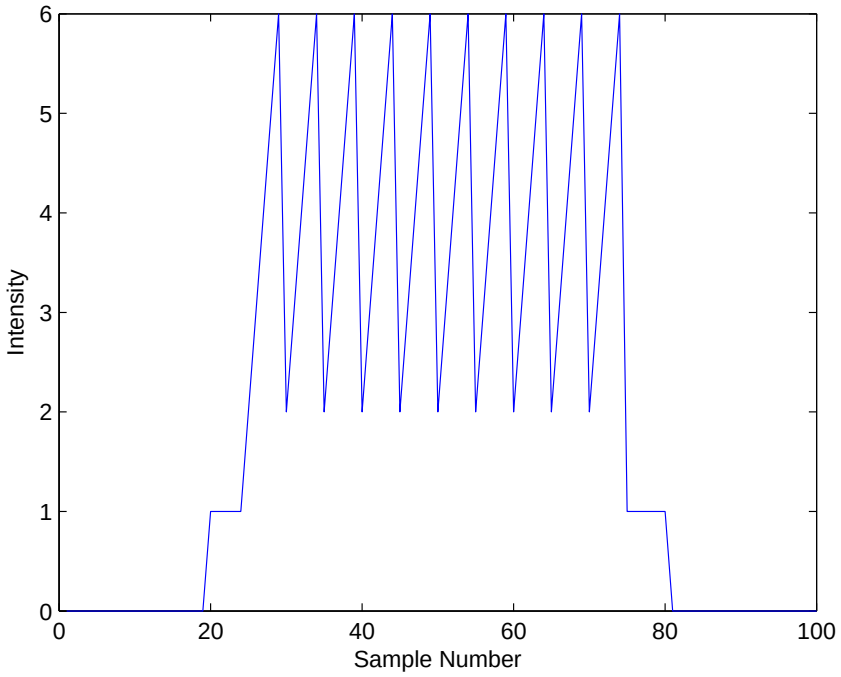

(a) Initial Signal

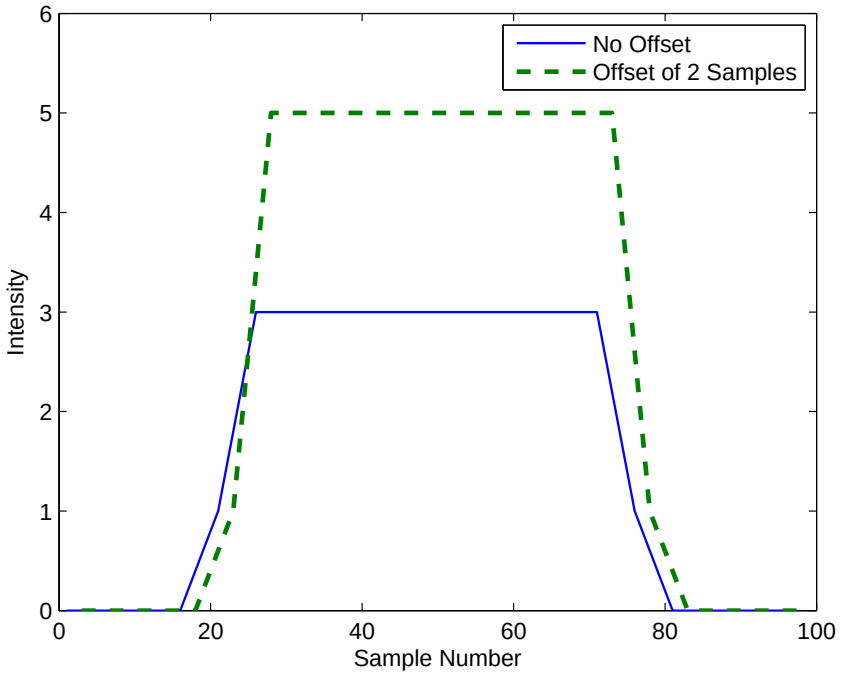

(b) Resampled Signal

Figure 4.2: Demonstration of the systematic error introduced when resampling a signal. The second figure plots every fifth sample of the first, with two different offsets giving completely different results; this is closely related to the suitability of a waveform shape for a linear interpolation based continous-phase model.

produce continuous amplitude and phase estimates, a signal return model contains continuous amplitude information in discrete, binned form. For many applications, a continuous range-measurement is required. We now introduce one possible approach to this problem using the assumption of component returns which are scaled and translated Dirac deltas. Diffuse-range component returns are modelled separately in Chapter 5.

Isoplanatic SST deconvolution approaches approximate the signal return model as the linear combination of a limited set of basis vectors, where these vectors are translated versions of an impulse response, in this case the reference waveform, ψ. The forward convolution can be modelled using a Toeplitz matrix, given by

$$H_\psi = \begin{pmatrix} \psi[0] & \psi[n-1] & \cdots & \psi[1] \\ \psi[1] & \psi[0] & \cdots & \psi[2] \\ \vdots & \vdots & \ddots & \vdots \\ \psi[n-1] & \psi[n-2] & \cdots & \psi[0] \end{pmatrix}, \tag{4.9}$$

where $\psi[i]$ is the ith of n discrete samples of the reference waveform. As with any sampling operation, these basis vectors suffer from aliasing of frequencies above the Nyquist frequency. The measurements of the correlation waveform, $\tilde{s}[i]$ are also discretised/quantised; however, because the measured correlation waveform is a continuous function of phase, the impact of aliasing changes with phase. An extreme example of this effect in the non-integrative/homodyne case is given in Fig. 4.2. Attempting to model one of the samples of the impulse response using the other would clearly result in misestimated amplitude. In the general case, both phase and amplitude error are possible.

In general, SST deconvolution methods can only model a continuous-phase Dirac delta type component return as the weighted sum of two adjacent basis vectors; in other words, as a linear interpolation. Thus the contribution of an individual component return to the correlation waveform is modelled by

$$f_\nu[i] = a\big(\psi[(i+\nu_\xi) \bmod n] + \lambda_\xi \Delta\psi[(i+\nu_\xi) \bmod n]\big), \tag{4.10}$$

where a is the estimated amplitude of the return, $\lambda_\xi \in [0,1)$ is the subsample location, $\nu_\xi \in \mathbb{Z}$ is the number of discrete samples to shift the basis vector and the discrete derivative of the basis vector is given by

$$\Delta\psi[i] = \psi[(i+1) \bmod n] - \psi[i]. \tag{4.11}$$

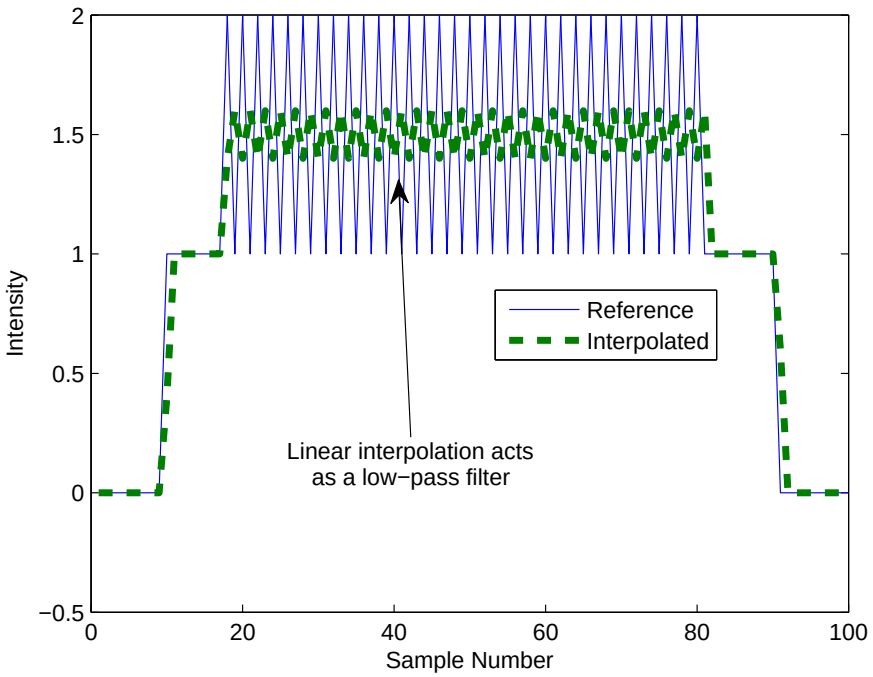

(a) Excessively Complicated Waveform

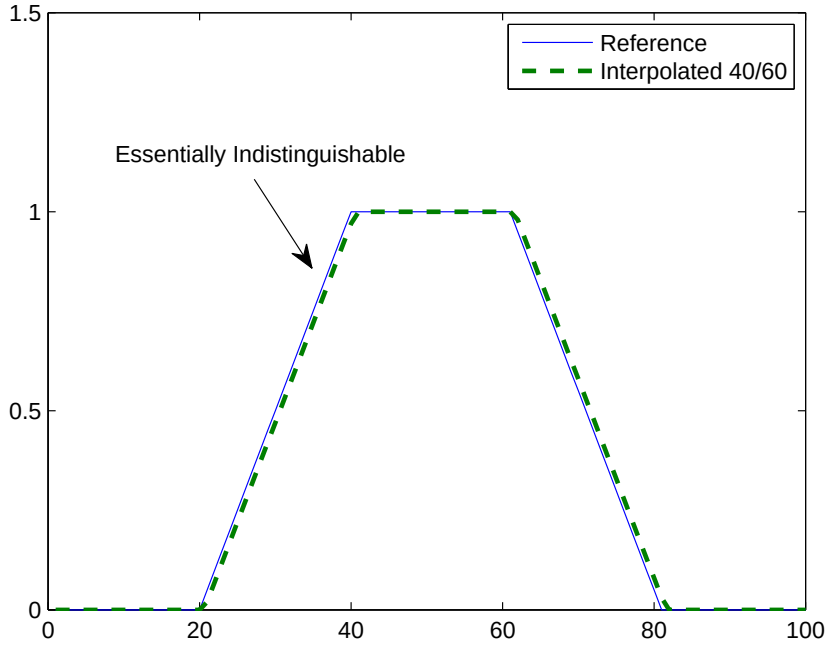

(b) Truncated-Triangle Waveform

Figure 4.3: Demonstration of linear interpolation between two waveform shape basis vectors intended as a continuous-phase model for the correlation waveform; waveforms with complicated detail are not well-modelled.

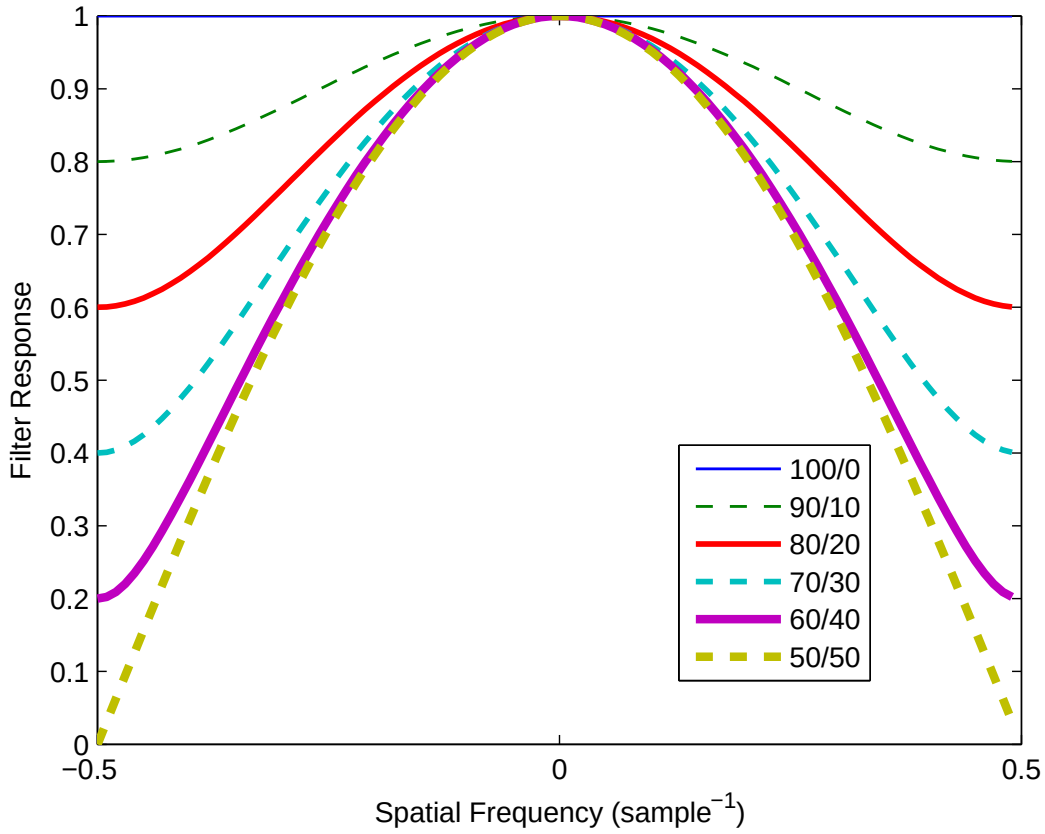

Figure 4.4: Frequency response of continuous-phase basis vector interpolation model for different intensity ratios ($\upsilon_\lambda[0]/\upsilon_\lambda[1]$).

The Fourier transform of the linear interpolation model, f_ν, is given by

$$F_\nu[u] = (e^{2\pi j u \nu_\xi} \Psi[u]) * \Upsilon_\lambda[u] \tag{4.12}$$

where the first term is given by the Fourier shift theorem and Υ_λ is the Fourier transform of the linear interpolation impulse response, υ_λ, given by

$$\upsilon_\lambda[0] = 1 - \lambda_\xi \tag{4.13}$$

$$\upsilon_\lambda[1] = \lambda_\xi \tag{4.14}$$

and

$$\Upsilon_\lambda[u] = (1 - \lambda_\xi) + \lambda_\xi \mathrm{e}^{2\pi j u/n}. \tag{4.15}$$

Analysing eqn. 4.15 shows significant attenuation for small numbers of phase steps and/or high spatial frequencies. This attenuation is a function of the subsample

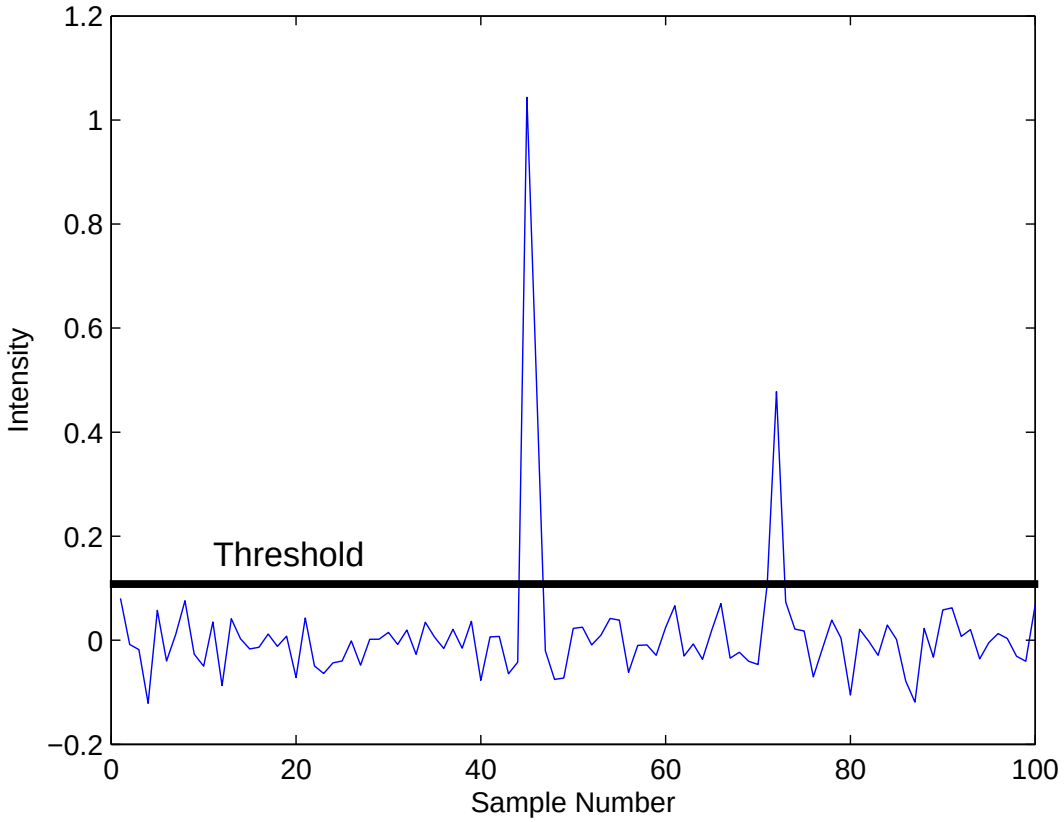

Figure 4.5: Setting a minimum intensity threshold for detection of a component return. (This figure is intended to be illustratory only, and should not be considered representative of the actual output from a SST deconvolution algorithm.)

location, as shown in Fig. 4.4. One disadvantage is the attenuation of spatial frequencies within the correlation waveform; conceivable approaches to prevent this would be to either sample the correlation waveform at a number of different sub-sample phase offsets or to model the waveform shape as a continuous function, rather than a piecewise interpolation.

There are two approaches to the practical determination of a, λ_ξ and ν_ξ: these are to directly fit eqn. 4.10 to the raw measured correlation waveform, which we discuss in Section 4.2; or to apply an off-the-shelf SST deconvolution algorithm, which fits the basis vectors to the data for each component return, and then post-process the data to determine the parameters for each component. We now address the latter approach.

As discussed in Chapter 3, the number of component returns is an ill-defined parameter. In order to separate out each component return we find all the groups of contiguous samples with intensities above a set threshold; for the experiments

presented here, the threshold was assumed to be 5% of total pixel intensity. The discrete Fourier transform bin corresponding to the negative fundamental spatial frequency was calculated of the region above the threshold, plus a single sample buffer on either side. In explicit mathematical form, that is

$$\xi = \sum_{i=0}^{n-1} e^{2\pi ji/n} f_\xi m_{\mathrm{x}}[i], \tag{4.16}$$

where f_ξ is the estimate of the signal return model produced by the SST deconvolution method and m_{x} is the mask for the specific component return in question.

While eqn. 4.10 models a continuous-phase return using a linear model, it does not directly determine phase and amplitude and must be fit to the measured data; the choice of algorithm for fitting the model can be considered separate from the design of the model itself, we discuss the fitting approaches separately. Assuming that a, λ_ξ and ν_ξ have been determined by some arbitrary algorithm which fits f_ν to f_ξ in a noise optimal manner, we can estimate the complex domain range measurement by

$$\xi = a\mathrm{e}^{2\pi j(\lambda_\xi+\nu_\xi)/(n\Psi[\frac{1}{n}])}. \tag{4.17}$$

The problem of identifying continuous ranges from the deconvolved SRM can be compared to the problem of producing continuous range data from range-gating systems. Centroiding, thresholding, simple maxima and second order methods have been applied to this problem (Wagner *et al.*, 2004). We use the masked Fourier transform in eqn. 4.17 because it produces results similar to centroiding but takes into account points that partially wrap around. There is also a certain elegance to this approach in that we are modelling a set of highly non-linear simultaneous equations of the form

$$\xi_l = \sum_{i=0}^{n-1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l-1}}, \tag{4.18}$$

using a simple linear model that produces deterministic estimates of values η_i, rather than an arbitrary local minimum.

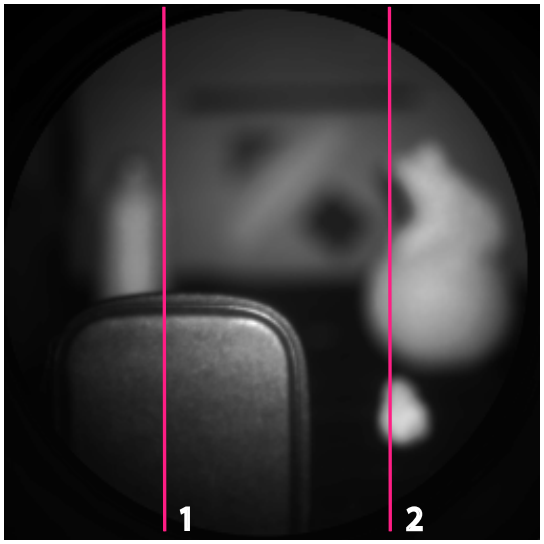

(a) Intensity Image of Scene

(b) Range Image (Standard Method)

(c) Brightest Return Range Value (LF)

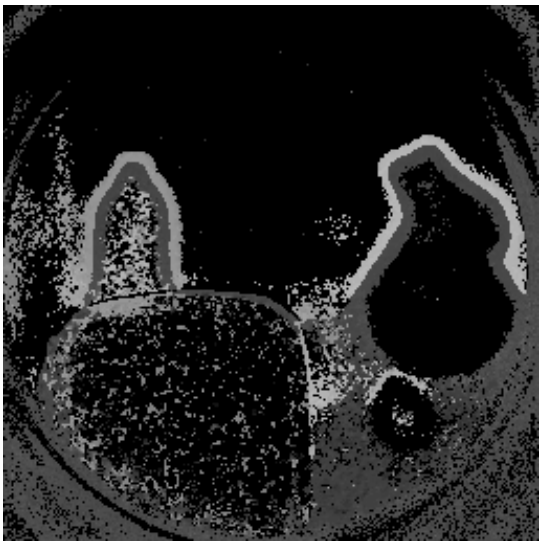

(d) Second Brightest Range Value (LF, black = none)

Figure 4.6: Range image data before and after processing using the Levy-Fullagar algorithm to separate out multiple components within each return using correlation waveform harmonics. White represents an object close to the camera, dark grey farther away. Note the light scattering outside the circular image intensifier and limited depth-of-field. Lines labelled '1' and '2' notate the location of slices illustrated in Fig. 4.7

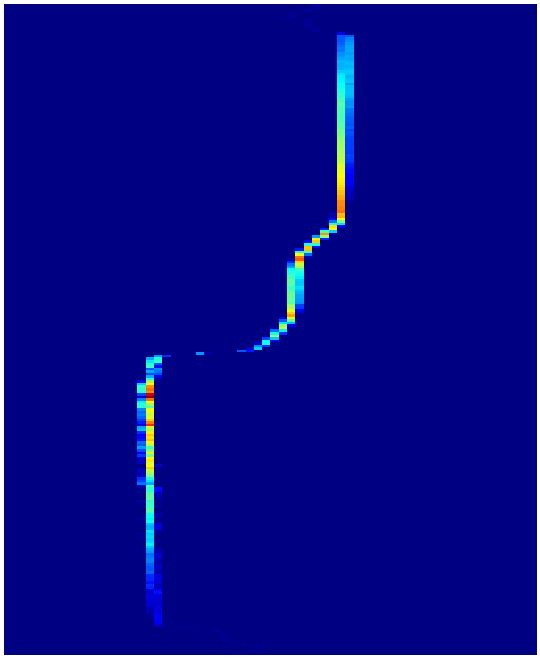

(a) Naïve Fourier Reconstruction (dataset 1)

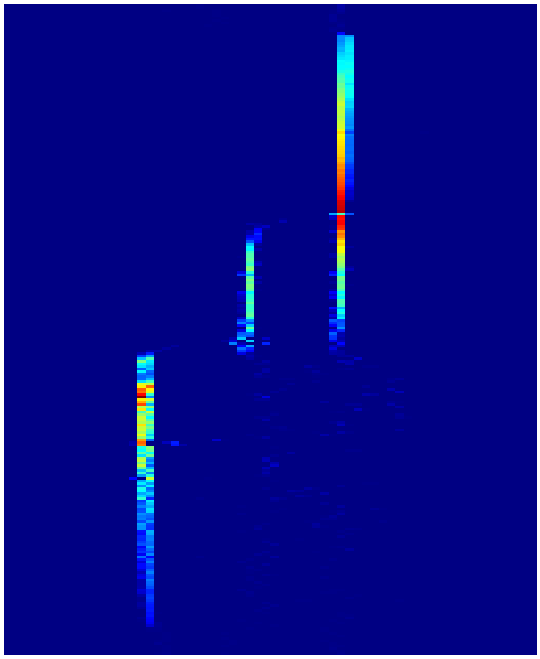

(b) Levy-Fullagar Reconstruction (dataset 1)

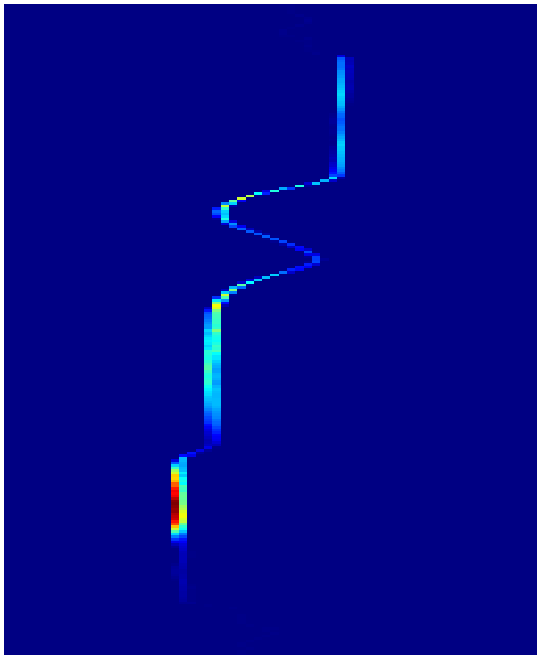

(c) Naïve Fourier Reconstruction (dataset 2)

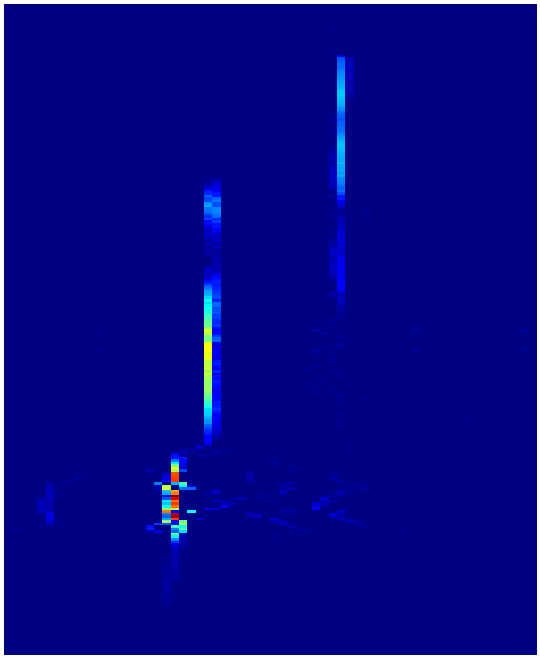

(d) Levy-Fullagar Reconstruction (with errors, dataset 2)

Figure 4.7: Vertical slices through a discrete signal return model estimated via the standard Fourier bin approach or the LF algorithm. The x axis represents range, the y axis represents vertical slice location – blue represents a low light level and red high. Any line parallel to the x axis that intersects more than one return in the recovered SRM is a mixed pixel in the standard version. The slices are labelled in Fig. 4.6.

4.1.4 Applying the Levy-Fullagar Algorithm to Component Return Separation

We now present a demonstration of the application of the Levy-Fullagar algorithm to the deconvolution of signal returns within each pixel using the UoW heterodyne lidar system. For the purposes of these experiments we utilised $\beta_{\mathrm{l}} = 4$ and the first eight harmonics of the correlation waveform, both sampled and modelled using 64 phase steps. In order to achieve as high a SNR as possible, 64 samples were averaged per phase step – additionally providing an estimate of the sample standard deviation, which was used to calculate σ_{l}^2. For these experiments, only one of the coaxial laser illumination sources was operational, which had a negative impact on the image SNR.

The example in Fig. 4.6 clearly shows the blurred object boundaries of the original range data and the improvement in Fig. 4.6c. Fig. 4.6d shows the range of the second brightest return, and regions where objects have blurred onto each other are clearly visible. The figure also shows light scattering, probably due to the relay lens coupling between the image intensifier and CCD, which results in changes to the measured ranges of dark objects.

A different way to look at the problem is by viewing slices through the SRM. Two slices through the Fig. 4.6 model are given in Figs. 4.7a to 4.7d. The second example shows a breakdown in the algorithm – this effect is known to occur in the brightest regions of images particularly when high microchannel plate voltages are used, and may be caused by the image intensifier.

In order to compare the RMS error for Fourier processing against the LF algorithm, we placed a flat board in front of the camera and fit a simple planar model to the surface; this was assumed to adequately approximate ground-truth for a noise analysis. In the the LF method RMS error versus the flat board model in the single return case was found to be smaller than the standard method RMS error (Table 4.1) despite the standard method data being the basis for the model. This shows that in the single return case, the LF method can decrease overall error by 30%. Fig. 4.8 illustrates the scene configuration utilised to analyse the two component return case; planar models were fit in a least squares manner to each of the surfaces and extrapolated over the mixed region between the two objects. Due to the nature of defocus, as the mixed region is traversed the relative intensities of the component returns change – this substantially impacts on the resultant overall RMS error after processing by the LF algorithm. Fig. 4.9 is a graph of RMS error as a function of the intensity ratio. The intensity ratio is the relative intensity of the returns composing each pixel (near object intensity/far object intensity). Error in

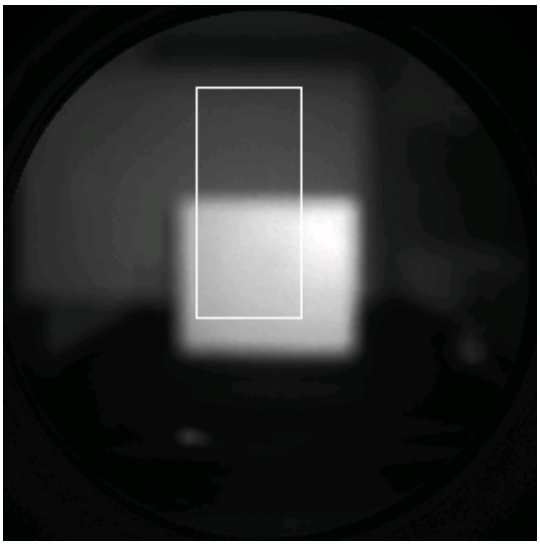

Figure 4.8: Amplitude image of the scene utilised for the analysis of the two component return RMS error as a function of intensity ratio. The rectangular region-of-interest shown in the figure was analysed by fitting a plane to each surface and extrapolating over the blurred region; this provided a suitable approximation to ground truth.

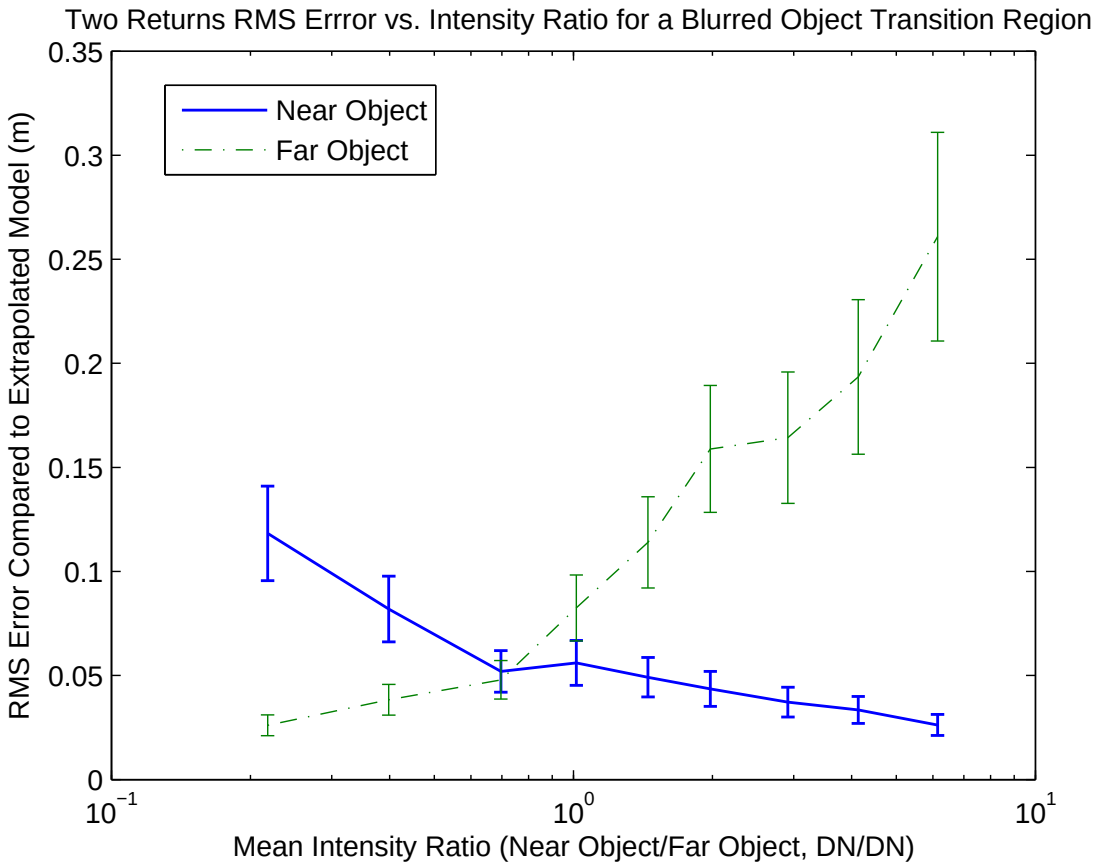

Figure 4.9: RMS error of recovered returns versus intensity ratio in the case of two returns

Table 4.1: Noise performance of the Levy-Fullagar algorithm

Single Return Results (RMS Error vs. Model)	
Standard Method (64 Beats)	0.0204m ±0.0005
Standard Method (1 Beat)	0.0580m ±0.0127
LF Method	0.0143m ±0.0020
LF Method (Ignr. False Ret.)	0.0105m ±0.0002
LF False Multiple Returns	0.61%
Two Returns Results (RMS Error vs. Model)	
Standard Method	N/A
LF Method	See Fig. 4.9
LF False Single Returns	1.15%

measuring the brightest return is fairly low, however the darker return appears to be negatively affected by the bounds inherent in the LF method.

4.2 Shape Fitting Methods

Whereas in Section 4.1 we determined the amplitudes and ranges to objects by post-processing data produced by SST deconvolution, in this section we directly model the waveform shape. In Section 3.1.3 we developed a general form for a truncated triangle waveform and in Section 4.1.3 we developed another linear model for a continuously phase-variable waveform; by combining these models with numerical optimisation methods we endeavour to estimate phase and amplitude.

The first model is a heterodyne variation of eqn. 3.22, implemented by integrating over each phase step rather than merely sampling at fixed intervals. In other words, for the reference waveform, ψ, defined by eqn. 3.22 and using heterodyning as defined by eqn. 3.34, we fit

$$f_\rho[i] = a \int_{-\phi_h/2}^{\phi_h/2} \psi(\phi_h i + \theta_h) \mathrm{d}\theta_h \tag{4.19}$$

to the measured correlation waveform, where ϕ_h is the phase step size.[1] The final model has five parameters, the phase θ and the amplitude, a, plus the three shape parameters $(\rho_h, \rho_l, \frac{c_\rho}{a_\rho})$. The shape parameters are separately calibrated for, so that only (a, θ) are determined by optimisation in a typical use-case. Despite the fact that the actual correlation waveform is not an ideal truncated-triangle, it is still a highly functional approach as long as the sensor/intensifier response approximates a rectangular function.

[1]For the sake of simplicity, ϕ_h is assumed to be identical to the domain of phases over which the heterodyne integration occurs, i.e., $\phi_h = \Delta\phi_d$.

The second model is given by eqn. 4.10; for convenience we call this the base-ratio model. This requires an explict reference waveform $\psi[i]$ and has two parameters fit at run-time: the amplitude, a, and the phase $\theta = 2\pi(\nu_\xi + \lambda_\xi)/n$.

In order to model a particular pixel, the models for each return within that pixel are summed together. Given the assumption of a specific number of component returns within a pixel we then determine the phases and amplitudes of those components by numerical maximisation of a likelihood function. In the case of Poisson distributed noise, which is the primary noise source for the image intensifier system, the log likelihood is described by

$$LL(\lambda) = -n\lambda + \log\lambda \sum_{i=0}^{n-1} k_i - \sum_{i=0}^{n-1} \log k_i!, \tag{4.20}$$

where λ is the underlying mean number of photons per sample and k_i is the ith sample of the measured data. Since the recorded data are not free parameters, the final term can be ignored. We implemented the numerical optimisation using Matlab 7.3's Sequential Quadratic Programming method for solving non-linear equations. Depending on the exact shape of the waveform, convergence to a global optimum is not guaranteed, thus the approach is non-deterministic in typical implementation. One approach to ensure determinacy is to brute-force search over the domain of ν_ξ; another approach, specific to the single return case, is to use a Fourier transform to calculate ν_ξ and use shape to determine only λ_ξ and a. We examine this latter approach in greater depth in Section 4.3.

4.2.1 Analysis of Shape Fitting Methods

These experiments were carried out on the University of Waikato full-field heterodyne ranger using a laser modulation frequency of 12 MHz and a camera frame-rate of 32 fps.

In order to use the truncated-triangle waveform model (eqn. 4.19) for deconvolution/phase identification it is necessary to determine the three shape parameters (p_h, p_l, $\frac{c}{a}$). In order to do this, a small region of the raw camera frames were averaged over a number of beat cycles to give the general shape of the waveform. Candidate pixels were required to be sufficiently bright to limit noise and the influence of scattered light, fairly close to the optical axis and not near the edge of an object. The parameters were then found via optimisation using the numerical method described above and this was repeated on an individual image basis for each experiment. The waveform ψ was estimated in a similar manner.

Single return precision was calculated by placing a flat board in front of the ranger and taking a series of range images. It was assumed for analysis purposes that each pixel contained only a single return. Each individual beat was considered to be a separate measurement. The methods under test were applied to the data over 10 beat cycles and a 50×50 pixel region of the board and the interbeat precision was estimated using the sample standard deviation. An estimate of the ideal range data was created by averaging the Fourier bin method data over time and then blurring it using a Gaussian blur of radius 5 pixels. The overall ranging error for a dataset was then estimated by measuring the RMS error against the ideal model. This allows error to be approximated without complicated calibration for effects such as irising. In the single return case, the Fourier bin values were used as seed values for the optimisation.

A scene was created where two boards were placed in front of each other so that they overlapped, the camera was then defocussed to produce a blurred region between the two boards and a large number of mixed pixels to work with, in a similar manner to the experiment in the previous section. Ten beat cycles were analysed in a similar manner to the single return case, except that for analysis purposes it was assumed that every pixel within this region contained two returns. In order to calculate the two return precision the source of each return was determined by a brute force search designed to minimise error – it was assumed that one source within each pixel belonged to the board in front, and one to the board in back. In order to identify the nature of any systematic error, the mean distance between returns was measured and compared to the actual mean distance between the two objects.

4.2.2 Waveform Models

Table 4.2 shows the parameters found in the two cases shown in Fig. 4.10. It was not possible to use the base-ratio method at 4 frames-per-beat because the sampled waveforms were not representative enough – typically they contained a significant constant offset. Attempting to apply the base-ratio method resulted in non-continuous range values and discrete levels in the range data. The same errors occur if the truncated-triangle parameters are fit to a poorly chosen region of the image where there is a large constant offset due to multiple returns. Strictly the four sample truncated-triangle parameter determination problem is under-defined, however the optimisation found very similar parameters to those of the 64 sample case. While we present some results using these estimated parameters, all the truncated-triangle model results using only four phase steps should be considered extremely

Table 4.2: Example fitted truncated-triangle model parameters. These are the parameters of the models shown in Fig. 4.10

	Shape			Scaling/Translation	
Frames-Per-Beat	p_h	p_l	$\frac{c}{a}$	θ (rad)	a (DN)
4	0.1459	0.02842	0	0.7626	16210
64	0.1692	0.03477	0.02398	3.3485	9445

Table 4.3: Interbeat precision and rms error against a reference model of phase measurements for different measurement methods in the single return case. Precision is 1σ precision, single optimisation attempt, FPB = frames-per-beat.

	Fourier-Bin		Trunc-Triangle		Base-Ratio	
PFB	Precision	RMS Error	Precision	RMS Error	Precision	RMS Error
4	0.0094	0.0193	0.0083	0.0348	N/A	N/A
8	0.0062	0.0163	0.0045	0.0119	0.0051	0.0117
16	0.0062	0.0168	0.0036	0.0086	0.0048	0.0124
32	0.0086	0.0138	0.0046	0.0093	0.0055	0.0086
64	0.0111	0.0123	0.0061	0.0086	0.0062	0.0075

suspect due to the potential for parameter misestimation. When the parameters are properly fit, the range measurements are approximately normally distributed; parameter misestimation results in easily recognisable non-standard distributions, often containing hard-limits.

4.2.3 Single Return Phase Measurement

The results from this experiment are presented in Table 4.3 and Fig. 4.11. The precision of the Fourier method decreases after 16 frames-per-beat – this is extremely surprising because it indicates that the extra light being integrated is not actually improving the signal to noise ratio at all, rather it appears to be adding noise. The most probable cause of this is the Direct Digital Synthesiser (DDS) which is used to produce the modulation signals. The DDS runs at 400MHz and uses an accumulator/lookup-table system to produce a roughly sinusoidal waveform which is converted to a rectangular wave via a comparator. Despite the insertion of a low-pass filter into the signal chain, jitter/instability in the rectangular wave output contributes noise.

It appears that under most circumstances the non-Fourier algorithms produce phase measurements with better precision and less overall error. An exception to this is the five frame-per-beat truncated-triangle case where the precision is quite good, but the overall error is quite high – possibly with a non-linear range response due to parameter misestimation.

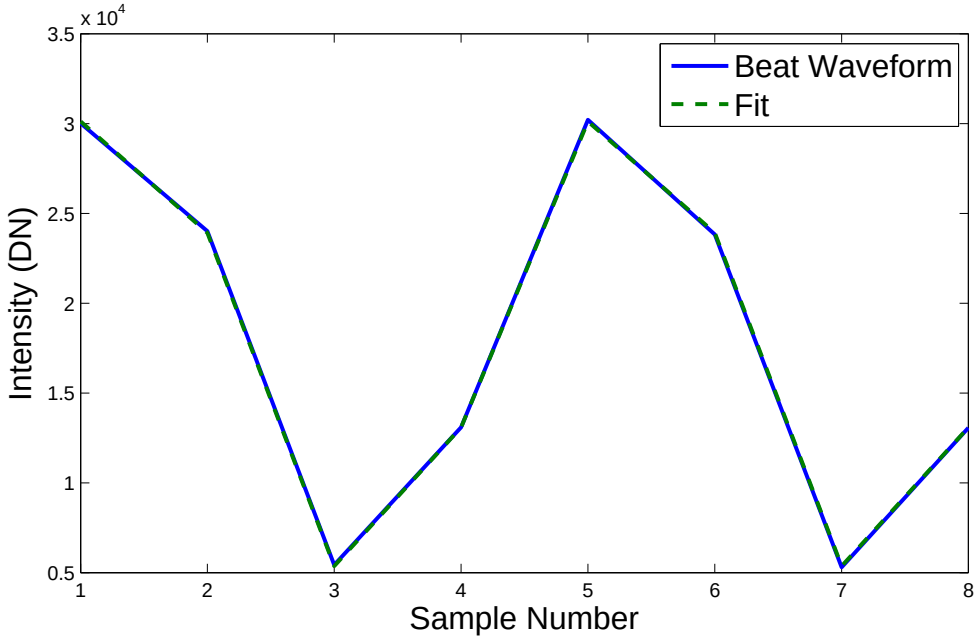

(a) 4 Frames-Per-Beat

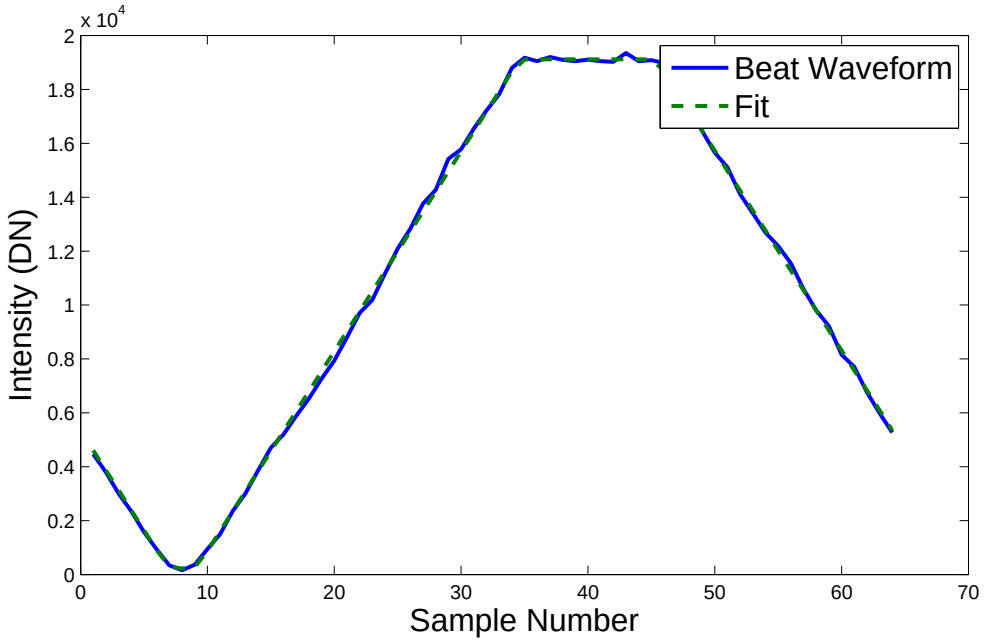

(b) 64 Frames-Per-Beat

Figure 4.10: Fitting a truncated triangle model to the reference waveform while varying the number of frames per beat; the waveforms are near indistinguishable.

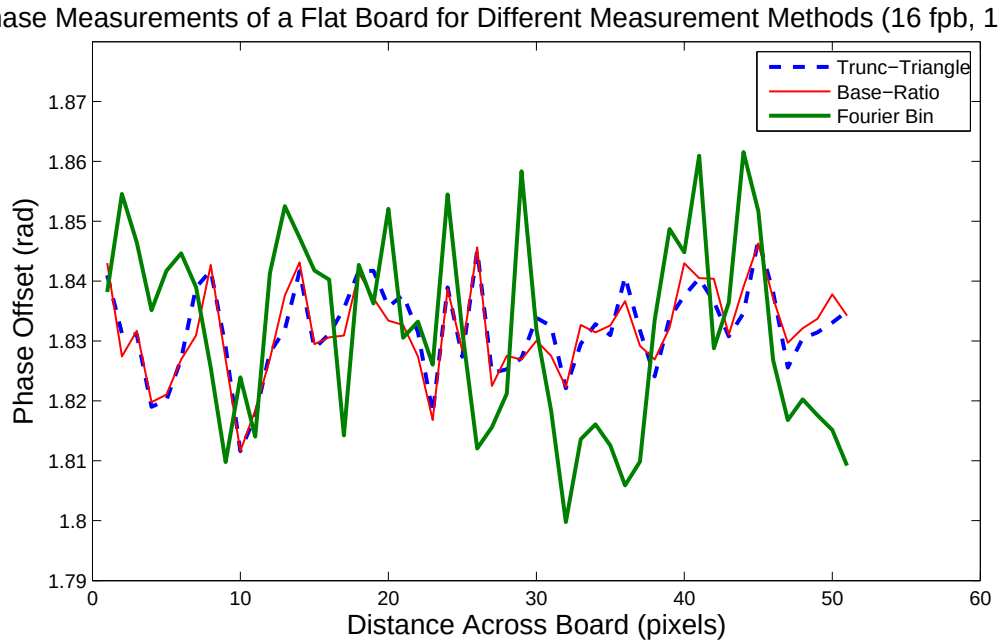

Figure 4.11: A slice through a range image of a flat board processed using different approaches. The base-ratio and trunc-triangle methods appear to substantially decrease random noise relative to the standard Fourier bin based method. The two shape based methods are highly correlated.

Table 4.4: The interbeat, 1σ precision of phase measurements for different measurement methods in the two returns case. PFB = frames-per-beat

	One Optimisation Attempt		Best of Five Attempts	
FPB	T-Triangle (rad)	Base-Ratio (rad)	T-Triangle (rad)	Base-Ratio (rad)
8	0.2457	0.3116	0.2137	0.1889
16	0.2715	0.4584	0.1842	0.1657
32	0.2404	0.4248	0.1425	0.0341
64	0.2849	0.3731	0.0225	0.0218

An optimal trade-off between precision and number of frames required appears to be between 8 and 16 frames-per-beat with the truncated-triangle method being slightly better than the base-ratio algorithm. One possibility is to average a smaller number of frames-per-beat over several beat cycles – this appears to be the only way to improve results after 16 frames-per-beat. A better base-ratio algorithm could trade-off complexity and precision, perhaps oversampling the waveform so that ψ is better representative of the underlying continuously variable waveform shape. One limitation of this analysis is that it does not take into account effects caused by linearity errors; given more time for analysis, this is the obvious next investigation.

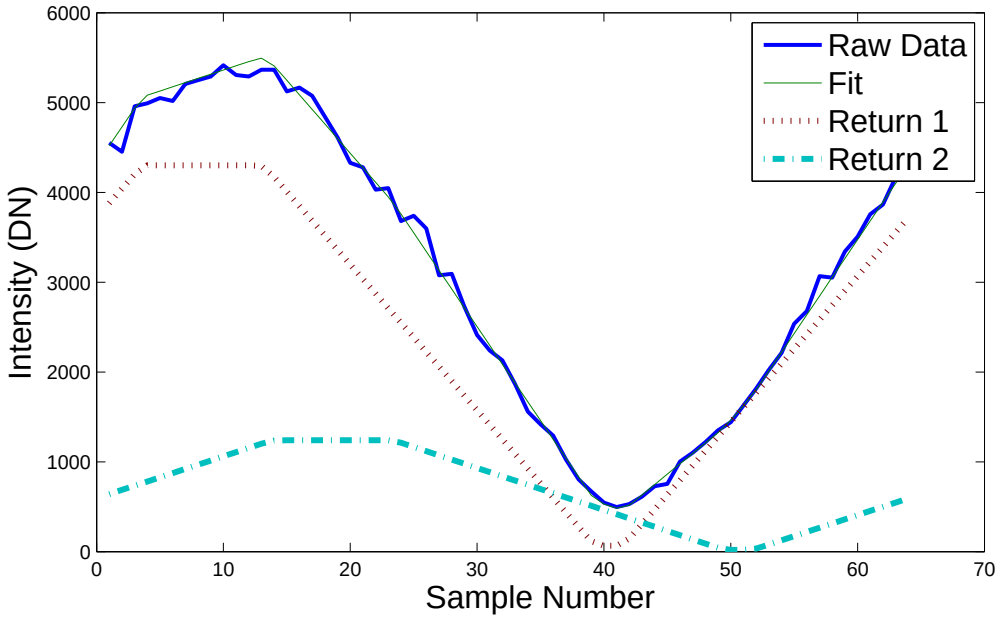

(a) Truncated-Triangle Method

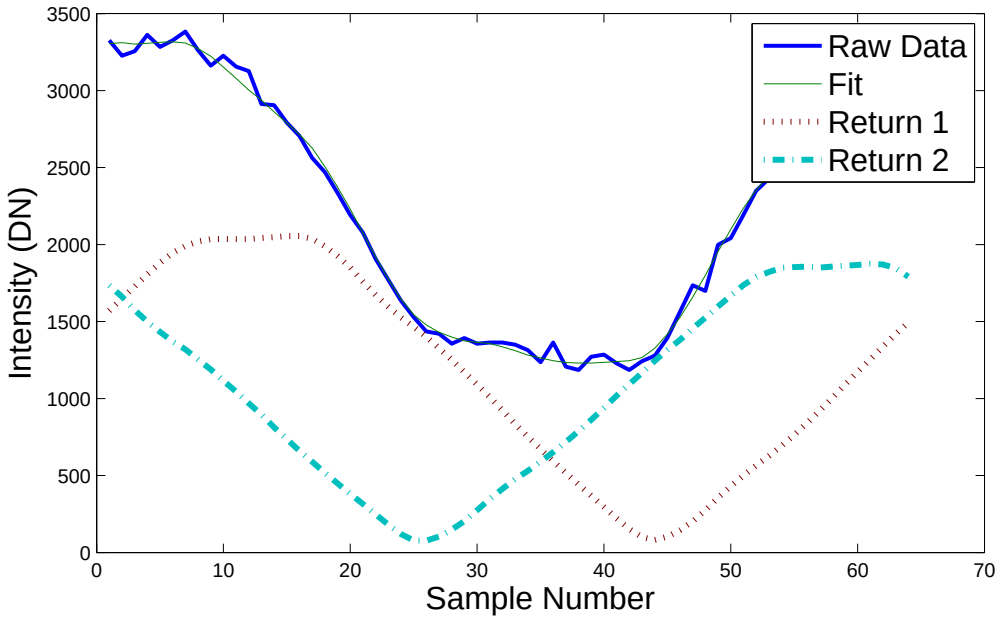

(b) Base-Ratio

Figure 4.12: Example waveform fits using different correlation waveform models using 64 phase steps.

Table 4.5: Mean estimated phase between returns in the two return case. FPB = frames-per-beat

	Best of Five Optimisation Attempts	
FPB	Trunc-Triangle (rad)	Base-Ratio (rad)
8	1.412	1.132
16	1.436	1.285
32	1.373	1.298
64	1.279	1.298
Actual	1.131	

4.2.4 Dual Return Phase Measurement

The results from the dual return experiment are shown in Tables 4.4 and 4.5 and graphical examples are shown in Fig. 4.12. The plotted results are the best of five optimisation attempts unless otherwise specified. The truncated-triangle and base-ratio methods both appear to perform reasonably well with multiple optimisation attempts and with a high number of frames-per-beat, however in general the precision is not very good. More detailed analysis of the raw data indicates that the majority of the measurements have a precision of around 0.02 radians, however there is a significant population of outliers that distort the metric. A good example is shown in Fig. 4.15 where multiple optimisation attempts have a huge impact on local minima. One possible solution to this may be to change from a 1D model to a 3D model and use spatial information to reduce the impact of the outliers.

Like other systems, such as Mure-Dubois and Hügli (2007a), the ranger suffers from issues with scattered light, particularly in dark regions, which is a source of subtly mixed pixels. This is contributed to by the use of a relay lens to couple the camera to the image intensifier. For example, in the case of Fig. 4.13 the mean light intensity across the entire image after calibration is 5.8×10^3 DN and the peak intensity is 4.5×10^4 DN. In theory, the regions outside the image intensifier should be completely black, but because of stray light the mean intensity in the corner regions of the image is 1.5×10^2 DN. If scattered light is integrated by a pixel it tends to result in an extra constant offset. Table 4.5 shows that there is a systematic error resulting in returns appearing farther apart than they actually are – while it is possible that this error is contributed to by stray light, this is a fairly consistent effect. One method for testing the ranger system in the two return case without scattering effects may be direct illumination of the image intensifier with several diffuse light sources at different ranges rather than using a physical scene. Another alternative might be a simulation based approach to determine if there are any other possible contributing factors to a systematic error.

(a) Raw Data Intensity

(b) First Return Intensity In Two Return Case

(c) Raw Data Phase

(d) First Return Phase In Two Return Case

Figure 4.13: A scene analysed using the base-ratio method at 64 frames-per-beat. Sharpening of the edges around objects is visible in 4.13d; the soft edges were due to mixed pixels. Figs. 4.14 and 4.15 are additional analyses of the same scene. Note: none of the intensity images are to the same scale.

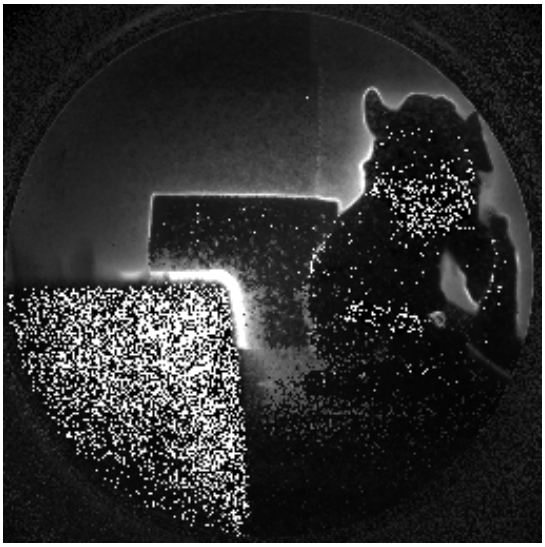

(a) Second Return Intensity In Two Return Case

(b) Log-Likelihood in Single Return Case, LL_1

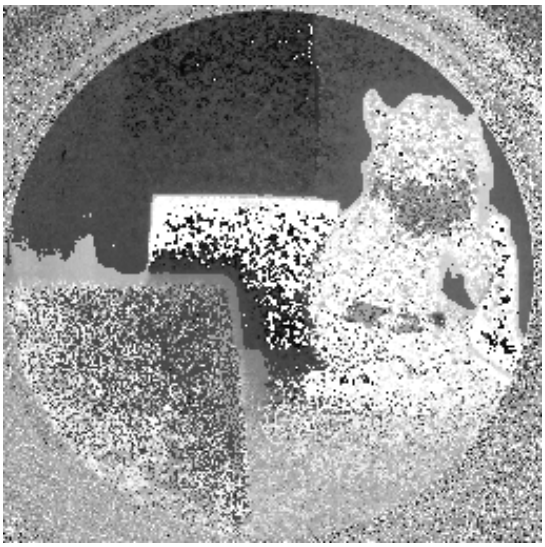

(c) Second Return Phase In Two Return Case

(d) Log-Likelihood Change Between One and Two Returns, ΔLL_{2-1}

Figure 4.14: A continued analysis of the scene from Fig. 4.13, giving second component returns and log-likelihoods. For 4.14d black represents values below zero, white the greatest positive ΔLL_{2-1}. Created using the base-ratio method and 64 frames-per-beat. Fig. 4.14a is deliberately saturated so as to bring out fine detail in the background. Note that the noisy regions in 4.14a mostly have a very low ΔLL_{2-1} in 4.14d. Figs. 4.13 and 4.15 are additional analyses of the same scene.

4.2.5 Estimating the Number of Returns

An example of the data output by the shape fitting methods is in Fig. 4.13. The challenge is to now identify the actual number of returns at each pixel. Note that the primary return/first return is the brightest return.

In an ideal situation there would be a simple local prior distribution describing the probability of a specific number of returns occurring within any pixel which could be incorporated into the optimisation process. Rather than use any specific prior, we instead use a simple method for global optimisation by prioritising returns which allows the user or possibly an advanced computer programme to determine the value of a single controlling parameter C_l.

The log-likelihood calculated as a part of the optimisation can then be used to determine the quality of a transition, namely, a change in state produced by incrementing the number of returns assumed to be within a pixel. The quality of the transition is defined by ΔLL_{2-1}, which is the change in log-likelihood caused by incrementing the assumed number of returns from one to two. The parameter, C_l, is the number of iterations that a greedy algorithm is allowed to execute. The algorithm starts by assuming that every pixel has a single return, and then each iteration finds the transition for which ΔLL_{2-1} is maximal and executes it. In this particular case the maximum number of returns is limited to two, however there is no intrinsic reason why the algorithm could not use a much larger number, or even start assuming no returns. A high ΔLL_{2-1} indicates that two returns fits the data better than one return. Given adequate fitting time, two returns should always be able to fit data at least as well as one return – however, we have found that a high ΔLL_{2-1} tends to indicate that LL_1 is a particularly poor fit and contained multiple returns. One might expect there to be problems with overfitting of the data, however in practice this does not appear to occur; this could potentially be due to the large size of the impulse response, which makes the method less sensitive to individual sample level fluctuations.

An example of the effect of C_l on a subregion of a range image and a demonstration of the impact of local minima is given in 4.15. The region within which multiple returns have been identified is very dark and thus much more susceptible to multiple returns from scattered light within the ranger. At the same time, the lack of intensity makes it harder to identify the range to the source of each return. Light from the object in front has been scattered onto the object at the back. The noise that is present in the images produced using only one optimisation attempt is due to local minima. In the $C_l = 400$ case only the clearly mixed pixels around the borders between objects are identified and separated out. As C_l increases, less

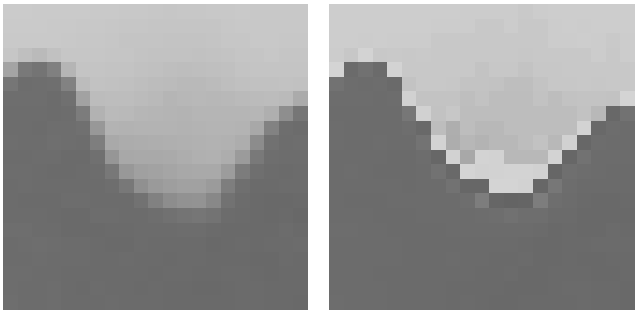

(a) Raw Phase

(b) Return 1 Phase at $C_l = 400$, Five Optimisation Attempts

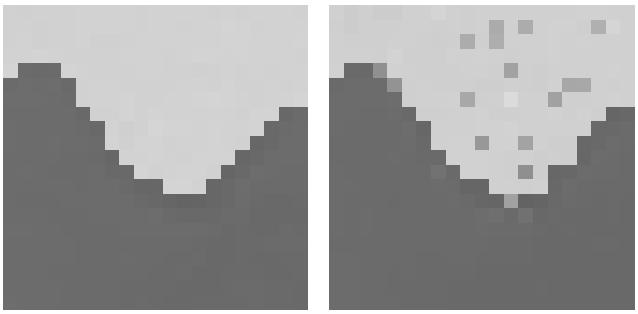

(c) Return 1 Phase at $C_l = 8000$, Five Optimisation Attempts

(d) Return 1 Phase at $C_l = 8000$, One Optimisation Attempt

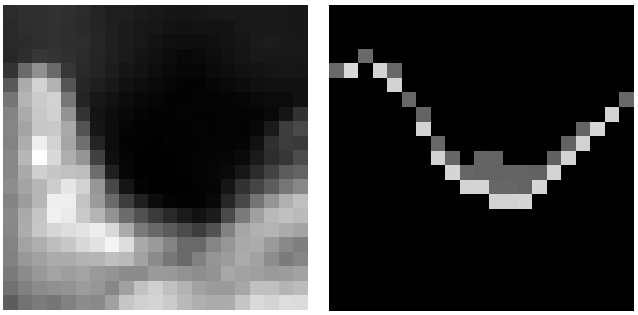

(e) Raw Intensity

(f) Return 2 Phase at $C_l = 400$, Five Optimisation Attempts

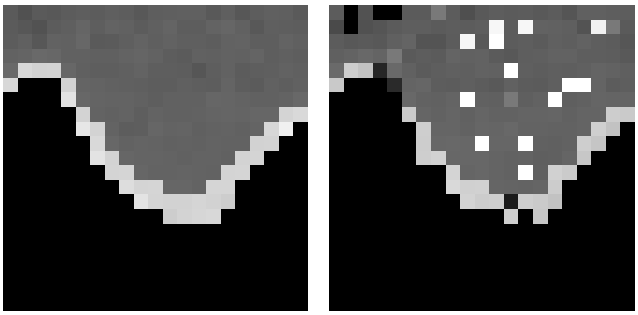

(g) Return 2 Phase at $C_l = 8000$, Five Optimisation Attempts

(h) Return 2 Phase at $C_l = 8000$, One Optimisation Attempt

Figure 4.15: An example of the impact that C_l and local minima have on an image using a subregion of scene 1 (from Fig. 4.13).

(a) First Return Intensity

(b) First Return Phase

(c) Second Return Intensity

(d) Second Return Phase

Figure 4.16: Recovered version of the scene from 4.13 using $C_l = 4000$. Mixed pixels around the edges of objects and some limited dark regions subject to scattered light have been recovered. Fig. 4.16b shows a noticeable transition between the one and two return regions due to the systematic overestimation of the phase difference between two returns. There is also an extremely dark object on the left which was not clearly recovered.

obviously mixed pixels containing scattered light are identified – eventually, if C_l is set too high, incorrect data is included in the recovered image. Notice that the C_l limiting was applied as a global optimisation, and this figure is just a subregion. A recovered version of the entire scene from Fig. 4.13 is given in Fig. 4.16.

4.3 Accelerating Waveform Shape Fitting

Section 4.2 developed two different shape models for the correlation waveform. Using optimisation these models were fit to measured correlation waveforms in order to determine the phase and intensity of each component return contributing to the signal. However, as a fundamental solution to the problem of mixed pixels/multipath interference this approach is inadequate as the presence of local minima decreases the quality of the recovered phase and amplitude estimates. Also, because the approaches require the use of numerical optimisation techniques, the computational cost is prohibitive. In Section 4.1 it was determined that use of the LF method improved the precision of range estimates, a similar effect was found with the shape fitting approaches in 4.2.

In this section we extend the base-ratio approach in a manner specific to the single component return case. By using the negative fundamental frequency bin of the Fourier transform we determine the coarse phase, ν_ξ, and then apply a weighted pseudoinverse to determine the fine phase, amplitude and an ambient light offset. The latter parameter is especially important for non-differential systems, where ambient light is unintentionally measured by the sensor, and to compensate for the much smaller temporal variation in sensor bias due to temperature. This parameter was not added to the earlier shape fitting models as it required the determination of an additional parameter by the numerical fitting algorithm.

4.3.1 Fitting An Updated Base-Ratio Model

The correlation waveform model from Eqn. 4.10 is rewritten as

$$f_\nu[i] = a(\psi[(i+\nu_\xi) \bmod n] + \lambda_\xi \Delta\psi[(i+\nu_\xi) \bmod n]) + \rho_\mathrm{I}, \tag{4.21}$$

where ρ_I is the contribution from ambient light. For a known coarse phase, ν_ξ, we can write f_ν as the linear combination of three basis vectors:

$$f_\nu \;=\; H_\nu \nu_\rho \tag{4.22}$$

$$= \begin{pmatrix} 1 & \psi[(\nu_\xi + 0) \bmod n] & \Delta\psi[(\nu_\xi + 0) \bmod n] \\ 1 & \psi[(\nu_\xi + 1) \bmod n] & \Delta\psi[(\nu_\xi + 1) \bmod n] \\ \vdots & \vdots & \vdots \\ 1 & \psi[(\nu_\xi + n - 1) \bmod n] & \Delta\psi[(\nu_\xi + n - 1) \bmod n] \end{pmatrix} \begin{pmatrix} \rho_{\mathrm{I}} \\ a \\ a\lambda_\xi \end{pmatrix}, \quad (4.23)$$

where ν_ρ is the vector of parameters and H_ν is the correlation waveform formation matrix transformation. As discussed in Section 3.4.1, the noise distribution for photon shot noise can be approximated as a Gaussian distribution where the variance is proportional to intensity. Thus the covariance matrix for the measured data, Σ_ν, is

$$\Sigma_\nu = \begin{pmatrix} \tilde{s}[0] & 0 & \cdots & 0 \\ 0 & \tilde{s}[1] & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \tilde{s}[n] \end{pmatrix}. \quad (4.24)$$

The pseudo-inverse of H in Eqn. 4.23 is equivalent to maximising the likelihood of ν_ρ given measured data $\tilde{s}$ and assuming a Gaussian error distribution. Solving the normal equations gives

$$\nu_\rho = (H_\nu^{\mathrm{T}} \Sigma_\nu^{-1} H_\nu)^{-1} H_\nu^{\mathrm{T}} \Sigma_\nu^{-1} \tilde{s}. \quad (4.25)$$

Solving this equation has a relatively trivial cost as opposed to the iterative numerical optimisation approach pursued in the previous section, however there is no ability to place constraints on the value of λ_ξ. Whereas in the previous section an explicit constraint was placed on the values of the fine phase, this is not possible for a pseudo-inverse. In some cases this can even result in negative values of λ_ξ, although this is infrequent. For relatively simple correlation waveform shapes this allows phase to be estimated even if the coarse phase is slightly misestimated, although it is subject to much greater potential for error. We now use this method to analyse the impact that illumination duty cycle has on ranging precision and overall error.

4.3.2 Analysing the Impact of Duty Cycle on Precision

A flat board was placed in front of the UoW range-imager and several sequences of beat cycles were captured at different laser modulation duty cycles at 20 MHz modulation frequency, 48 samples per beat, 1 beat per second. The correlation waveform was sampled and then convolutionally blurred in order to simulate a worst

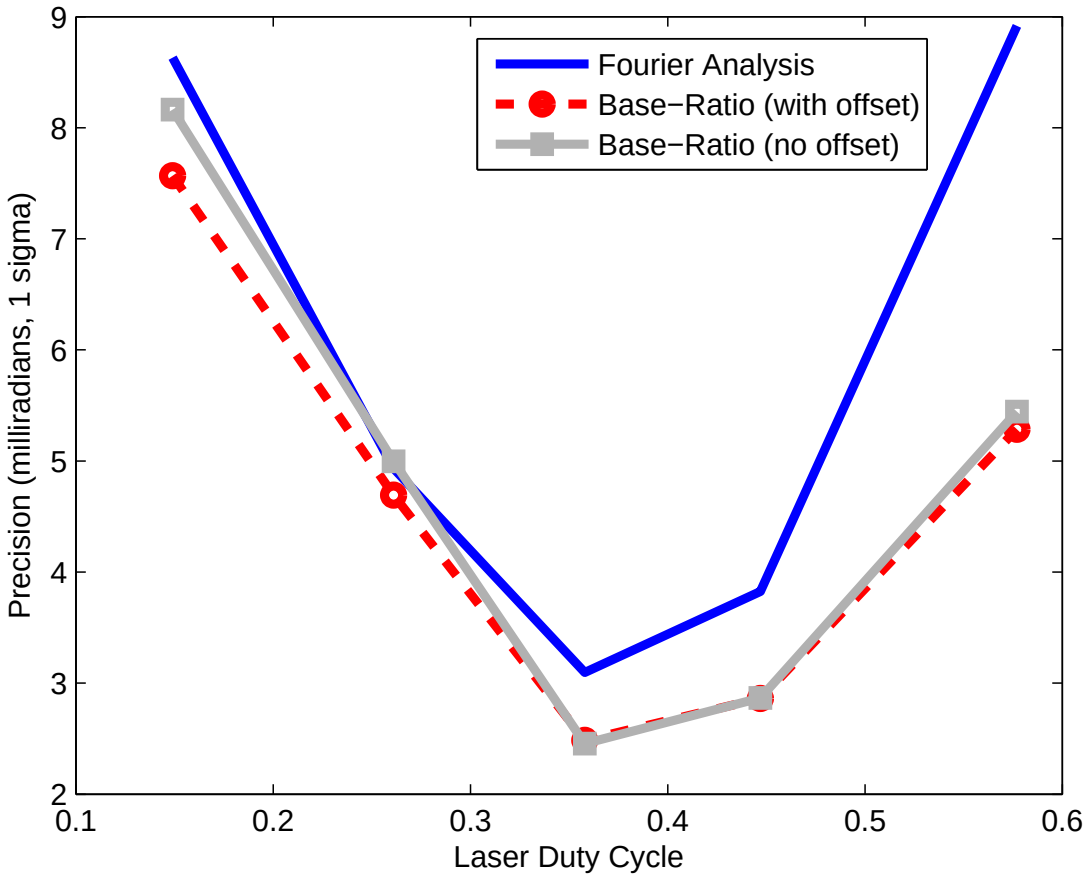

Figure 4.17: Phase measurement precision as a function of duty cycle and estimation algorithm

case interpolation scenario. Range measurements were produced by post-processing these data using the Fourier analysis method, Base-Ratio method and the new ML method. Precision was estimated by taking the sample standard deviation of each pixel's range measurements over time. Since the number of frames per cycle was sufficiently high to avoid aliasing, a range model was created by taking the Fourier analysis data, averaging it over time and then blurring it in the complex domain using a $\sigma = 5$ pixels Gaussian blur. This model was then considered to be an accurate estimate of the actual range. The overall error was then calculated as the RMS error of the range measurements versus this model. Fig. 4.19 was produced by taking the sampled 35.8% duty cycle reference waveform, adding varying amounts of Poisson distributed noise and then taking range measurements via the Fourier and new ML approach.

Precision results in Fig. 4.17 show a clear improvement in precision in the case of high duty cycles, with little improvement at low duty cycles. A contributing factor

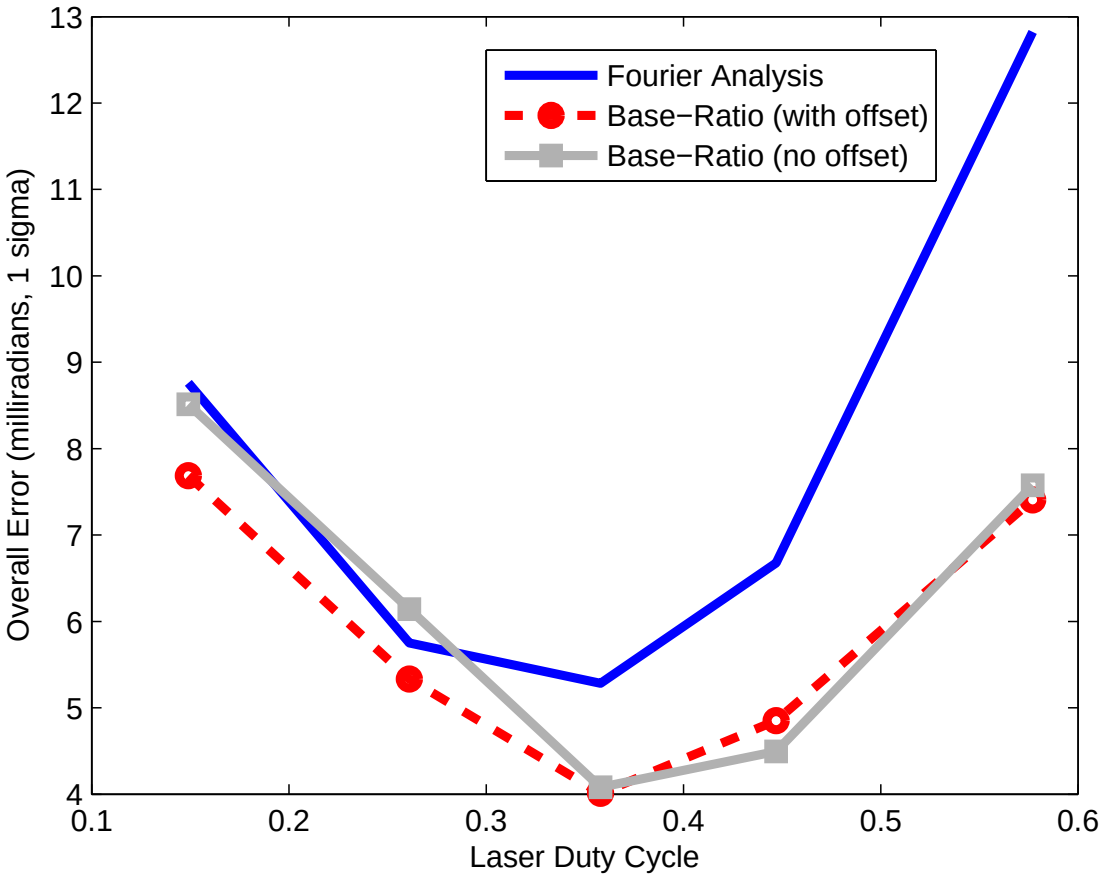

Figure 4.18: Overall phase measurement error as a function of duty cycle and estimation algorithm (including systematic and random components)

is the more complicated shape of the lower duty cycle waveforms. Out of the five different duty cycles sampled, a 35.8% duty cycle was clearly the best in all cases. Comparing overall error (Fig. 4.18) to precision indicates an increase in systematic error components as the duty cycles increase. It is not due to simple offsets in the data and may be related to noise properties.

For the Fourier analysis method, precision is inversely proportional to the square root of the SNR. If Poisson distributed shot noise is the only noise source, then the noise power is proportional to the total integrated light intensity. The signal strength is the absolute value of the fundamental Fourier bin, hence

$$precision \propto \frac{\sqrt{\sum_{i=0}^{n-1} \tilde{s}[i]}}{|\sum_{i=0}^{n-1} \tilde{s}[i] e^{-2j\pi \frac{i}{n}}|}. \tag{4.26}$$

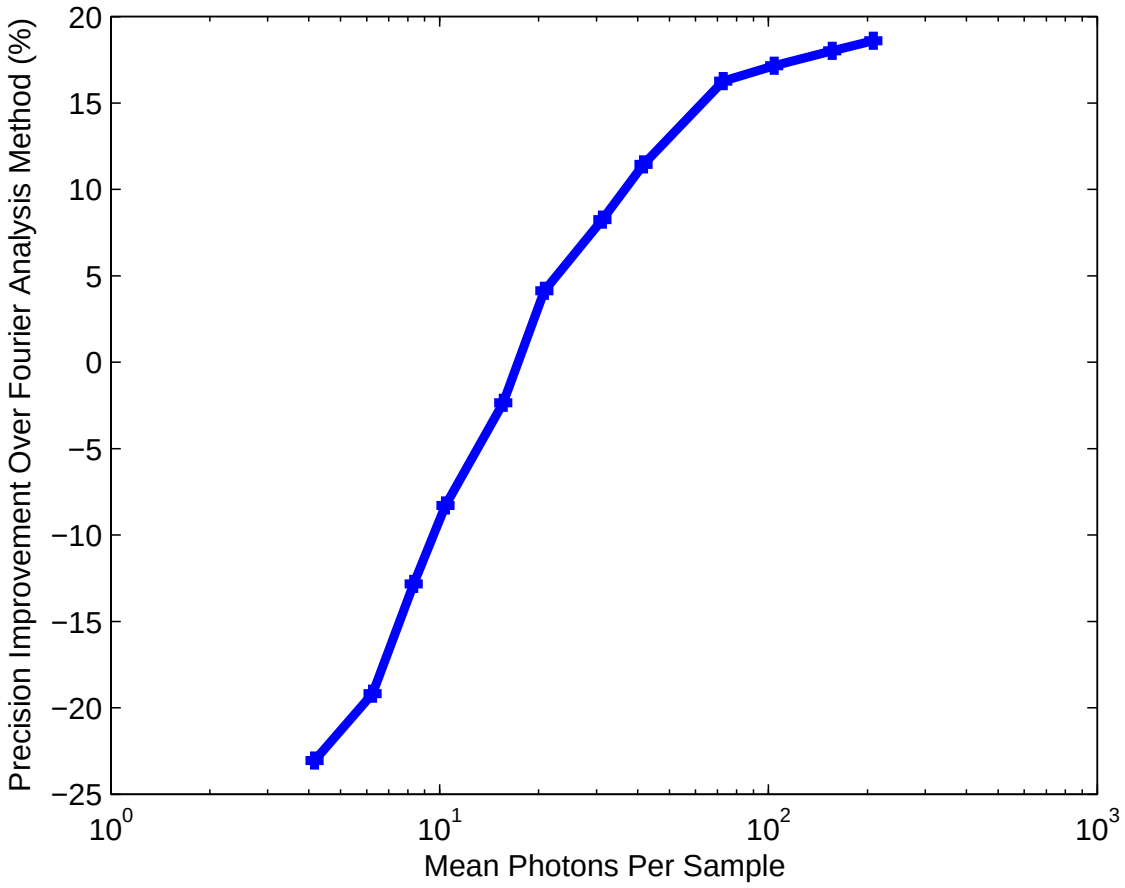

Figure 4.19: Improvement in phase precision versus mean number of photons per sample for the new method versus Fourier analysis. Results generated via simulation from sampled 35.8% laser duty cycle correlation waveform.

However, since we are limited to rectangular modulation,

$$precision \propto \frac{1}{\sqrt{x}\ \mathrm{sinc}(x)}, \tag{4.27}$$

where $\{x \in \mathbb{R}; 0 \leq x \leq 1\}$ is the duty cycle. This is a convex function over $[0, 1]$, with a global minimum at a 37.1% duty cycle that very closely matches the curve of the results in Fig. 4.17. Since this relationship is separable, this means that in the Fourier analysis case 37.1% is the optimum duty cycle for both the laser duty cycle and intensifier duty under all circumstances. This relationship does not appear to have been previous mentioned in the literature.

However, this does not necessarily hold when other noise sources are taken into account and when the new method is analysed. From an intuitive perspective on overall phase information content, the 44.7% duty cycle sample might have been

expected to have the greatest overall phase information content as the sloped region could be considered to be larger than any other waveform[2]. Correct determination of the factors that influence the precision of the new method requires further analysis, although it appears to follow the general trend of the precision of the Fourier analysis method.

A simulation (Fig. 4.19) shows how the amount of light collected affects the precision improvement. Low mean Poisson distributions are leptokurtic, thus poorly modelled by Gaussian distributions – hence below about 20 photons per sample, there is no precision improvement.

In general, the algorithm substantially improves overall RMS error and precision across a range of laser duty cycles versus the standard Fourier analysis method, also performing slightly better than the previous Base-Ratio approach without offset compensation. The measurement precision depends upon the laser duty cycle, which can be optimised based on the parameters of the intensifier modulation using a simple relationship. Because the new method does not require numerical optimisation, it is over two orders of magnitude faster than the original Base-Ratio fitting approach.

4.4 Looking for Better Approaches

In this chapter we described several possible new approaches to separating out mixed pixels and/or improving the precision of single component returns using the harmonic content of the correlation waveform. Of these only one method, using a weighted pseudo-inverse on a linear interpolation model, was plausible for a practical on-line system due to computational complexity. Even so, this method only improved the precision of range measurements, it did not separate out multiple components within a mixed pixel.

Fundamentally, all the approaches in this chapter are inadequate because an AMCW system is not designed for harmonic content in the correlation waveform and designing a system to optimise harmonic content results in a range-gating system, which is not a new invention. Because of noise limitations, in subsequent chapters we concentrate on sequences of measurements at different modulation frequencies. Solving for the component returns within a correlation waveform with harmonic content is mathematically equivalent to solving for the component returns

[2]It seems strange for large horizontal regions of the waveform to contribute significantly to improving estimates of fine phase, given that translation has little or no impact on these values. One might expect only regions which change when the waveform is translated would contribute to phase precision, ergo the larger the sloped portion of the waveform, the better the phase precision.

by taking measurements at different modulation frequencies using only the information contained in the fundamental frequency of the correlation waveform, except for resulting in a much higher SNR for any fixed measurement interval. While it is possible to synthesise a correlation waveform from measurements captured in this manner and then apply the methods presented here, the fact that it is possible to separate out two components using the correlation waveform harmonics suggests that there is a more fundamental mathematical solution: reduced to the simplest case, that involves taking two measurements at different modulation frequencies and applying these to determination of the amplitude and phase of two component returns within each pixel.

In the next chapter we introduce two different closed-form solutions for determination of the component returns within a mixed pixel. The first relies on four measurements and uses Cauchy distributions to model diffuse-range returns, such as those produced by volumetric scattering, the second uses attenuation ratio polynomials and requires two measurements of the attenuation ratio at different modulation frequencies, plus one standard complex domain range measurement.

Chapter 5

Closed-Form Approaches To Mixed Pixel Separation

In this chapter polynomial models for the multiple component return problem in both the point-return and diffuse-range cases are presented. We develop simple solutions for several two component return cases and explain the difficulties involved in solving additional cases. Unless specified otherwise, this treatment ignores the impact of aliasing and other systematic errors.

The first model addresses the diffuse-range case: this corresponds to backscattering sources which are distributed over a variety of ranges from the camera, such as volumetric scattering by fog and hair. The scattering of light in clouds is commonly addressed in doppler lidar for analysis and prediction of weather (Klett, 1985; Eloranta, 1998), but the methods are of little relevance to full-field AMCW lidar, which is primarily suited to short-range measurements. It is important to note that the term diffuse-range refers to closely located backscattering sources integrated by a single pixel: thus reflectivity which is diffuse over range (not to be confused with diffuse reflectivity or Lambertian reflectance).

In contrast to Chapter 4 – which developed waveform shape fitting methods in the spatial domain using numerical optimisation – this chapter presents closed-form solutions in the Fourier domain, which allows practical real-time implementation. While Chapter 4 used correlation waveform harmonics, as opposed to multiple separate measurements at different modulation frequencies, both methods implicitly sample exactly the same data. Theoretically, the approaches from Chapter 4 could be applied to a waveform synthesised from explicit measurements at different modulation frequencies, in the same manner as approaches detailed in this chapter could potentially be applied to implicit measurements of different spatial frequencies extracted from the harmonics of the correlation waveform at a single modulation

frequency.

Section 5.1 presents three different parametisations of the first model, although a closed-form inverse is only determined for the Cauchy distribution model of component returns[1]. The closed-form solution requires three complex domain measurements at relative frequencies of $1:2:3$ and one measurement of the total integrated intensity of the signal returns[2]. The solution is thus over-determined, taking a seven-dimensional input and returning a six-dimensional output (amplitude, phase and range-spread parameter for each of two component returns).

The second model assumes the point-return/Dirac delta model for component returns and determines the phase and amplitude relationship between the component returns using measurements of the attenuation ratio at two different modulation frequencies. Given at least one complex domain range measurement, it is then possible to determine candidate solutions for the absolute phase and amplitude of the component returns. Depending on the particular frequency ratio, there may be one or more possible candidate solutions. Whereas the first model attempts to determine the model parameters directly from complex domain range measurements, the second model requires knowledge of the attenuation ratio, which is not able to be directly measured in most systems. Typically, the calculation of the attenuation ratio requires a complex domain range measurement and a measurement of the total integrated intensity of the signal returns. As a result, despite the fact that the attenuation ratio method appears to be exactly-defined in the $2:1$ frequency ratio case, practical implementation requires an additional measurement. We have been unable to find any closed-form effectively exactly-determined approaches. Chapter 7 includes a brief comparison of the performance of the attenuation ratio method to other approaches.

5.1 Modelling Diffuse-Range Returns

In Chapter 4 component returns in signal return models were modelled as Dirac delta functions versus range. In reality, however, a component return may be spread over a region rather than an infinitesimal point. In this section we develop four polynomial models for component returns spread over a region of space: a Dirac delta model, a Uniform model, a Gaussian distribution model and a Cauchy distribution model.

[1]Strictly, there are four parametisations, but the Cauchy model is a simple, backwards compatible extension of the Dirac delta model.

[2]This is equivalent to sampling the zeroth spatial frequency of the signal return model. A direct measurement is difficult due to offsets from ambient light and temporal changes in measurement bias. The best approach is to approximate the total integrated intensity by the modulus of a measurement at an extremely low frequency, such as 10kHZ.

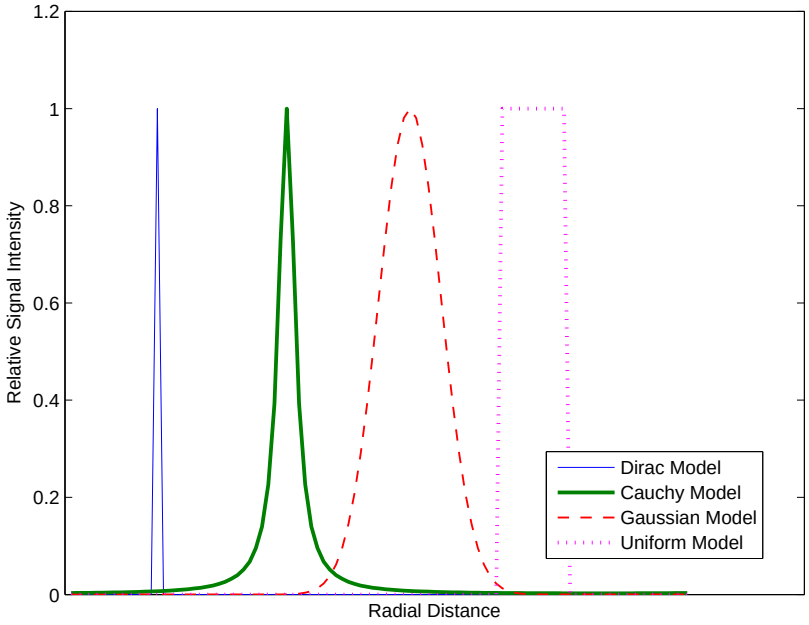

Figure 5.1: Candidate diffuse-range return models

Table 5.1: Diffuse-range return models at relative frequency r_l.

		Measurement Representations	
Distribution	Symbol	Cartesian	Polar
Dirac	$\delta(x)$	$\frac{\eta^{r_l}}{\lvert\eta\rvert^{r_l-1}}$	$a\nu^{r_l}$
Uniform	$\mathcal{U}(-\frac{x}{2w}, \frac{x}{2w})$	$\frac{\mathrm{sinc}(r_l w)}{\mathrm{sinc}(w)} \frac{\eta^{r_l}}{\lvert\eta\rvert^{r_l-1}}$	$\frac{\mathrm{sinc}(r_l w)}{\mathrm{sinc}(w)} a\nu^{r_l}$
Gaussian	$\mathcal{N}(0, \sigma^2)$	$\nu_{\mathcal{N}}^{r_l^2-1} \frac{\eta^{r_l}}{\lvert\eta\rvert^{r_l-1}}$	$a\nu^{r_l} \lvert\nu\rvert^{r_l^2-r_l-1}$
Cauchy	$\mathcal{L}(0, \lambda_{\mathcal{C}})$	$\nu_{\mathcal{L}}^{\lvert r_l\rvert-1} \frac{\eta^{r_l}}{\lvert\eta\rvert^{r_l-1}}$	$a\nu^{r_l}$ (+ve r_l) or $a(\nu^*)^{\lvert r_l\rvert}$ (-ve r_l)

An example of each of the four different models is given in Fig. 5.1 – this graph shows the measured intensity versus range for a single pixel.

5.1.1 Dirac Model

Fig. 5.1 and Table 5.1 show representations of four different diffuse-range return models. The most elementary model is the Dirac delta model, which assumes that the signal return model is a sparse spike train. In the measurement domain, this means that as relative frequency changes, phase varies but the modulus of the mea-

surement intensity is conserved.

In the general case, a measurement, ξ_l, at a relative frequency r_l can be notated as

$$\xi_l = \sum_{i=0}^{n-1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l-1}}. \tag{5.1}$$

Expressing each return in a Cartesian form,

$$\eta_i = R_i + jI_i, \tag{5.2}$$

where $R_i, I_i \in \mathbb{R}$, gives

$$\xi_l = \sum_{i=0}^{n-1} \frac{(R_i + jI_i)^{r_l}}{(R_i^2 + I_i^2)^{(r_l-1)/2}}. \tag{5.3}$$

This can be converted to polynomial form by rearranging to remove the fractional powers, with higher relative frequencies and more component returns resulting in increased polynomial order.

The two return case for positive, odd values of the relative frequency, $r_l \in \mathbb{Z}$, expands to

$$\xi_l S_0^{(r_l-1)/2} S_1^{(r_l-1)/2} - U_0^{r_l} S_1^{(r_l-1)/2} - U_1^{r_l} S_0^{(r_l-1)/2} = 0, \tag{5.4}$$

where $U_i = R_i + jI_i$ and $S_i = R_i^2 + I_i^2$. This polynomial is r_lth order over each of R_0, R_1, I_0 and I_1, and if necessary can be written explicitly as a separate real and a separate imaginary equation. Using the phase symmetry of measurements around a relative frequency of zero[3], Eqn. 5.4 can be rewritten for the negative, odd case as

$$\xi_l^* S_0^{-(r_l+1)/2} S_1^{-(r_l+1)/2} - (U_0^*)^{-r_l} S_1^{-(r_l+1)/2} - (U_1^*)^{-r_l} S_0^{-(r_l+1)/2} = 0, \tag{5.5}$$

where $U_i^* = R_i - jI_i$.

For positive, even relative frequencies we find

$$(\xi_l^2 S_0^{r_l-1} S_1^{r_l-1} - U_0^{2r_l} S_1^{r-1} - U_1^{2r_l} S_0^{r-1})^2 - U_0^{2r_l} U_1^{2r_l} S_0^{r-1} S_1^{r-1} = 0, \tag{5.6}$$

which is $4r_l$th order. Negative, even relative frequencies can be rewritten in the same manner as the previous case.

The final case is the zeroth spatial frequency, given by

$$(\xi_l^2 - S_0 - S_1)^2 - S_0 S_1 = 0, \tag{5.7}$$

[3]The negative frequency bins of the Fourier transform of a real function are complex conjugates of the equivalent positive bins. For a signal return model $f_\xi(r)$, this implies that $F_\xi(-u) = F_\xi(u)^*$.

which is fourth order.

Combinations of measurements of different spatial frequencies allow the separation of multiple component returns within a signal return model; in an ideal world, it would be possible to combine these systems of polynomial equations and determine a closed-form inverse. However, most combinations of measurements sufficient to separate out two or more component returns result in at least one polynomial of degree five or above with no obvious factorisation. Not only is there no general solution for the roots of polynomial of fifth order and above (due to the Abel-Ruffini theorem), but the solution also requires the application of a non-linear constraint – that of being a real root, not a complex root.

The only spatial frequencies with less than fifth order polynomials are 0, ± 1, and ± 3. The zeroth spatial frequency only has half the information content of a non-zero spatial frequency and measurements of positive and negative spatial frequencies are complex conjugates of each other. In order to solve the two component problem, two non-zero independent measurements are required. For example, if a measurement is taken at a relative frequency of one and another at two, there are two possible solutions corresponding to the two permutations of the same component returns. Ignoring positivity, there is only one frequency ratio that results in polynomials of less than fifth order – that of relative frequencies three and one. This gives the following simultaneous equations:

$$\xi_0 - U_0 - U_1 = 0 \tag{5.8}$$

$$\xi_1 S_0 S_1 - U_0^3 S_1 - U_1^3 S_0 = 0. \tag{5.9}$$

Substituting Eqn. 5.8 into Eqn. 5.9 gives

$$\begin{aligned} \xi_1(R_0^2 + I_0^2)((\pi_\mathrm{R} - R_0)^2 + (\pi_\mathrm{I} - I_0)^2) - (\xi_0 - R_0 - jI_0)^3(R_0^2 + I_0^2) \\ -(R_0 + jI_0)^3((\pi_\mathrm{R} - R_0)^2 + (\pi_\mathrm{I} - I_0)^2) = 0 \end{aligned} \tag{5.10}$$

where $\pi_\mathrm{R} = \Re(\xi_0)$ and $\pi_\mathrm{I} = \Im(\xi_0)$. This is fifth order and does not have a trivial factorisation, therefore it does not appear to lead to a plausible closed-form solution.

An alternative representation is to break each return into an amplitude component, $a_i \in \mathbb{R}^+$, and a phase component $|\nu| = 1$, where

$$\eta_i = a_i \nu_i. \tag{5.11}$$

This gives a particularly simple polynomial for each measurement, that is

$$\xi_l = \sum_{i=0}^{n-1} a_i \nu_i^{r_l}. \tag{5.12}$$

There are several plausible approaches to using this notation, one is to attempt to solve it using a minimum number of simultaneous measurements and non-linear constraints on the amplitude and phase variables – making these constraints explicit by breaking up the phase component into real and imaginary parts results in a set of simultaneous polynomials that can be rearranged to give the Cartesian polynomial form explicated earlier. Another approach is ignore these constraints and utilise a deliberately overdetermined situation in order to utilise a simple inverse – the main constraint being that this requires a large increase in the number of measurements required for multiple component separation, which may not always be possible due to device frequency limitations and the extra time and space required. In order to determine the phase and amplitude of two component returns, four measurements are required. We now provide a solution for the specific case of sampling relative frequencies $r_l \in \{0, 1, 2, 3\}$.

The fundamental equations are given by

$$\xi_0 = a_0 + a_1 \tag{5.13}$$
$$\xi_1 = a_0\nu_0 + a_1\nu_1 \tag{5.14}$$
$$\xi_2 = a_0\nu_0^2 + a_1\nu_1^2 \tag{5.15}$$
$$\xi_3 = a_0\nu_0^3 + a_1\nu_1^3. \tag{5.16}$$

From Eqn. 5.13,

$$a_1 = \xi_0 - a_0. \tag{5.17}$$

Substituting Eqn. 5.17 into Eqn. 5.14 gives

$$a_0 = \frac{\xi_0\nu_1 - xi_1}{\nu_1 - \nu_0}, \tag{5.18}$$

which combined with Eqn. 5.15 allows the determination that

$$\nu_1 = \frac{\xi_1\nu_0 - \xi_2}{\xi_0\nu_0 - \xi_1}. \tag{5.19}$$

Substituting these equations into Eqn. 5.16 gives a polynomial equation, namely,

$$\alpha\nu_0^2 + \beta\nu_0 + \gamma = 0, \tag{5.20}$$

where $\alpha = \xi_0\xi_2 - \xi_1^2$, $\beta = \xi_1\xi_2 - \xi_0\xi_3$ and $\gamma = \xi_1\xi_3 - \xi_2^2$. This can be trivially solved to give two different solutions corresponding to the two different permutations of the same two component returns. This is a nice solution in that there is a direct inverse, but is problematic in a practical sense because it requires four separate measurements when two is enough to fully constrain the problem. One mitigating aspect, is that by measuring to what extent the positive real number constraints on $\{R_0, R_1, I_0, I_1\}$ are broken, may give a measure of how much error there is likely to be in the component return estimates.

5.1.2 Uniform Model

If the component return is assumed to be uniformly spread over a region with a depth of w, then the return can be considered to be given by a translated rect function. Taking the Fourier transform of a rect function gives a sinc function in the Fourier domain, viz

$$\mathcal{F}\left\{\text{rect}\left(\frac{x}{w}\right)\right\} = \text{sinc}(uw). \tag{5.21}$$

Given that a component return, η, is notated at a relative frequency of one, a uniform model is otherwise identical to a Dirac model apart from a relative attenuation constant given by

$$\rho_{\text{U}} = \frac{\text{sinc}(r_l w)}{\text{sinc}(w)} \tag{5.22}$$

such that

$$\xi_l = \rho_{\text{U}} \sum_{i=0}^{n-1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l - 1}}. \tag{5.23}$$

In the single component return case, the width w can be determined by dividing the moduli of measurements at two appropriately chosen frequencies. There are several issues, including that there does not appear to be a simple inverse – perhaps suggesting a practical implementation utilising a LUT. Additionally, the inverse is a multivalued function; in practice, it is reasonable to always assume that the smallest solution is correct, because the width of the return is necessarily going to be substantially smaller than the measurement ambiguity intervals.

Unfortunately, a uniform model does not lend itself to easy inversion. Whereas in Section 5.1.1 there was a simple closed-form inverse for the two component return case, in the two component return case with two unknown width parameters

there are now six unknown parameters and equations which involve transcendental functions. Consequently, we do not address this model further.

5.1.3 Gaussian Model

A Gaussian model is a very natural model for component returns that originate from a region of space. Gaussian mixture models are distributions composed of the weighted sum of multiple Gaussian distributions and commonly occur in statistics (Bishop, 2006). These models are commonly fit to measurements in pulsed radar and lidar (Hofton *et al.*, 2000) in order to determine the range and amplitude of component returns. These systems typically involve direct sampling in the spatial domain, whereas AMCW lidar relies on sparse sampling directly in the Fourier domain and common numerical fitting approaches such as Expectation Maximisation (Bishop, 2006) are not guaranteed to converge to the correct answer due to local minima.

The zero centred Gaussian distribution with a variance of σ^2 is given by

$$\mathcal{N}(0,\sigma^2) = \frac{1}{\sqrt{2\pi}\sigma}\mathrm{e}^{-\frac{x}{2\sigma^2}}. \tag{5.24}$$

Taking the Fourier transform gives

$$\mathcal{F}\{\mathcal{N}(0,\sigma^2)\} = \mathrm{e}^{-2\pi^2u^2\sigma^2}. \tag{5.25}$$

The attenuation constant encodes the rate at which the specific Gaussian distribution attenuates modulation frequencies as modulation frequency increases. From the preceeding equation, we can write the attenuation constant as

$$\nu_{\mathcal{N}} = \mathrm{e}^{-2\pi^2\sigma^2}. \tag{5.26}$$

Using the attenuation constant we can reexpress the Gaussian model from Eqn. 5.1 as

$$\xi_l = \nu_{\mathcal{N}}^{r_l^2-1}\sum_{i=0}^{n-1}\frac{\eta_i^{r_l}}{|\eta_i|^{r_l-1}}. \tag{5.27}$$

In the single return case, the attenuation constant (and thus width) can be estimated from two measurements, ξ_0 and ξ_1, at relative frequencies, r_0 and r_1 by

$$\nu_{\mathcal{N}} = \sqrt[r_1^2-r_0^2]{\frac{|\xi_1|}{|\xi_0|}}. \tag{5.28}$$

Attempting to notate this model in polynomial form results in very complicated polynomials which need to be solved simultaneously. In Cartesian form, the polynomials are of the same order as in the Dirac delta case (see Section 5.1.1). In polar form, odd relative frequencies result in $(3|r|-1)/2$th order polynomials, while even relative frequencies give $2(3|r|-1)$th order polynomials. As a result, a Gaussian model looks unsuitable for closed-form inversion and we do not analyse it further.

5.1.4 Cauchy Model

In Section 5.1.1 we developed a polar form for representing measurements of a scene sampled at different relative frequencies – as a result, we unintentionally developed a model which allows both the representation of a diffuse-range component return and direct closed-form inversion. We now show how the Dirac delta model can be extended to handle diffuse-range component returns.

In statistics, the PDF of a zero centred Cauchy-Lorentz distribution is given by

$$f_{\mathcal{C}}(x|0,\lambda_{\mathcal{C}}) = \frac{\lambda_{\mathcal{C}}}{\pi(x^2+\lambda_{\mathcal{C}})}. \tag{5.29}$$

The Fourier transformation of the Cauchy distribution is given by a zero centred Laplace distribution, viz

$$\mathcal{F}\{f_{\mathcal{C}}(x|0,\lambda_{\mathcal{C}})\} = f_{\mathcal{L}}(u|0,\lambda_{\mathcal{C}}) = \frac{1}{2\lambda_{\mathcal{C}}}\mathrm{e}^{-\frac{|u|}{\lambda_{\mathcal{C}}}}. \tag{5.30}$$

If we assume that the spatial distribution of each component return is given by a Cauchy distribution, then we can model the relationship between measurements at different positive relative frequencies using Eqn. 5.12, that is

$$\xi_l = \sum_{i=0}^{n-1} a_i \nu_i^{r_l}. \tag{5.31}$$

For negative frequencies we introduce a minor extension to handle the conjugate relationship with positive frequencies:

$$\xi_l^* = \sum_{i=0}^{n-1} a_i \nu_i^{|r_l|}. \tag{5.32}$$

Whereas in Section 5.1.1 it was assumed that $|\nu_i| = 1$, we now use the modulus of ν_i to encode the attenutating properties of the Cauchy distribution in the Fourier

domain. In other words, let ν_i be defined as

$$\nu_i = \frac{\eta}{|\eta|} \mathrm{e}^{-1/\lambda_{\mathcal{C}}}, \tag{5.33}$$

thus the relationship between the modulii of measurements of a component return at relative frequencies r_0 and r_1 are given by

$$\frac{f_{\mathcal{L}}(r_1|0,\lambda_{\mathcal{C}})}{f_{\mathcal{L}}(r_0|0,\lambda_{\mathcal{C}})} = \frac{|\nu_i|^{|r_1|}}{|\nu_i|^{|r_0|}}. \tag{5.34}$$

It can be clearly seen that the Dirac delta model is merely a degenerate case of the Cauchy model, when $\lambda_{\mathcal{C}} \to 0$. Taking into account the conjugate relationship for negative relative frequencies, the inversion formula from Eqns. 5.17 to 5.20 can be used to find the tuple (a_0, a_1, ν_0, ν_1). Whereas for a Dirac delta model, three complex and one positive real domain measurements are required in order to determine four positive real values, in this case three complex and one positive real domain measurements are required in order to determine two real and two complex domain values. This is significantly less overdetermined, thus more efficient. One major limitation to this approach is that it is unable to handle more than two component returns because the polynomials become too large for closed-form inverses; solving for a greater number of component returns ultimately requires a numerical/iterative or lookup table approach.

5.2 Attenuation Ratio Polynomials

In the previous section we introduced the idea of modelling range-diffuse measurements using simultaneous multivariate polynomial equations; while in a number of overdetermined cases there are closed-form solutions, generally the polynomial equations involved are too complicated and high order for a closed-form solution. This section introduces a different representation for Dirac delta type point-returns using the attenuation ratio. This allows the relationship between component returns to be analysed independently of the absolute phase and amplitude; as a result, the polynomial order can be reduced. In order to introduce the attenuation ratio, we first develop a new concept of complex domain range measurement normalisation.

5.2.1 Normalisation of Range Measurements

This section introduces a new transfomation referred to as complex domain measurement normalisation, or normalisation for short. Often, normalisation of a complex

number implies division of one complex number by another, however in the case of complex domain range measurements there is a more useful type of normalisation. There are two equivalent types of normalisation developed here, normalisation of measurements and normalisation of component returns.

An example of the more intuitative component return normalisation is given in Fig. 5.2. The technique works by scaling phase and amplitude so that the closest component has a fixed intensity of 1 and a distance of 0. As a result, the intensity and phase of the normalised second component completely encodes the relative intensity and distance of the two components.

Measurement normalisation works in the same way, by normalising one measurement to one such that the other measurement encodes the relationship between the two measurements. In previous work the divison of one complex domain range measurement by another has been used for the synthesis of synthetic wavelength measurements in order to extend the ambiguity interval of an AMCW lidar or interferometer system (Dändliker *et al.*, 1998). The measurement normalisation technique developed here differs in that the complex domain range measurements are converted to the same relative frequency before division. Thus complex domain measurement normalisation works by factoring in the difference in modulation frequency so that the normalised measurement represents the relative amplitude and distance, although the distance may not always have a direct physical connotation. The remainder of this section develops a mathematical model for the normalisation concept.

We use the same phasor notation utilised earlier in the thesis for each component return but now analyse it in more detail. Each return within an arbitrary pixel is notated as a phasor $\eta_i \in \mathbb{C}$ as if it were a pixel with a single return measured at a particular spatial frequency $v = -\frac{2}{\lambda}$. Let a phasor, η_i, be defined by

$$\eta_i = a_i \mathrm{e}^{-2\pi j v d_i} = a_i \mathrm{e}^{j\theta_i}, \tag{5.35}$$

where $a_i \in \mathbb{R}^+$ is the intensity and the range from the camera, $d_i \in \mathbb{R}^+$, is represented as a phase angle $\theta_i \in [0, 2\pi)$. An ideal, unperturbed AMCW lidar measurement of a particular pixel at a spatial frequency vr_l, where $r_l \in \mathbb{Z}$ is the relative spatial frequency, is given by

$$\xi_l = F_\xi(vr_l) = \sum_{i=0}^{n-1} a_i \mathrm{e}^{jr_l\theta_i} = \sum_{i=0}^{n-1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l-1}}, \tag{5.36}$$

where n is the number of component returns. In cases other than $r_l = 1$, the equation

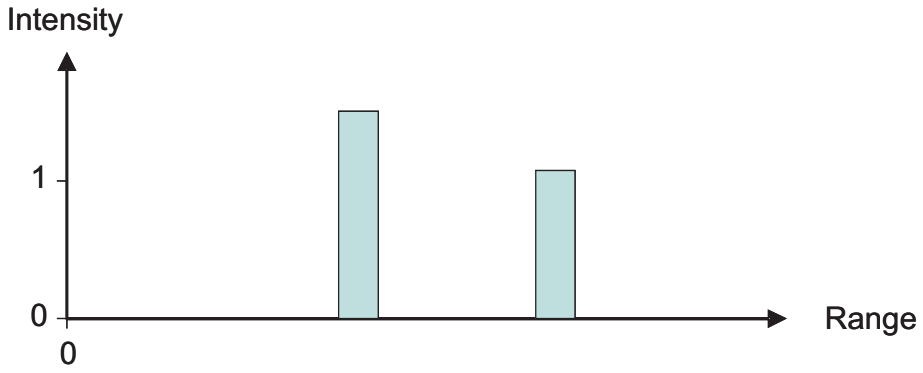

(a) Before Normalisation by the First Component Return

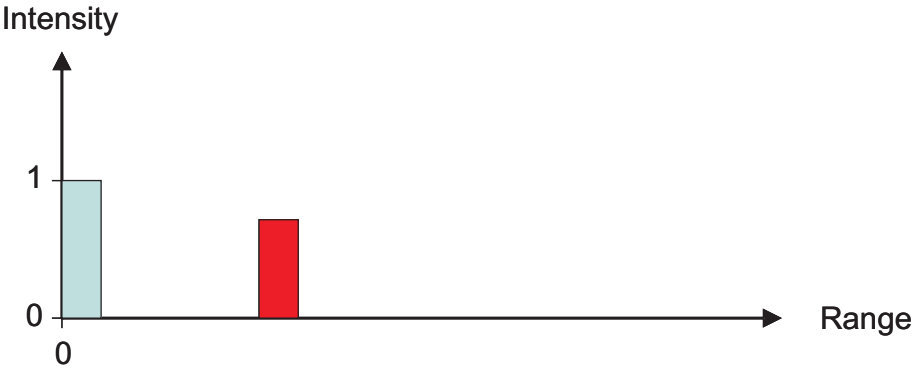

(b) After Normalisation by the First Component Return

Figure 5.2: Graphical illustration of the normalisation of component returns. Afterwards, the red component completely encodes the intensity and distance relationship between the two component returns. Measurement normalisation can be considered in an analogous manner, although it is more difficult to physicaly visualise due to the different spatial frequencies.

is non-holomorphic in terms of the input phasors η_i as the complex modulus operator does not not satisfy the Cauchy-Riemann theorem.

A sequence of measurements at different relative frequencies forms a set of simultaneous equations relating the set of component returns $\mathrm{H} = \{\eta_0, \eta_1, ..., \eta_{n-1}\}$ to the set of measurements $\Xi = \{\xi_0, \xi_1, ..., \xi_{N-1}\}$. For convenience, let

$$\Xi = R(\mathrm{H}) \tag{5.37}$$

be the forward relationship.

Multiplication of a component return η_i by a complex coefficient is equivalent to scaling in intensity and translating the component return in range. For a scaling coefficient $\alpha \in \mathbb{R}^+$ and a distance $d \in \mathbb{R}$, the translated return $\hat{\eta}_i$ is given by

$$\hat{\eta_i} = \mu_{\mathrm{n}} \eta_i, \tag{5.38}$$

where the complex coefficient, μ_{n}, is defined by

$$\mu_{\mathrm{n}} = \alpha \mathrm{e}^{-2\pi j v d} = \alpha \mathrm{e}^{j\phi}, \tag{5.39}$$

where ϕ is the phase shift equivalent to translation by d. Multiplying the set of all component returns H by μ_{n} is equivalent to multiplying each range measurement ξ_l by a complex coefficient μ_l, such that

$$\begin{aligned} \hat{\xi}_l &= \sum_{i=0}^{n-1} \frac{(\mu_{\mathrm{n}}\eta_i)^{r_l}}{|\mu_{\mathrm{n}}\eta_i|^{r_l-1}} && (5.40)\\ &= \sum_{i=0}^{n-1} \alpha a_i \mathrm{e}^{j r_l(\theta_i+\phi)} && (5.41)\\ &= \mu_l \sum_{i=0}^{n-1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l-1}}, && (5.42) \end{aligned}$$

where μ_l is defined for a relative frequency of r_l by

$$\mu_l = \alpha \mathrm{e}^{-2\pi j r_l v d} = \alpha \mathrm{e}^{j r_l \phi} = \frac{\mu_{\mathrm{n}}^{r_l}}{|\mu_{\mathrm{n}}|^{r_l-1}}. \tag{5.43}$$

In order to maintain consistency between measurements, the entire set of measurements Ξ must be transformed at once. Use of this transformation allows normalisation of H relative to a single component return or of Ξ relative to a single measurement.

For example, one can model the relationship between component returns in the two return case by normalising the value of the first component return, η_0, to a value of 1. This is equivalent to setting $\mu_{\mathrm{n}} = \eta_0^{-1}$, giving

$$\hat{\mathrm{H}} = \mu_{\mathrm{n}}(\eta_0, \eta_1) = \left(1, \frac{\eta_1}{\eta_0}\right) = (1, b\mathrm{e}^{j\theta}), \tag{5.44}$$

where $b = a_1/a_0$ is the relative magnitude of the two component returns and $\theta = \theta_1 - \theta_0$ is the relative phase. The normalised measurement is found by

$$\begin{aligned} \hat{\xi}_l &= \mu_l \sum_{i=0}^{1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l-1}} = \mu_l \sum_{i=0}^{1} a_i \mathrm{e}^{j r_l \theta_i} && (5.45)\\ &= \frac{a_0 \mathrm{e}^{j r_l \theta_0}}{a_0 \mathrm{e}^{j r_l \theta_0}} + \frac{a_1 \mathrm{e}^{j r_l \theta_1}}{a_0 \mathrm{e}^{j r_l \theta_0}} && (5.46)\\ &= \Lambda_{\mathrm{f}}(b, r_l\theta) && (5.47) \end{aligned}$$

where the function $\Lambda_{\mathrm{f}}(c, \phi_{\mathrm{f}})$ is defined in order to assist in modelling the relationship

between the two component returns, as

$$\Lambda_{\mathrm{f}}(c, \phi_{\mathrm{f}}) = 1 + c\mathrm{e}^{j\phi_{\mathrm{f}}}. \tag{5.48}$$

For $b \in [0, 1]$, $\Lambda_{\mathrm{f}}(b, \theta)$ can be considered the perturbation of the primary component return by the second. By denormalising Eqn. 5.47, any two component measurement can be modelled by

$$\xi_l = \frac{\eta_0^{r_l}}{|\eta_0|^{r_l - 1}} \hat{\xi}_l = a_0 \mathrm{e}^{j r_l \theta_0} \Lambda_{\mathrm{f}}(b, r_l \theta). \tag{5.49}$$

5.2.2 The Attenuation Ratio

Let the attenuation ratio τ_l at a particular modulation frequency $v r_l$ be

$$\tau_l = \frac{|\xi_l|}{\kappa} = \frac{\left| \sum_{i=0}^{n-1} \frac{\eta_i^{r_l}}{|\eta_i|^{r_l - 1}} \right|}{\sum_{i=0}^{n-1} |\eta_i|}, \tag{5.50}$$

where κ is the total integrated intensity, namely the sum of the modulii of the component returns. Potentially, this can be approximated by taking the modulus of a measurement at an extremely low modulation frequency. In the two return case we can model this using Eqn. 5.49 by

$$\tau_l = \frac{a_0 |\Lambda_{\mathrm{f}}(b, r_l \theta)|}{a_0 (1 + b)} = \frac{|\Lambda_{\mathrm{f}}(b, r_l \theta)|}{1 + b}. \tag{5.51}$$

The attenuation ratio is the ratio of the net active intensity to the sum total of the active intensity of the component returns and indicates how much cancellation has occurred between component returns for a particular pixel.

A geometric interpretation of Eqn. 5.48 combined with the law of cosines gives

$$|\Lambda_{\mathrm{f}}(c, \phi_{\mathrm{f}})|^2 = 1 + c^2 + 2c \cos(\phi_{\mathrm{f}}). \tag{5.52}$$

Assume two measurements, ξ_0 and ξ_1, at relative frequencies, r_0 and r_1, with known attenuation ratios, τ_0 and τ_1. Solving Eqn. 5.51 to find $|\Lambda_{\mathrm{f}}(b, \theta)|$ and combining with Eqn. 5.52 gives

$$|\Lambda_{\mathrm{f}}(b, r_l \theta)|^2 = 1 + b^2 + 2b \cos(r_l \theta) = \tau_l^2 (1 + b)^2, \tag{5.53}$$

hence

$$b^2 + b \frac{2 \cos(r_l \theta) - 2\tau_l^2}{1 - \tau_l^2} + 1 = 0. \tag{5.54}$$

This is the attenuation ratio polynomial for a measurement at a particular relative

frequency r_l. Given a sequence of attenuation ratios at different spatial frequencies, a set of polynomials is formed which can be simultaneously solved in the same manner as the Cartesian and polar models for diffuse-range returns discussed earlier in this chapter. We now address the solution of these polynomial systems.

5.2.3 Relative Phase Polynomials

Eqn. 5.54 is separately valid for each measurement, thus

$$\frac{2\cos(r_0\theta) - 2\tau_0^2}{1-\tau_0^2} = \frac{2\cos(r_1\theta) - 2\tau_1^2}{1-\tau_1^2}, \tag{5.55}$$

or

$$\alpha\cos(r_1\theta) - \beta\cos(r_0\theta) + \gamma = 0, \tag{5.56}$$

where $\alpha = 1-\tau_0^2$, $\beta = 1-\tau_1^2$ and $\gamma = \tau_0^2 - \tau_1^2$. Eqn. 5.56 rewritten in terms of $\cos(\theta)$ using Chebyshev polynomials gives the attenuation ratio relative phase polynomial

$$\alpha T_{r_1}(\cos(\theta)) - \beta T_{r_0}(\cos(\theta)) + \gamma = 0, \tag{5.57}$$

where $T_n(x)$ is a Chebyshev polynomial of the first kind, defined by the trigonometric identity

$$T_n(x) = \cos(n\cos^{-1}(x)). \tag{5.58}$$

Chebyshev polynomials are equivalent to a cosine basis resampled using the function $\sqrt{1-x^2}$ and are commonly used for polynomial fitting because they offer a good approximation to the min-max polynomial and avoid Runge's Phenomenon (Press *et al.*, 2007). The first four Chebyshev polynomials are given by

$$T_1(x) = x, \tag{5.59}$$
$$T_2(x) = 2x^2 - 1, \tag{5.60}$$
$$T_3(x) = 4x^3 - 3x, \tag{5.61}$$
$$T_4(x) = 8x^4 - 8x^2 + 1, \tag{5.62}$$

By rewriting Eqn. 5.56 using Chebyshev polynomials, a seemingly difficult equation has been reduced to a relatively low order polynomial.

5.2.4 Determination of Relative Phase

The order of the polynomial given by Eqn. 5.57 is given by $\max(r_1, r_0)$. As the order increases, closed-form solutions become less useful and eventually impossible. The simplest possible case is for a frequency ratio of two to one.

Given $r_1 = 2$ and $r_0 = 1$, Eqn. 5.57 becomes

$$2\alpha\cos^2(\theta) - \beta\cos(\theta) - \alpha + \gamma = 0, \tag{5.63}$$

which has the solutions

$$\cos(\theta) \in \left\{1, -\frac{2\alpha - \beta}{2\alpha}\right\}. \tag{5.64}$$

Closed-form solutions for other frequencies are given in Table 5.2. The solutions for $r_1 = 4$ are not given, due to their extreme size, which limits their usefulness. Every equation has α in the denominator resulting in a substantial increase in noise sensitivity as α approaches zero, which occurs in regions with small perturbations. The relative phase and amplitude as a function of the attenuation ratio at different frequency ratios are plotted in Figs. 5.3 and 5.4. Both 3 : 2 and 3 : 1 result in multiple solutions, with the number of solutions varying as a function of attenuation ratio.

Table 5.2 gives known closed-form solutions for different frequency ratios. In addition to the explicitly listed solutions in the table, all relative phase polynomials have $\cos(\theta) = 1$ as an additional solution. This is because the solution equations in Table 5.2 become undefined as $\alpha, \beta \to 0$, which only occurs when there is a single component return.

Table 5.2: Known closed-form solutions for relative phase polynomials

r_1	r_0	Cosine of Relative Phase (Solutions for $\cos(\theta)$, excluding Unity)
2	1	$-(2\alpha - \beta)/(2\alpha)$
3	1	$-(\alpha \pm \sqrt{\alpha\beta})/(2\alpha)$
3	2	$-(2\alpha - \beta \pm \sqrt{\beta(4\alpha + \beta)})/(4\alpha)$
4	x	(Too Big to Print)

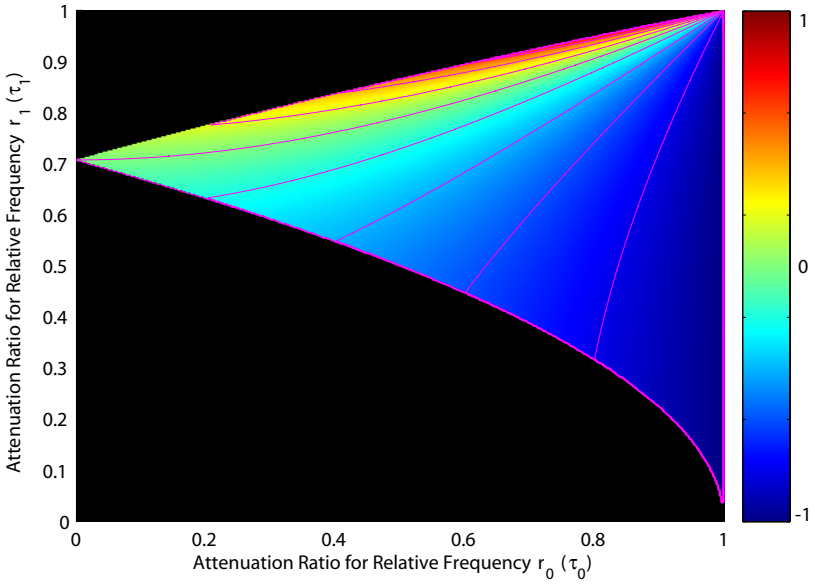

(a) Two to One

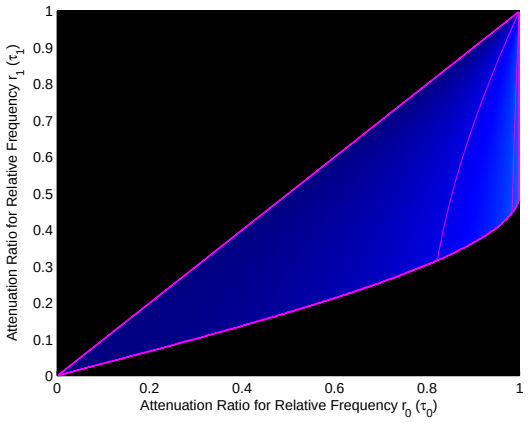

(b) Three to One (Solution One)

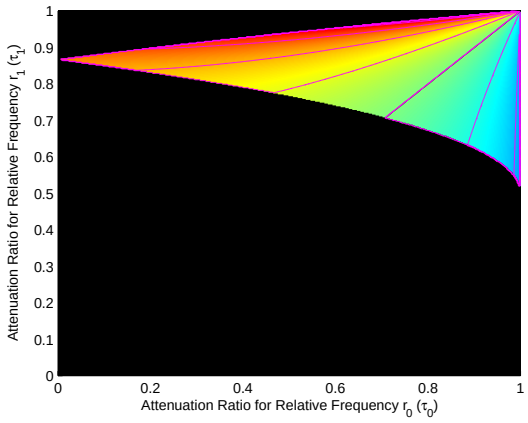

(c) Three to One (Solution Two)

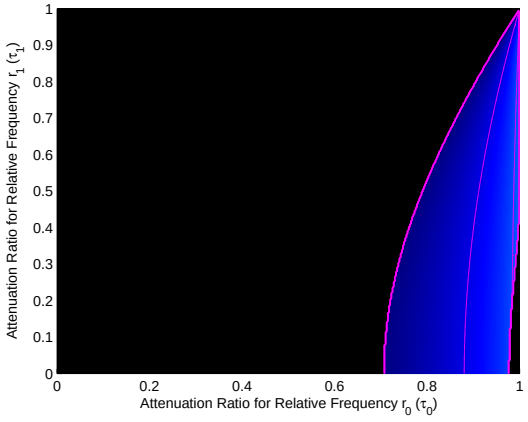

(d) Three to Two (Solution One)

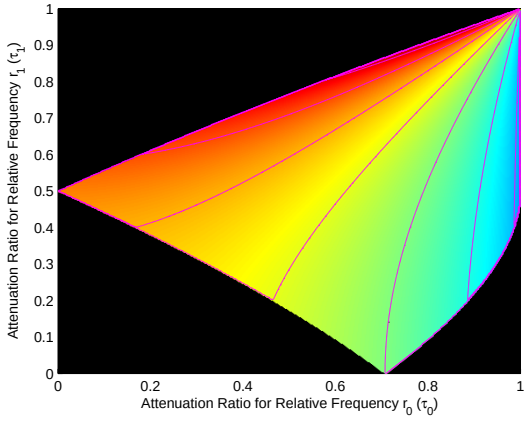

(e) Three to Two (Solution Two)

Figure 5.3: The cosine of relative phase plotted against the attenuation ratio tuple (τ_0, τ_1) for different frequency ratios. Black represents no valid solution, blue $= -1$ and red $= 1$. Each contour line represents an increment of 0.2.

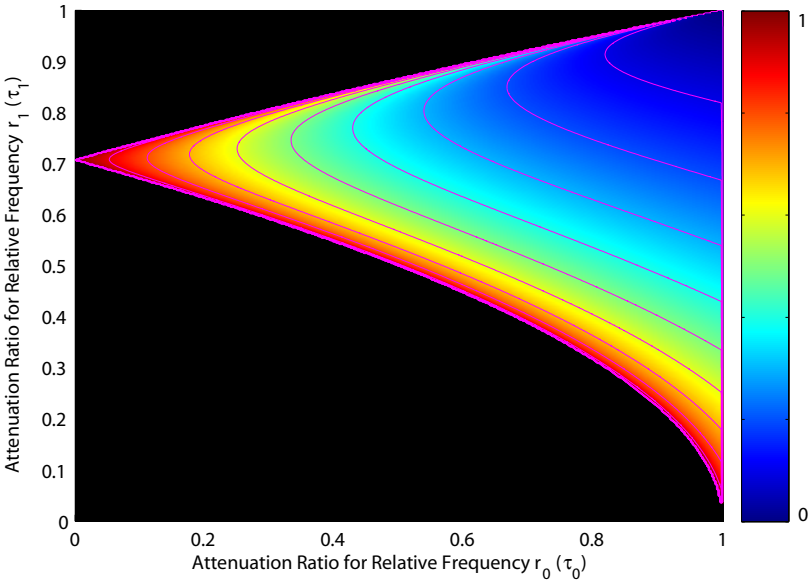

(a) Two to One

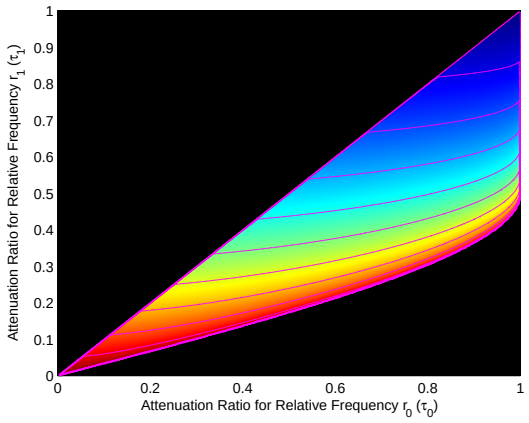

(b) Three to One (Solution One)

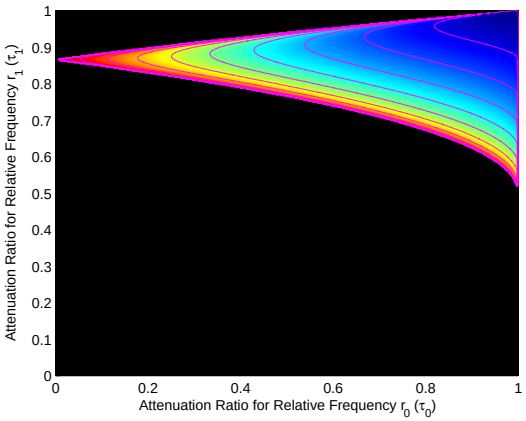

(c) Three to One (Solution Two)

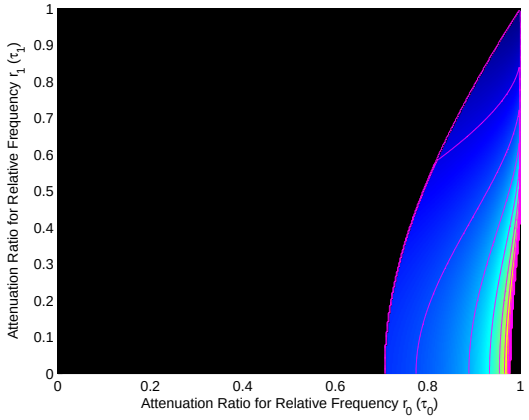

(d) Three to Two (Solution One)

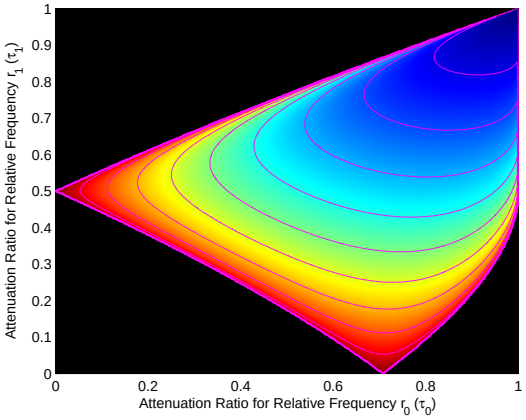

(e) Three to Two (Solution Two)

Figure 5.4: Relative amplitude plotted against the attenuation ratio tuple (τ_0, τ_1) for different frequency ratios. Black represents no valid solution, blue = 0 and red = 1. Each contour line represents an increment of 0.1.

5.2.5 Determining Relative Intensity

The relative intensity, b, is determined from Eqn. 5.54 by substituting the second solution from Eqn. 5.64, giving

$$b = \frac{\tau_0^2 - \cos(\theta) \pm \sqrt{(1-\cos(\theta))(2\tau_0^2 - \cos(\theta) - 1)}}{1-\tau_0^2} \tag{5.65}$$

For valid (τ_0, τ_1), the positive square root variation gives $b \geq 1$ and the negative variation gives $b \in [0, 1]$ – the alternative values for b being reciprocals of each other. At most there are two unique candidate solution tuples (b, θ) and $(b, -\theta)$, since $(b^{-1}, -\theta)$ and (b^{-1}, θ) correspond to alternative normalisations of the same component returns.

5.2.6 Determining Absolute Phase and Amplitude

Given a known relationship between the two component returns, we now determine absolute phase and intensity. This can be considered a denormalisation operation in the case where the unnormalised sum of the component returns and the normalised individual components are known. The component returns can be found for a candidate solution tuple (b, θ) by

$$\eta_0 = \zeta \mathrm{e}^{2\pi j m / r_0} \tag{5.66}$$

$$\eta_1 = b\eta_0 \mathrm{e}^{j\theta} \tag{5.67}$$

where $\zeta \in \mathbb{C}$ is the undisambiguated[4] first component return calculated by

$$\zeta = \frac{\xi_0^{\frac{1}{r_0}}}{|\xi_0|^{\frac{1}{r_0}-1}} \cdot \frac{|\Lambda_\mathrm{f}(b, r_0\theta)|^{\frac{1}{r_0}-1}}{\Lambda_\mathrm{f}(b, r_0\theta)^{\frac{1}{r_0}}} \tag{5.68}$$

and the disambiguation constant, $m \in \mathbb{Z}$, is determined by

$$\mathrm{e}^{2\pi j m \frac{r_1}{r_0}} = \frac{|\zeta|^{r_1-1}\xi_1}{\zeta^{r_1}\Lambda_\mathrm{f}(b, r_1\theta)}. \tag{5.69}$$

Valid candidate solution tuples can be discriminated from invalid using m, which is an integer for a valid solution in the absence of noise. In the presence of noise it may merely be close to an integer.

When this algorithm is implemented on real data, a substantial reduction in error can be achieved by calculating (τ_0, τ_1) for each pixel and then constraining

[4]Before removal of the cyclic phase ambiguity in the component return (disambiguation).

each value to a valid bound (e.g. $\tau_0 \in [0, 1)$) before any further processing. This reduces the incidence of noise corrupted data values, that otherwise result in highly erroneous estimates of component returns.

5.3 Summary/Analysis of Closed-form Methods

In this chapter the concept of amplitude and phase normalisation of complex domain range measurements was explained. A number of models were developed for non-point-source backscattering sources, where the backscattered intensity is spread over a region of range values; it was found that a closed-form Cauchy distribution model for range-diffuse sources naturally falls out of attempts to model point-source backscattering sources. Another closed-form solution was determined using attenuation ratios, where Chebyshev polynomials were found to enable the notation of the problem as the solution of low-order simultaneous polynomials. Due to time constraints, analysis of these methods was not completed during the course of the PhD[5]

In the next chapter, the phase and amplitude normalisation processes discussed here are applied to the generation of a coefficient called the characteristic measurement, which encodes the relationship between component returns, irrespective of absolute intensity and range. Using the characteristic measurement it is possible to reduce the dimensionality of the inverse problem, albeit no further closed-form solutions are developed in this thesis. The characteristic measurement also enables the calculation of a large number of useful bounds on phase perturbation, relative intensity and other interesting parameters of the component returns. Unlike the methods briefly developed in this chapter, a full analysis is performed using real and simulated data in Chapter 7.

[5]Real data results and some performance information is available in the paper Godbaz *et al.* (2012a), published several months after initial submission of the thesis. Both methods were demonstrated working on real data as well as simulated data, although the performance of the attenuation ratio was only barely operable given the noise levels of the real range-images. In general, all methods struggle with resolving two components which are located very close to each other – rather uniquely, the attenuation ratio method was found to additionally struggle with separating out two component returns near 180 degrees out-of-phase with each other (This was true at a frequency ratio of two-to-one; whether this holds at other frequency ratios remains undetermined.) The Cauchy algorithm was found to perform substantially better than the attenuation ratio, a simple generalisation of the algorithm was also determined, which allows the use of any sequence of frequency ratios $\{r_\mathrm{b}, r_\mathrm{b}+1, r_\mathrm{b}+2, r_\mathrm{b}+3\}$, where the integer base frequency is given by $r_\mathrm{b} \geq 0$. One issue awaiting resolution is false detection of diffuse-range returns.

Chapter 6

Mixed Pixel Separation and Bounding Using The Characteristic Measurement

This chapter introduces a new concept called the characteristic measurement. The characteristic measurement encodes the relative phase and intensity of the component returns in the multiple return separation problem. Use of the characteristic measurement allows the reduction of the dimensionality of the component return separation problem. Analysis of the phase perturbation problem using the characteristic measurement additionally allows closed-form bounds to be placed on the phase and amplitude of the underlying component returns.

Section 6.1 identifies a projection mapping, which is utilised for graphing the behaviour of functions over the entire complex plane. This is followed by the development of useful mathematical properties of phase perturbation and the characteristic measurement – these properties are then used to find bounding methods for the phase and intensity of the underlying component returns.

Whereas Chapter 5 gave two different approaches to closed-form inversion of the two component return mixed pixel separation problem, there appears to be no simple closed-form inverse in the two complex domain measurement, two non-diffuse-range component return case. This is despite the problem being exactly-determined and solvable using numerical methods. In the naïve case, a four dimensional lookup table is impracticable due to memory constraints, hence the importance of measurement normalisation and the concept of the characteristic measurement, which reduces the inverse problem to two dimensions.

6.1 Visualising the Entire Complex Plane

In order to visualise the behaviour of functions over the full gamut of possible input values it is necessary to derive a projection mapping from the entire complex plane, including infinity, to a finite subset of the complex plane. For the figures presented herin, the complex plane is projected onto the Riemann Sphere using Inverse Stereographic projection, followed by an Azimuthal Equidistant projection back onto a subset of the complex plane. This section explain how this combined projection is calculated and how the resultant function plots should be interpreted.

6.1.1 Inverse Stereographic Projection onto the Riemann Sphere

Inverse Stereographic projection is a conformal transformation, preserving relative angles, but is non-area preserving. It is one of the most natural methods for projecting a plane onto the surface of a sphere.

Given a unit sphere centred at $(1, 1, 1)$, the projection of a point on the complex plane $(\Re(w), \Im(w), 0)$ onto the sphere is given by the point on the surface of the sphere intersected by a line from $(0, 0, 2)$ to $(\Re(w), \Im(w), 0)$. If we notate the location on the sphere in Spherical Coordinates relative to the centre of the sphere, (r, ϕ, θ), where ϕ is azimuthal angle, θ is zenith angle and r is radial distance, then

$$r = 1 \tag{6.1}$$

$$\phi = \arg(w) \tag{6.2}$$

$$\theta = \frac{\pi}{2} - \text{atan2}(2|w|, |w|^2 - 1). \tag{6.3}$$

6.1.2 Azimuthal Equidistant Projection onto the Complex Plane

The Azimuthal Equidistant Projection (Tobler, 1962) is a synthetic projection which conserves directionality and distance from the origin. It is non-conformal and non-area preserving. For this particular application, we align the origin of the projection with the south pole of the Riemann Sphere.

For a point $(\phi, \theta, 1)$ on the Riemann Sphere, the projected point on the complex plane is given by

$$p(w) = \left(\theta + \frac{\pi}{2}\right) \mathrm{e}^{j\phi}. \tag{6.4}$$

If the output range is scaled so that $|p(w)| \in [0, 2]$ then the complete projection

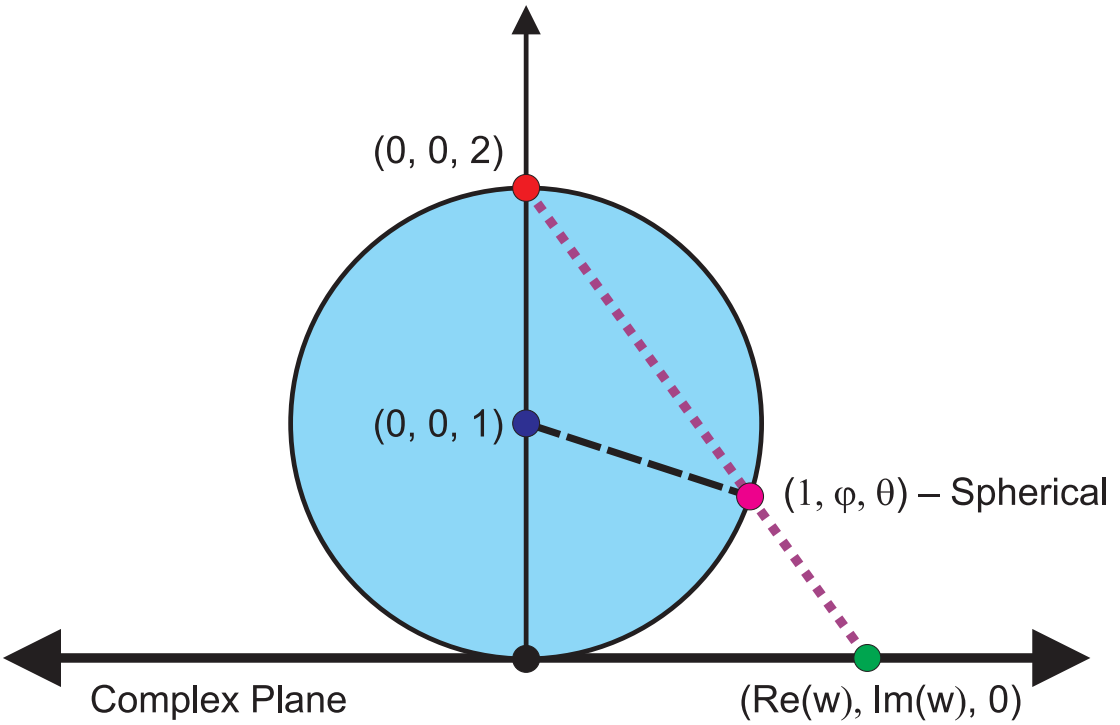

Figure 6.1: Inverse stereographic projection of the complex plane onto the Riemann Sphere. The green point on the complex plane, w, is projected onto the unit sphere centred at $(0, 0, 1)$. This gives the magenta projected point at $(1, \phi, \theta)$ in spherical coordinates centred on $(0, 0, 1)$.

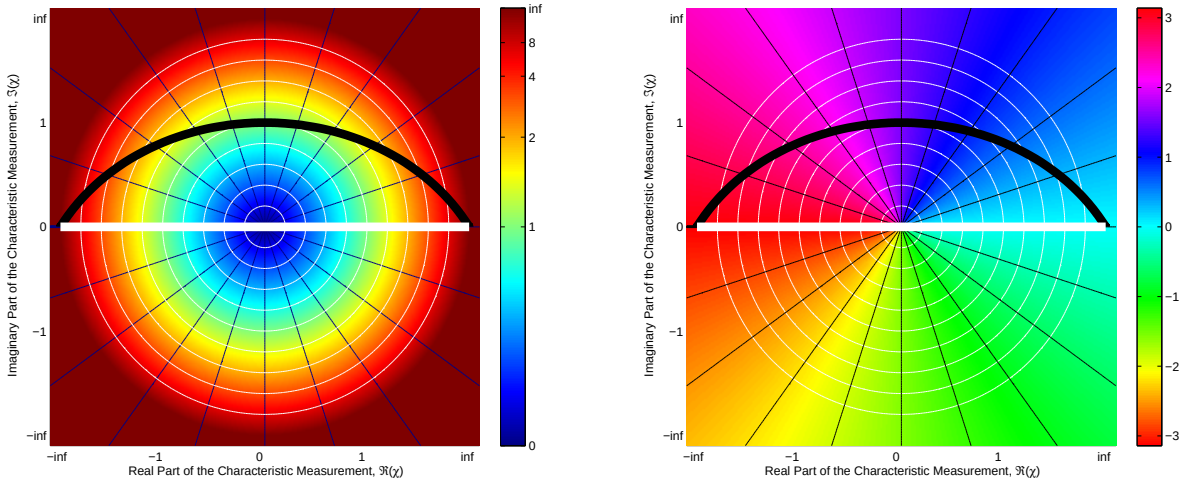

(a) $|\chi|$ (b) $\arg(\chi)$

Figure 6.2: Illustration of the combined projection designed primarily for visualisation of functions of the characteristic measurement.

function can be written as:

$$p(w) = \frac{2w}{\pi|w|}\left(\pi - \text{atan2}(2|w|, |w|^2 - 1)\right). \tag{6.5}$$

In addition to projecting the complex plane, this function is used to compress the dynamic range of values on the real line.

6.1.3 Interpreting Projected Functions

The combined projection given by Eqn. 6.5 is non-conformal and non-area preserving, however it is continuously differentiable and visually pleasing. Two examples are given in Fig. 6.2.

Radial distance from the origin is proportional to zenith angle on the sphere, while angle around the origin is proportional to azimuthal angle. Due to the radial nature of the projection, the values specified on the real axis only hold at $\Im(\chi) = 0$ and for the imaginary axis at $\Re(\chi) = 0$. Any line that does not pass through the origin becomes an arc when projected. The white line is a projection of the line $1x$, whereas the black arc is a projection of the line $1x + 1i$; $x \in \mathbb{R}$ is a free parameter. Values at infinity are continued out to the border of the image. This projection is particularly useful in that it differentiates between values at infinity with different phase angles, which are of great importance in our later analysis.

6.2 The Characteristic Measurement

Section 5.2.1 developed a method of normalisation that can be applied to either measurements or component returns. By applying this method to a tuple of range measurements, where one of the range measurements is normalised to unity, the remaining non-unity values encode the phase and amplitude relationship between the component returns – which have specific useful properties. The term characteristic measurement is introduced to cover these non-unity normalised values; in the two measurement case, a single complex coefficient. We now provide a precise mathematical definition.

6.2.1 An Explicit Definition of the Characteristic Measurement

Given two measurements $\Xi = \{\xi_0, \xi_1\}$ at relative frequencies r_0 and r_1, normalising ξ_0 to 1, via Eqn. 5.43

$$\mu_0 = \frac{1}{\xi_0} \quad \Rightarrow \quad \mu_{\mathrm{n}} = \frac{|\xi_0|^{1/r_0 - 1}}{\xi_0^{1/r_0}}. \tag{6.6}$$

Using the definition of normalisation from Eqns. 5.40 to 5.43, the normalised measurements ($\hat{\Xi}$) are determined by

$$\begin{aligned} \hat{\Xi} &= \left(\xi_0 \frac{|\xi_0|^{r_0/r_0 - 1}}{\xi_0^{r_0/r_0}}, \xi_1 \frac{|\xi_0|^{r_1/r_0 - 1}}{\xi_0^{r_1/r_0}} \right), & (6.7) \\ &= (1, \chi), & (6.8) \end{aligned}$$

where

$$\chi = \frac{\xi_1 |\xi_0|^{r_1/r_0 - 1}}{\xi_0^{r_1/r_0}}. \tag{6.9}$$

The value χ defines the relationship between the two measurements, referred to as the characteristic measurement. This can be extended to higher dimensional cases by allowing χ to be a vector. The characteristic measurement is global translation and scaling invariant; it only encodes the relationship between the component returns. In other words, it can be expressed solely as a function of relative intensity, b, and relative phase, θ. In polar constraints form, this can be shown by substituting Eqn. 5.49 for each measurement

$$\begin{aligned} |\chi| &= \frac{|\xi_1|}{|\xi_0|} = \frac{a_0 |\Lambda_{\mathrm{f}}(b, r_1\theta)|}{a_0 |\Lambda_{\mathrm{f}}(b, r_0\theta)|} = \frac{|\Lambda_{\mathrm{f}}(b, r_1\theta)|}{|\Lambda_{\mathrm{f}}(b, r_0\theta)|} & (6.10) \\ \arg(\chi) &= \arg(\xi_1) - \frac{r_1}{r_0} \arg(\xi_0) & (6.11) \\ &= \arg\left(\mathrm{e}^{jr_1\theta_0} \Lambda_{\mathrm{f}}(b, r_1\theta) \right) - \frac{r_1}{r_0} \arg\left(\mathrm{e}^{jr_0\theta_0} \Lambda_{\mathrm{f}}(b, r_0\theta) \right) & (6.12) \\ &= \arg\left(\Lambda_{\mathrm{f}}(b, r_1\theta) \right) - \frac{r_1}{r_0} \arg\left(\Lambda_{\mathrm{f}}(b, r_0\theta) \right). & (6.13) \end{aligned}$$

6.2.2 Applying the Characteristic Measurement to the Multiple Return Separation Problem

Assuming the existence of an inverse $\mathrm{H} = R^{-1}(\Xi)$ that determines the component returns from the set of recorded measurements, then the component returns, H, may

be determined from the normalised measurements by

$$\hat{\mathrm{H}} = R^{-1}(\hat{\Xi}) = \hat{R}^{-1}(\chi) \tag{6.14}$$

$$\mathrm{H} = \frac{\hat{\mathrm{H}}}{\mu_{\mathrm{n}}}. \tag{6.15}$$

In this manner the domain of the function $R^{-1}(\Xi)$ has been reduced, giving the dimensionality reduced inverse $\hat{R}^{-1}(\chi)$.

The characteristic measurement has a certain similarity to the synthetic measurements produced by the synthetic wavelength technique, which was discussed in Chapter 2 along with other phase standard unwrapping techniques. Both the characteristic measurement and synthetic wavelength technique involve synthesising a measurement by subtracting the phase of one measurement from the phase of another. However, the characteristic measurement weights each phase according to the relative frequency in addition to encoding the amplitude relationship.

In general, phase unwrapping is implicit in the multiple component return problem for the majority of two-measurement frequency ratios. Whereas in a typical phase unwrapping situation, two measurements at different modulation frequencies will result in a single unambiguous range, in practice the multiple component separation problem results in multiple solutions, thus retaining ambiguity. For example, the complex domain measurements produced by a single component return measured at two different modulation frequencies with a ratio of $3:2$ happens to be identical to the complex domain range measurements produced by two additional, completely different combinations of two non-zero amplitude component returns. While the single component return case is generally a more probable underlying situation, there is no trivial way to resolve this ambiguity. This is discussed further in Chapter 7.

6.3 Mathematical Properties of the Characteristic Measurement

There are a number of different bounds that can be placed on properties of the component returns using the attenuation ratio and characteristic measurement. We now start to develop these methods; by derivation where possible, otherwise using a more emperical approach.

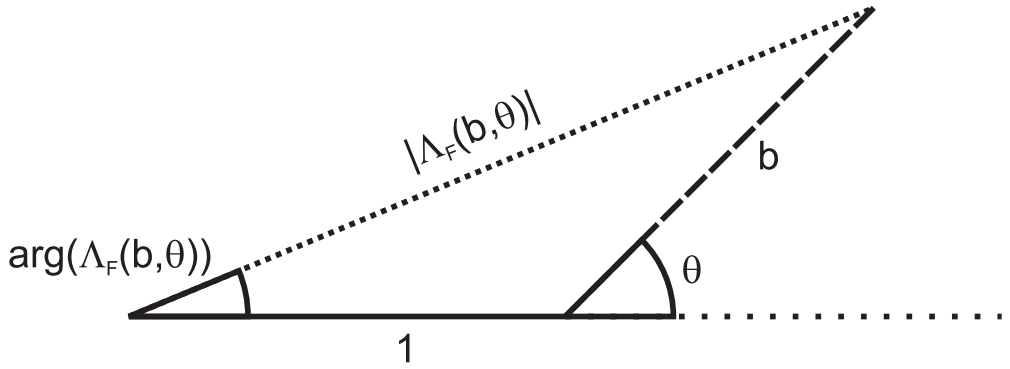

Figure 6.3: Illustration of the geometric meaning of $\Lambda_{\mathrm{f}}(b, \theta)$.

6.3.1 Properties of Phase Perturbation

In Eqn. 5.48 we introduced the use of the function $\Lambda_{\mathrm{f}}(b, \theta)$, which represents the phase perturbation introduced by a second component return, with a relative amplitude of b and relative phase of θ, on a primary component return. This was achieved by normalising the first component return to a value of unity using the approach notated in Eqns. 5.38 to 5.43.

In order to understand the properties of the characteristic measurement, we first need to understand the properties of $\Lambda_{\mathrm{f}}(b, \theta)$ on which the characteristic measurement is built. A graphical representation of this model for phase perturbation is given in Fig. 6.3

Fig. 6.4 shows the phase and amplitude response of $\Lambda_{\mathrm{f}}(b, \theta)$. The figure shows a symmetry around $\theta = \pi$ – changing the sign of the relative phase is equivalent to taking the complex conjugate, i.e.

$$\begin{aligned} \Lambda_{\mathrm{f}}(b, -\theta)^* &= (1 + b\mathrm{e}^{-j\theta})^* && (6.16) \\ &= 1 + b\mathrm{e}^{j\theta} && (6.17) \\ &= \Lambda_{\mathrm{f}}(b, \theta) && (6.18) \end{aligned}$$

There is also a simple relationship between positive and negative values of b, namely

$$\begin{aligned} \Lambda_{\mathrm{f}}(-b, \theta + \pi) &= 1 - b\mathrm{e}^{j\theta}\mathrm{e}^{j\pi} && (6.19) \\ &= 1 + b\mathrm{e}^{j\theta} && (6.20) \\ &= \Lambda_{\mathrm{f}}(b, \theta). && (6.21) \end{aligned}$$

The final relationship is equivalent to changing which component return is being

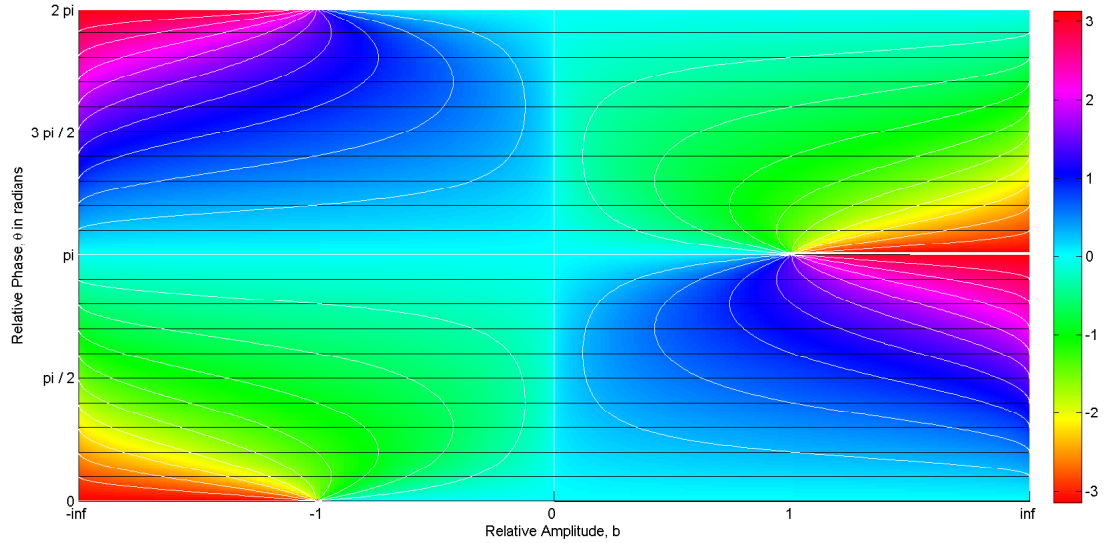

(a) $|\Lambda_f(b,\theta)|$

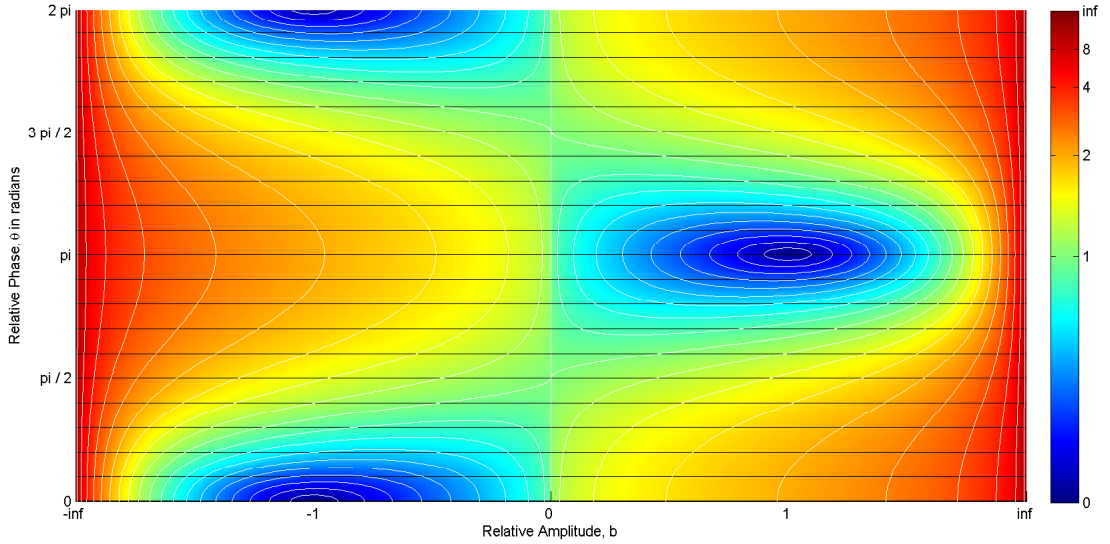

(b) $\arg(\Lambda_f(b,\theta))$

Figure 6.4: Plots of the phase and magnitude of $\Lambda_f(b,\theta)$.

implicitly normalised to one,

$$be^{j\theta}\Lambda_f(b^{-1},-\theta) = be^{j\theta}(1+b^{-1}e^{-j\theta}) \quad (6.22)$$
$$= be^{j\theta}+1 \quad (6.23)$$
$$= \Lambda_f(b,\theta). \quad (6.24)$$

This final relationship can be decomposed into a phase symmetry,

$$\arg(\Lambda_f(b,\theta)) = \arg(\Lambda_f(b^{-1},\theta)), \quad (6.25)$$

and an amplitude relationship,

$$|\Lambda_{\mathrm{f}}(b,\theta)| = |b\Lambda_{\mathrm{f}}(b^{-1},\theta)|. \tag{6.26}$$

There is also a degenerate case when $b = 0$ or $\theta = 0$ – the other parameter can freely vary without changing either the phase or amplitude of the output, that is

$$\Lambda_{\mathrm{f}}(0,\theta) = \Lambda_{\mathrm{f}}(b,0) = 0. \tag{6.27}$$

This corresponds to the case in which there is no phase perturbation, which while ideal, is not particularly interesting. In contrast, the situation where $b = 1$ and $\theta = \pi$ is a highly interesting condition where there is complete cancellation and the slightest perturbation can have an enormous impact on $\arg(\Lambda_{\mathrm{f}}(b,\theta))$. This particular degenerate case is the only situation in which amplitude can reach zero, which later in this chapter becomes important for bounding the extrema of the characteristic measurement. Tracing a path around the zero in an anticlockwise direction, it can be seen that the phase goes through a complete 2π phase cycle in the positive direction.

Plotting the phase and amplitude of Λ_{f} on more traditional graphs gives Fig. 6.5. In general, there are three different broad classes of situation. In the first case, where $b \in (0,1]$, the phase perturbation increases as θ moves away from zero, reaches a maxima and then returns to zero at $\theta = \pi$ before repeating an inverted variation of the same pattern. Taking the derivative of $\Lambda_{\mathrm{f}}(b,\theta)$ with respect to θ gives

$$\frac{d(\arg(\Lambda_{\mathrm{f}}(b,\theta)))}{d\theta} = \frac{b(\cos(\theta)+b)}{1+b^2+2b\cos(\theta)}. \tag{6.28}$$

Setting this to zero gives

$$\theta = \pm\cos^{-1}(-b), \tag{6.29}$$

which gives the relative phase corresponding to maximal phase perturbation for values of $b \in (0,1]$. For $b > 1$, maximal phase perturbation always occurs at $\theta = \pi$. For $b \in (0,1]$, the maximal phase perturbation achieved can be found by substituting $\cos^{-1}(-b)$ back into $\Lambda_{\mathrm{f}}(b,\theta)$, giving

$$\begin{aligned}
\arg(\Lambda_{\mathrm{f}}(b,\cos^{-1}(-b))) &= \tan^{-1}\left(\frac{b\sqrt{1-b^2}}{1-b^2}\right) && (6.30)\\
&= \tan^{-1}\left(\frac{b}{\sqrt{1-b^2}}\right) && (6.31)\\
&= \sin^{-1}(b). && (6.32)
\end{aligned}$$

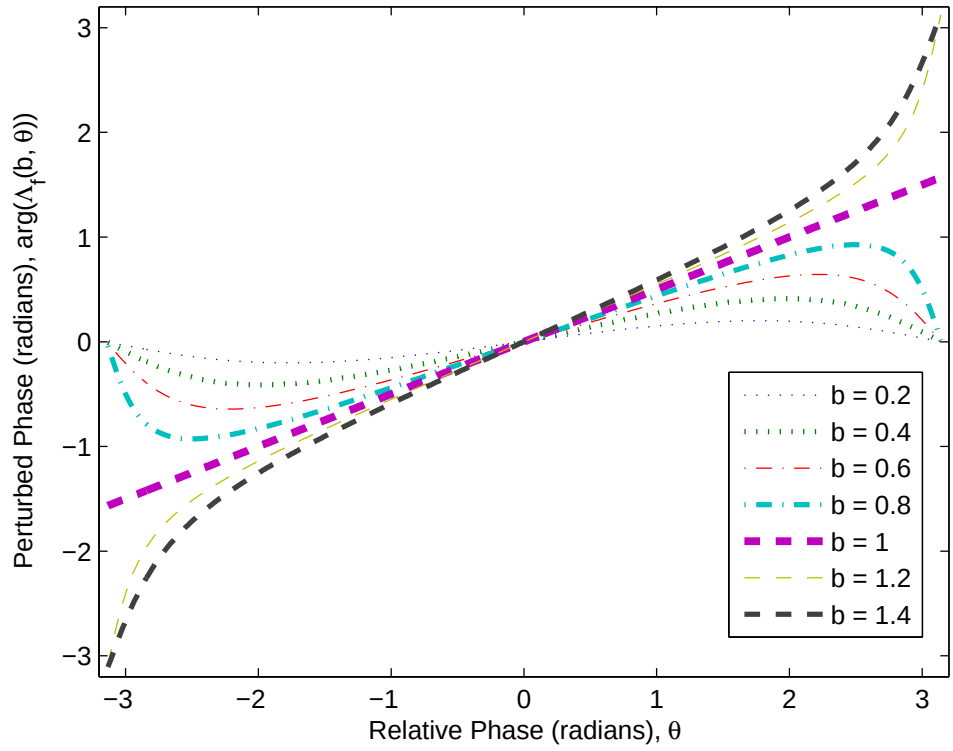

(a) $\arg(\Lambda_{\mathrm{f}}(c,\theta))$

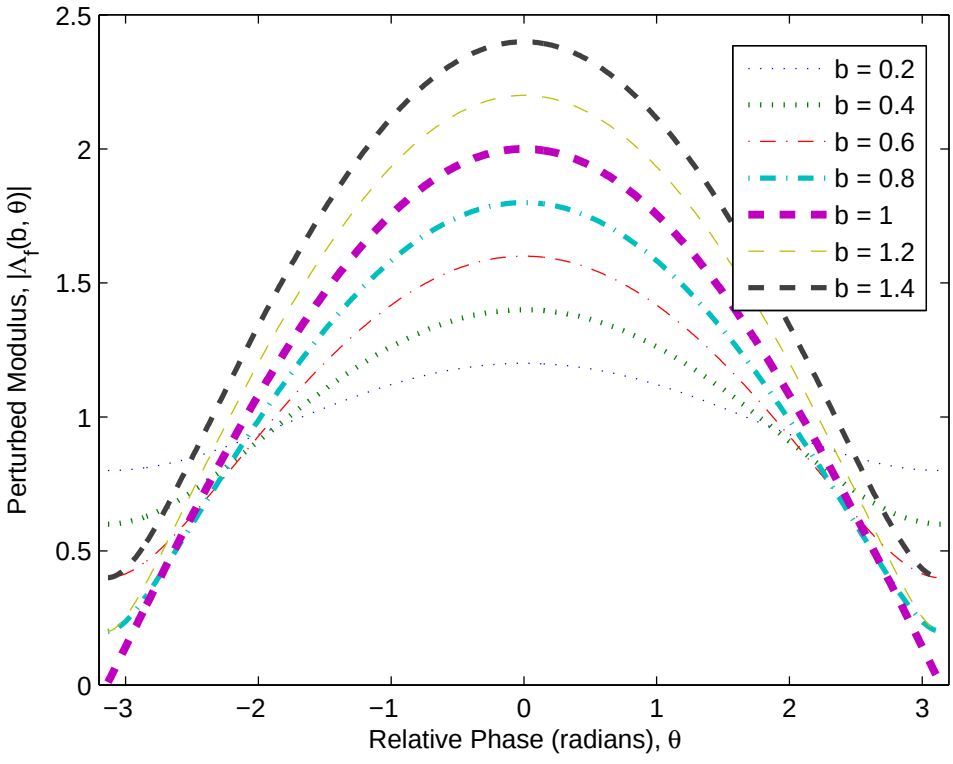

(b) $|\Lambda_{\mathrm{f}}(c,\theta)|$

Figure 6.5: The phase and amplitude response of $\Lambda_{\mathrm{f}}(c,\theta)$

In the most extreme variation of this case, where $b = 0$, there is no phase perturbation whatsoever and the response is flat. Because the second component is smaller than the first (ie. $b < 1$), $\theta = \pi$ corresponds to maximal intensity perturbation and no phase perturbation. The only other stationary point is $\theta = 0$, which corresponds to maximal intensity.

The second case is for $b = 1$. As b approaches one, the phase response for $\phi_\mathrm{f} \leq \cos^{-1}(-b)$ approaches linear with a slope of $\frac{1}{2}$, with a phase reversal near $\theta = \pi$ which becomes increasingly compact. At $b = 1$, $\arg(\Lambda_\mathrm{f}(1, \theta))$ becomes a discontinuous linear function with a slope of $\frac{1}{2}$. Given that the primary component return and the secondary component return are of equal length, with their sum they form an isosceles triangle. Utilising Fig. 6.3, it can be seen that θ is supplementary with one of the internal angles of the isosceles triangle. From the triangle rule this gives

$$2\arg(\Lambda_\mathrm{f}(1, \theta)) = \pi - (\pi - \theta) \tag{6.33}$$

$$\Rightarrow \quad \arg(\Lambda_\mathrm{f}(1, \theta)) = \frac{\theta}{2}. \tag{6.34}$$

At $\theta = \pi$ the amplitude is zero and the phase is undefined.

The third case is when the second component return is brighter than the first component return ($b > 1$); thus $\arg(\Lambda_\mathrm{f}(b, \theta))$ utilises the full dynamic range of output phase values by sweeping across $(-\pi, \pi]$. As $b \to \infty$, the phase response becomes linear again, that is

$$\lim_{b \to \infty} \arg(\Lambda_\mathrm{f}(b, \theta)) = \theta. \tag{6.35}$$

If the relative intensity of two component returns and $\arg(\Lambda(b, \theta))$ are known, it is possible to determine the relative phase θ. Starting from the logarithmic definition of $\tan^{-1}(x)$ as

$$\phi = \tan^{-1}(x) = -\frac{j}{2} \ln\left(\frac{1 + jx}{1 - jx}\right), \tag{6.36}$$

$\arg(\Lambda_\mathrm{f}(b, \theta))$ can be written as

$$\mathrm{e}^{2j\arg(\Lambda_\mathrm{f}(b,\theta))} = \frac{1 + j\frac{b\sin(\theta)}{1+b\cos(\theta)}}{1 - j\frac{b\sin(\theta)}{1+b\cos(\theta)}} \tag{6.37}$$

$$= \frac{1 + b\mathrm{e}^{j\theta}}{1 + b\mathrm{e}^{-j\theta}} \tag{6.38}$$

$$= \frac{1 + b\omega_\Lambda}{1 + b\omega_\Lambda^{-1}}, \tag{6.39}$$

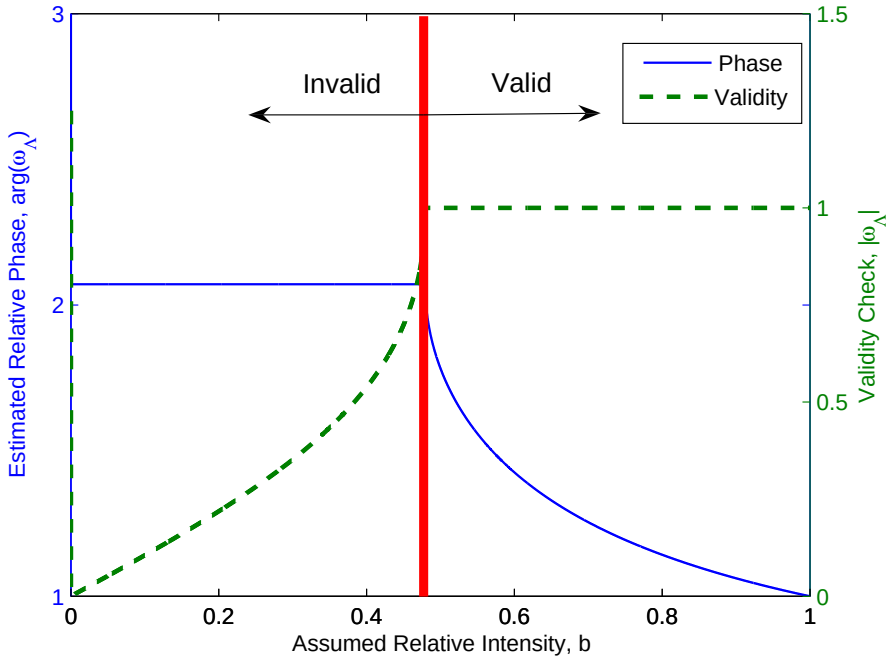

Figure 6.6: Estimating the phase of the secondary component return as a function of assumed relative intensity. Only values where $|\omega_\Lambda| = 1$ are valid, all other values correspond to physically impossible phenomena.

where $\omega_\Lambda = \mathrm{e}^{j\theta}$. Inverting this gives

$$\omega_\Lambda = \frac{\omega_\phi - 1 + \sqrt{4b^2\omega_\phi + \omega_\phi^2 - 2\omega_\phi + 1}}{2b} \tag{6.40}$$

where $\omega_\phi = e^{2j\arg(\Lambda_\mathrm{f}(b,\theta))}$. Note that the other solution to the quadratic equation implicit in eqn. 6.39 is incorrect for our purposes. We now explain the behaviour of this inverse function using an example.

Fig. 6.6 shows the behaviour of eqn. 6.40 given the assumption that $\omega_\phi = \mathrm{e}^{j}$. The minimum value of $\arg(\omega_\Lambda)$ occurs at $b = 1$, where

$$\theta_\mathrm{f} = \frac{\theta}{2} = \frac{\arg(\omega_\phi)}{2}. \tag{6.41}$$

The relationship agrees with Eqn. 6.34 and is utilised in Section 6.4.6, Eqn. 6.161 as the basis for a lower bound on θ; for a known θ_f, no value of θ lower than this is possible. The maximum possible valid value of $\arg(\omega_\Lambda)$ occurs at

$$b = \sin\left(\frac{\arg(\omega_\Lambda)}{2}\right). \tag{6.42}$$

This relationship agrees with Eqn. 6.32 and is used in Section 6.4.4, Eqn. 6.131 as the basis for a lower bound on b. No valid values are possible below this point as is demonstrated in Fig. 6.39 by analysing the behaviour of $|\omega_\Lambda|$. For valid values of b, ω_Λ is confined to the unit circle; values of $b < \sin(1/2)$, $|\omega_\Lambda| < 1$ indicate the invalidity of b. The relative phase corresponding to the minimum possible value of relative intensity corresponds to the maximum valid relative phase – as θ increases above this point, it is no longer physically possible for the required amount of phase perturbation to be achieved.

6.3.2 The Forward Transformation

Let the utility function, $\Lambda_\chi(b, \theta)$, represent the forward transformation given by

$$\Lambda_\chi(b, \theta) = \frac{\Lambda_\mathrm{f}(b, r_1\theta)|\Lambda_\mathrm{f}(b, r_0\theta)|^{(r_1/r_0)-1}}{\Lambda_\mathrm{f}(b, r_0\theta)^{r_1/r_0}}. \tag{6.43}$$

To ease analysis, we make the assumption that r_0 and r_1 are coprime, that is, they have no common factors other than unity. Cursory examination of the above equation shows that if r_0 is greater than one, then $\Lambda_\chi(b, \theta)$ is a multivalued function with at least r_0 branches due to the complex root. In other words, if $\chi \in \mathbb{C}$ is a valid output value from $\Lambda_\chi(b, \theta)$ then the complete set of output values is given by

$$X = \{\chi \mathrm{e}^0, \chi \mathrm{e}^{2\pi/r_0}, \ldots \chi \mathrm{e}^{(n-1)2\pi/r_0} : n \in \mathbb{Z}\}. \tag{6.44}$$

Plots of this function for several different frequency ratios are given in Figs. 6.7 and 6.8. For the sake of notational simplicity, we typically avoid explicitly notating the rotational ambiguity.

The remainder of this section explains the behaviour of this function, why it behaves the way it does and how this allows bounds to be placed on the input parameters.

6.3.3 Symmetries and the Modulus of the Characteristic Measurement

It is immediately apparent from Fig. 6.7 that there are similar symmetries in $\Lambda_\chi(b, \theta)$ to those of $\Lambda_\mathrm{f}(b, \theta)$ described in Section 6.3.1. In other words, the following hold

$$\Lambda_\chi(b, \theta) = \Lambda_\chi(b, -\theta)^* \tag{6.45}$$

$$\Lambda_\chi(b, \theta) = \Lambda_\chi(-b, \theta + \pi) \tag{6.46}$$

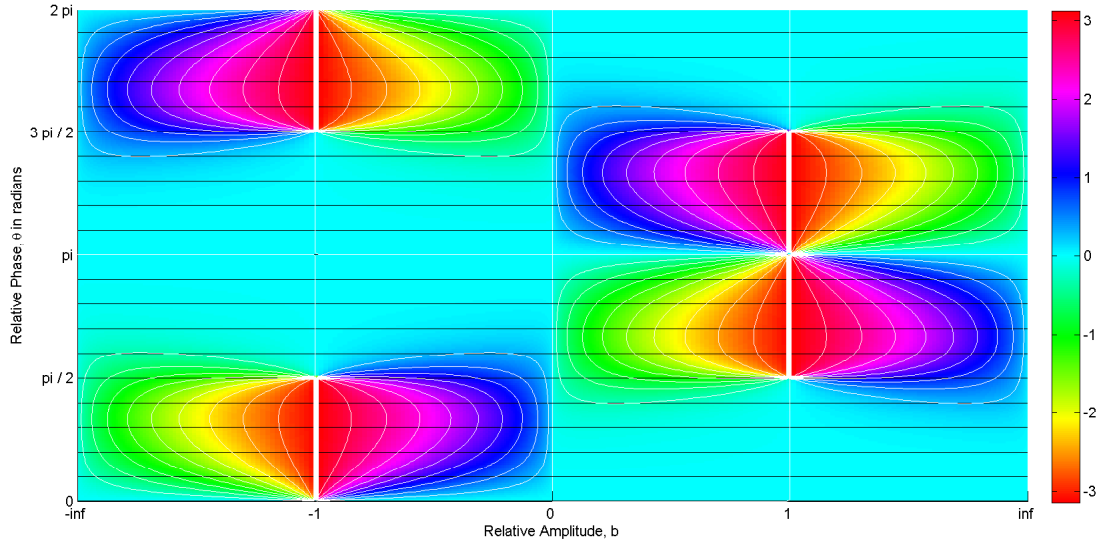

(a) $\arg(\Lambda_\chi(b,\theta))$ at 2:1

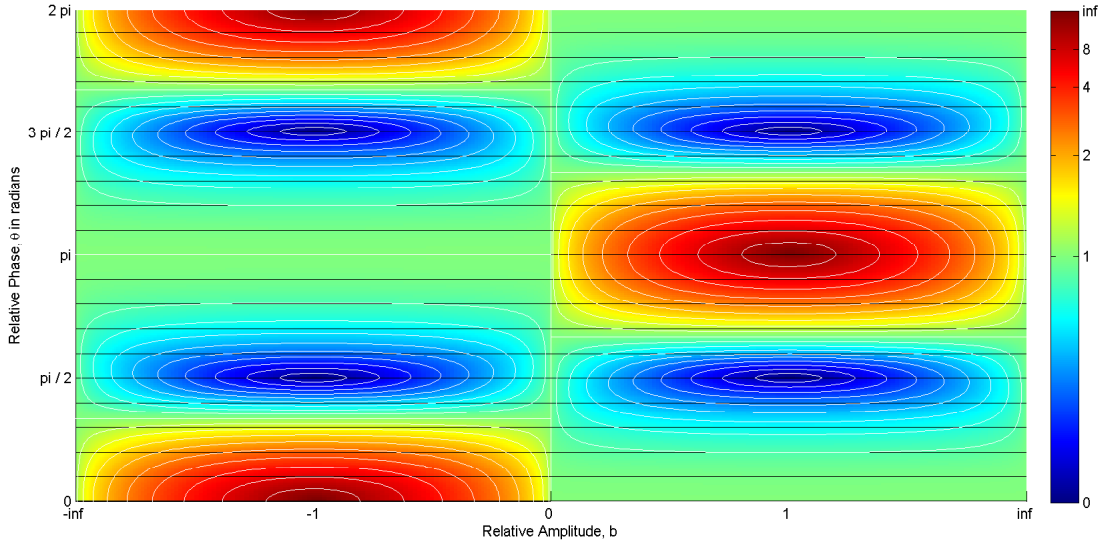

(b) $|\Lambda_\chi(b,\theta)|$ at 2:1

Figure 6.7: Plots of the phase and magnitude of $\Lambda_\chi(b,\theta)$.

But rather than the symmetry expressed in Eqn. 6.24, Λ_f satisfies

$$\Lambda_\chi(b,\theta) \quad = \quad \Lambda_\chi(b^{-1},-\theta). \tag{6.47}$$

In other words, there is a conjugate reflection across a line at $b=1$ because $(1, b\mathrm{e}^{j\theta})$ and $(b^{-1}, \mathrm{e}^{-j\theta})$ encode the same relative phase and relative magnitude. From this perspective, all possible two component relationships can be encoded by the tuple (b,θ) where $b \in [0,1]$ and $\theta \in [0, 2\pi)$ – the majority of further analysis we restrict to this domain.

Fig. 6.7 shows that the most extreme modulii, in other words the modulii which

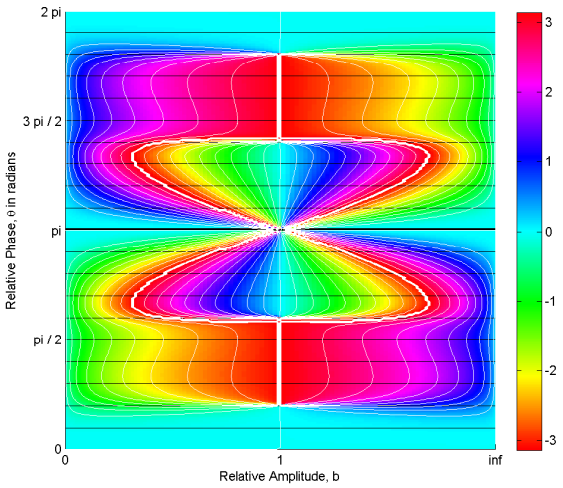

(a) $\arg(\Lambda_\chi(b,\theta))$ at 5:1

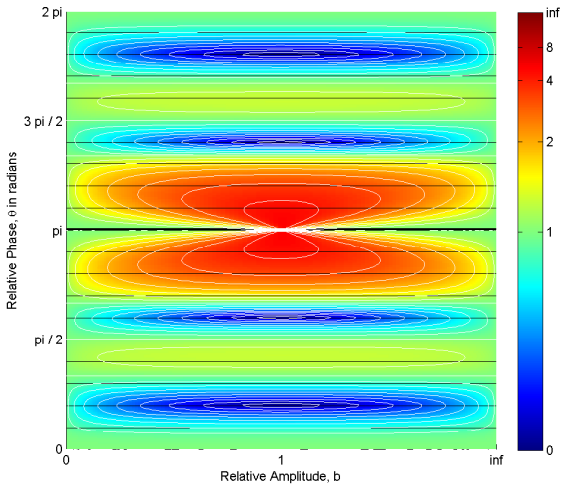

(b) $|\Lambda_\chi(b,\theta)|$ at 5:1

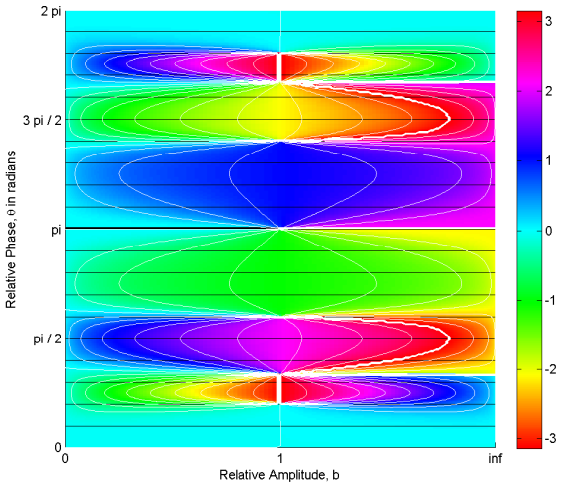

(c) $\arg(\Lambda_\chi(b,\theta))$ at 5:3

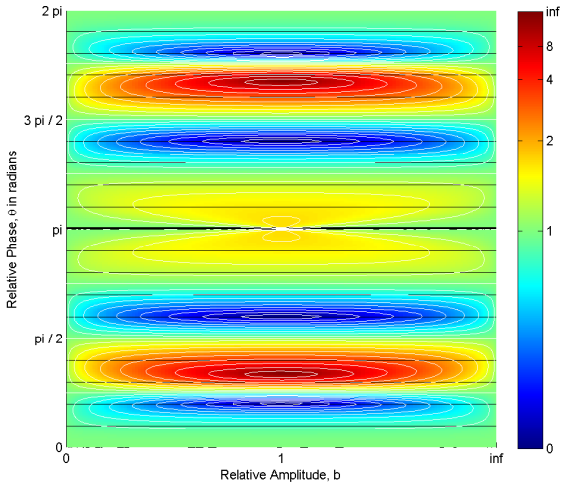

(d) $|\Lambda_\chi(b,\theta)|$ at 5:3

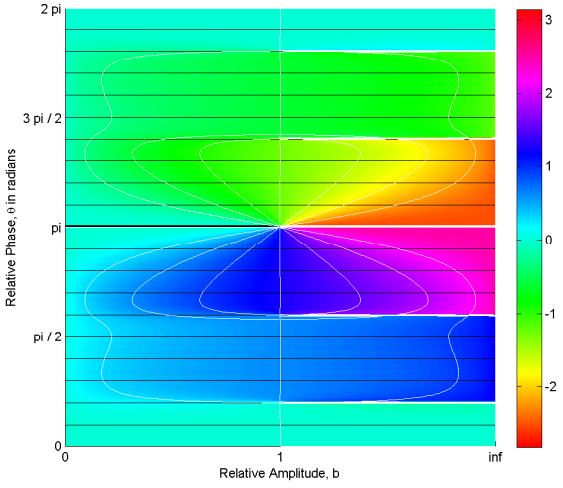

(e) $\arg(\Lambda_\chi(b,\theta))$ at 1:5

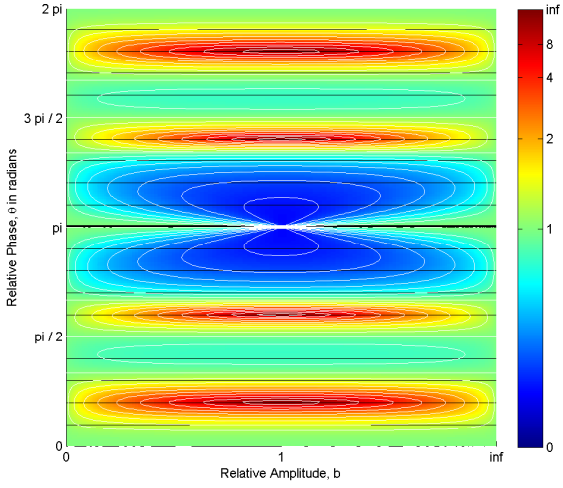

(f) $|\Lambda_\chi(b,\theta)|$ at 1:5

Figure 6.8: Plots of the phase and magnitude of $\Lambda_\chi(b,\theta)$ at different frequency ratios.

have diverged farthest from a value of unity, all occur on the line given by $b = 1$. This can be illustrated by calculating the stationary points of the modulus squared versus relative intensity, b. From Eqn. 6.10 we have

$$|\Lambda_\chi(b,\theta)|^2 = \frac{|\Lambda_\mathrm{f}(b,r_1\theta)|^2}{|\Lambda_\mathrm{f}(b,r_0\theta)|^2}. \tag{6.48}$$

Differentiation with respect to b gives

$$\frac{\partial|\Lambda_\chi(b,\theta)|^2}{\partial b} = \frac{\left(\frac{\partial|\Lambda_\mathrm{f}(b,r_1\theta)|^2}{\partial b}\right)|\Lambda_\mathrm{f}(b,r_0\theta)|^2 - \left(\frac{\partial|\Lambda_\mathrm{f}(b,r_0\theta)|^2}{\partial b}\right)|\Lambda_\mathrm{f}(b,r_1\theta)|^2}{|\Lambda_\mathrm{f}(b,r_0\theta)|^4}. \tag{6.49}$$

Given that

$$|\Lambda_\mathrm{f}(b,r_l\theta)|^2 = 1 + b^2 + 2b\cos(r_l\theta) \tag{6.50}$$

$$\frac{\partial|\Lambda_\mathrm{f}(b,r_l\theta)|^2}{\partial b} = 2(\cos(r_l\theta) + b), \tag{6.51}$$

then

$$\begin{aligned}\frac{\partial|\Lambda_\chi(b,\theta)|^2}{\partial b} &= 2\frac{(\cos(r_1\theta)+b)(1+b^2+2b\cos(r_0\theta))}{|\Lambda(b,r_0\theta)|^4} \\ &\quad -2\frac{(\cos(r_0\theta)+b)(1+b^2+2b\cos(r_1\theta))}{|\Lambda(b,r_0\theta)|^4} \qquad (6.52)\\ &= 2\frac{(b+1)(b-1)(\cos(r_0\theta)-\cos(r_1\theta))}{|\Lambda_\mathrm{f}(b,r_0\theta)|^4} \qquad (6.53)\end{aligned}$$

Setting this to zero, gives a stationary point at $b = 1$ and another at $b = -1$, which is the image of the first stationary point using the identity in Eqn. 6.46. This is either a maximum or minimum depending on the sign of

$$(\cos(r_0\theta) - \cos(r_1\theta))(\cos(r_0\theta)b^2 + 2b + \cos(r_0\theta))), \tag{6.54}$$

with postive values corresponding to mimima, negative to maxima. Ultimately, $b \to 1$ bounds the most extreme values of χ in terms of distance from the unit circle. Another less obvious stationary point is given by $\cos(r_0\theta) - \cos(r_1\theta) = 0$, which corresponds to the unit circle, as indicated by

$$|\Lambda_\chi(b,\theta)|^2 = \tfrac{|\Lambda_\mathrm{f}(b,r_1\theta)|^2}{|\Lambda_\mathrm{f}(b,r_0\theta)|^2} = 1 \tag{6.55}$$

$$\Rightarrow \quad |\Lambda(b,r_1\theta)|^2 - |\Lambda_\mathrm{f}(b,r_0\theta)|^2 = 0 \tag{6.56}$$

$$\Rightarrow \quad \cos(r_0\theta) - \cos(r_1\theta) = 0. \tag{6.57}$$

Although not a true stationary point *per se*, as $b \to \infty$ the derivative also goes to zero.

Fig. 6.9a shows how the value of b changes the squared modulus of the characteristic measurement. At $b = 0$, the squared modulus is always unity – this situation corresponds to the situation where there is only a single component return. As the value of b increases towards 1, the squared modulus approaches maximal distance from unity.

A degeneracy occurs when there is only a single component return, not only just when $b = 0$[1]:

$$\Lambda_\chi(0, \theta) = \lim_{b \to \infty} \Lambda_\chi(b, \theta) = \Lambda_\chi(b, 0) = \mathrm{e}^{2\pi j m / r_0}, \tag{6.58}$$

where $m \in \mathbb{Z}$, which are the r_0th roots of unity. By measuring the distance from these roots, an indication of how mixed a measurement is can be achieved; this is elaborated upon, for a specific frequency ratio of $2:1$ in Section 6.4.6.

An understanding of the bounding case at $b = 1$ is best achieved by analysing the underlying numerator and denominator in the formatory equation Eqn. 6.48. In particular, there are a number of special cases: if $|\Lambda_\mathrm{f}(b, r_0\theta)|^2 = 0$, then there is a pole; if $|\Lambda(b, r_1\theta)|^2 = 0$, then there is a zero; that is, unless both conditions hold, in which case there is a singularity with a value of

$$\lim_{\theta \to \pi} |\Lambda_\chi(1, \theta)| = \lim_{\theta \to \pi} \sqrt{\frac{1 + \cos(r_1\theta)}{1 + \cos(r_0\theta)}} \tag{6.59}$$

$$= \frac{r_1}{r_0}. \tag{6.60}$$

All of these special cases are demonstrated in Fig. 6.9b.

In the same manner that the $b = 1$ case bounds the modulus of the characteristic measurement, it also defines the behaviour of the phase of the characteristic measurement. Looking at Fig. 6.7a, we can see that tracing a contour around the pole at $(b, \theta) = (1, \pi)$ in the anticlockwise direction passes through a phase shift of -4π. On the other hand, the zero at $(b, \theta) = (1, \pi/2)$ passes through a phase shift of 2π.

As a general rule, each zero corresponds to a phase shift of 2π, each pole $-(r_1/r_0)\pi$ and each singularity $2\pi(1 - 1/r_0)$. If at any point the phase shift around a point is a non-integer multiple of 2π, then that point is a branch point, the anchor for a branch cut. For example, in Fig. 6.8e there are four poles and one singularity on $b = 1$, each of which is a branch point. In this particular figure, each branch point

[1] Albeit, this is primarily a notational issue.

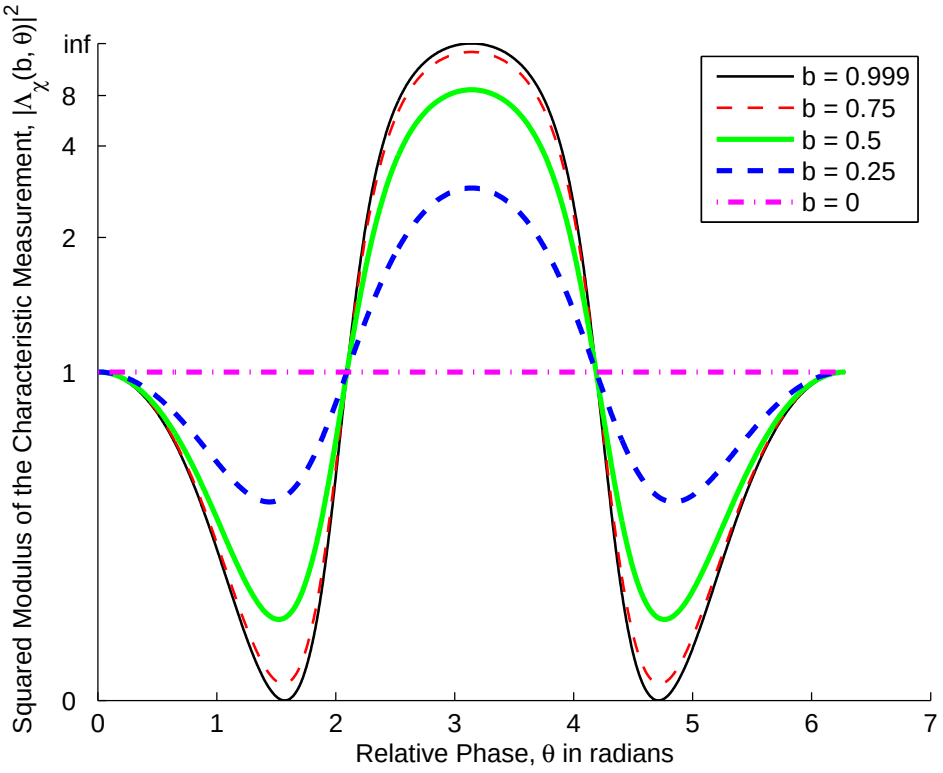

(a) Varying b at 2 : 1

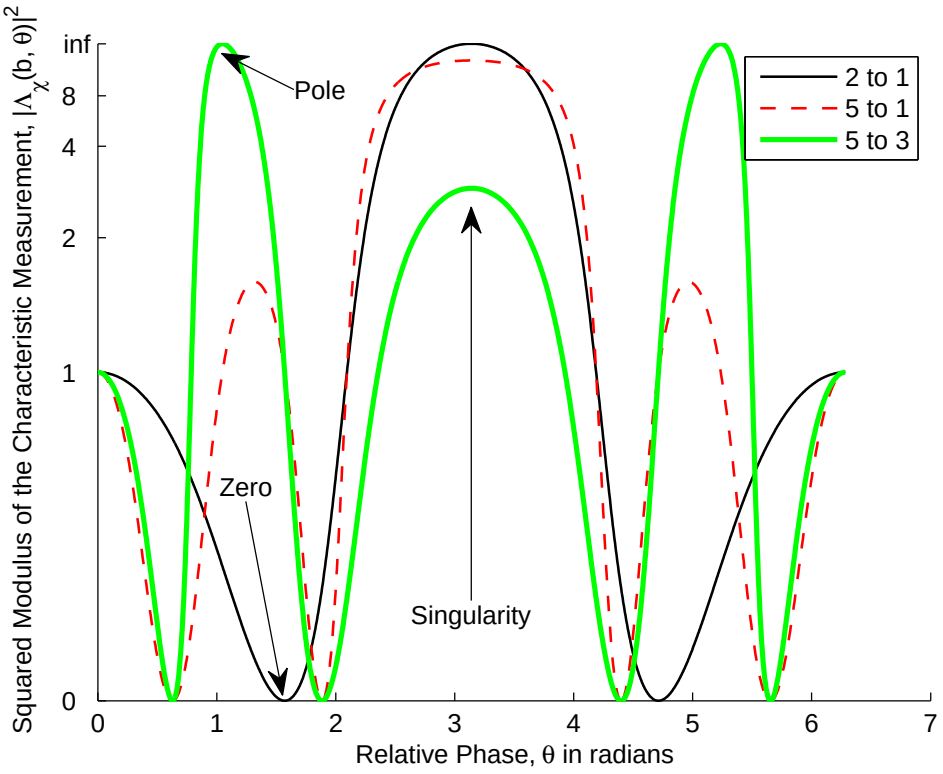

(b) Bounding the Squared Modulus at Different Relative Frequencies ($b \to 1$)

Figure 6.9: Analysing the squared modulus of the characteristic measurement.

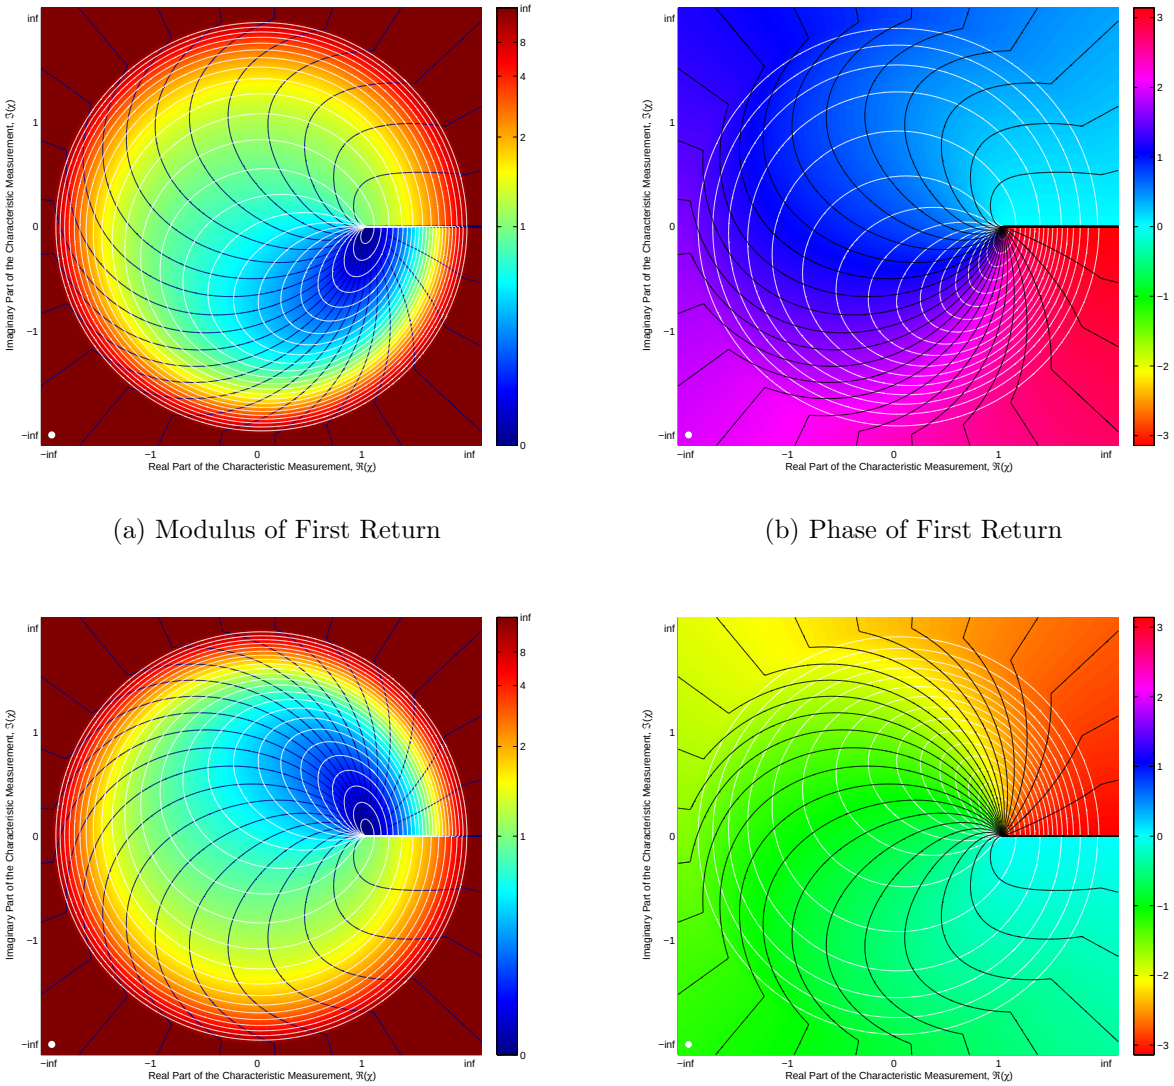

(a) Modulus of First Return

(b) Phase of First Return

(c) Modulus of Second Return

(d) Phase of Second Return

Figure 6.10: The inverse mapping in the 2 : 1 frequency ratio case.

has been used as the anchor for a branch cut extending out to infinity. Because the normalisation of the measurement at a relative frequency of one by a measurement at a relative frequency of five introduces a $2\pi/5$ phase ambiguity, Fig. 6.8e only shows a fifth of the underlying solution; each section being a phase shifted image of the others, as given in Section 6.3.2. This is a strong rationale for always normalising the highest relative frequency by the lowest relative frequency. For example, Fig. 6.8a does not have any branch cuts while Fig. 6.8e does.

Figure 6.11: Phase of the inverse mapping in the 3 : 2 frequency ratio case. Layout: Top row, primary return; bottom row, secondary return; Left column, first solution; middle column, second solution; right column, third solution.

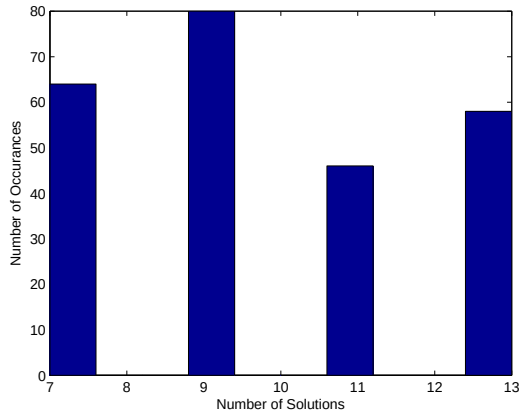

(a) Monte-Carlo Simulation at 7 : 2

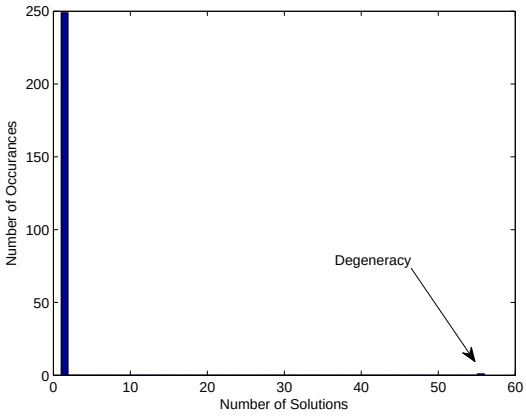

(b) Monte-Carlo Simulation at 3 : 1

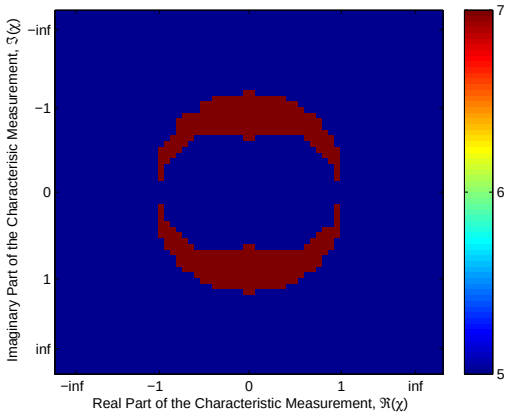

(c) Number of Solutions vs. χ at 5 : 2

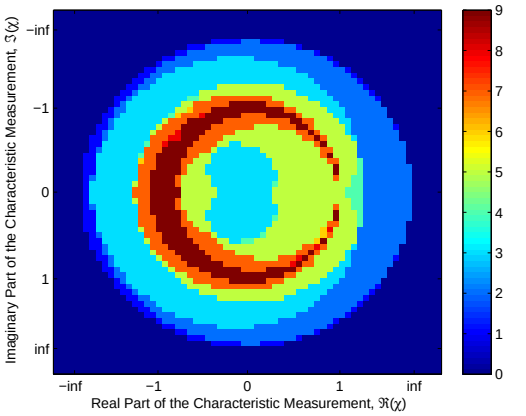

(d) Number of Solutions vs. χ at 7 : 1

Figure 6.12: Variation in the number of solutions for different frequency ratios.

6.3.4 Understanding the Inverse

Thus far we have discussed the forward problem: the formation of a set of measurements from a set of component returns. More important, for a practical situation, is the ability to find the component returns from a set of measurements. Ideally, it would be possible to determine a closed-form inverse to do this. As yet, this has not been achieved; the lack of success does not indicate that it is impossible, merely that it is a non-trivial[2] exercise. However, even without a closed-form inverse it is possible to determine the component returns using numerical methods, or lookup table based approaches. Despite the lack of a full closed-form inverse, it is still possible to analyse the behaviour of the forward function so as to place close-form bounds on the properties of the underlying component returns; bounds such as maximum phase perturbation and maximum total integrated intensity. We now analyse the

[2]If it were trivial, one would hope that four years of research would have found a solution.

Table 6.1: Number of returns versus relative frequency.

Frequency Ratios		Number of Solutions	
r_1	r_0	Min	Max
2	1	1	1
3	1	1	1
4	1	2	4
5	1	1	4
6	1	3	9
7	1	2	9
8	1	4	16
9	1	2	16
10	1	5	25
11	1	3	25
12	1	6	36
3	2	3	3
5	2	5	7
7	2	7	13
4	3	6	6
5	3	5	6
7	3	7	11
5	4	10	10
6	5	15	15
7	6	21	21
Apparent Relationships			
n	1	?	$\lfloor n/2 \rfloor^2$
n	$n-1$	$n(n-1)/2$	$n(n-1)/2$

properties of the inverse.

Using numerical optimisation it is possible to plot the inverse function. Fig. 6.10 gives the inverse function for a frequency ratio of $2:1$, which is the simplest case. For a given value of χ, the modulus of the first component return is given by a specific location on Fig. 6.10a, the phase by the same location on Fig. 6.10 and the equivalent for the second component return using Figs. 6.10c and d. Together, the first and second component returns form a single continuous manifold $H^{-1}(\chi)$. If this function was complex differentiable versus χ, then it would be a Riemann Surface, but it is not. In this particular case, the first component return is the conjugate reflection of the second, although this does not necessarily hold for other measurement frequency ratios. The most important feature is the branch point at $\chi = 1$; this branch point corresponds to the point at which there is only a single component return. A branch cut from $\chi = 1$ out to infinity forms the boundary at which the solution surfaces of the first and second component returns are stitched

together. Even without a closed-form inverse, intuitively it appears that unwrapping around the branch point is likely to provide fundamental insight into the behaviour of the function – an approach which we continue in Section 6.4.6. This is related to the Cartesian layout for the figures: we have chosen to use a projected Cartesian mapping for the figures, partly because it highlights certain properties – such as the potential for unwrapping around the branch points – which would be less conspicuous in a zero centred polar layout.

Depending on the particular frequency ratio, the inverse function, $\mathrm{H}^{-1}(\chi)$, is a multivalued function. An example of this is given in Fig. 6.11, which gives a plot of the phase of the inverse function for a frequency ratio of $3:2$. As predicted by Eqn. 6.58, there are two points corresponding to single component returns (being $\chi = \pm 1$), each of which is a branch point plotted with a branch cut to infinity. In Section 6.3.2, it was discussed how the phase ambiguity introduced in $\arg(\chi)$ results in multiple solutions. Instinctively, one might assume that this would mean that a frequency ratio of $3:2$ would have two possible solutions – however, this is not true. Analysing the inverse function closely shows that there are two solutions that are conjugate reflections of each other, and one additional solution for which the upper two quadrants are a conjugate reflection of the lower two quadrants, thus the 2-fold rotational symmetry implicit in Eqn. 6.44 for $r_0 = 2$ is retained, but more than two solutions occur. The unexpected number of component returns prompts a more detailed analysis of the number of solutions across a wider variety of frequency ratios.

In order to analyse the propensity for multiple solutions, a Monte-Carlo simulation using randomly generated measurement phasors was run. For each frequency ratio 250 random measurement phasors were generated and then 400 randomly initialised optimisation attempts were run using a least squares cost function. Each solution was recorded and solutions within a threshold Euclidian distance of each other were considered to constitute a single solution. This was necessary since there is a certain amount of error in a numerically optimised value. For the purposes of this method, the minimum distance between the best permutation of the underlying complex phasors was utilised as a cost function. A larger simulation was considered, but was ultimately time-prohibitive due to the limitations of the optimisation method; obviously there is no guarantee that all the solutions were found. The prior distribution used for phasor generation is discussed in Section 7.1.

Some example results from the Monte-Carlo simulation are presented in Fig. 6.12. Fig. 6.12a shows the frequency of occurance of different numbers of solutions at a frequency ratio of $7:2$; note that only odd solution counts occur. Fig. 6.12 also

includes a graph at $3:1$. This particular experiment shows that the Monte-Carlo simulation has generated a degenerate case: a case where there are multiple possible different component returns that map to the same value of the characteristic measurement. The degenerate case corresponds to $\theta \approx \pi$, which results in an indeterminate relative intensity (as explained more deeply in Section 6.3.6). Earlier in this section we described the behaviour of the inverse in the case of a frequency ratio of $3:2$, for which case there are three different possible solutions. While it appears reasonable for there to be multiple possible solutions for a particular value of χ – particularly in cases where $r_0 > 1$ – it is not immediately obvious why there are different number of possible component returns for different values of χ. Figs. 6.12c and d show how the number of solutions changes over the domain of χ; note that the greatest number of solutions always occurs near the unit circle.

A more complete set of results from the Monte-Carlo simulation are presented in Table 6.1. It appears that the number of solutions for a frequency ratio of the form $n:(n-1)$ is always given by $n(n-1)/2$ and that the maximum number of solutions for a frequency ratio of $n:1$ is given by $\lfloor n/2 \rfloor^2$. Conventionally, frequency ratios like $5:4$ are used with methods like the synthetic wavelength in order to maximise the disambiguated range, unfortunately the higher the values of r_1 and r_0, the greater the number of possible solutions. Past a certain point, this becomes a major limiting factor as in most applications there is a requirement for determination of the correct solution. This issue is discussed in more detail in Section 7.2.

One of the interesting aspects of $n:1$ ratios is that if $\lfloor n/2 \rfloor$ is even, then ignoring special cases with an infintesimal probability, the number of solutions for any particular value of χ is always even. In addition, the number of solutions is always odd for odd values of $\lfloor n/2 \rfloor$. One might be inclined to wonder what differences there are between ratios of, say $8:1$ and $9:1$, which both have up to 16 possible solutions, with only an even number of solutions for any particular value of χ. One difference is as to whether the valid solution-space for the characteristic measurement extends to infinity. Given $r_0 = 1$, there is only one possible value of θ for which there is the potential for $|\chi| \to \infty$, that is $\theta \approx \pi$. If r_1 is even, then the solution-space extends to infinity; if r_1 is odd, then there is a singularity instead and the solution-space does not extend to infinity.

6.3.5 Asymptotes and Folds

We now discuss the topology of the inverse by introducing folds and asymptotes in the context of the characteristic measurement solution manifold. Earlier in this chapter we described the inverse solution as being a single continuous manifold from

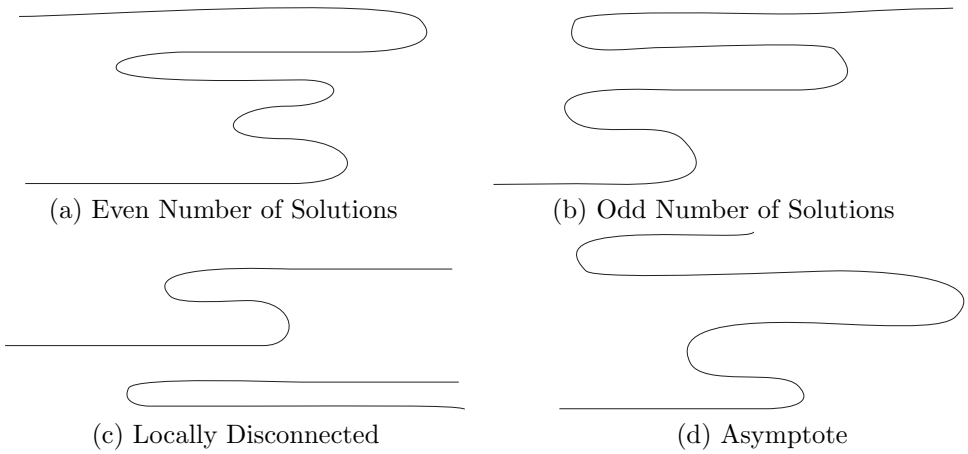

(a) Even Number of Solutions (b) Odd Number of Solutions

(c) Locally Disconnected (d) Asymptote

Figure 6.13: A simple model of the inverse manifold.

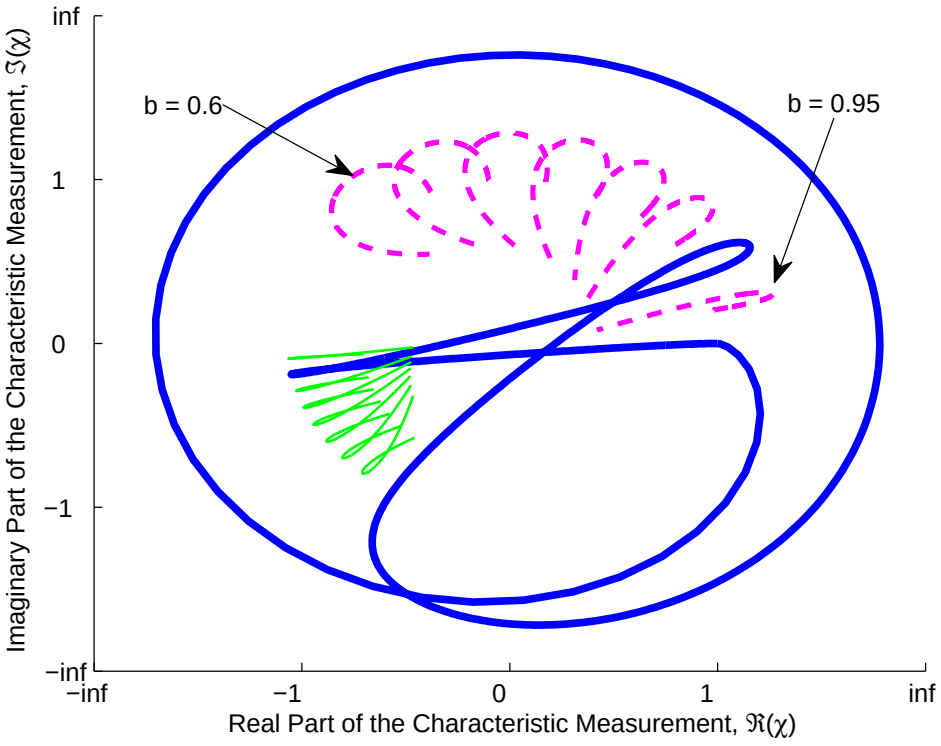

Figure 6.14: The formation of folds: the characteristic measurement as a function of b and θ at a frequency ratio of $7:1$. The thick line represents the values of χ produced for $b = 0.9$, $\theta \in [0, \pi]$. The dashed lines represent $b \in [0.6, 0.95]$ (increments of 0.05), $\theta \in [\pi/5, 2\pi/5]$. The thin lines represent the same as the dashed, but for $\theta \in [\pi/2, 7\pi/10]$. See also Fig. 6.12d.

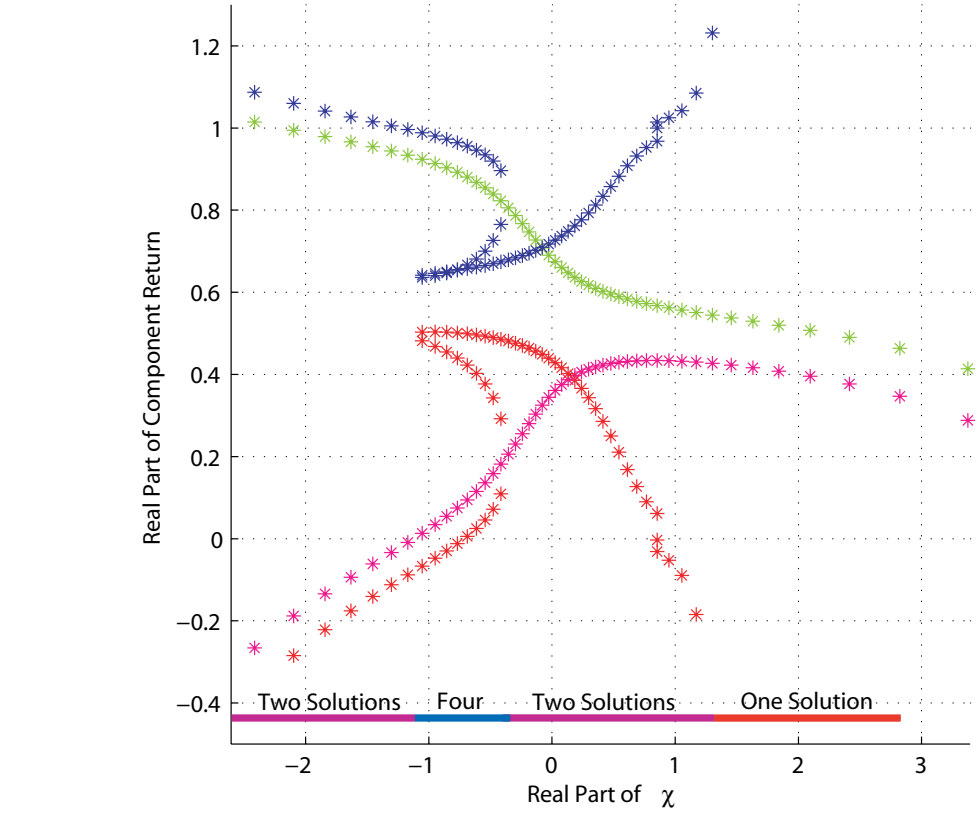

(a)

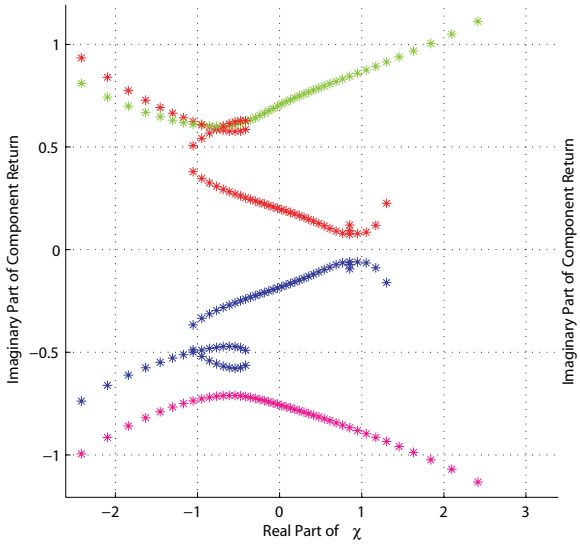

(b)

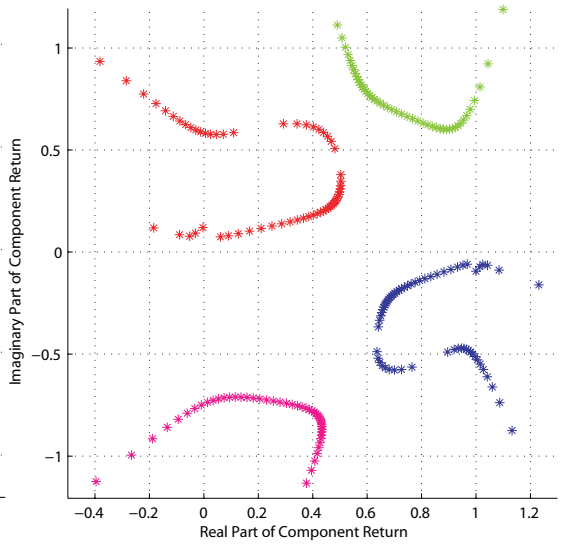

(c)

Figure 6.15: Solutions to the two component return problem for fixed values of $\Im(\chi)$ at 5 : 1. The solutions form a single connected surface in complex space; as the surface folds back over itself, the number of solutions for a particular value of the characteristic measurement changes. The blue and red points correspond to one particular possible solution and the lime and magenta points to another. Note that for a particular value of $\Re(\chi)$, the sum of the real parts of the component returns for a solution always add to unity and the sum of the imaginary parts always add to zero. See also, Fig. 6.16

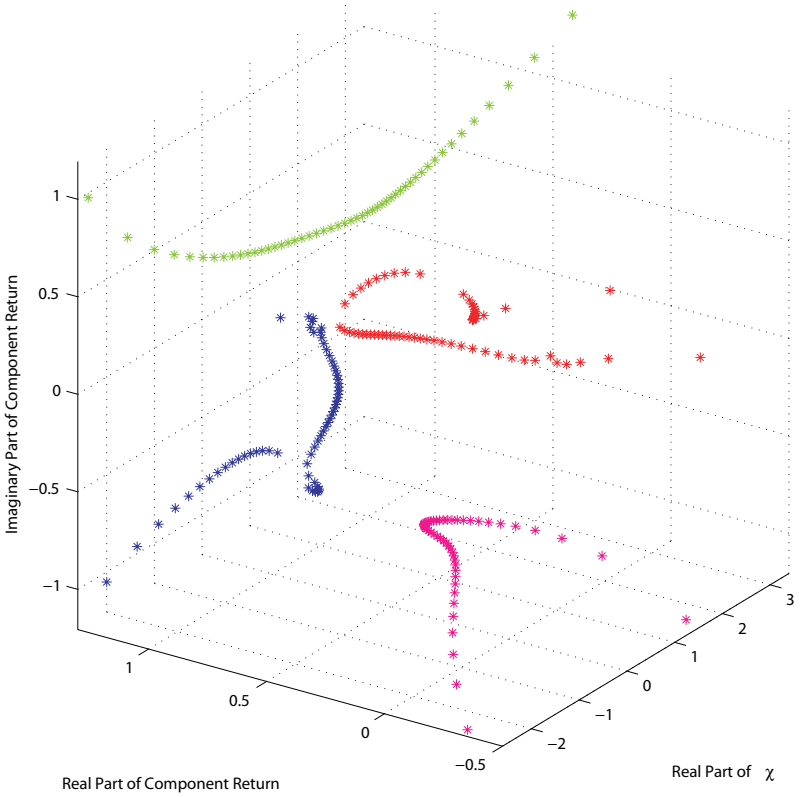

Figure 6.16: A three dimensional visualisation of solutions to the two component return problem for a fixed value of $\Im(\chi)$. See also, Fig. 6.15.

which all possible solutions and both component returns can be inferred. One way to visualise the phenomenon of multiple solutions is to use the metaphor of origami. When folding a piece of paper, there are certain physical constraints which must be fulfilled[3] Figs. 6.13a and b show how the number of solutions can vary over different regions, while only having either odd or even numbers of solutions. Fig. 6.13c shows a situation where the manifold is not connected locally, although the inverse is always globally connected in practice.

If we consider the forward problem, as the values of b and θ change, the values progressively sweep across the surface of the sheet. Depending on the frequency ratio, at certain values of θ the direction of travel suddenly changes. An example is given in Fig. 6.14. Determination of fold location is discussed in Section 6.4.3[4].

Another example of a fold is shown in Figs. 6.15 and 6.16. At a frequency ratio of $5:1$, there are either one, two or four solutions for any particular value of the characteristic measurement. Fig. 6.15 illustrates how the solution set given by the blue and red datapoints folds back on itself so as to result in a varying number of

[3]While the specific constraints are not necessarily the same, it is a useful analogy.

[4]The problem is non-trivial as the values of the θ for which this occurs are a function of the value of b.

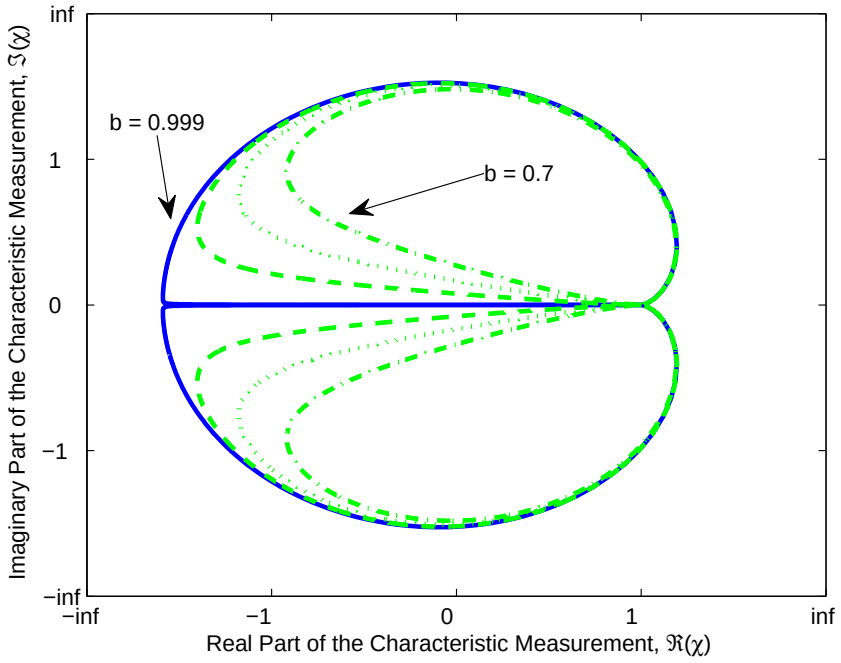

Figure 6.17: The formation of asymptotes: the characteristic measurement as a function of b and θ at a frequency ratio of $3:1$. Lines are plotted for $b \in \{0.7, 0.8, 0.9, 0.999\}$ and $\theta \in [-\pi, \pi]$.

component returns. The lime and magenta datapoints appear disconnected from the blue and red datapoints, but all values are actually part of the same continuous manifold.

While folds occur in most, if not all non-$n:(n-1)$ frequency ratio cases, asymptotes occur more rarely. Asymptotes occur when in the inverse, non-infinite values of the characteristic measurement map to normalised component returns for which the modulus asymptotically approaches infinity. Typically, normalised component returns only approach infinity as the characteristic measurement approaches infinity. The asymptotes only occur near singularities of $|\chi|$, which solely occur at $\theta = \pi$ as $b \to 1$, when both r_0 and r_1 are odd. The example given in Fig. 6.17 occurs at a frequency ratio of $3:1$. The full inverse manifold for $3:1$ is plotted in Fig. 6.18.

6.3.6 The Unit Circle

In Fig. 6.9b there are a number of specific relative phases (values of θ) at which the modulus of the characteristic measurement equals unity; apart from $b = 0$, these are the only cases in which this occurs. These conditions correspond to the situation where the range of χ is restricted to the unit circle. Ignoring $b = 0$, for a particular

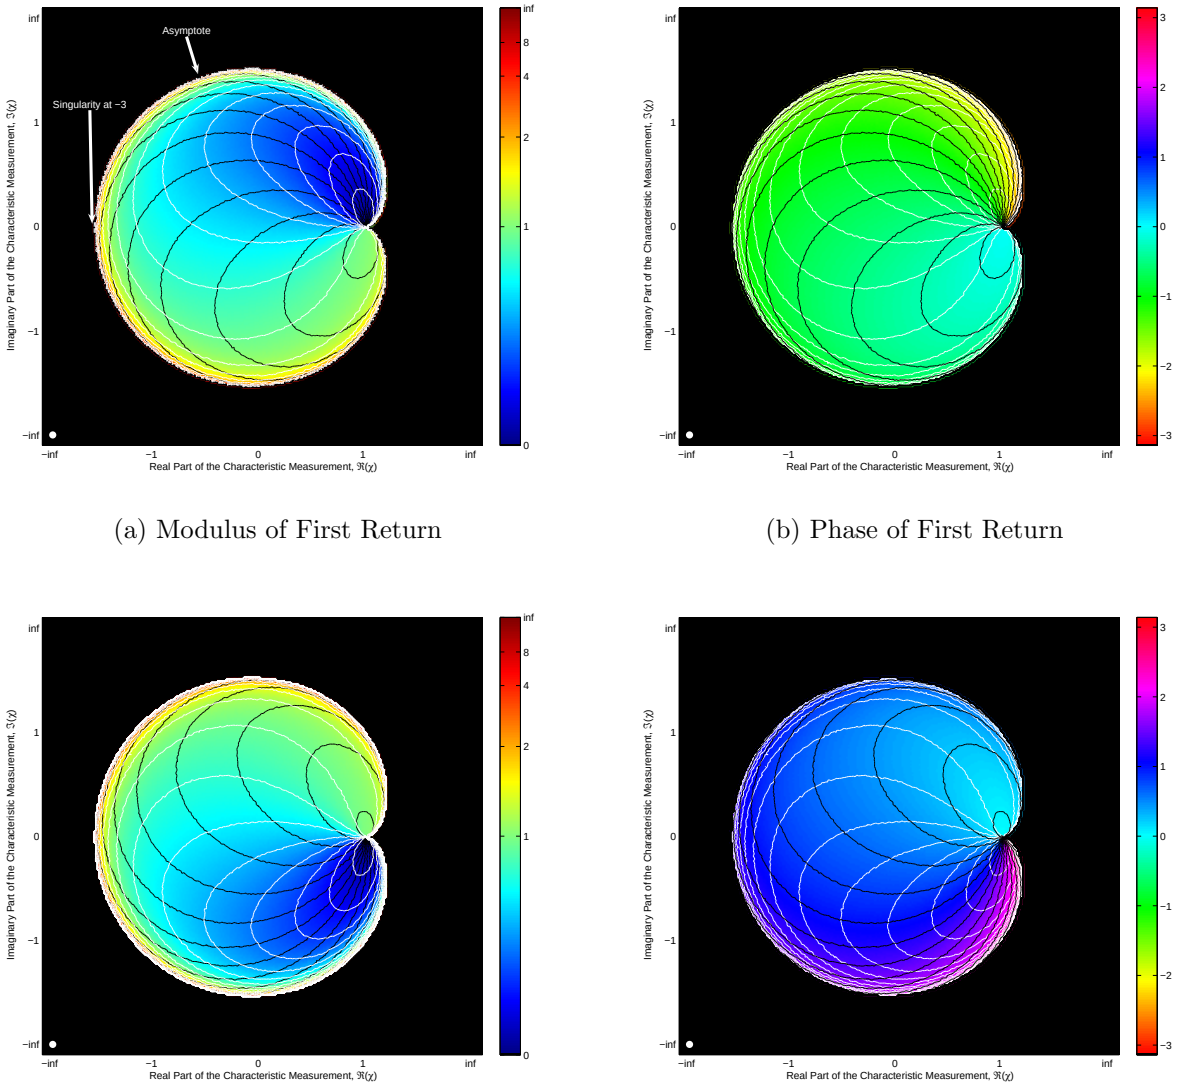

(a) Modulus of First Return

(b) Phase of First Return

(c) Modulus of Second Return

(d) Phase of Second Return

Figure 6.18: The inverse mapping in the 3 : 1 frequency ratio case.

frequency ratio $r_1 : r_0$, these points are given by the solutions to Eqn. 6.57, that is

$$\cos(r_0\theta) - \cos(r_1\theta) = 0 \tag{6.61}$$

$$T_{r_0}(\cos(\theta)) - T_{r_1}(\cos(\theta)) = 0. \tag{6.62}$$

A number of solutions are given in Table 6.2.

There are two types of solution for θ, the first type maps all values of b onto unity – thus is degenerate. The second maps values of b onto the unit circle, where the phase of χ is a function of b. We now show why the first type maps onto unity by modelling the phase of the characteristic measurement. The first situation occurs

Table 6.2: Values of θ for which $|\Lambda_\chi(b,\theta)| = 1$, for selected frequency ratios.

Frequencies		Solutions for $\pm\theta$	
r_1	r_0	Maps to Unity	Maps to the Unit Circle
2	1	0	$2\pi/3$
3	1	0, π	$\pi/2$
4	1	0, $2\pi/3$	$2\pi/5$, ~ 2.513
5	1	0, $\pi/2$, π	$\pi/3$, $2\pi/3$
3	2	0	$2\pi/5$, ~ 2.513
5	2	0, $2\pi/3$	~ 0.898, ~ 1.795, ~ 2.693
4	3	0	~ 0.898, ~ 1.795, ~ 2.693

when $\arg(\chi) = 0$; we now derive a new expression for the phase component of the characteristic measurement in order to reason about this condition.

Using the logarithmic expression of the phase of Λ_f from Eqn. 6.38 we rewrite the phase of the characteristic measurement (Eqn. 6.11) as

$$(1 + b\mathrm{e}^{jr_1\theta})^{r_0}(1 + b\mathrm{e}^{-jr_0\theta})^{r_1} - \mathrm{e}^{2jr_0\arg(\chi)}(1 + b\mathrm{e}^{jr_0\theta})^{r_1}(1 + b\mathrm{e}^{-jr_1\theta})^{r_0} = 0, \qquad (6.63)$$

which when multiplied by $\mathrm{e}^{j\max(r_1,r_0)\theta}$, gives a polynomial over b and $\mathrm{e}^{j\theta}$ encoding the phase relationship; albeit, a major limitation is that the definition of the tangent operation intrinsically introduces a π radian ambiguity.

If $\arg(\chi) = 0$, then

$$(1 + bc^{r_1})^{r_0}(1 + bc^{-r_0})^{r_1} - (1 + bc^{r_0})^{r_1}(1 + bc^{-r_1})^{r_0} = 0, \qquad (6.64)$$

where $c = \mathrm{e}^{j\theta}$. There are a number of trivial solutions to this problem, including $\theta = 0$ and $b \in \{0, 1\}$. The solution to this equation which we are most interested in is

$$\mathrm{e}^{jr_0\theta} - \mathrm{e}^{jr_1\theta} = 0. \qquad (6.65)$$

Or, expressed differently:

$$\cos(r_0\theta) - \cos(r_1\theta) = 0 \quad \wedge \quad \sin(r_0\theta) - \sin(r_1\theta) = 0. \qquad (6.66)$$

This equation adds an additional constraint to Eqn. 6.57. Any value of θ which satisfies Eqn. 6.66 always results in $\chi = 1$ and thus belongs to the first, unity mapping class (as opposed to the second, unit circle mapping class.) From a practical standpoint this means that for certain measurement frequency ratios, there are non-zero phase relationships which are degenerate, thus leading to inseparable component re-

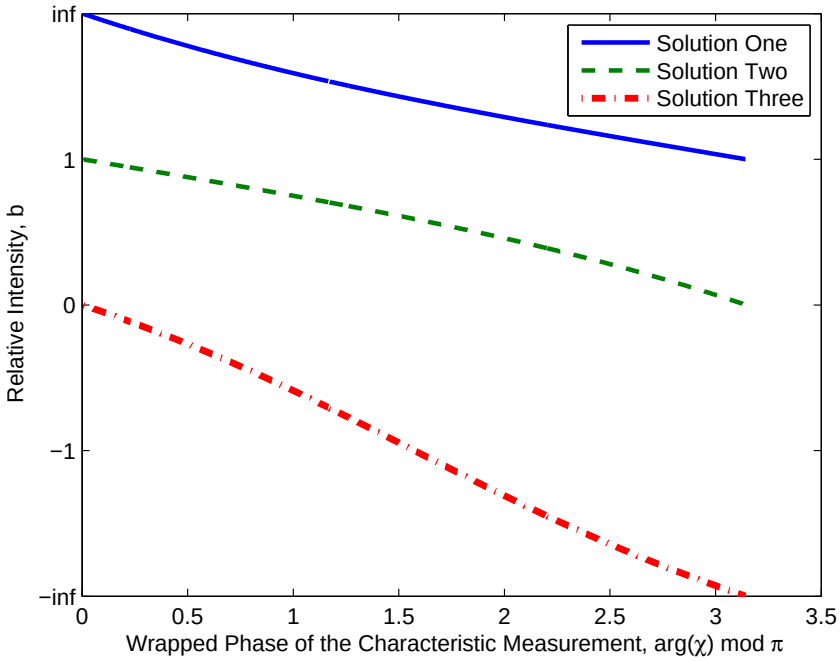

Figure 6.19: Determination of the relative intensity of the component returns, b, in the special case where the characteristic measurement is confined to the unit circle for a frequency ratio of $2:1$.

turns for any characteristic measurement based method or any other method which *implicitly* uses the characteristic measurement. In other words, any inversion of two measurements at a ratio of, say, $5:1$ will suffer from the same fundamental problem. However, an important outcome is that $n:(n-1)$ ratios do not appear to suffer from the malady; in order to ensure invertability, ratios like $5:1$ ought to be avoided in practical situations.

For the second class, which maps to the unit circle, the phase of the characteristic measurement can be considered to be a function of b. For example, at a frequency ratio of $2:1$ Eqn. 6.63 reduces to

$$(1+bc^2)(c+b)^2 - \mathrm{e}^{2j\arg(\chi)}(1+bc)^2(c^2+b) = 0, \tag{6.67}$$

which is cubic over b, thus allowing calculation of a closed form inverse. Assuming $\theta = 2\pi/3$, we find the relationship given in Fig. 6.19. Because of the π radian ambiguity, all values of the characteristic measurement are mapped to $[0,\pi]$ – whereas solutions one and three actually correspond to $\arg(\chi) \in [0,\pi]$, solution two corresponds to $\arg(\chi) \in [-\pi,0]$. This ambiguity must be resolved by another method, like a reference forward transformation.

We have now shown that starting from $\chi = 1$, perturbing the phase of the

characteristic measurement is equivalent to changing the relative amplitude of the component returns. As it turns out, perturbing the modulus of the characteristic measurement is equivalent to perturbing the relative phase of the component returns. We now start applying this understanding of the behaviour of the characteristic measurement to the bounding problem.

6.4 Bounding Component Parameters

In this section we develop bounds for individual parameters using single dimensions of the two dimensional characteristic measurement – using either the modulus or the phase. We also introduce some bounds using attenuation ratios, which are later applied to the characteristic measurement. Finally, we develop compound bounds – by using knowledge of a constraint it is often possible to infer other constraints, and thus improve our understanding of the underlying relationships.

6.4.1 Bounding Phase Using $|\chi|^2$

We now analyse the special cases corresponding to maximal phase perturbation in order to determine bounds on . We start by analysing the stationary points of the phase of the characteristic measurement versus b and θ.

The phase of $\Lambda_\chi(b,\theta)$ can be written in terms of the phase of $\Lambda_\mathrm{f}(b,\theta)$, viz

$$\arg(\Lambda_\chi(b,\theta)) \;=\; \arg(\Lambda_\mathrm{f}(b,r_1\theta)) - \frac{r_1}{r_0}\arg(\Lambda_\mathrm{f}(b,r_0\theta)). \tag{6.68}$$

Taking the derivative versus θ and utilising Eqn. 6.28 for the derivative of $\Lambda_\mathrm{f}(b,\theta)$ gives

$$\frac{\partial\arg(\Lambda_\chi(b,\theta))}{\partial\theta} \;=\; \frac{\partial\arg(\Lambda_\mathrm{f}(b,r_1\theta))}{\partial\theta} - \frac{r_1}{r_0}\frac{\partial\arg(\Lambda_\mathrm{f}(b,r_0\theta))}{\partial\theta} \tag{6.69}$$

$$\;=\; r_1\left(\frac{b(\cos(r_1\theta)+b)}{1+b^2+2b\cos(r_1\theta)} - \frac{b(\cos(r_0\theta)+b)}{1+b^2+2b\cos(r_0\theta)}\right). \tag{6.70}$$

Setting this to zero gives the polynomial equation,

$$\begin{aligned} &b(r_0(T_{r_1}(\cos(\theta))+b)(1+b^2+2bT_{r_0}(\cos(\theta))) \\ &-r_1(T_{r_0}(\cos(\theta))+b)(1+b^2+2bT_{r_1}(\cos(\theta)))) \;=\; 0, \end{aligned} \tag{6.71}$$

which factorises to

$$br_1(b+1)(b-1)(T_{r_0}(\cos(\theta)) - T_{r_1}(\cos(\theta))) \;=\; 0, \tag{6.72}$$

giving $b \in \{0, -1, 1\}$ or the unit circle. The first case, $b = 0$, corresponds to minimal phase perturbation; the second and third cases, $b \in \{-1, 1\}$, correspond to maximal phase perturbation. The unit circle is a bit more complicated; substituting $\cos(r_1\theta) = \cos(r_0\theta)$ into Eqn. 6.70, gives

$$\frac{\partial \arg(\Lambda_\chi(b,\theta))}{\partial\theta} = r_1\left(\frac{b(\cos(r_0\theta)+b)}{1+b^2+2b\cos(r_0\theta)} - \frac{b(\cos(r_0\theta)+b)}{1+b^2+2b\cos(r_0\theta)}\right), \quad (6.73)$$

therefore

$$\frac{\partial \arg(\Lambda_\chi(b,\theta))^2}{\partial^2\theta} = 0. \quad (6.74)$$

Thus the unit circle corresponds to a point of inflection and is not useful for bounding relative phase, θ, although it is used below for bounding other parameters.

In the maximal phase perturbation case, corresponding to $b = 1$, the modulus squared of the characteristic measurement becomes

$$|\Lambda_\chi(1,\theta)|^2 = \frac{1+T_{r_1}(\cos(\theta))}{1+T_{r_0}(\cos(\theta))}, \quad (6.75)$$

which is a rational function that can be inverted for specific frequency ratios. In particular, the inverses for frequency ratios of $2:1$ and $5:1$ are now explained in the following.

For a frequency ratio of $2:1$, Eqn. 6.75 reduces to

$$|\chi|^2 = \frac{1+T_2(\cos(\theta))}{1+T_1(\cos(\theta))} \quad (6.76)$$

$$= \frac{1+(2\cos^2(\theta)-1)}{1+\cos(\theta)} \quad (6.77)$$

$$= \frac{2\cos^2(\theta)}{1+\cos(\theta)} \quad (6.78)$$

$$\Rightarrow \cos(\theta) = \frac{|\chi|^2 \pm \sqrt{|\chi|^4+8|\chi|^2}}{4}. \quad (6.79)$$

A plot of Eqn. 6.79 is given in Fig. 6.20a. The positive variation only gives valid results for $|\chi| \in [0, 1]$ and the negative variation is clearly a case of more extreme perturbation. Alas, unless $b = 1$, this does not allow us to place a bound of any sort on θ – just because it does not correspond to the maximum perturbation case, does not make any particular value of θ impossible. On the other hand, we do know that this corresponds to the maximum perturbation case, which means that by assuming that $b = 1$ holds, an upper bound can be placed on the modulus of θ_f. Following

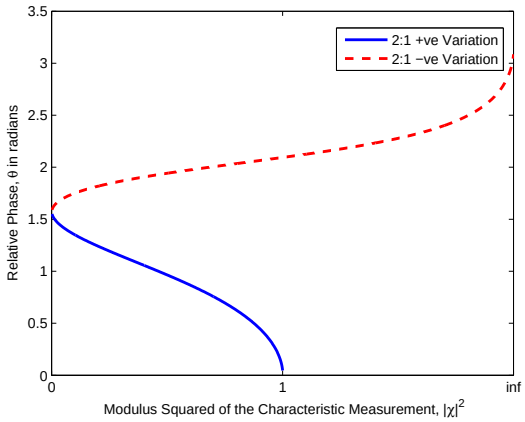

(a) θ at 2:1

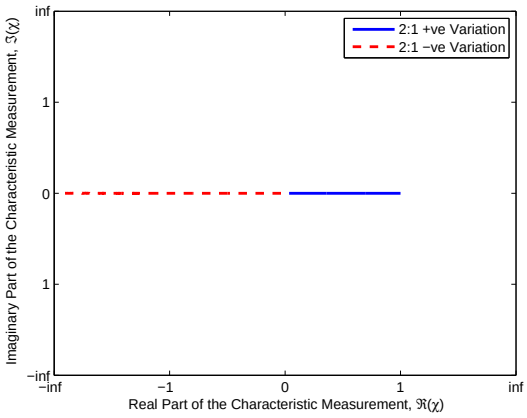

(b) χ at 2:1

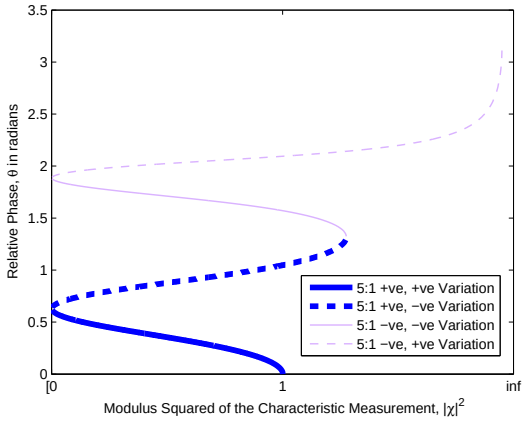

(c) θ at 5:1

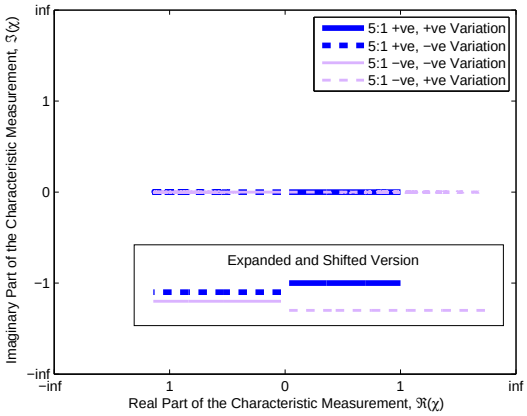

(d) θ at 5:1

Figure 6.20: Plot of the special case corresponding to maximal phase perturbation of the primary component return

this logic, we arrive at

$$|\theta_{\mathrm{f}}| \leq \frac{1}{2}\cos^{-1}\left(\frac{|\chi|^2 - \sqrt{|\chi|^4 + 8|\chi|^2}}{4}\right). \tag{6.80}$$

Unfortunately, this is a somewhat loose bound in practice.

The equivalent at 5:1 gives

$$\cos(\theta) \;=\; \frac{1 \pm \sqrt{5 \pm 4|\chi|}}{4}. \tag{6.81}$$

Plotting θ as a fuction of $|\chi|^2$ gives Fig. 6.20c. In order to understand the physical significance, we substitute the value of θ into $\Lambda_\chi(1,\theta)$ to determine the resultant full values of the characteristic measurement. These are shown in Fig. 6.20d.

From Fig. 6.20c, it is obvious that the -ve, +ve variation results in the most extreme value of θ, therefore minimum $\cos(\theta)$. A bound on phase perturbation of the primary component return is therefore given by

$$|\theta_{\mathrm{f}}| \leq \frac{1}{2}\cos^{-1}\left(\frac{1-\sqrt{5+4|\chi|}}{4}\right). \tag{6.82}$$

Fig. 6.21 illustrates several different bounds that are developed in this chapter, in particular Fig. 6.21a plots the bound from Eqn. 6.82. While many of the bounds are quite loose, there is a lot of potential for incorporating the information into processing algorithms where the necessity for numerical optimisation routines or lookup tables is prohibitive.

6.4.2 Bounding Phase Using $\arg(\chi)$

This section develops bounds on phase perturbation using the phase of the characteristic measurement. Bounds using the phase information from the characteristic measurement become substantially more difficult as $r_0 > 1$, because there are multiple possible bounds available, depending on which solution is chosen. We start by deriving a lower bound on the value of θ. This is followed by an investigation of the impact of the modulus of the characteristic measurement on the phase, which leads to the development of some ad-hoc bounds on phase perturbation for a relative frequency of $2:1$.

Given a particular value of $\arg(\chi)$, a lower bound can be placed on θ. The phase of the characteristic measurement is given by

$$\arg(\chi) \;=\; \arg(\Lambda_{\mathrm{f}}(b, r_1\theta)) - \frac{r_1}{r_0}\arg(\Lambda_{\mathrm{f}}(b, r_0\theta)) + \frac{2\pi}{r_0}m, \tag{6.83}$$

where $m \in \mathbb{Z}$ is an arbitrary constant. In the most efficient case, there is no cancellation between the first and second terms, in other words,

$$\mathrm{sgn}(\arg(\Lambda_{\mathrm{f}}(b, r_1\theta))) = -\mathrm{sgn}(\arg(\Lambda_{\mathrm{f}}(b, r_0\theta))). \tag{6.84}$$

If $r_1 > r_0$, for small values of θ the r_1 term has a greater magnitude than the r_0 term, that is

$$|\arg(\Lambda_{\mathrm{f}}(b, r_1\theta))| > |\arg(\Lambda_{\mathrm{f}}(b, r_0\theta))|. \tag{6.85}$$

Values of θ for which this does not hold correspond to cases for which there is a

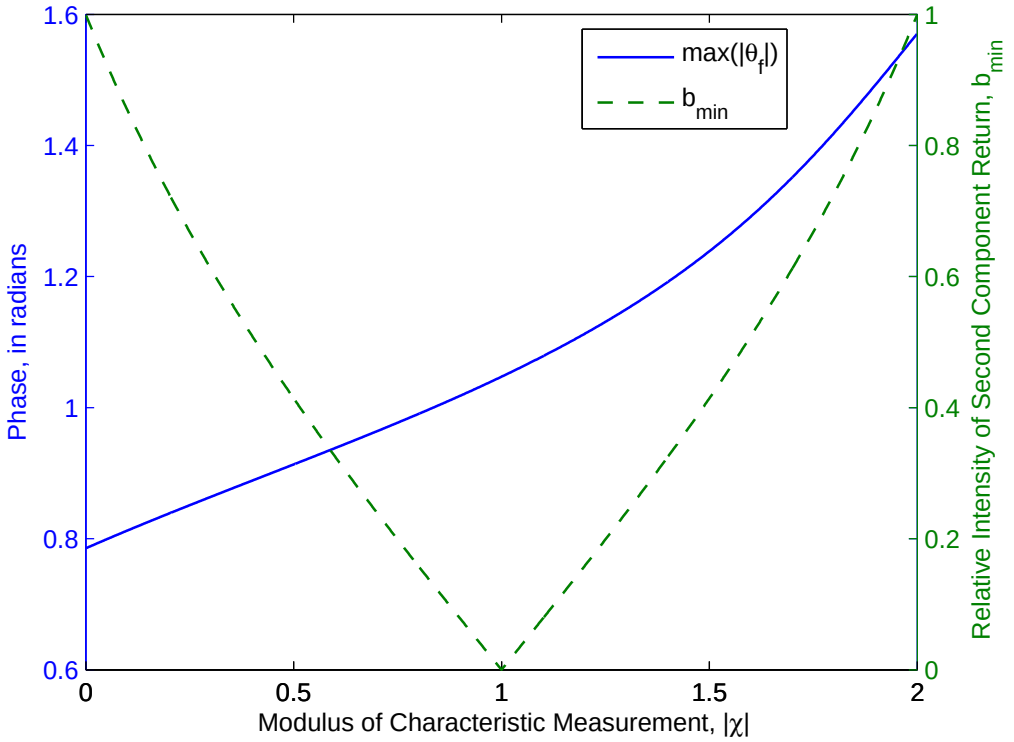

(a) Modulus of Characteristic Measurement

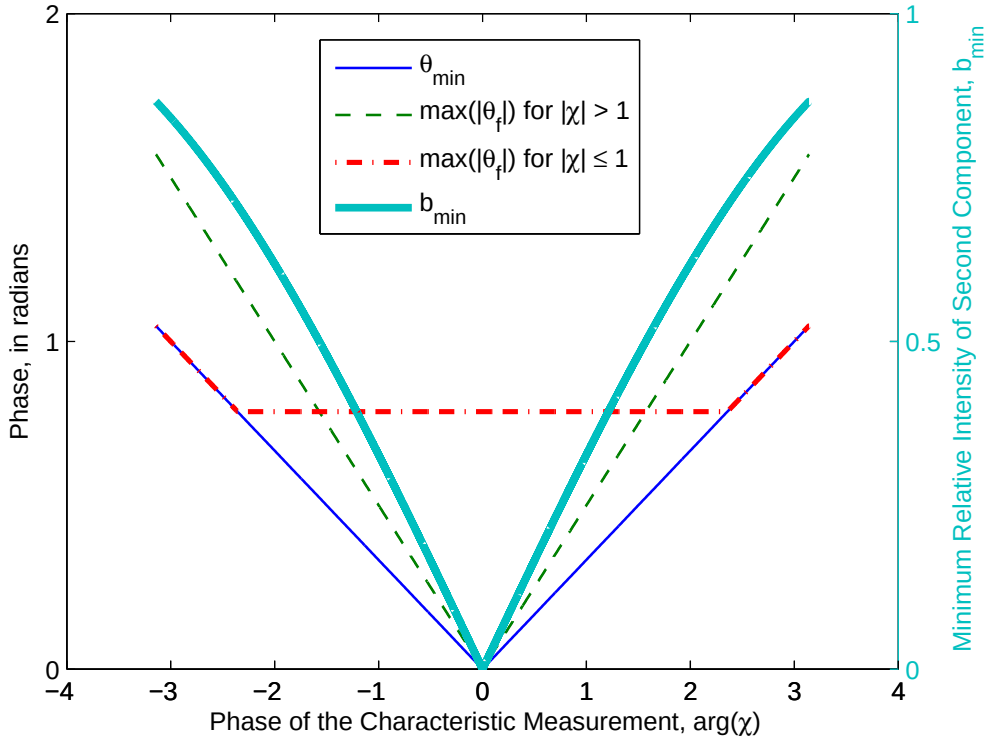

(b) Phase of Characteristic Measurement

Figure 6.21: Bounds on phase and relative intensity using the characteristic measurement assuming a frequency ratio of 2 : 1.

smaller value of θ with the same value of $|\Lambda_{\mathrm{f}}(b, r_1\theta)|$ which is the true bounding case. This allows us to write

$$\left|\arg(\chi) - \frac{2\pi}{r_0}m\right| \leq \left(1 + \frac{r_1}{r_0}\right)|\arg(\Lambda_{\mathrm{f}}(b, r_1\theta))|\,. \tag{6.86}$$

If $|\arg(\chi)| > \pi/2$ then Eqn. 6.86 accepts a looser than necessary bound for the sake of simplicity.

In order to facilitate a practical bound on relative phase, a specific value needs to be determined for m. In an ideal situation, m would be fully constrained by other knowledge such as measurements at other frequencies or spatial connectivity. Without this information, m may take any integer value, therefore must be chosen to minimally constrain any derived bounds. In this particular case, this constraint amounts to choosing m such that

$$\forall_{n\in\mathbb{Z}}\left|\arg(\chi) - \frac{2\pi}{r_0}m\right| \leq \left|\arg(\chi) + \frac{2\pi}{r_0}n\right|. \tag{6.87}$$

Recall that maximum phase perturbation of a primary component return occurs when the secondary component return is the same brightness as the primary component, that is, when $b = 1$. From Eqn. 6.34, in this circumstance

$$\arg(\Lambda_{\mathrm{f}}(1, \theta)) \;=\; \frac{\theta}{2}, \tag{6.88}$$

where $\theta \in [-\pi, \pi]$. By combining Eqns. 6.86 and 6.88, the following minimum bound on θ arises:

$$\frac{\arg(\chi) - \frac{2\pi}{r_0}m}{1 + \frac{r_1}{r_0}} \leq \frac{r_1}{2}\theta_{\min}. \tag{6.89}$$

For $2\!:\!1$, $\mathrm{sgn}(\arg(\chi)) = -\mathrm{sgn}(\theta)$ also holds. In practice, this bound is primarily useful when $\arg(\chi)$ is significantly non-zero.

In Section 6.4.1, we analysed the derivative of the phase of the characteristic measurement as a function of θ, we now analyse the remaining derivative of the phase of the characteristic measurement so that we can explain how the modulus of the characteristic measurement impacts on relative phase. Taking the derivative of Eqn. 6.68 versus b gives

$$\frac{\partial\arg(\Lambda_\chi(b,\theta)}{\partial b} \;=\; \frac{\partial\arg(\Lambda_{\mathrm{f}}(b, r_1\theta))}{\partial\theta} - \frac{r_1}{r_0}\frac{\partial\arg(\Lambda_{\mathrm{f}}(b, r_0\theta))}{\partial\theta} \tag{6.90}$$

$$\;=\; \frac{\sin(r_1\theta)}{1 + b^2 + 2b\cos(r_1\theta)} - \frac{r_1}{r_0}\frac{\sin(r_0\theta)}{1 + b^2 + 2b\cos(r_0\theta)}. \tag{6.91}$$

Setting this to zero and rearranging gives

$$r_0 \sin(r_1\theta)(1 + b^2 + 2b\cos(r_0\theta)) - r_1 \sin(r_0\theta)(1 + b^2 + 2b\cos(r_1\theta)) \quad = \quad 0. \tag{6.92}$$

For a particular value of θ, one can determine b via the quadratic equation. Only a small percentage of values of θ correspond to valid, non-complex values for b. Unlike the other stationary points of the characteristic measurement, which have clear impacts on the shape of the manifold or behaviour of the multiple return problem, as yet we have been unable to interpret the implication of the this equation. The only other non-trivial set of stationary points is given by $\partial|\Lambda_\chi(b,\theta)|^2/\partial\theta = 0$ and clearly correlates to folds in the solution manifold, whereas $\partial|\Lambda_\chi(b,\theta)|^2/\partial b = 0$ and $\partial\arg(\Lambda_\chi(b,\theta))/\partial\theta = 0$ correspond to specific fixed values of b or θ. Two example plots of these stationary points are given in Figs. 6.23a and 6.23b. In the five to one ratio case, the stationary points do not appear to correlate with any specific known behaviour. However, in the three to two case the stationary points of $\partial\arg(\Lambda_\chi(b,\theta))/\partial b = 0$ appear to be a subset of the stationary points of $\partial|\Lambda_\chi(b,\theta)|^2/\partial\theta = 0$. This is particularly interesting, because if the function were holomorphic one would expect the set of stationary points to be identical for both.

The square of the modulus of the characteristic measurement is an indicator of which measurement, either that at r_0 or that at r_1, has the greatest impact on the phase of the characteristic measurement. From Eqn. 6.10,

$$|\chi|^2(1 + b^2 + 2b\cos(r_0\theta)) \quad = \quad 1 + b^2 + 2b\cos(r_1\theta). \tag{6.93}$$

The derivative of $\arg(\chi)$ versus θ was given in Eqn. 6.70 as

$$\frac{\partial\arg(\Lambda_\chi(b,\theta))}{\partial\theta} \quad = \quad r_1\left(\frac{b(\cos(r_1\theta)+b)}{1+b^2+2b\cos(r_1\theta)} - \frac{b(\cos(r_0\theta)+b)}{1+b^2+2b\cos(r_0\theta)}\right). \tag{6.94}$$

By substituting Eqn. 6.93 into Eqns. 6.91 and 6.94 gives,

$$\frac{\partial\arg(\Lambda_\chi(b,\theta))}{\partial\theta} \quad \propto \quad |\chi|^{-2}(\cos(r_1\theta)+b) - (\cos(r_0\theta)+b) \tag{6.95}$$

$$\frac{\partial\arg(\Lambda_\chi(b,\theta)}{\partial b} \quad \propto \quad |\chi|^{-2}\sin(r_1\theta) - \frac{r_1}{r_0}\sin(r_0\theta). \tag{6.96}$$

In other words, as $|\chi| \to \infty$, the phase of the characteristic measurement is dominated by the phase of the measurement at r_0 and as $|\chi| \to 0$, the phase of the characteristic measurement is dominated by the phase of the measurement at r_1. Consider now the case $|\chi| \to \infty$. In order for $|\chi| \to \infty$, it is necessary for $|\Lambda_f(b, r_0\theta)|^2 \to 0$

while $|\Lambda_{\mathrm{f}}(b, r_1\theta)|^2$ remains non-zero. In other words

$$|\Lambda_{\mathrm{f}}(b, r_0\theta)|^2 \;=\; 1 + b^2 + b\cos(r_0\theta) \;\approx\; 0 \tag{6.97}$$

$$\Rightarrow \quad (r_0\theta) \bmod 2\pi \;\approx\; \pi. \tag{6.98}$$

Given infinitisimal perturbations, ϵ_θ in the relative phase and ϵ_{b} in the relative intensity, $\arg(\chi)$ becomes

$$\arg(\chi) \;=\; \arg(\Lambda_{\mathrm{f}}(1-\epsilon_{\mathrm{b}}, r_1(\theta+\epsilon_\theta))) - \frac{r_1}{r_0}\arg(\Lambda_{\mathrm{f}}(1-\epsilon_{\mathrm{b}}, r_0(\theta+\epsilon_\theta))). \tag{6.99}$$

Given $|\epsilon_\theta| \ll \theta$ and $|\epsilon_{\mathrm{b}}| \ll 1$ and the relationships in Eqns. 6.95 and 6.96 we can regard the first term as a constant ϕ_{r_1}, with the second term as the only function of ϵ_θ and ϵ_{b}, viz

$$\arg(\chi) \;=\; \phi_{r_1} - \frac{r_1}{r_0}\arg(\Lambda_{\mathrm{f}}(1-\epsilon_{\mathrm{b}}, r_0(\theta+\epsilon_\theta))) \tag{6.100}$$

$$\;=\; \phi_{r_1} - \frac{r_1}{r_0}\theta_{\mathrm{f}} \tag{6.101}$$

$$\Rightarrow \quad |\theta_f| \;\leq\; \frac{r_0}{r_1}|\arg(\chi) - \phi_{r_1}| \;\wedge\; \mathrm{sgn}(\theta_{\mathrm{f}}) = -\mathrm{sgn}(\arg(\chi)) \tag{6.102}$$

where θ_{f} is the phase perturbation of the primary component return notated at a relative frequency of r_0. In addition to issues regarding the determination of ϕ_{r_1}, it is not immediately obvious how to utilise this equation as a bound. We can address these bounds by chosing a frequency ratio of $2:1$, which immediately results in $\phi_{r_1} = 0$, a useful disambiguation.

It is important to determine when Eqn. 6.102 holds true. We know that it is necessary for the second term from Eqns. 6.95 and 6.96 to dominate, it is a reasonable assumption that this requires, at a bare minimum, that $|\chi| \geq 1$. In order to resolve this, we rely upon an empirical approach. Fig. 6.22a shows a plot of the value of θ as a function of χ and Fig. 6.22b shows a plot of the value of θ_{f} as a function of χ. The black circle corresponds to $|\chi| \to \infty$, the bounding case; and the white circle corresponds to the unit circle, which provides another special case to utilise as a bound.

Several interesting inferences can be made from Fig. 6.22. From the first plot, it appears that

$$\left(|\chi| > 1 \Rightarrow |\theta| > \frac{2\pi}{3}\right) \wedge \left(|\chi| < 1 \Rightarrow |\theta| < \frac{2\pi}{3}\right) \tag{6.103}$$

which is a simple binary constraint. From the second plot, we can also infer that

Eqn. 6.102 holds for all values of $|\chi| > 1$. The equation also holds for the two leftmost quadrants of $|\chi| < 1$, but is broken as a constraint in the top right and bottom right quadrants. A practical bound requires an approach to deal with these regions of the solution-space. We now approach this problem using the discovery that the values of θ_f around the unit circle are equally spaced.

In Section 6.3.6, we noted that the unit circle corresponded to certain fixed values of θ. For a ratio of $2:1$, $\theta = \pm 2\pi/3$ and varying b sweeps around the unit circle. In the positive case we write $\arg(\chi)$ as

$$\begin{aligned}
\arg(\chi) &= \arg\left(\frac{\Lambda_\mathrm{f}(b, 4\pi/3)}{\Lambda_\mathrm{f}(b, 2\pi/3)^2}\right) && (6.104)\\
&= \arg(\Lambda_\mathrm{f}(b, 4\pi/3)) - 2\arg(\Lambda_\mathrm{f}(b, 2\pi/3)) && (6.105)\\
&= \arg(\Lambda_\mathrm{f}(b, -2\pi/3)) - 2\arg(\Lambda_\mathrm{f}(b, 2\pi/3)) && (6.106)\\
&= -3\arg(\Lambda_\mathrm{f}(b, 2\pi/3)) = -3\theta_\mathrm{f}. && (6.107)
\end{aligned}$$

This fundamental result is unimpacted by changing the sign of θ. Supplementary analysis reveals this to be a specialised form of a general solution. Recall that the constraints from Section 6.3.6, relative phases that map to unity are given by

$$\mathrm{e}^{r_1\theta} - \mathrm{e}^{r_0\theta} = 0. \quad (6.108)$$

Also recall that $\cos(r_1\theta) = \cos(r_0\theta) \Leftrightarrow |\chi| = 1$. In order for χ to be non-unity, but still map to the unit circle it is necessary that

$$\mathrm{e}^{r_1\theta} - \mathrm{e}^{-r_0\theta} = 0. \quad (6.109)$$

Using Eqn. 6.109 allows Eqn. 6.107 to be generalised to

$$\begin{aligned}
\arg(\chi) &= \arg\left(\frac{1 + b\mathrm{e}^{jr_1\theta}}{(1 + b\mathrm{e}^{jr_0\theta})^{r_1/r_0}}\right) && (6.110)\\
&= \arg\left(\frac{1 + b\mathrm{e}^{-jr_0\theta}}{(1 + b\mathrm{e}^{jr_0\theta})^{r_1/r_0}}\right) && (6.111)\\
&= \arg\left((1 + b\mathrm{e}^{jr_0\theta})^{-r_1/r_0 - 1}\right) && (6.112)\\
&= -\left(1 + \frac{r_1}{r_0}\right)\theta_\mathrm{f} + \frac{2\pi}{r_0}m, && (6.113)
\end{aligned}$$

where $m \in \mathbb{Z}$ is an arbitrary constant. Let us resume developing the specific bound at $2:1$. Within the unit circle, $3|\theta_\mathrm{f}| \leq |\arg(\chi)|$ is tighter than $2|\theta_\mathrm{f}| \leq |\arg(\chi)|$, but does not always hold. Referring again to Fig. 6.22b, we can ameliorate this using

the ad-hoc inclusion of the special case at $|\chi| = 0$. This gives

$$|\Lambda_{\mathrm{f}}(b, 2\theta)|^2 \;=\; 1 + b^2 + 2b\cos(2\theta) \;=\; 0 \tag{6.114}$$

$$\Rightarrow \qquad \theta = \tfrac{\pi}{2} \;\wedge\; b = 1 \tag{6.115}$$

$$\Rightarrow \qquad \theta_{\mathrm{f}} = \tfrac{\pi}{4}. \tag{6.116}$$

Thus for $|\chi| < 1$, phase perturbation is bounded by

$$|\theta_{\mathrm{f}}| \leq \max\left(\frac{\pi}{4}, \frac{|\arg(\chi)|}{3}\right). \tag{6.117}$$

The bounds from Eqns. 6.102 and 6.117 on phase perturbation are plotted in Fig. 6.21b; however, the raw bound relationships do not give an intuitive understanding of how loose or tight the bounds are in practical usage – this is discussed in the next chapter.

6.4.3 Bounding Relative Intensity Using $|\chi|^2$

We now focus on bounding relative intensity, b, given the modulus of the characteristic measurement. We start by considering the stationary points of $|\Lambda_\chi(b, \theta|^2$ versus θ.

The condition for extrema with respect to b of the characteristic measurement, Eqn. 6.92, is difficult to manipulate because it contains multiple sine terms using different frequency ratios. While it is possible to rewrite some particular cases using Chebyshev polynomials of the second kind, detailed analysis is not particularly valuable. In Section 6.3.5, we described how the solution manifold is folded back over itself and how this results in multiple solutions for specific values of χ and associated measurement frequency ratios. This set of stationary points corresponds to the location of folds in the manifold, which is illustrated in Fig. 6.23. Fig. 6.23a shows the set of stationary points for a frequency ratio of $3:2$; recall that this frequency ratio has three different possible solutions – as a result there are no folds. The stationary points still exist, but are not correlated with fold lines. In Fig. 6.23b, the set of stationary points is plotted for a frequency ratio of $5:1$. This particular frequency has any of either zero, one, two or four solutions for any particular value of χ. In this figure, the stationary points are clearly correlated with folds. Note that there is a small section missing from one of the fold lines on the left; this is an artefact of the method used to find the stationary points. We can conclude that the stationary points are necessary, but not sufficient, for the presence of a fold.

We now focus on the bounding problem. The values of θ corresponding to

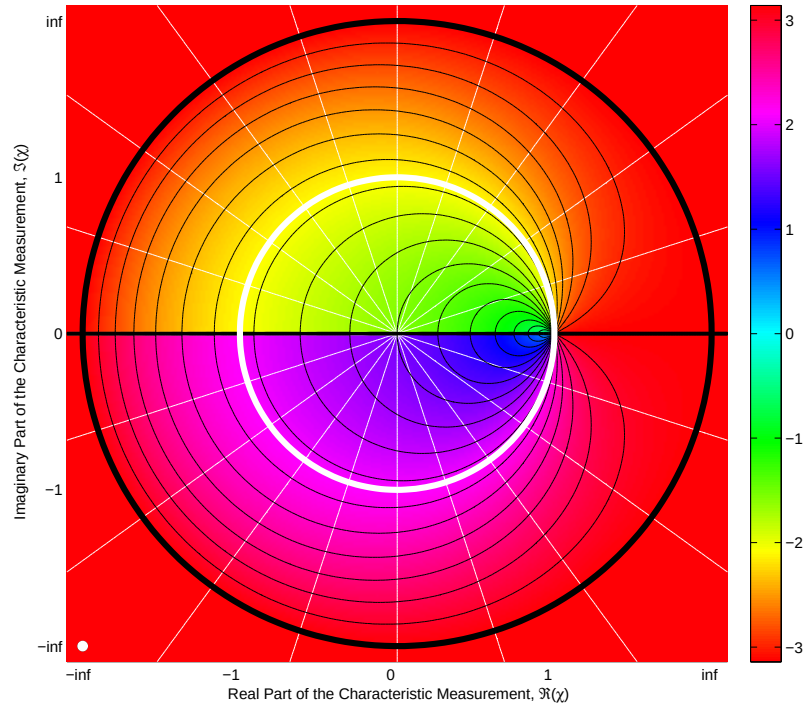

(a) θ vs. $\arg(\chi)$ (Black Contours)

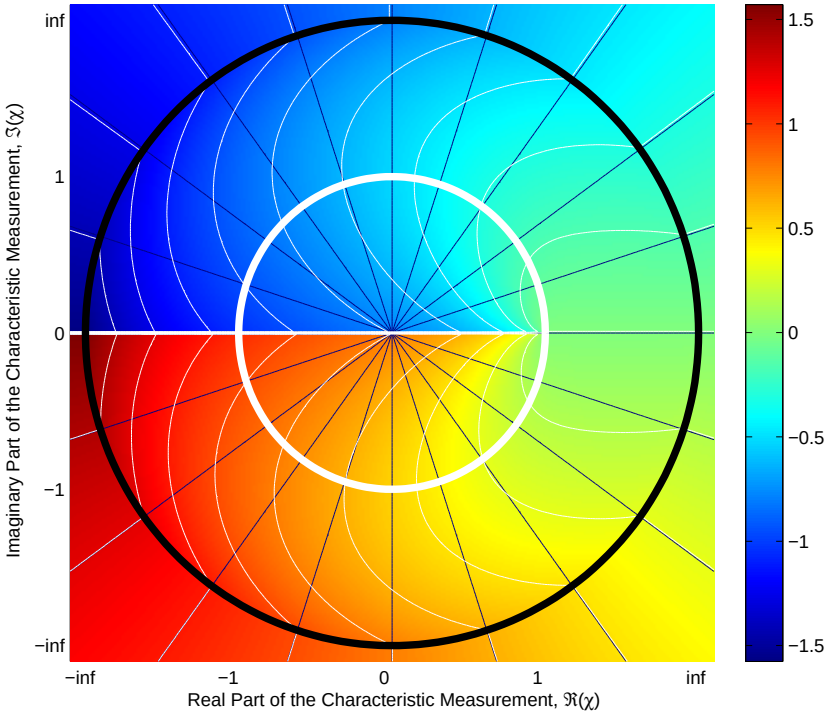

(b) θ_{f} vs. $\arg(\chi)$ (Blue Contours)

Figure 6.22: Parameter relationships with $\arg(\chi)$.

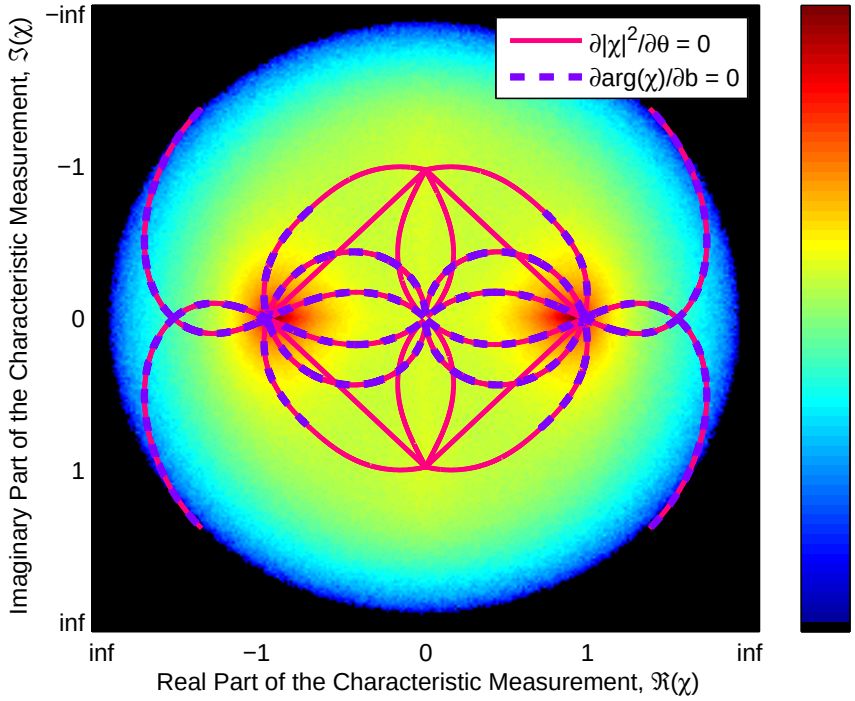

(a) 3 : 2

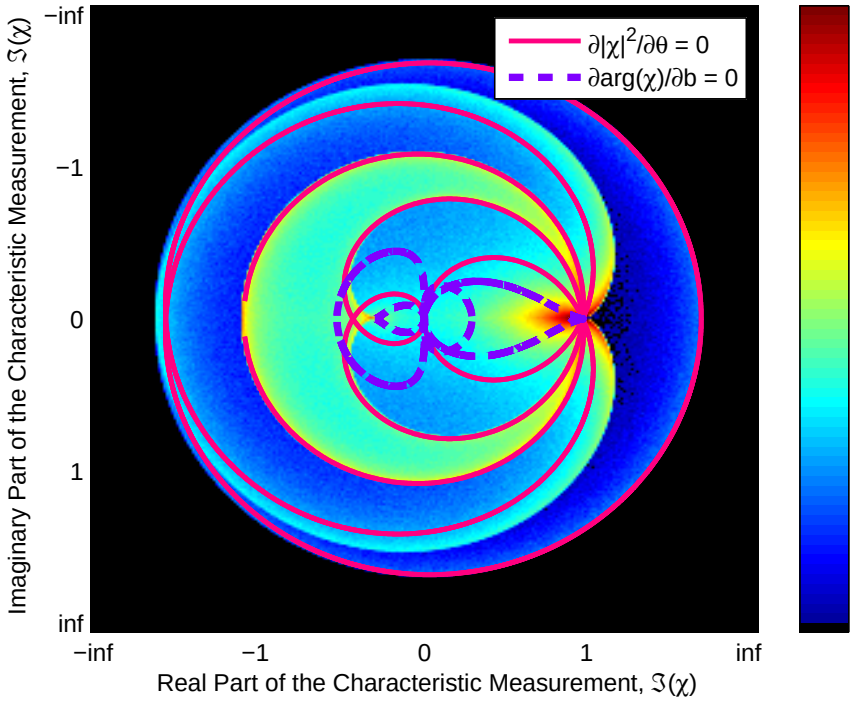

(b) 5 : 1

Figure 6.23: A plot of the stationary points corresponding to $\partial|\Lambda_\chi(b,\theta)|^2/\partial\theta = 0$ and $\partial\arg(\Lambda_\chi(b,\theta))/\partial b = 0$.

maximum and minimum values of $|\chi|^2$ may be determined by explicitly enumerating over the set of relative phases corresponding to zeros, poles and singularities, given by

$$L_\theta = \{\theta\colon |\Lambda_{\mathrm{f}}(1, r_1\theta)|^2 = 0 \vee |\Lambda_{\mathrm{f}}(1, r_0\theta)|^2 = 0\}. \tag{6.118}$$

We know from Eqn. 6.53 that the most extreme values of the characteristic measurement occur when $b = 1$, thus in order to find the values of $\theta \in L_\theta$ that correspond to maximum and minimum values of $|\chi|^2$ we take the equation for the square of the modulus of the characteristic measurement and substitute $b = 1$, giving

$$|\Lambda_\chi(1,\theta)|^2 = \frac{1 + b^2 + 2b\cos(r_1\theta)}{1 + b^2 + 2b\cos(r_0\theta)} \tag{6.119}$$

$$= \frac{1 + \cos(r_1\theta)}{1 + \cos(r_0\theta)}. \tag{6.120}$$

In order to determine the maximum bounding case, we need to find the value of $\theta \in L_\theta$ for which $|\Lambda_\chi(1,\theta)|^2$ is largest. On its own, this is inadequate due to differences between poles – just because $|\Lambda_\chi(1,\theta)|^2 = \infty$ for two different poles does not imply that the modulii are the same for values of b less than unity. In other words, while the denominator may dominate at $b = 1$, for $b \neq 1$ the numerator is of equal importance. When comparing two poles, the pole with the largest numerator bounds the other. An equivalent comparison can be implemented for the bounding case for $|\chi|^2 < 1$. Subsequently, the bounding cases, $\theta \in \{\theta_{\mathrm{low}}, \theta_{\mathrm{high}}\}$, are given by the solution to

$$\forall_{\phi\in(L_\theta\setminus\{\theta_{\mathrm{low}}\})} \begin{pmatrix} & \dfrac{1+\cos(r_1\theta_{\mathrm{low}})}{1+\cos(r_0\theta_{\mathrm{low}})} & < & \dfrac{1+\cos(r_1\phi)}{1+\cos(r_0\phi)} & \\ \vee & |\Lambda_{\mathrm{f}}(1, r_1\theta_{\mathrm{low}})|^2 & = & |\Lambda_{\mathrm{f}}(1, r_1\phi)|^2 & = 0 \\ \wedge & |\Lambda_{\mathrm{f}}(1, r_0\theta_{\mathrm{low}})|^2 & \geq & |\Lambda_{\mathrm{f}}(1, r_0\phi)|^2 & \end{pmatrix}$$

$$\wedge \quad \forall_{\phi\in(L_\theta\setminus\{\theta_{\mathrm{high}}\})} \begin{pmatrix} & \dfrac{1+\cos(r_1\theta_{\mathrm{high}})}{1+\cos(r_0\theta_{\mathrm{high}})} & > & \dfrac{1+\cos(r_1\phi)}{1+\cos(r_0\phi)} & \\ \vee & |\Lambda_{\mathrm{f}}(1, r_0\theta_{\mathrm{high}})|^2 & = & |\Lambda_{\mathrm{f}}(1, r_0\phi)|^2 & = 0 \\ \wedge & |\Lambda_{\mathrm{f}}(1, r_1\theta_{\mathrm{high}})|^2 & \geq & |\Lambda_{\mathrm{f}}(1, r_1\phi)|^2 & \end{pmatrix}, \tag{6.121}$$

where θ_{min} corresponds to the bounding case for values of the modulus of the characteristic measurement below unity, and θ_{min} is the bounding case for values of the modulus of the characteristic measurement above unity. As an important sidenote, for $|\chi|^2 = 1$ we cannot bound b in any manner whatsoever; as we know from Section 6.3.6, b can take any value so long as θ is confined to one of a limit set of values specific to the measurement frequency ratio. As it turns out, though, there is no

need to explicitly handle this as it is implicit in the high and low bounds notated below.

For any value of the characteristic measurement,

$$\begin{aligned} &\left(|\chi|^2 \geq 1 \;\Rightarrow\; |\chi|^2 \;\leq\; \frac{1+b^2+2b\cos(r_1\theta_{\text{high}})}{1+b^2+2b\cos(r_0\theta_{\text{high}})}\right) \\ \vee\; &\left(|\chi|^2 \leq 1 \;\Rightarrow\; |\chi|^2 \;\geq\; \frac{1+b^2+2b\cos(r_1\theta_{\text{low}})}{1+b^2+2b\cos(r_0\theta_{\text{low}})}\right). \end{aligned} \tag{6.122}$$

The modulus of the characteristic measurement can therefore be rewritten as

$$|\chi|^2 = \frac{1+b_{\text{min}}^2+2b_{\text{min}}\cos(r_1\theta_{\text{b}})}{1+b_{\text{min}}^2+2b_{\text{min}}\cos(r_0\theta_{\text{b}})}, \tag{6.123}$$

where

$$\theta_{\text{b}} \;=\; \begin{cases} \theta_{\text{high}}, & \text{if } |\chi| > 1 \\ \theta_{\text{low}}, & \text{if } |\chi| \leq 1 \end{cases}. \tag{6.124}$$

Eqn. 6.123 then becomes

$$(|\chi|^2-1)b^2 + 2(\cos(r_1\theta_{\text{b}})|\chi|^2 - \cos(r_0\theta_{\text{b}}))b + |\chi|^2 - 1 = 0. \tag{6.125}$$

For frequency ratios where the range of valid values of the modulus of the characteristic measurement is given by $|\chi| \in [0,\infty]$, the solutions for b_{min} are quite simple. The only cases for which the solutions are more complicated occur when $r_0 = 1$ and $r_1/2 \notin \mathbb{Z}$ or vice-versa. If $|\chi| \in [0,\infty]$, then

$$(|\chi| > 1 \;\Rightarrow\; \cos(r_1\theta_{\text{b}}) = 1 \wedge \cos(r_0\theta_{\text{b}}) = -1) \tag{6.126}$$

$$\wedge\;\; (|\chi| < 1 \;\Rightarrow\; \cos(r_1\theta_{\text{b}}) = -1 \wedge \cos(r_0\theta_{\text{b}}) = 1) \tag{6.127}$$

which implies that for $b \in [0,1]$

$$b_{\text{min}} \;=\; \begin{cases} \frac{|\chi|-1}{|\chi|+1} & \text{if } |\chi| > 1 \\ -\frac{|\chi|-1}{|\chi|+1} & \text{if } |\chi| \leq 1 \end{cases} \tag{6.128}$$

This bound always holds irrespective of the frequency ratio; however, for frequency ratios such as 3 : 1, there are more complex, tighter bounds available taking into account the lack of any poles in the modulus of the characteristic measurement. This bound is plotted in Fig. 6.21a for the two-to-one frequency ratio case.

6.4.4 Bounding Relative Intensity Using $\arg(\chi)$

From Eqn. 6.32, we know that for any particular relative intensity b, where the relative phase, θ, is a free variable, the maximum phase perturbation of the unit phasor is given by

$$\arg(\Lambda_{\mathrm{f}}(b, \cos^{-1}(b))) = \sin^{-1}(b). \tag{6.129}$$

We can now apply this to the phase of the characteristic measurement. Using the same assumptions as Eqn. 6.86, the minimum value of b occurs when

$$|\arg(\chi) - \frac{2\pi}{r_0}m| = \left(1 + \frac{r_1}{r_0}\right)\sin^{-1}(b), \tag{6.130}$$

thus

$$b_{\min} = \sin\left(\frac{r_0\left|\arg(\chi) - \frac{2\pi}{r_0}m\right|}{r_0 + r_1}\right), \tag{6.131}$$

where m is given by Eqn. 6.87. This bound is plotted in the two-to-one frequency ratio case in Fig. 6.21.

6.4.5 Bounding Using Attenuation Ratios

The concept of the attenuation ratio is closely related to the characteristic measurement: the attenuation ratio is the modulus of the characteristic measurement calculated from measurements at $r_1 \in \mathbb{Z}\ \{0\}$ and $r_0 = 0$. Separation of component returns using attenuation ratio polynomials was discussed in Section 5.2. We now demonstrate how a measurement of the attenuation ratio can be used to place bounds on minimum relative intensity, minimum relative phase and maximum phase perturbation.

Combining Eqns. 5.48, 5.51 and the law of cosines allows us to write the cosine of the phase perturbation of the primary component return, $\cos(\theta_{\mathrm{f}})$, as a linear function of relative intensity, b, viz

$$b^2 = 1 + \tau_l^2(1+b)^2 - 2\tau_l(1+b)\cos(\theta_{\mathrm{f}}), \tag{6.132}$$

simplifying to

$$2\cos(\theta_{\mathrm{f}}) = b\left(\tau_l - \frac{1}{\tau_l}\right) + \tau_l + \frac{1}{\tau_l}. \tag{6.133}$$

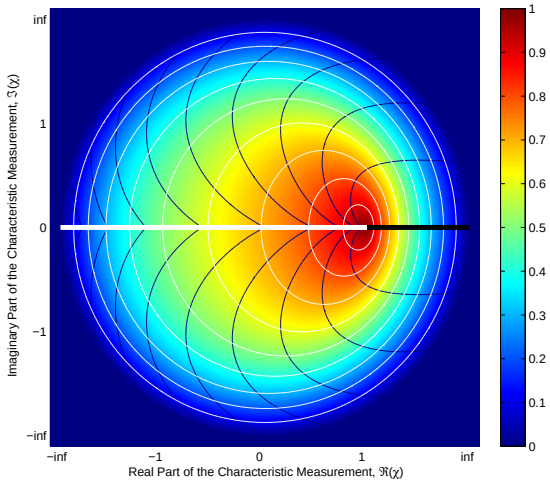

(a) τ_0 vs. θ_f(Blue Contours)

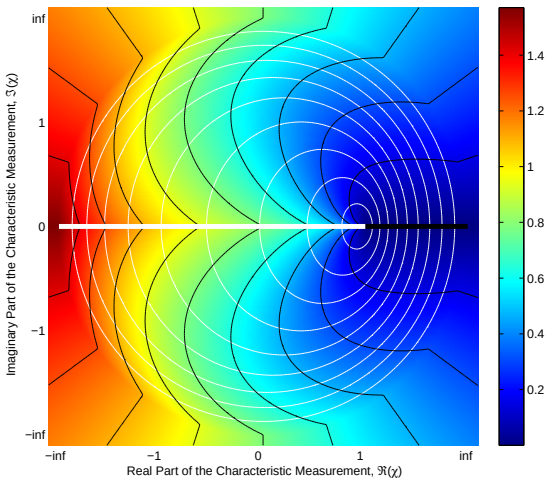

(b) θ_f vs. τ_0 (White Contours)

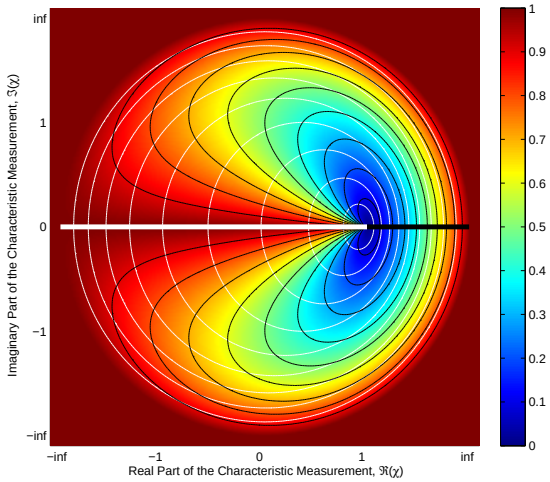

(c) b vs. τ_0 (White Contours)

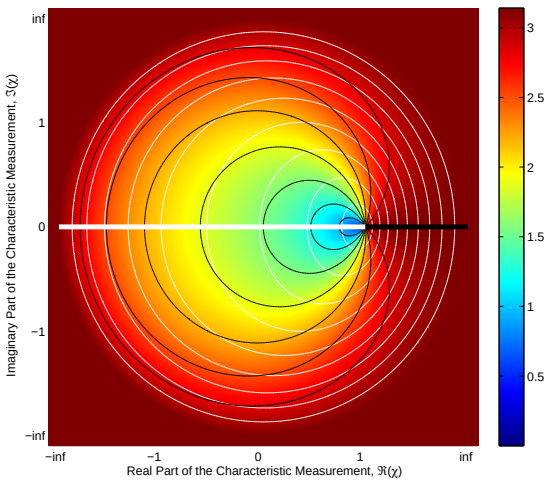

(d) θ vs. τ_0 (White Contours)

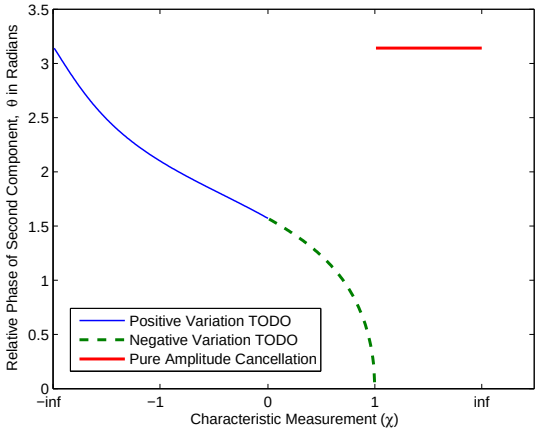

(e) Relative Phase vs. the Characteristic Measurement

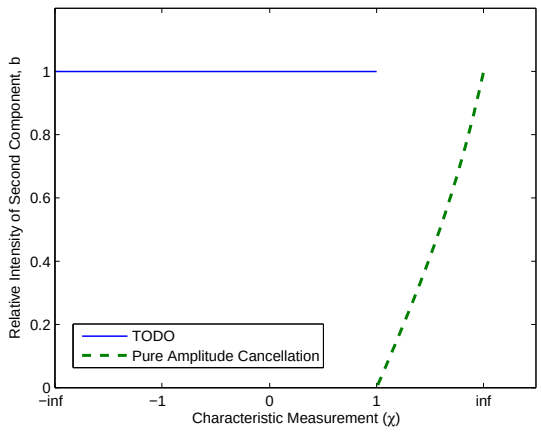

(f) Relative Intensity vs. the Characteristic Measurement

Figure 6.24: Phase, intensity and attenuation ratio relationships at a relative frequency of 2 : 1.

We assume that the returns are ordered such that $b \in [0,1]$ and θ_f corresponds to the phase perturbation of the primary component return notated at the same relative frequency as the attenuation ratio τ_l.

Setting $\frac{d|\Lambda_\mathrm{f}(b,\theta)|^2}{d\theta} = 0$ gives two stationary points. The maximum at $\theta = 0$ corresponds to the degenerate case where there is only one component return and b can take any positive value. The minimum at $\theta = \pi$ corresponds to pure active intensity cancellation with no phase perturbation, resulting in Eqn. 5.51 simplifying to

$$\tau_l = \frac{1-b}{1+b}. \tag{6.134}$$

If $\tau_l \neq 1$ then partial active intensity cancellation is occurring and since $\theta = \pi$ is the most efficient case in terms of active intensity cancellation, this implies a minimum bound on the value of b, namely

$$b_{min} = \frac{1-\tau_l}{1+\tau_l}. \tag{6.135}$$

A bound can be placed on the phase perturbation of the primary component return θ_f by finding the minimum value of $\cos(\theta_\mathrm{f})$ such that $b \in [0,1]$, giving

$$|\theta_\mathrm{f}| \leq \cos^{-1}(\tau_l). \tag{6.136}$$

Because Eqn. 6.133 is linear over b, this does not correspond to a stationary point of any sort. We conclude that maximal phase perturbation for a particular value of τ_l occurs when there are two returns each with the same active brightness. This bound still holds, even if there are additional component returns or an overestimate of total integrated intensity that has not been accounted for as long as the primary component return composes 50% or greater of the total integrated intensity. However in these cases the estimate of b_{min} increases, which is erroneous.

Fig. 6.24 shows the relationship between phase, intensity and the attenuation ratio at a frequency ratio of 2 : 1. Figs. 6.24a and b show how τ_0 can be used to bound phase perturbation at a relative frequency of one – the maximum is given by the white line and the minimum by the black. Fig. 6.24c shows the relationship between τ_0 and b, given a value of $|\chi - 1|$ the minimum possible value of b occurs for real values of the characteristic measurement above unity. Figs. 6.24e and f show plots along the real line, where the bounding cases occur.

The final bound – on θ using the attenuation ratio – is derived in the next section in Eqns. 6.158 through 6.160. It is given by

$$\theta \;\geq\; \cos^{-1}(2\tau_l^2 - 1). \tag{6.137}$$

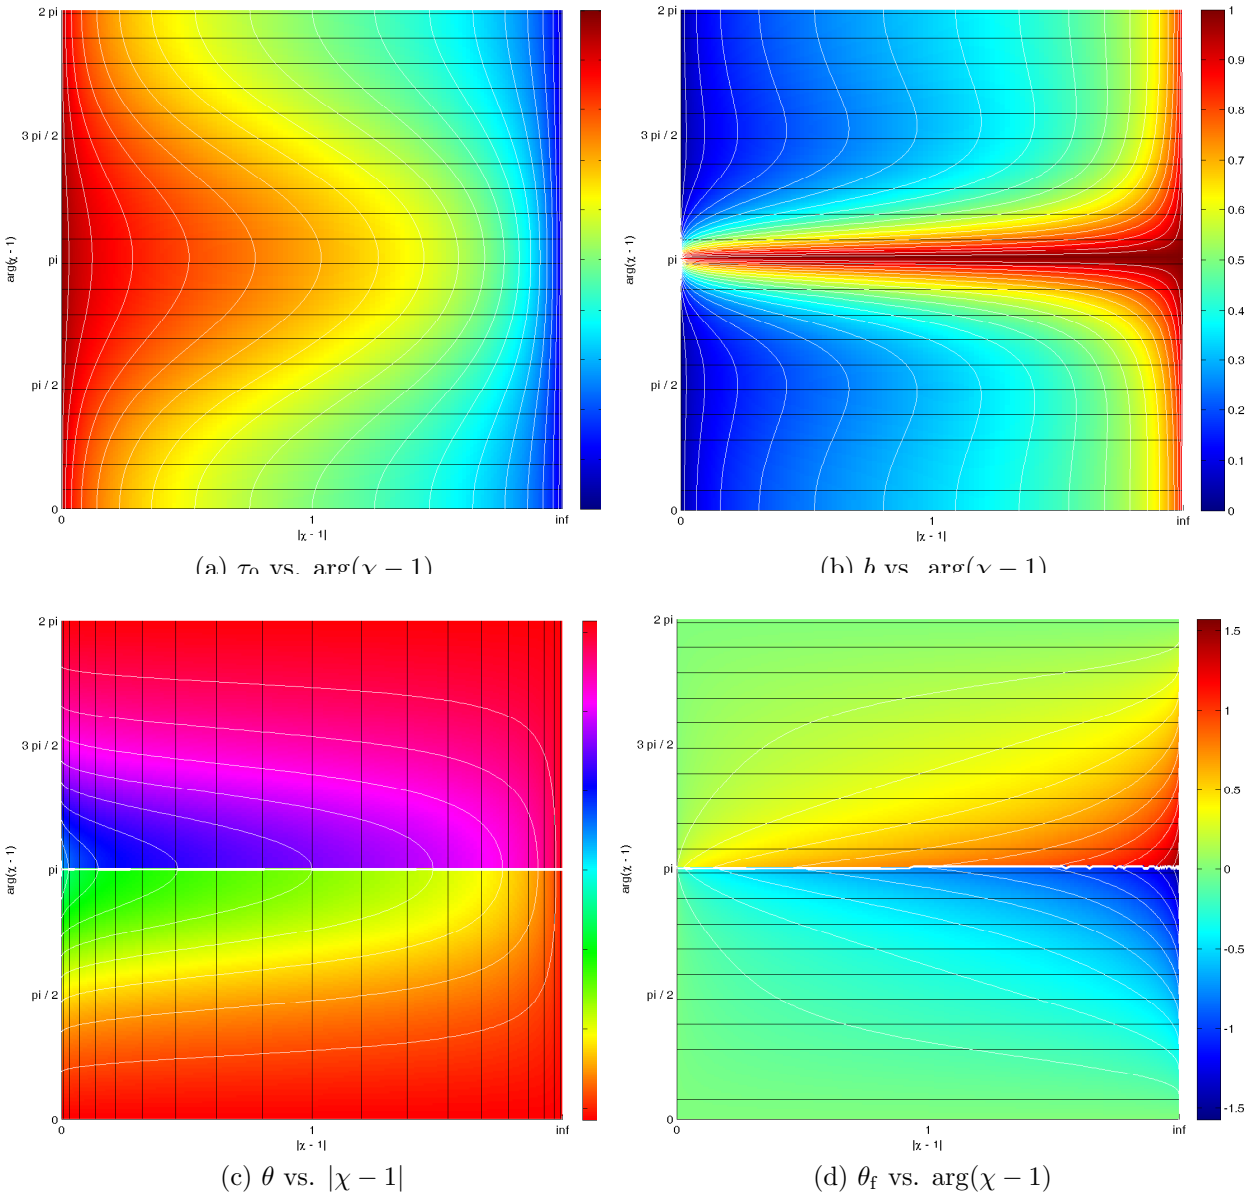

(a) τ_0 vs. $\arg(\chi - 1)$

(b) b vs. $\arg(\chi - 1)$

(c) θ vs. $|\chi - 1|$

(d) θ_f vs. $\arg(\chi - 1)$

Figure 6.25: Descriptive parameters of the mixed pixel problem at a frequency ratio of 2 : 1, unwrapped around the branch point at $\chi = 1$.

A graphical summary of the bounds available from measurements of the attenuation ratio is provided in Fig. 6.26.

6.4.6 Bounds Around $\chi = 1$

For measurement ratios where $r_0 = 1$ a remarkable transformation is possible by unwrapping around the singular branch point at $\chi = 1$ using a polar mapping. In this section we focus on the specific frequency ratio of 2 : 1, although it is possible to analyse other specific frequency ratios in a similar manner. We continue, using a

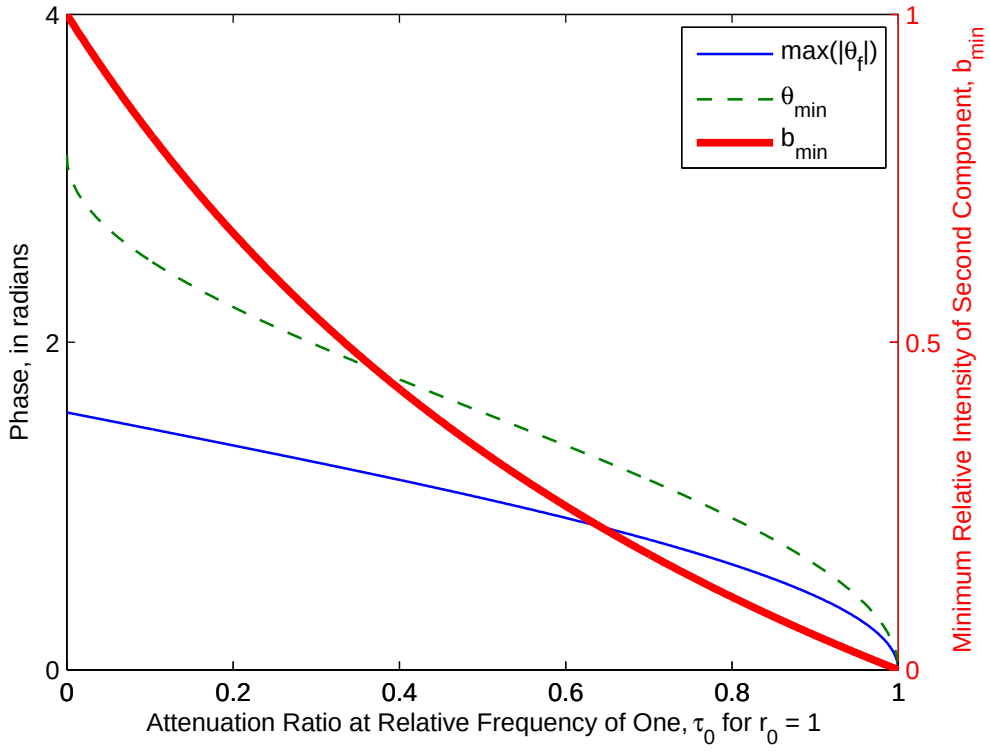

(a) Attenuation Ratio

Figure 6.26: Bounds on phase and relative intensity using the attenuation ratio.

Table 6.3: Parameter bounding conditions

	Bound			
Value	$\arg(\chi-1)=0$	$\arg(\chi-1)=\pi$	$\lvert\chi-1\rvert=0$	$\lvert\chi-1\rvert\to\infty$
τ_0	min	max	max	min
τ_1	max	min	N/A	N/A
$\lvert\theta\rvert$	max	min	min	max
$\lvert\theta_\mathrm{f}\rvert$	min	max	min	max
$\lvert\hat{\eta}_0\rvert+\lvert\hat{\eta}_1\rvert$	max	min	min	max
b	N/A	max	min	max

mixed theoretical and empirical approach, to develop a series of bounding methods based around unwrapping the branch point at unity.

The best way to appreciate the profound results from this unwrapping is by plotting θ, b and other descriptive parameters in this new, unwrapped mapping. Fig. 6.25 plots τ_0, b, θ and θ_f in this new translated polar form. Whereas in the zero centred polar and simple Cartesian mappings it is difficult to derive simple bounds for many parameters, this mapping allows a large number of easily contrived bounds. For comparison, see Cartesian plots in Figs. 6.22 and 6.24.

A list of specific bounding cases is given in Table 6.3. For a particular value of

$\arg(\chi - 1)$, the value of a parameter, say b, can be bounded using the $|\chi - 1| = 0$ and $|\chi - 1| \to \infty$ columns. In this particular case, for a specific value of $\arg(\chi - 1)$, the minimal value of b occurs when $|\chi - 1| = 0$ and the maximal value occurs when $|\chi - 1| \to \infty$. By explicitly determining the four bounding special cases, simple single input bounds can be placed on all the parameters in the table. An interesting aspect of this exercise is that the bounding cases often correspond to the same bounding cases discussed earlier, such as the maximal perturbation case, except that the manner in which these bounds are applied to the data is altered. Despite the elegance of this projection, some bounds – such as placing a minimum on b – which were possible in a zero centred polar layout are no longer possible. The information presented in Table 6.3 was generated by manual inspection of a numerically generated model of the inverse manifold. Formal proofs of these bounds are difficult due to the complicated nature of the derivatives of $|\chi - 1|$ and $\arg(\chi - 1)$; this is left as future work.

The first bounding case occurs when $\arg(\chi - 1) = 0$. This corresponds to perfect amplitude cancellation – in other words $\theta = \pi$, $b \in [0, 1]$. This means that

$$\chi = \frac{\Lambda_{\mathrm{f}}(b, 2\theta)|\Lambda_{\mathrm{f}}(b, \theta)|}{\Lambda_{\mathrm{f}}(b, \theta)} \tag{6.138}$$

$$= \frac{1 + b}{1 - b} \tag{6.139}$$

$$= \frac{\tau_1}{\tau_0} \tag{6.140}$$

which gives $\chi \in [1, \infty]$. No cancellation occurs in the measurement at a relative frequency of two, due to the relative phase of 2π in the measurements at that frequency – this gives $\tau_1 = 1$. The immediate consequence is that

$$\chi = \frac{\tau_1}{\tau_0} \tag{6.141}$$

$$\Rightarrow \quad |\chi - 1| + 1 = \frac{1}{\tau_0}. \tag{6.142}$$

These values are plotted in the superunity region of Figs. 6.24e and f. The bounds that can be derived in this case using the empirical relationships from Table 6.3 include:

$$\tau_0 \geq \frac{1}{1 + |\chi - 1|} \tag{6.143}$$

$$\tau_1 \leq 1 \tag{6.144}$$

$$|\hat{\eta}_0| + |\hat{\eta}_1| \leq 1 + |\chi + 1| \tag{6.145}$$

and

$$|\theta_{\mathrm{f}}| \geq 0 \tag{6.146}$$

$$|\theta| \leq \pi. \tag{6.147}$$

For example, Eqn. 6.143 is a consequence of Eqn. 6.142 given that $\arg(\chi - 1) = 0$ corresponds to the minimum value of τ_0 (from Table 6.3.) Some of the above bounds are obvious, some are inutile, but others are of great importance – particularly the first. We do not write any further bounds of the form of Eqn. 6.145 as bounds on the normalised total integrated intensity of the component returns can be inferred from bounds on τ_0.

Using the minimum bound on τ_0, we can place a maximum bound on phase perturbation at a relative frequency of one using one of the attenuation ratio relationships discussed earlier. Using Eqn. 6.136 we bound phase perturbation using $|\chi - 1|$ by

$$|\theta_{\mathrm{f}}| \leq \cos^{-1}(\tau_0) \tag{6.148}$$

$$\Rightarrow \quad |\theta_{\mathrm{f}}| \leq \cos^{-1}(\min(\tau_0)) \tag{6.149}$$

$$\leq \sec^{-1}(1 + |\chi - 1|). \tag{6.150}$$

This is one of the more useful, and natural, phase perturbation bounds. Useful in that it performs much better than most other phase bounding methods and natural in that it is based upon the distance of a value from the branch point at $\chi = 1$. Since the branch point corresponds to the special case where there is only a single component return, this can be institutively interpreted as inferring the maximum possible phase perturbation using the distance in normalised relative intensity/phase space from the aforementioned single component return case.

The second bounding case occurs when $\arg(\chi - 1) = \pi$. This corresponds to the maximal phase perturbation case which we discussed in Section 6.4.1. We can write the characteristic measurement as

$$\chi = \frac{(1 + \mathrm{e}^{2j\theta})|1 + \mathrm{e}^{j\theta}|}{1 + \mathrm{e}^{j\theta}} \tag{6.151}$$

This leads to the glaringly obvious conclusion that

$$b \leq 1, \tag{6.152}$$

but also to some less obvious ones about phase. In section 6.4.1 we inverted the max-

imal phase perturbation case using $|\chi|$ to place a bound on θ_{f}. Rewriting Eqn. 6.79 to use $|\chi - 1|$ instead of $|\chi|$ leads to

$$\cos(\theta_{\min}) = \begin{cases} \frac{(1-|\chi-1|)^2+\sqrt{(1-|\chi-1|)^4+8(1-|\chi-1|)^2}}{4}, & |\chi-1| \leq 1 \\ \frac{(1-|\chi-1|)^2-\sqrt{(1-|\chi-1|)^4+8(1-|\chi-1|)^2}}{4}, & |\chi-1| > 1 \end{cases}. \quad (6.153)$$

Using Table 6.3 we conclude that

$$\tau_0 \leq \frac{|\Lambda_{\mathrm{f}}(1, \theta_{\min})|}{2} \quad (6.154)$$

$$\tau_1 \geq \frac{|\Lambda_{\mathrm{f}}(1, 2\theta_{\min})|}{2} \quad (6.155)$$

$$|\theta| \geq \theta_{\min} \quad (6.156)$$

$$|\theta_{\mathrm{f}}| \leq \frac{\theta_{\min}}{2}. \quad (6.157)$$

We can also create a new general bound by combining Eqns. 6.154 and 6.156. Inverting the former and combining with the latter, we find

$$\tau_0 \leq \frac{|\Lambda(1, \theta_{\min})|}{2} \quad (6.158)$$

$$\leq \frac{\sqrt{2 + 2\cos(\theta_{\min})}}{2} \quad (6.159)$$

$$\Rightarrow \quad \theta \geq \cos^{-1}(2\tau_0^2 - 1). \quad (6.160)$$

This is the third bound mentioned in Section 6.4.5. Therefore, if we can measure the attenuation ratio we can place a minimum bound on the relative phase of the two component returns. In a similar manner, combining Eqns. 6.156 and 6.157 gives the bound

$$\frac{|\theta|}{2} \geq |\theta_{\mathrm{f}}|. \quad (6.161)$$

The third translated polar bounding case occurs when $|\chi - 1| = 0$. This configuration corresponds to the perfect single component return case where

$$b = 0 \vee \theta = 0. \quad (6.162)$$

As a general rule, this bounding case is not of practical use. We notate these derived bounds for the sake of comprehensiveness:

$$\tau_0 \leq 1 \quad (6.163)$$

$$|\theta| \geq 0 \quad (6.164)$$

$$|\theta_{\mathrm{f}}| \geq 0 \tag{6.165}$$

$$b \geq 0. \tag{6.166}$$

The final bounding case is when $|\chi - 1| \to \infty$; $\tau_0 = 0 \wedge \tau_1 \neq 0$ is necessary and sufficient for this to occur. From Section 6.4.2 we know that this corresponds to the case where $\arg(\chi)$ is completely dominated by the r_0 term. From Eqn. 6.102 we know that this bounding case corresponds to the circumstance given by

$$\theta \approx \pi \wedge \theta_{\mathrm{f}} = -\frac{\arg(\chi)}{2} \wedge b \approx 1. \tag{6.167}$$

Given that

$$\lim_{\gamma \to \infty} \arg(\gamma \mathrm{e}^{j\phi} - 1) = \phi, \tag{6.168}$$

we can rewrite Eqn. 6.168 using $\arg(\chi - 1)$ instead, giving the conditions at $|\chi - 1| \to \infty$ as

$$\theta \approx \pi \wedge \theta_{\mathrm{f}} = -\frac{\arg(\chi - 1)}{2} \wedge b \approx 1. \tag{6.169}$$

Using Table 6.3, we find the following generally applicable bounds:

$$\tau_0 \geq 0 \tag{6.170}$$

$$|\theta| \leq \pi \tag{6.171}$$

$$|\theta_{\mathrm{f}}| \leq \frac{\arg(\chi - 1)}{2} \wedge \mathrm{sgn}(\theta_{\mathrm{f}}) = -\mathrm{sgn}(\arg(\chi - 1)) \tag{6.172}$$

$$b \leq 1. \tag{6.173}$$

The most important bounds are summarised in Figs. 6.27 and 6.28.

As a final bound we present a fully empirically derived lower bound on total integrated intensity. For any particular value of $|\chi - 1|$, the minimum bounding case is non-trivial and difficult to determine. However, by analysing numerically generated minimum values it is possible to put an approximate lower bound on total integrated intensity. Fig. 6.29 plots the normalised total integrated intensity divided by $|\chi - 1|$ for this minimum case. This value is always between (approximately) 0.65 and 1. From this, a very rough bound could be written as

$$|\hat{\eta}_0| + |\hat{\eta}_1| \geq 0.65|\chi - 1|, \tag{6.174}$$

however, this bound is relatively loose. By fitting a rational function with cubic numerator and denominator to the curve, it is possible to very closely match the

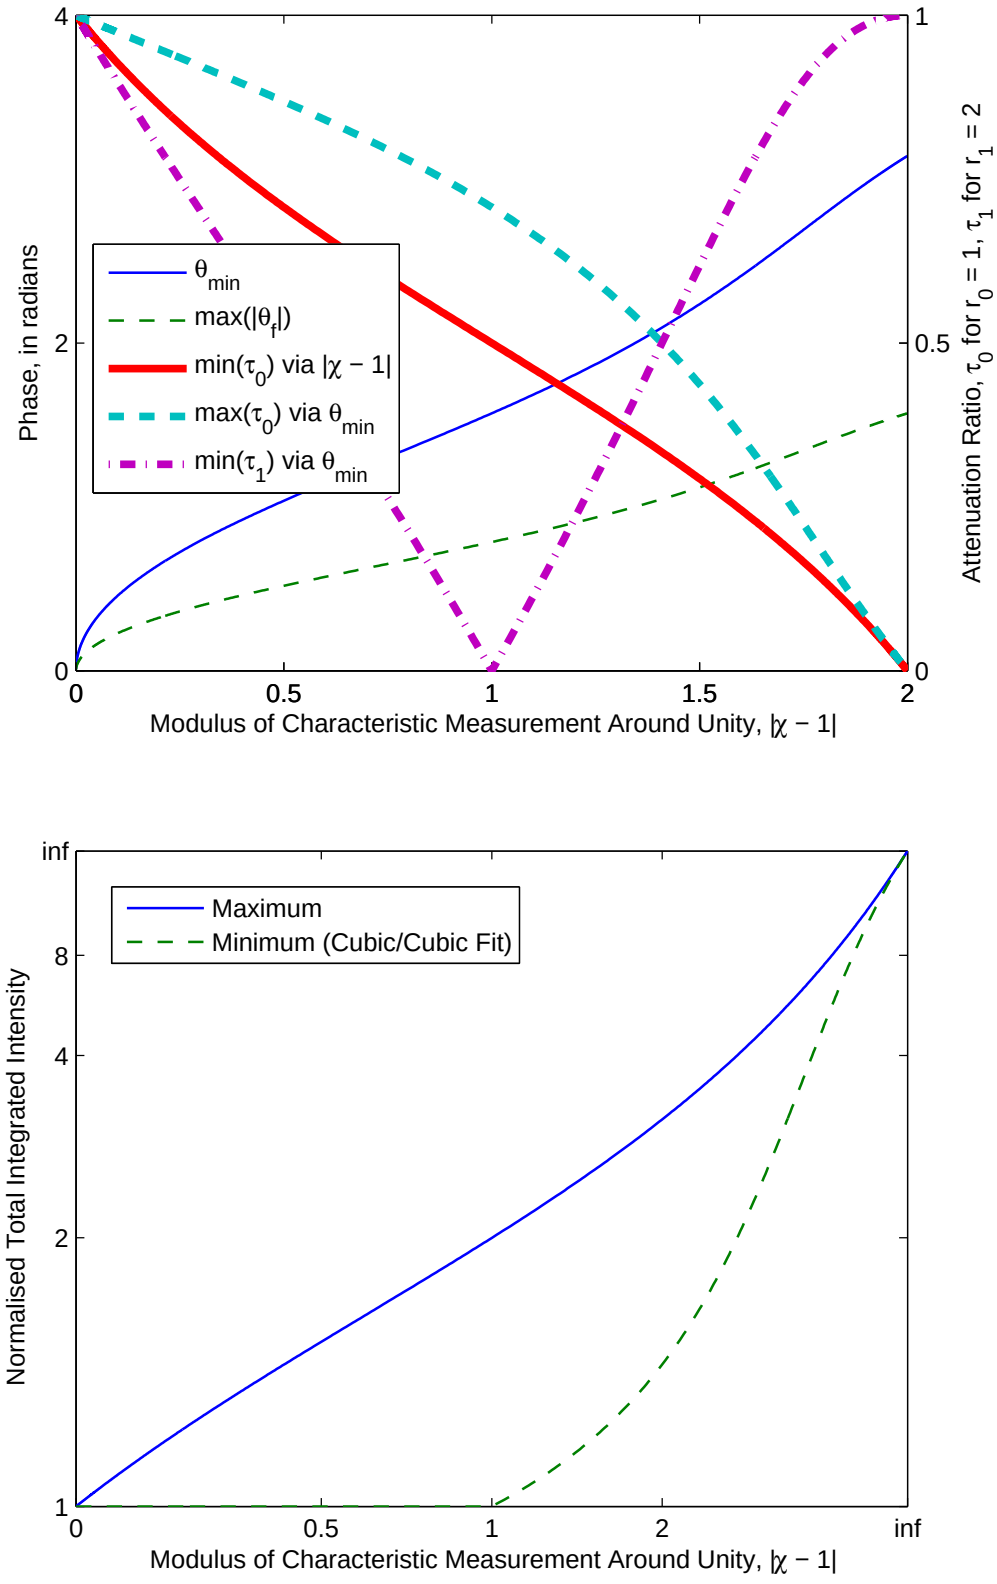

Figure 6.27: Bounds on phase, attenuation ratios and normalised total integrated intensity using the modulus of the characteristic measurement around unity at a frequency ratio of 2 : 1.

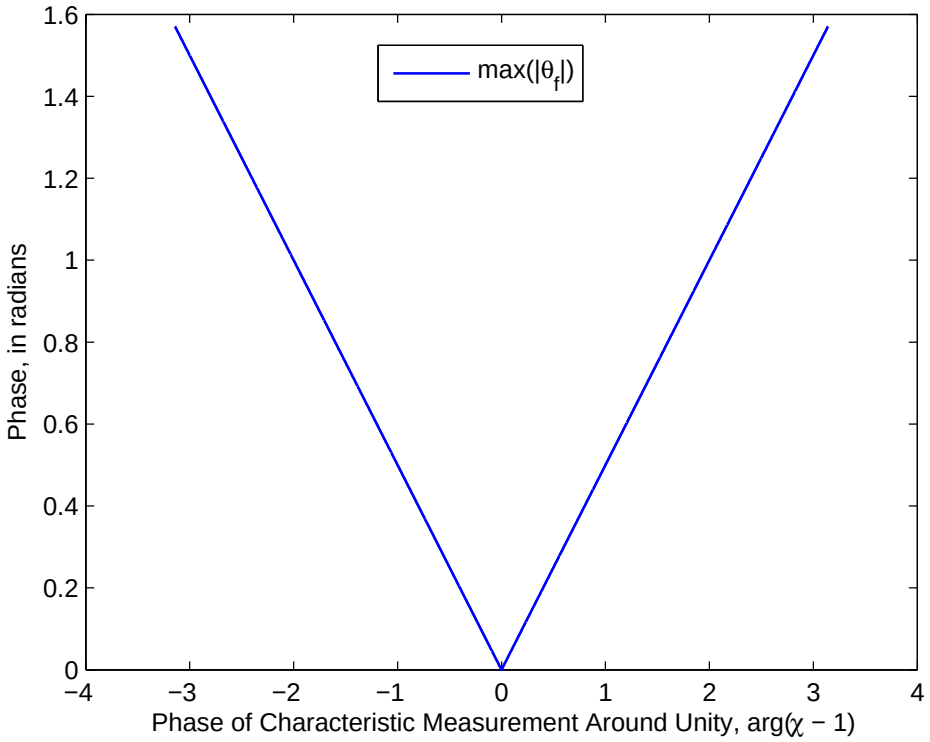

(a) Phase of Characteristic Measurement Around Unity

Figure 6.28: Bound on phase perturbation using the characteristic measurement assuming a frequency ratio of 2 : 1.

raw data in Fig. 6.29. The bound arrived at for $|\chi - 1| \geq 1$ is

$$|\hat{\eta}_0| + |\hat{\eta}_1| \geq |\chi - 1| \frac{0.9979|\chi - 1|^3 - 1.55|\chi - 1|^2 + 6.25|\chi - 1| + 0.7143}{|\chi - 1|^3 + 0.3301|\chi - 1|^2 + 6.811|\chi - 1| - 1.731}, \quad (6.175)$$

for $|\chi - 1|$ we assume $|\hat{\eta}_0| + |\hat{\eta}_1| > 1$.

6.4.7 Using the Bounding Methods

In this section we developed a large number of bounds on properties of the component returns from measurements of the characteristic measurement or attenuation ratio. While these bounds have intellectual interest, they are also practically valuable; in general, the most valuable bounds are phase perturbation maxima, i.e. a maximum value for $|\theta_f|$ at some frequency r_0. If we can assume that the brightest component return is the 'correct' return (for example, if there is slight intra-camera scattering induced multipath interference), then we can state hard limits on the possible true range values for each measurement because the phase of the brightest return is $\arg(\xi_0) \pm \max(|\theta_f|)$. Use of the other bounds is more application specific:

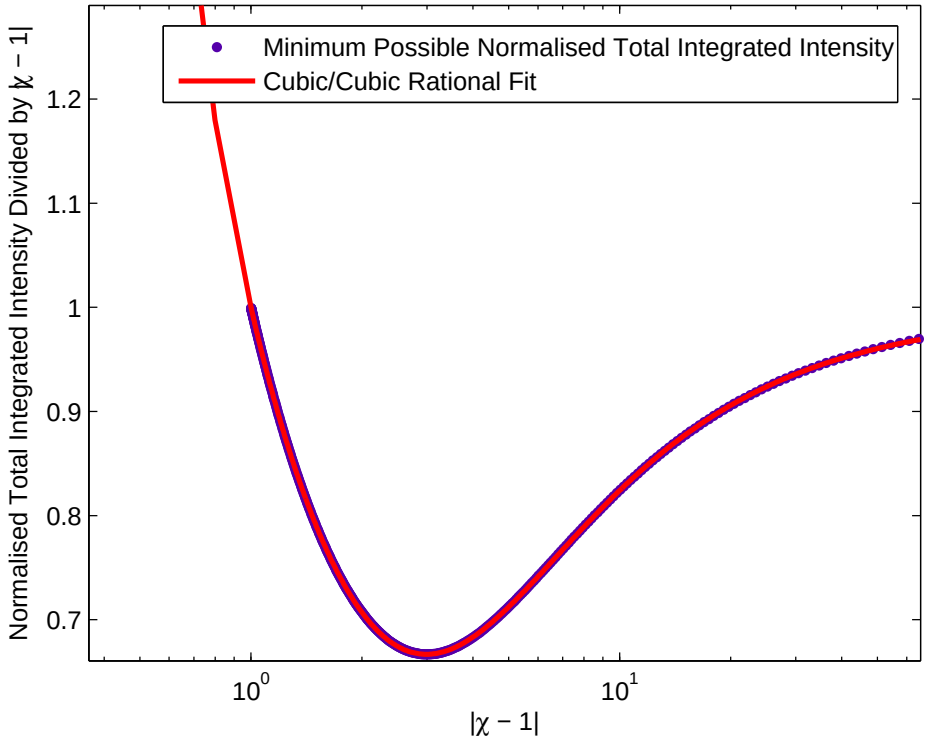

Figure 6.29: Cubic/cubic rational empiricial minimum bound on normalised total integrated intensity, $\hat{\kappa}$

for example, it might be useful to know what the phase and intensity relationship between the component returns is. While a minimum on relative phase is not particularly useful for detection of mixed pixels, estimates of $\min(b)$ are useful indicators of the level of mixing; one simple approach to detection of mixed pixels is to place a threshold on $\min(b)$, above which the pixel is assumed to be mixed. Not all bounds are useful in practice; as we determine in the next chapter, some bounding methods produce bounds which are so loose as to be useful. Some bounds have little direct practical value; it appears that bounds on attenuation ratios are too esoteric to be useful for anything other than the development of other bounds (on the other hand, the derived bounds are quite useful.)

6.5 Summary

This chapter introduced the new concept of the characteristic measurement, utilising phase and amplitude normalisation of measurements in order to simplify the mixed pixel/multipath interference problem. In addition to allowing the dimensionality of the inverse to be reduced, allowing easier implementation, the characteristic mea-

surement also makes it easier to reason about the properties of phase perturbation, allowing the derivation of new bounding methods. Subsequently, we developed a library of bounding methods and a number of general techniques for determining bounding methods for the characteristic measurement at any given frequency ratio. One challenging piece of future work is to extend the bounding methods developed for the two-to-one frequency ratio based around the branch point at $\chi = 1$ to other frequency ratios.

In the next chapter we develop a general approach to numerically testing these methods using Monte-Carlo simulations and then test the methods on real and simulated data in order to demonstrate their operability.

Chapter 7

Characterising the Performance of Characteristic Measurement Based Approaches

In Chapters 5 and 6 we developed a series of methods for determining the phase and intensity of component returns within a pixel. In addition, Chapter 6 developed a series of methods for bounding characteristics like phase perturbation, relative intensity and phase in closed-form using the characteristic measurement. In this chapter the methods of Chapters 5 and 6 are evaluated.

First off, the assumptions made in the simulations are outlined, then the characteristic measurement lookup table method is evaluated. This includes a comparison to attenuation ratio based inversion, a series of experiments on real-data and systematic evaluation of the bounding methods from Chapter 6.

7.1 Modelling Component Returns

In order to determine the performance of various multiple component return separation algorithms it is necessary to make assumptions about the relative intensity and phase of the component returns that are generated. Here the assumptions and noise statistics utilised in the simulations are outlined.

There are many models for natural images. Previous models for images have used priors such as Gaussian derivative priors (Levin *et al.*, 2007), long-tailed priors (Ruderman and Bialek, 1994) and priors over wavelet decomposition coefficients (Wainwright and Simoncelli, 2000). Other work has analysed the spatial frequency distribution, often considered to be proportional to $1/f^2$, where f is the spatial frequency (Van Der Schaaf and Van Hateren, 1996). Because we are only interested

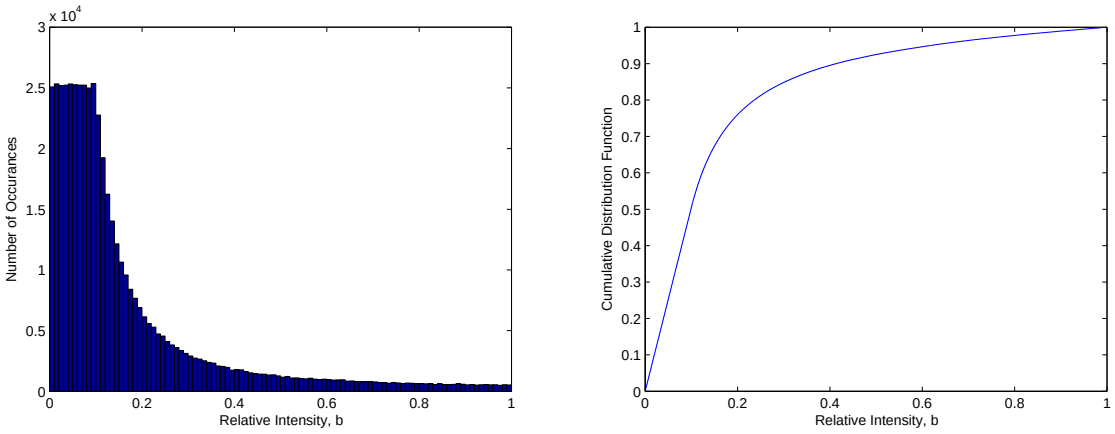

(a) Histogram (b) Cumulative Distribution Function

Figure 7.1: The prior distribution assumed for relative intensity (similar to the uniform ratio distribution)

in a single pixel at a time, these previous analyses are of less value. Whereas natural image models are typically interested in the marginal distribution of pixel intensities and the relationships between the intensities of adjacent pixels, we are interested in the relationships between the intensities of different components within a pixel. These components within a pixel may be from any number of sources, including intra-scene reflections (multipath interference), intra-camera scattering and mixed pixels around the edges of objects. Each of these situations can be modelled separately; for example, a model for mixed pixels could potentially include effects such as occlusion. However, with added complexity comes more potential for fundamental mistakes.

The simplest possible assumptions for the distribution of intensity values are the Gaussian and uniform distributions, the choice being arbitrary. Given this, let the distribution of relative intensity, b be given by

$$a_0 \sim \mathcal{U}(0,1) \tag{7.1}$$

$$a_1 \sim \mathcal{U}(0,b_{\mathrm{M}}) \tag{7.2}$$

$$b = \frac{\min(a_0,a_1)}{\max(a_0,a_1)}, \tag{7.3}$$

where a_0 and a_1 are the intensities of the component returns and b_{M} is approximately the median of b. As $b_{\mathrm{M}} \to 0$, the distribution of b asymptotically approaches a scaled uniform ratio distribution. For our experiments we arbitrarily chose a value of $b_{\mathrm{M}} = 0.1$. In practice, this is a value which varies with the physical characteristics of any one scene. Fig. 7.1a shows a histogram of the distribution of b in this one

particular case. Alternative approaches, which were decided against, include: to use a Gaussian ratio distribution, with the modulus of each return being sampled from a Gaussian distribution; or a Rayleigh ratio distribution, with each component return being sampled from a zero centred, symmetric complex Gaussian distribution.

While it is possible to plot relative intensity using a normal histogram, most of the parameters used in this section are long-tailed; for example, the phase error in separated component returns. Conventional histograms are not suited to data with a very large dynamic range, which necessitates a logarithmic scale. In order to plot these data, we utilise plots of the Cumulative Distribution Function (CDF) – for example, Fig. 7.1 gives the CDF of b.

Assumptions about relative phase are even more challenging to rationalise than assumptions about relative intensity, partly because it is so particularly challenging to define precisely what a component return is in real data. If a single pixel is integrating over a large region and the target is a sloped board, do we always define it as a single component return? What if we have a continuous surface with concavities and holes in it? In order to avoid these questions, we make the assumption of a uniform prior for relative phase. That is,

$$\theta \sim \mathcal{U}(-\pi, \pi). \tag{7.4}$$

Chapter 3 provides a lengthy mathematical description of the formation process of component returns. There are multiple noise sources including both zero centred Gaussian and Poisson components. One approach to the noise simulation problem would be to model everything at the lowest level possible, including factors such as the offset from ambient light, modulation depth and the relationship between the measured brightness of ambient light and the drop-off in the brightness of the active illumination. As layer upon layer of detail is added, there is a loss of clarity in the results; it is necessary to indulge in some simplification to simplify reasoning and evaluation, otherwise it can be a struggle to draw precise conclusions.

An extremely simple noise model is utilised for our simulations: ignoring ambient light and modulation depth, and assuming a fixed SNR for each simulation, where the SNR is defined in terms of the total signal power over all the component returns we are measuring. In other words

$$\sigma^2 = \frac{a_0^2 + a_1^2}{SNR} \tag{7.5}$$

$$\epsilon_x \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2), \tag{7.6}$$

where σ^2 is the variance of the simulated measurement noise, SNR is the signal to

noise ratio, ϵ_x is the noise applied to the xth complex domain range measurement. The latter equation can be interpreted in terms of the distributions of the real and imaginary components of ϵ_x, giving

$$\Re(\epsilon_x) \sim \mathcal{N}\left(0, \frac{1}{2}\sigma^2\right) \tag{7.7}$$

$$\Im(\epsilon_x) \sim \mathcal{N}\left(0, \frac{1}{2}\sigma^2\right). \tag{7.8}$$

7.2 Inversion Using the Characteristic Measurement

Chapter 6 developed the concept of the characteristic measurement. In particular, Section 6.2.2 showed that using the characteristic measurement it is possible to reduce the dimensionality of the two measurement, two component return problem to a two dimensional inverse problem, rather than a four dimensional inverse problem. A number of different bounding methods were developed; no closed-form inverse was developed. However, because of the useful nature of the two measurement, two component problem, it is necessary to develop a practicable solution method.

One approach to the inverse problem is to apply a numerical optimisation method, such as the Newton-Raphson or Levenberg-Marquardt algorithms (Press *et al.*, 2007), in order to find a solution. The primary problem with using numerical methods to solve a problem like this is that the computational complexity is high, resulting in poor performance; in practice, so poor that it is only possible to apply to data offline. In order for a method to be practical, it is necessary for that method to be usable in real-time. This requires a quickly evaluable model of the inverse function.

7.2.1 Approaches to Modelling the Inverse Function

There are a number of plausible approaches to modelling the inverse including explicit lookup tables and mathematical approximations including polynomial models. Explicit lookup tables are implemented by sampling the output of a function across a set of selected values from the domain of the function. Whenever an input value is given, the values in the lookup table that are closest to the input value are used to calculate an estimate of the correct output value. Generally either interpolation or a nearest neighbour approach is used to calculate the output value. Common interpolation techniques include linear intepolation, polynomial interpolation, cubic

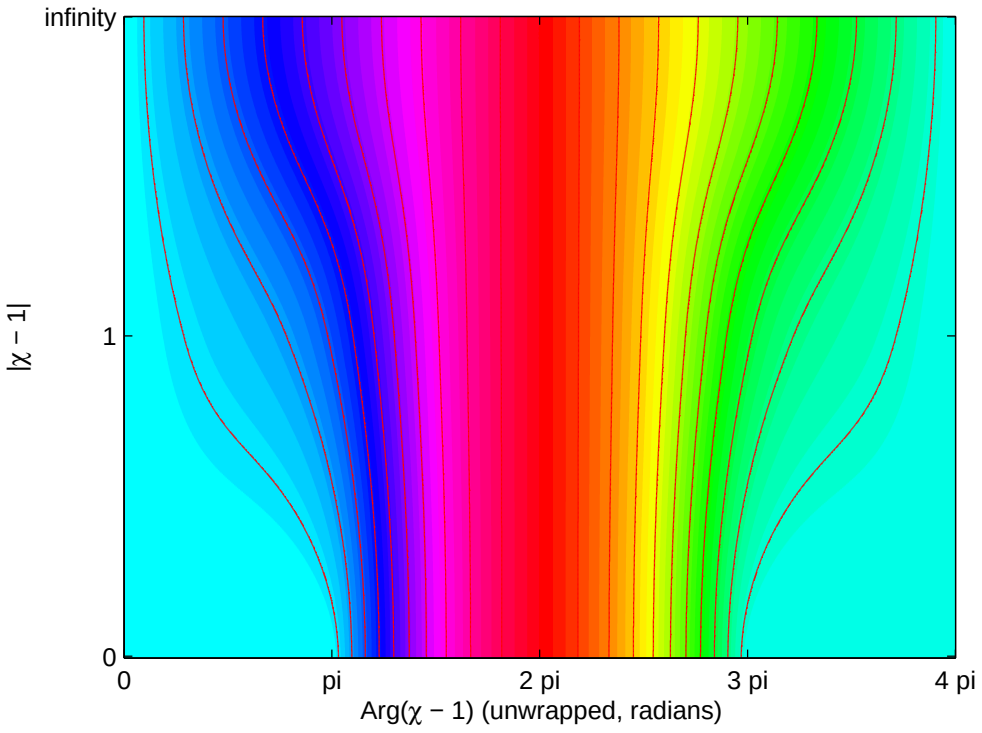

(a) Normalised Phase

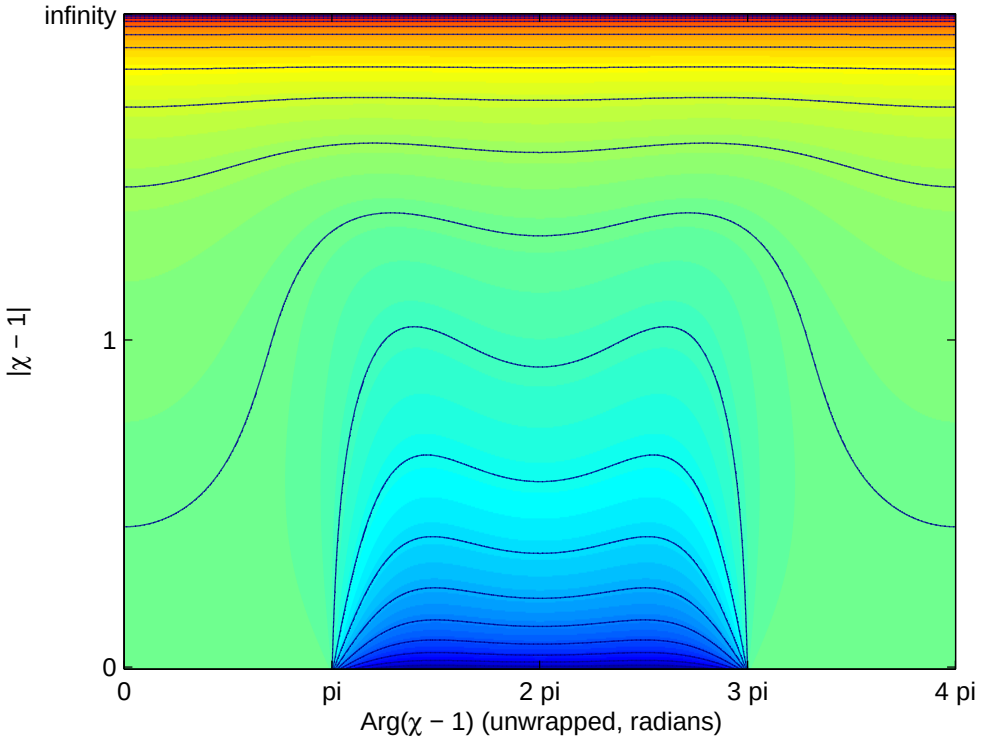

(b) Normalised Intensity

Figure 7.2: The lookup table for the inverse two measurement, two component return problem at a relative frequency of 2 : 1 using a translated polar mapping

splines and radial basis functions (Press *et al.*, 2007).

Experiments with modelling the behaviour of the inverse function using polynomials and transcendental functions were generally unsuccessful; not because of the impossibility, but rather the lack of any elegant model for the inverse. It was necessary for models to be of fairly high order and – from an aesthetic point of view – rather ugly in order to provide high quality results. As a result this approach has been discarded as unsuitable, instead utilising a lookup table.

Modelling the inverse function using a lookup table can be quite a complicated endeavour due to the nature of the connectivity of the inverse function manifold. Fig. 6.11 showed the inverse function in the case of a frequency ratio of $3:2$; if utilising a Cartesian mapping of the same form used in that figure, then it is necessary to incorporate a connectivity table in order to interpolate between points. For example, if we have a point right on the edge of a branch cut – how do we know from which solution to take the points with which to calculate the interpolated value? This requires knowledge of which points are adjacent to each other on the solution manifold, hence a connectivity table. The problem is that, especially around the r_0th roots of unity, the shape of the table can make interpolation difficult or meaningless. While not a generally solution to this problem, for ratios of $2:1$ and $3:1$ using a translated polar mapping allows the entire table, including both component returns, to be unwrapped onto a single continuous sheet.

Section 6.4.6 discussed the use of $|\chi - 1|$ and $\arg(\chi - 1)$ for bounding various properties of the underlying component returns. Fig. 6.25 provides graphs of these properties. An obvious extension of this method is to take the $2:1$ inverse manifold from Fig. 6.10 and unwrap around $\chi = 1$ in a similar manner. Whereas the bounding methods modelled the phase around $\chi = 1$ as $\arg(\chi - 1) \in [0, 2\pi]$, we now extend this representation. If we encode both the first and second component returns on the same manifold, we can use $\arg(\chi - 1) + 2\pi m \in [0, 4\pi]$, where $m \in [0, 1]$ is the component return number, as the first dimension of the lookup table. An example lookup table is given in Fig. 7.2; this is best illustrated by comparing it directly to the $2:1$ LUT in Fig. 6.10. There are some significant benefits to this lookup table layout over a simple Cartesian layout, the primary one being the lack of necessity for any connectivity table and the general smoothness of the table. This smoothness property allows computationally simple bilinear interpolation to provide sufficient performance for a practical implementation.

Our implementation uses uniform sampling of $\arg(\chi - 1) + 2\pi m$ over the entire range from 0 to 4π and a special compressed projection of $|\chi - 1|$; this is different from the projection utilised in Chapter 6 for Cartesian models of infinite manifolds

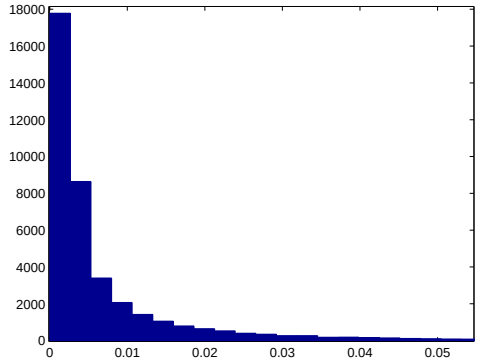

(a) Phase Error – Higher Resolution, High Phase Error Values Truncated

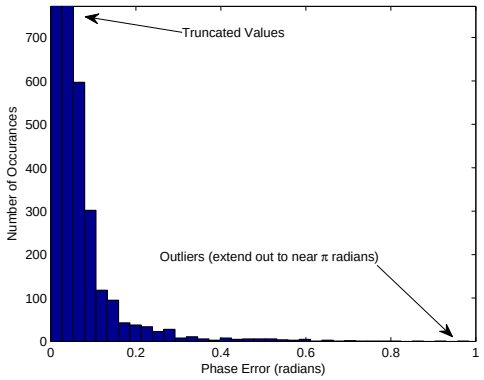

(b) Phase Error – Lower Resolution, High Occurances Truncated

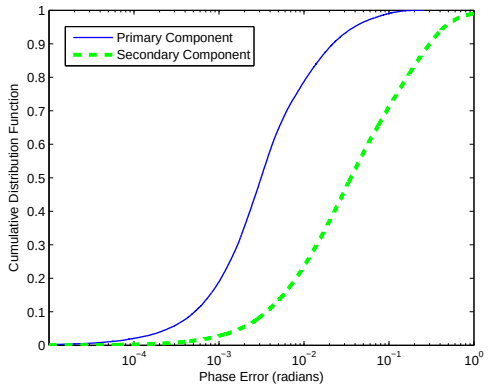

(c) CDF of Phase Error

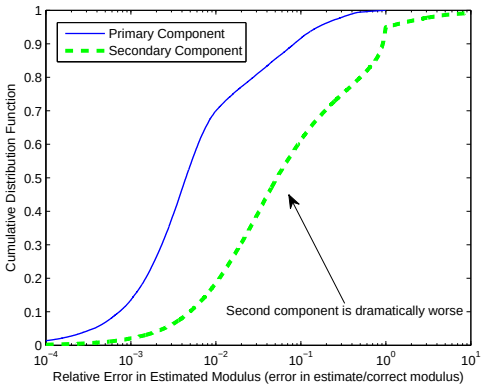

(d) CDF of Modulus Error

Figure 7.3: Histograms of phase error in primary return after processing at a frequency ratio of 2 : 1 and CDFs of phase error and relative modulus error for both the primary and secondary component returns. A Monte Carlo simulation of 500000 runs at a SNR of 25000 : 1.

and plotting graphs from zero to infinity. The projection operation given in Eqn. 6.5 tends to sample particularly large values in such a manner that the values cannot be adequately interpolated because adjacent values increase in magnitude too rapidly near infinity. This is ameliorated by defining an alternative projection operation, which is utilised for the $|\chi - 1|$ dimension of the lookup table. This projection operation is given by

$$p(w) \quad = \quad \begin{cases} 2 - \frac{1}{\sqrt{w}} & w \geq 1 \\ \sqrt{w} & w < 1 \end{cases} \tag{7.9}$$

where $p(|\chi - 1|) \in [0, 2]$ is the corresponding location on the second lookup table axis. As a result, this function increases the sampling near the high and low ex-

tremes of the $|\chi - 1|$ dimension, relative to Eqn. 6.5. The $p(|\chi - 1|)$ dimension uses equally spaced samples and the underlying normalised component returns are calculated by bilinear interpolation of complex domain values. The interpolated values are then denormalised using Eqns. 6.14 and 6.15 to determine the true component returns. For the result presented in this thesis, the lookup tables were produced using numerical least squares methods and the assemblage of the component returns into a single continuous manifold was partly achieved through manual intervention; in particular there were certain regions near $|\chi - 1| \to 0$ which were particularly difficult to optimise, due to high sensitivity to the numerical precision of 8-byte ('double' precision) floating point values and optimisation termination conditions.

7.2.2 The Impact of Frequency Ratio/Utilisation

Section 7.2.1 developed a translated polar lookup table method specific to frequency ratios of 2 : 1 and 3 : 1. However, it is possible to implement other frequency ratios. For simulated data and offline processing it is possible to use direct optimisation of a dataset to determine the component returns. This section analyses the relative performance of different relative frequencies in order to understand which frequency ratios give the best results and why.

The first case analysed is when the frequency ratio is 2 : 1. The histograms given in Fig. 7.3a and 7.3b were produced using data simulated by the method of Section 7.1. Fig. 7.3a truncates high phase error values and Fig. 7.3b truncates high occurances in order to display as much of the curve as possible. These graphs show how extreme the dynamic range of phase error is for the inverse problem; the distribution of phase error is extremely skewed and almost decays in an exponential manner. As discussed above in Section 7.1, this provides the rationale behind the use of Cumulative Distribution Functions to display the nature of the error distribution. Related to this behaviour, common metrics like RMS error implicitly assume the existence of a Gaussian distribution; if the error is not Gaussian distributed, then the outliers are far more heavily weighted than they ought to be. This has the potential to give error statistics that are unrepresentative and seriously overestimate the error in the average instance.

Fig. 7.3c gives the CDF for phase error for both the primary and secondary component returns. The phase error in the secondary component is at least an order of magnitude worse than in the primary component return. Also, even with quite a low noise level (SNR of 25000 : 1), approximately 10% of values have greater than 10 milliradians of phase error. This is due to the unique nature of the problem, in particular contributed to by the phase relationship between component returns;

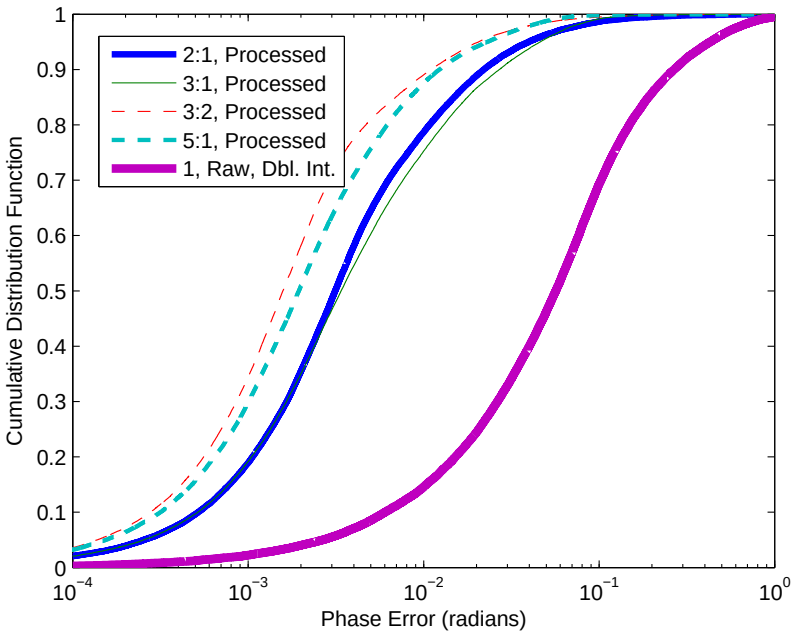

(a) SNR = 25000 : 1

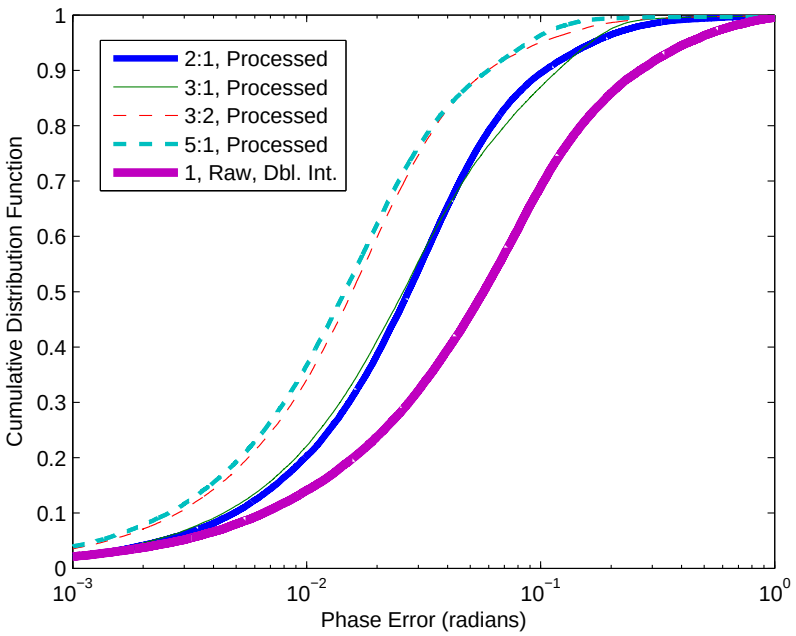

(b) SNR = 200 : 1

Figure 7.4: CDFs of overall phase error in the separated primary return at different frequency ratios using optimisation assuming total integration time

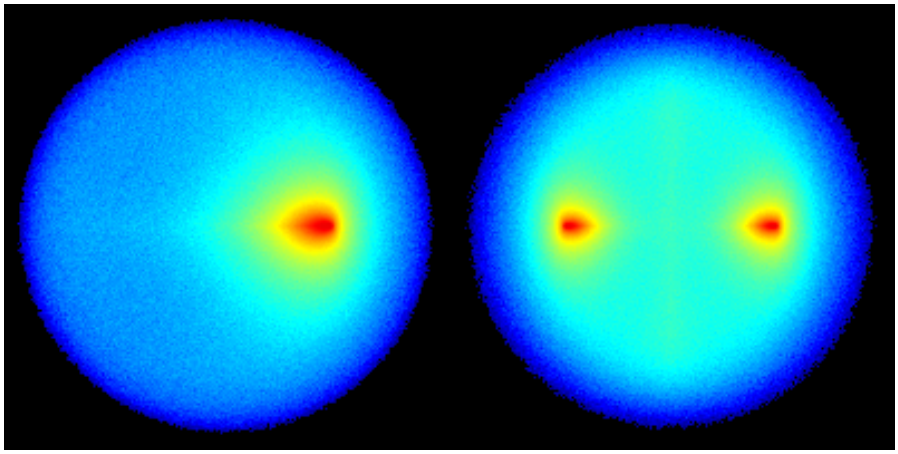

(a) 2:1 (b) 3:2

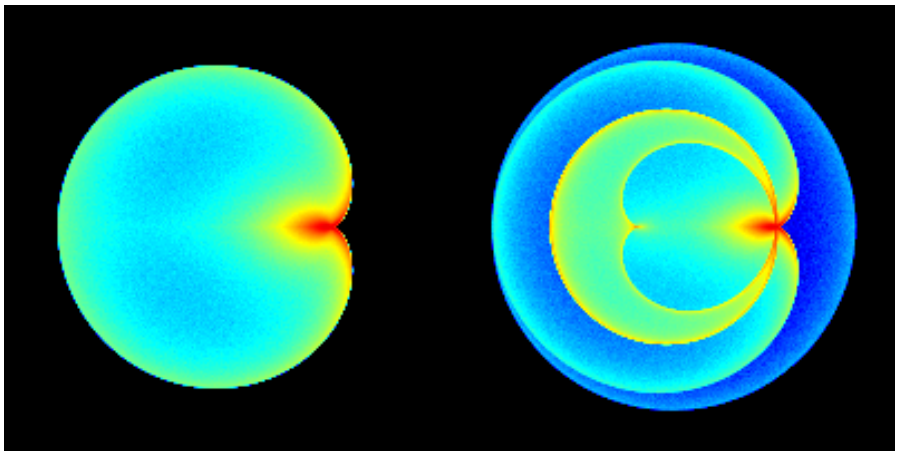

(c) 3:1 (d) 5:1

Figure 7.5: Lookup table utilisation versus frequency ratio. Red correponds to high utilisation, dark blue to low utilisation.

this is described in greater detail in Section 7.2.3.

Assuming a constant total integration time, and ignoring readout times and other temporal inefficiencies, we simulated the formation of measurements at different frequency ratios to give the CDF curves in Figs. 7.4a and 7.4b. In order to provide a fair reference, measurements were also simulated at a relative frequency of unity. If we give the total noise power corresponding to the error in a single measurement of a two frequency capture as σ^2, then we can write the total noise power in the unity relative frequency measurement as $2\sigma^2$. In Figs. 7.4a and 7.4b, this single frequency reference is notated as '1, Raw, Dbl. Int.' in order to conserve figure space.

Fig. 7.4a was simulated using an assumption of a signal to noise ratio of 25000 : 1. The reference curve is clearly far to the right (greater noise levels) of the processed data at any relative frequency, meaning a substantial improvement in noise levels of

at least an order of magnitude. Comparing the results at different frequency ratios, we can see that 2 : 1 and 3 : 1 give roughly equivalent performance, but with a very slightly different shape to the CDF curves; this is as a direct result of the differing shapes of the inverse manifolds.

One particularly useful way to understand the precision implications of the inverse manifold is to plot the lookup table utilisation density. This can be considered to be counting how many times each LUT index value is accessed. Plots of the utilisation density in the 2 : 1 and 3 : 1 are given in Figs. 7.5a and 7.5c – the scale is logarithmic from dark blue to red, where black corresponds to the lack of a valid solution. Given the assumption of $b_{\mathrm{M}} = 0.1$, there are a large number of points clustered around the r_0th roots of unity that correspond to near-single component returns. Regarding the plot for 2 : 1, there is only one feature of interest: the high usage region near unity. Likewise for 3 : 2, where there are high utilisation regions around $\sqrt{1} = \pm 1$. We now address the more complicated and interesting cases of 3 : 1 and 5 : 1.

Fig. 7.5c shows a utilisation plot for a frequency ratio of 3 : 1. In addition to the clustering around unity, there is an increase in utilisation around the border of the valid region of the inverse. This border corresponds to the asymptote discussed in Section 6.3.5.

If the input component return values for the forward problem (Eqn. 6.43) were uniformly distributed, one could draw conclusions about the derivatives of the forward problem relative to component return tuples, (η_0, η_1). A high utilisation density for a particular region of the inverse problem LUT means that the modulus of the derivatives in this region is relatively low. If the modulus of the derivatives in a region is low, a small difference in the value of the characteristic measurement may correspond to a large change in the component return tuple. On the other hand, if there is a low utilisation density and there is a relatively uniform prior, it is likely that at least some of the derivatives have a large magnitude – in other words, for a given difference in the characteristic measurement, there is a greater change in the corresponding component return tuple than in a region with high utilisation. Given the naïve assumption that the error distribution in measurements of χ is constant across all possible values of χ, we can infer that regions of high utilisation result tend to correspond to high noise sensitivity in the inverse function and regions of low utilisation correspond to low noise sensitivity in the inverse function. Thus the shape of the inverse manifold can be seen to result in variation in the shape of the overall phase error CDF for different frequency ratios.

The most extreme utilisation table plotted herein is Fig. 7.5d, which corresponds

to a frequency ratio of $5:1$. As relative phase and intensity changes, the inverse manifold folds back over itself, resulting in either one, two or four solutions for any particular value of the characteristic measurement. In addition to the high usage near $\chi = 1$, there are a series of curved features near the unit circle; these correspond to the folds present in the manifold. Near folds, the magnitudes of the derivatives of χ with respect to b and θ decrease, this appears to result in an increase in utilisation and a corresponding decrease in relative phase and amplitude precision in the case of the practical inverse problem. In addition to the multiple solution problem, this is a disincentive to using $5:1$ and other similar ratios (in particular, any ratios where both frequencies are either odd, or both frequencies are even.)

We now attempt to ameliorate a major systematic bias in Fig. 7.4. As discussed in more depth in Section 7.2.6 below, there is no guaranteed manner in which to determine the correct component return. In order to avoid testing the phase error due to the algorithm to determine the correct solution, rather than inherent to a particular frequency ratio, we utilised the naïve approach of assuming that the choice of solution was given by the solution that minimised the measured overall phase error versus the ideal, unprocessed data from which the simulated range measurements at different relative frequencies were generated. This results in a decrease in the estimated error which may not be truely representative of performance at that particular relative frequency, but there does not appear to be a viable alternative approach available.

We can illustrate the impact that this approach could potentially have on phase error by analysing the value given by

$$Q_m \quad = \quad \mathbb{E}[\min(X_i) \,|\, Y_i \sim \mathcal{N}(0, \mathcal{I}_m) \wedge \forall_{k \in [0,m-1]} X_i[k] = |Y_i[k]| \,]. \qquad (7.10)$$

The value Q_m corresponds to the expected output of a function that finds the element with the minimum absolute value from a set of m samples sampled from the same unity variance, zero centred Gaussian distribution. From the point of view of the multiple solution problem, this is equivalent to making the assumption that the phase error for any particular solution is given by the absolute value of a sample from a unity variance, zero centred Gaussian distribution. While this is related to the extreme value distribution, it is not the same because the extreme value distribution deals with the probability distribution of the maxima of a Gaussian distribution, as opposed to the value closest to zero. Fig. 7.6 shows how the distribution of $\min(X_i)$ changes with the value of m; as m increases, the distribution becomes progressively more long-tailed.

The relative impact on estimated phase error of choosing the best of m solutions,

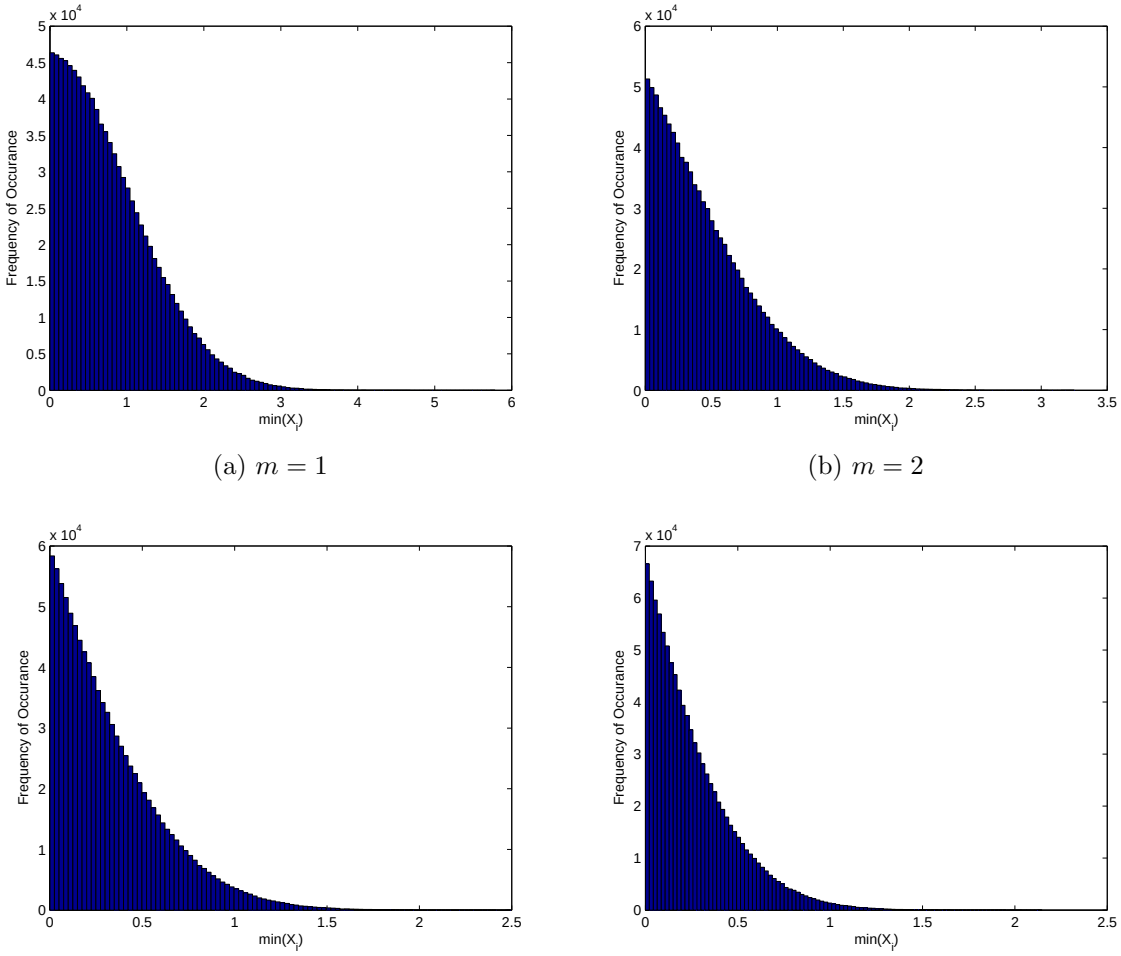

(a) $m = 1$ (b) $m = 2$

(c) $m = 3$ (d) $m = 4$

Figure 7.6: The distribution of $\min(X_i)$ for various values of m.

as opposed to only having a single possible solution can be expressed as an overall phase error scaling factor:

$$c_\epsilon = \frac{Q_m}{Q_1}. \tag{7.11}$$

Table 7.1 gives values of Q_m and c_ϵ for not only a Gaussian distribution, but also using the assumptions of uniform or exponential distributions for Y. This allows us to compare how the minimum scales with m and also to what extent the type of distribution influences this scaling. The table shows that distributions with greater kurtosis have the lowest values of c_ϵ and distributions with less kurtosis have the highest values of c_ϵ. Fig. 7.7 shows the distribution of the expected phase error if the correct solution is unknown and chosen randomly based upon the expected probability of each solution as a function of the characteristic measurement assuming

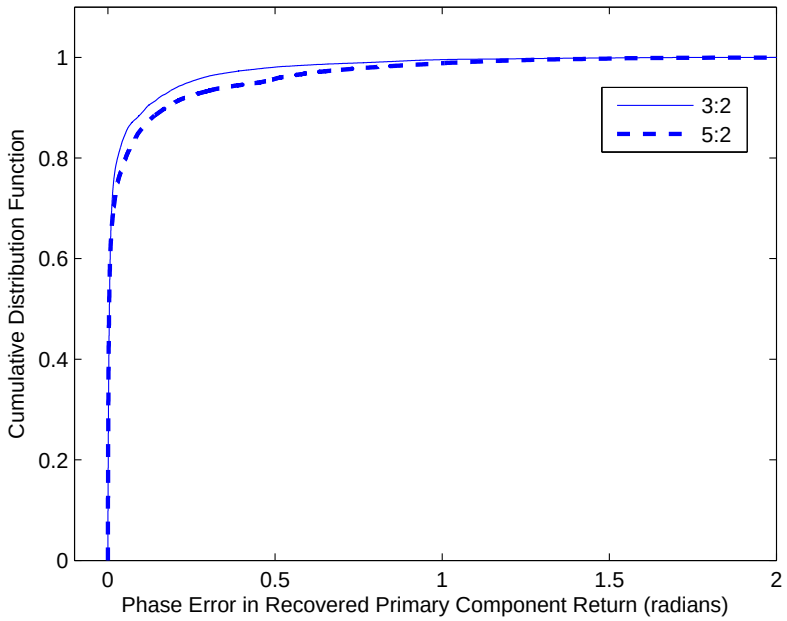

Figure 7.7: CDFs of phase error when the solution is chosen randomly using the marginal probabilities associated with the prior distribution from Section 7.1 for each value of the characteristic measurement, via Monte-Carlo simulation

the prior from Section 7.1. This was achieved using a Monte-Carlo simulation. The resulting error distribution for a frequency ratio of $3:2$ is highly leptokurtic and relatively well modelled by an exponential distribution. Note that 50% of the values are correct, therefore have an error of zero, so it is only appropriate to model values above this point as exponential – if anything, this makes the distribution even more kurtotic. The distribution for $5:2$, while similar, suffers from some slight discontinuities related to varying number of solutions across the inverse manifold.

In Fig. 7.4 we plotted the CDF curves for different frequency ratios. Since a frequency ratio of $3:2$ results in three solutions we can apply an extremely naïve correction factor by dividing phase error by c_ϵ for different distributions at $m = 3$. The results are plotted in Fig. 7.8. While this does not constitute absolute proof, the figure indicates that there is a high level of probability that a frequency ratio of $3:2$ does not result in significantly better phase error than $2:1$. Given that $3:2$ requires an algorithm to determine which solution is correct, frequency ratios of $2:1$ and $3:1$ appear to be obviously superior choices.

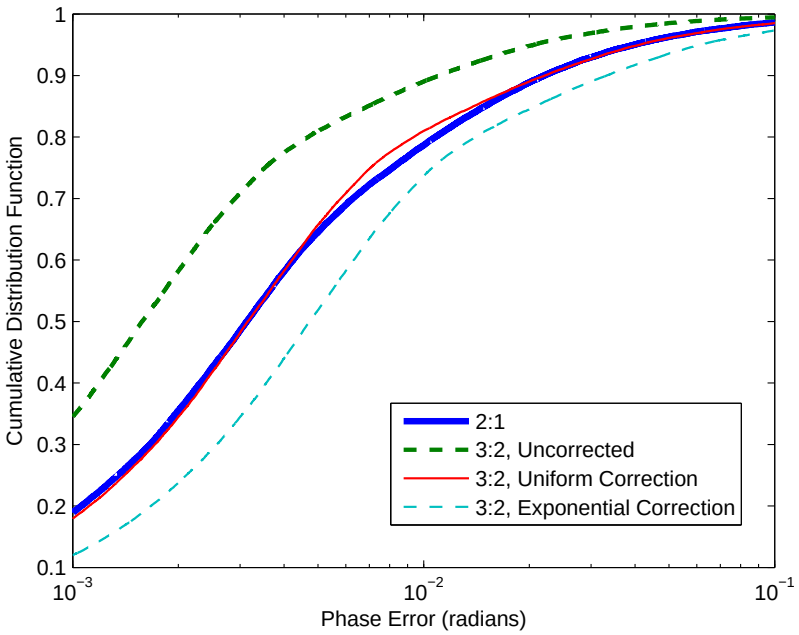

Figure 7.8: CDFs of phase error at a SNR of 25000 : 1, using correction factors c_ϵ from Table 7.1. This graph uses the same underlying data as Fig. 7.4.

Table 7.1: Impact of choosing the minimum absolute value from a set of m samples from a zero centred, unity variance, independent Gaussian distribution versus a uniform distribution over $[0, 1]$ or an exponential distribution with a rate parameter $\lambda = 1$.

	Gaussian		Uniform		Exponential	
m	Q_m	c_ϵ	Q_m	c_ϵ	Q_m	c_ϵ
1	0.798	1	1/2	1	1	1
2	0.467	0.586	1/3	2/3	1/2	1/2
3	0.335	0.420	1/4	1/2	1/3	1/3
4	0.262	0.328	1/5	2/5	1/4	1/4

7.2.3 The Impact of Relative Intensity and Phase

We now address the impact that the underlying relative phase, θ and relative intensity, b, have on precision, accuracy and overall phase error. Using fixed values of SNR and b, Figs. 7.9, 7.10 and 7.11 are plots of systematic error, precision and overall RMS error as a function of θ.

Analysing the systematic error in the unprocessed, single measurement case case in Fig. 7.9 we can see that up to a certain point, the greater the relative phase, the greater the systematic error – the perturbation having the same sign as the relative phase of the smaller component return. As relative intensity increases, the systematic phase perturbation increases and SNR has no discernible impact. There

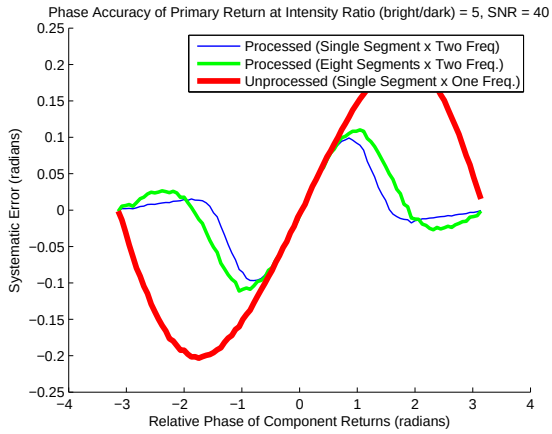

(a) $b^{-1} = 5$, SNR $= 40$

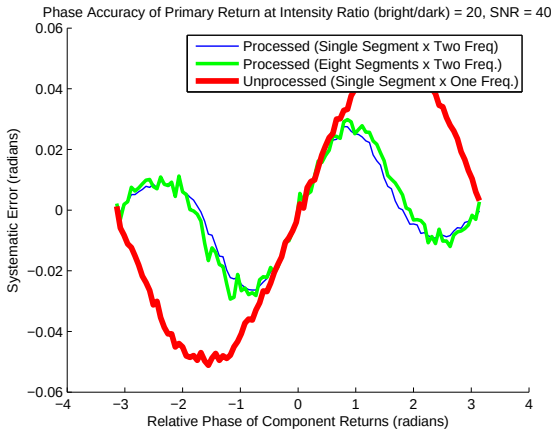

(b) $b^{-1} = 20$, SNR $= 40$

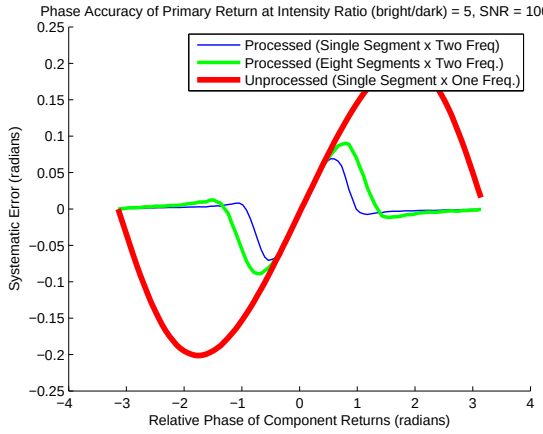

(c) $b^{-1} = 5$, SNR $= 1000$

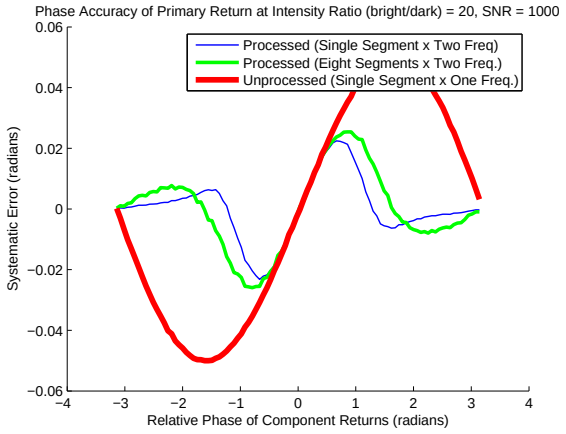

(d) $b^{-1} = 20$, SNR $= 1000$

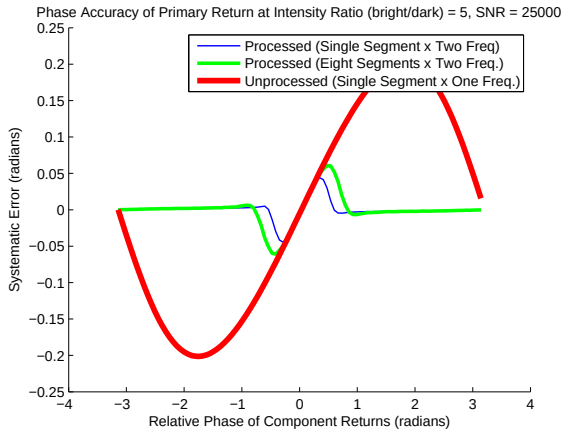

(e) $b^{-1} = 5$, SNR $= 25000$

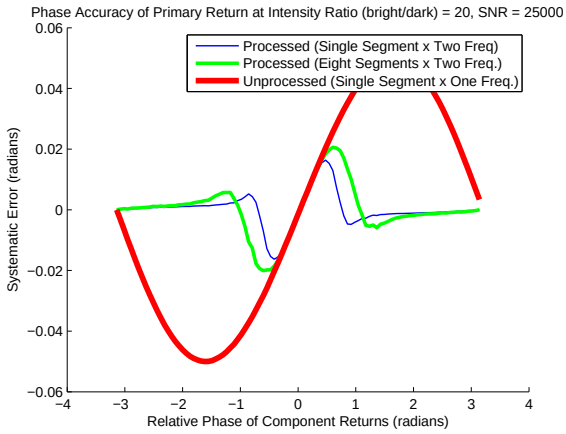

(f) $b^{-1} = 20$, SNR $= 25000$

Figure 7.9: Accuracy versus phase relationship for different SNRs and relative intensities.

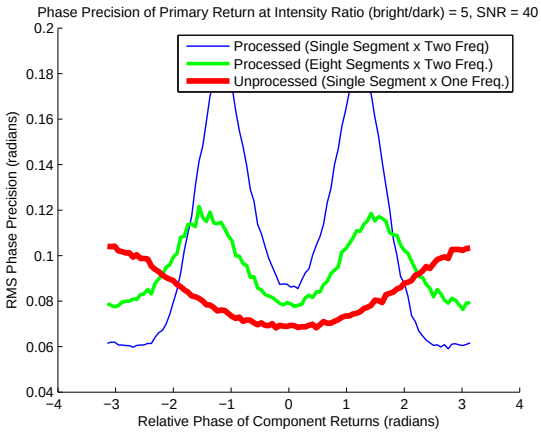

(a) $b^{-1} = 5$, SNR $= 40$

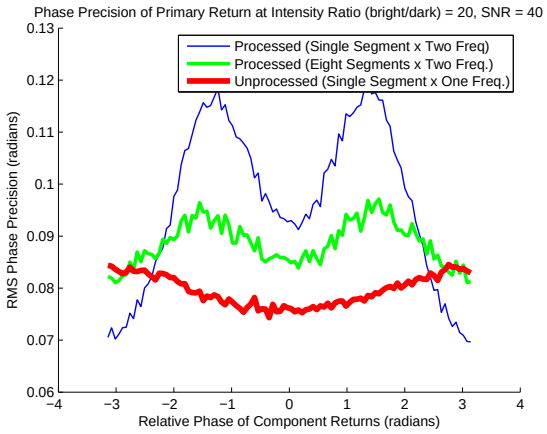

(b) $b^{-1} = 20$, SNR $= 40$

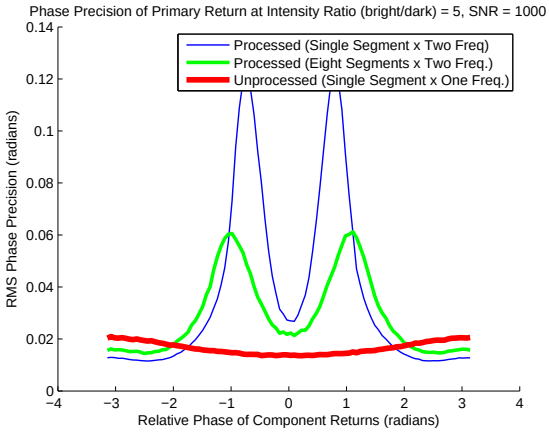

(c) $b^{-1} = 5$, SNR $= 1000$

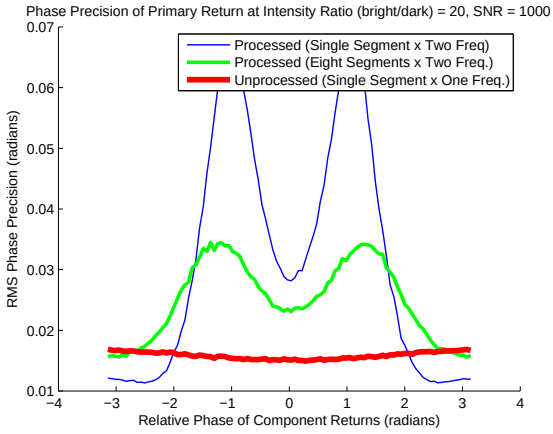

(d) $b^{-1} = 20$, SNR $= 1000$

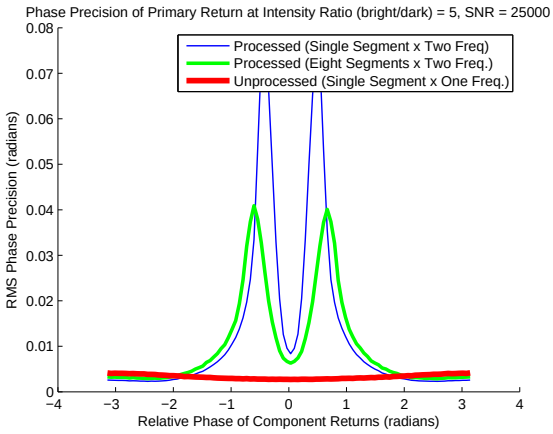

(e) $b^{-1} = 5$, SNR $= 25000$

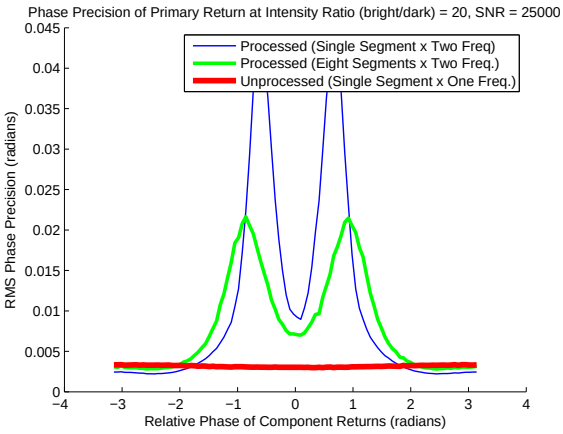

(f) $b^{-1} = 20$, SNR $= 25000$

Figure 7.10: Precision versus phase relationship for different SNRs and relative intensities.

are two processed curves plotted on each subfigure, each one corresponding to a slightly different utilisation of the same overall sensor integration time. The first, in blue, cuts the integration time into two measurements, one at twice the modulation frequency of the other. The second makes eight independent measurements at one frequency followed by eight independent measurements at twice the modulation frequency. In the first case, the measurement is processed to give a single result; in the second case, each of the eight segments is combined with a segment from the other frequency to produce eight independent measurements which are then processed separately and the mean measurement calculated post-processing. So the blue line can be considered to correspond to averaging before processing, and the green line to averaging after processing. Comparing the two cases indicates that averaging before processing results in a significant improvement in systematic phase error, whereas averaging after processing results in an error distribution that is largely a scaled up (worse) version of the blue line.

For each combination of averaging method, relative intensity and SNR there is a particular relative phase at which the processed and unprocessed curves start to diverge. For example, in the case of $b^{-1} = 5$, SNR $= 40$, that point is near$\theta = 0.5$ radians. Below this relative phase, the algorithm does not result in improved phase precision; however, it does not result in worse phase precision either. Fig. 7.9 shows that after the processed data diverges from the unprocessed data, the systematic error eventually decreases and crosses zero, with a slight pseudooscillatory overshoot. The location of this zero crossing changes with averaging method, relative intensity and SNR.

In order to understand the systematic errors introduced by the processing method, we consider the expected value of a linear function of a random normal variable:

$$\alpha = \mathbb{E}[x \mid x \sim \mathcal{N}(0, \sigma^2)]. \tag{7.12}$$

Due to symmetry, $\alpha = 0$, irrespective of the value of σ^2. Now let f_{Q} be the arbitrary function

$$f_{\mathrm{Q}}(x) \quad = \begin{cases} x & x < 0 \\ x^3 & x \geq 0 \end{cases} . \tag{7.13}$$

The expected value of $f_{\mathrm{Q}}(x)$,

$$\beta = \mathbb{E}[f_{\mathrm{Q}}(x) \mid x \sim \mathcal{N}(0, \sigma^2)], \tag{7.14}$$

β, is a function of σ^2 because $f_{\mathrm{Q}}(x)$ is no-longer symmetric around zero. This is

illustrated in Fig. 7.12. The inverse problem $\mathrm{H}^{-1}(\chi)$ can be considered in a similar manner to f_{Q} – because of the shape of the inverse manifold, noise in the measured data can result in systematic perturbations in estimates of the component returns. Since these perturbations are predictable, given prior knowledge of the SNR of a measurement – for example, knowledge of the total integrated intensity and ambient light offset – a third dimension can potentially be added to a lookup table system in order to correct for the systematic offset. We now analyse the impact of relative phase and intensity on phase precision.

Fig. 7.10 is a plot of phase precision versus relative phase for different SNRs and relative intensities. Phase precision in the unprocessed case is best when $\theta = 0$; as θ increases, progressively more cancellation of the complex domain measurements occurs, resulting in a poorer signal intensity. Note that this assumes that the SNR is measured relative to the total integrated intensity of the components, rather than the summed measurements. Processed data, on the other hand, has highest precision near $\theta = \pi$, slightly worse precision at $\theta = 0$ and horrific precision in a region just to either side of $\theta = 0$. This poor precision region corresponds to a region just outside where the processed and unprocessed systematic errors diverged (as shown in Fig. 7.9). This is at the point where the algorithm starts actually separating out component returns. It appears that initially, the algorithm is particularly sensitive to slight perturbations in the parent data. Comparing these random error peaks across various values of b^{-1} and SNR clearly indicates that the worse the SNR or the darker the component return, the bigger the peaks and the farther away from $\theta = 0$ they lie. In the same manner that Fig. 7.9 compared averaging before and after the data had been processed, the blue line corresponds to data which has been averaged before processing, and the green line corresponds to data which has been averaged after processing. Comparing the two clearly shows that post-process averaging results in a very significant decrease in peak precision at the cost of shifting the peak precision slightly farther away from $\theta = 0$, and of increasing the minimum possible systematic error (for some particular value of θ.) We now analyse overall RMS phase error.

Fig. 7.11 is a plot of overall RMS phase error versus θ for different values of b^{-1} and SNR – this shows how precision and accuracy combine to impact the final processed measurements. We can see that there is a particular threshold at which the processed RMS error diverges from the unprocessed RMS error – this does not correspond to exactly the same relative phase as for the divergence of processed and unprocessed accuracy (as in Fig. 7.9), instead it corresponds to a point slightly farther away from $\theta = 0$. As SNR and relative intensity increase, this threshold

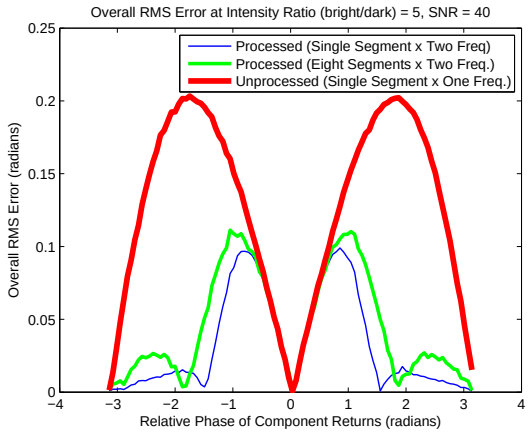

(a) $b^{-1} = 5$, SNR $= 40$

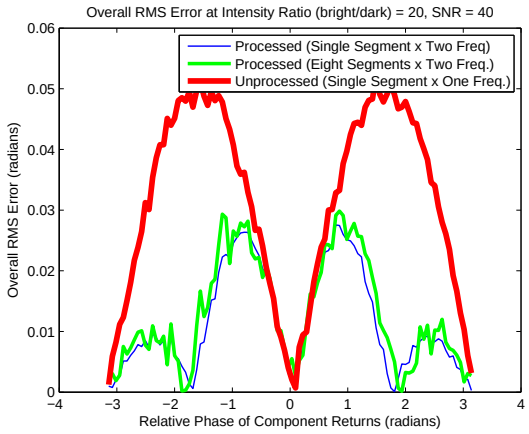

(b) $b^{-1} = 20$, SNR $= 40$

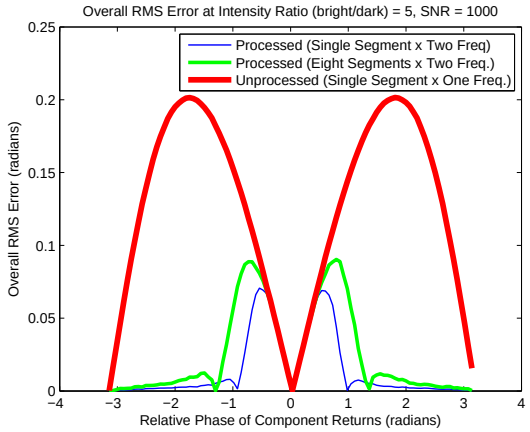

(c) $b^{-1} = 5$, SNR $= 1000$

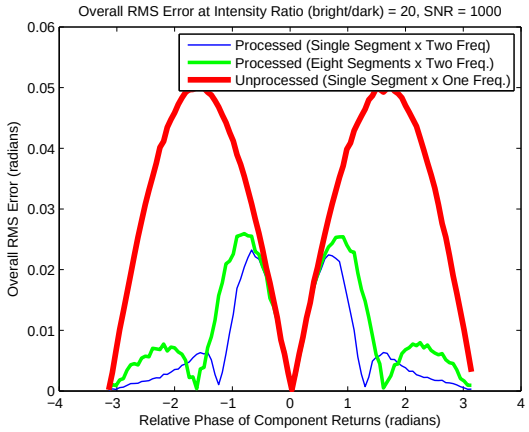

(d) $b^{-1} = 20$, SNR $= 1000$

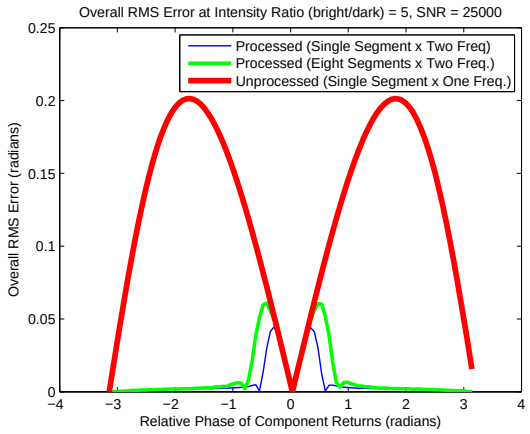

(e) $b^{-1} = 5$, SNR $= 25000$

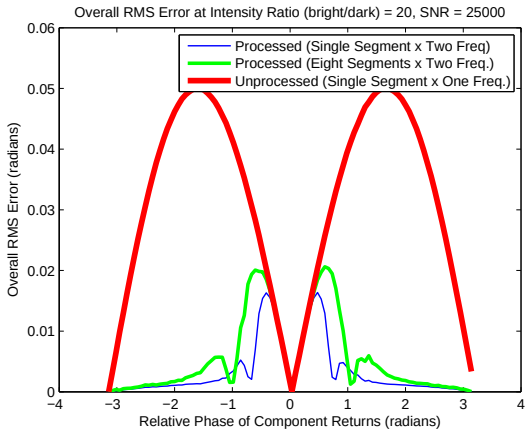

(f) $b^{-1} = 20$, SNR $= 25000$

Figure 7.11: Overall RMS error versus phase relationship for different SNRs and relative intensities.

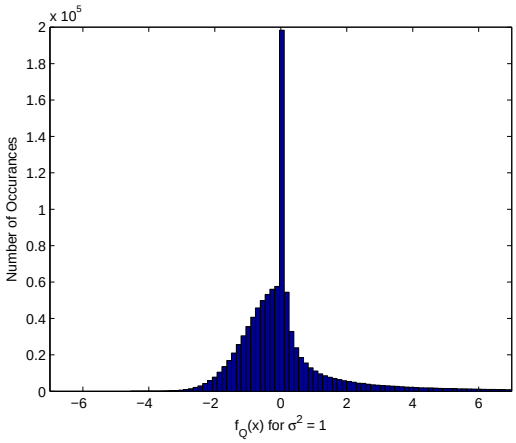

(a) $\sigma^2 = 1$

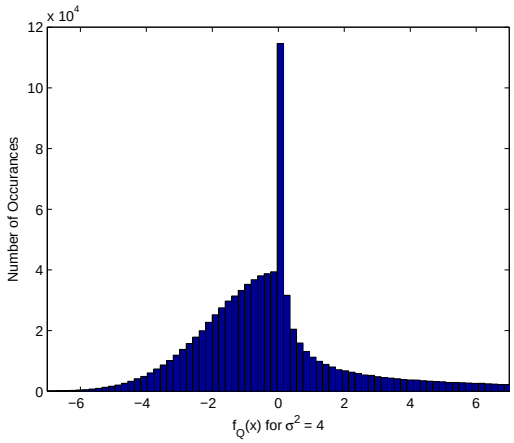

(b) $\sigma^2 = 4$

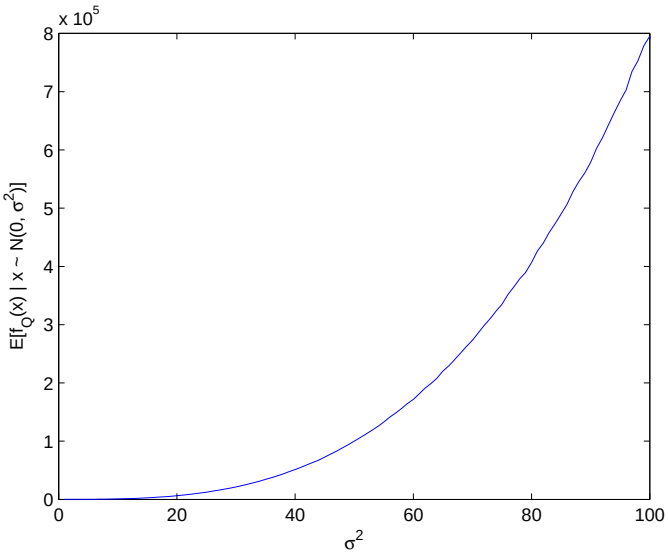

(c) $m = 4$

Figure 7.12: The distribution of $f_{\mathrm{Q}}(x)$ as a function of σ^2. Histograms generated using a 1000000 sample Monte-Carlo simulation.

decreases and peak processed RMS error decreases; as SNR and relative intensity decrease, the threshold increases and peak processed RMS error increases. This threshold is a combination of the impacts of accuracy and phase; random phase error peaks just outside the processed/unprocessed accuracy divergence threshold so as to result in a continuous curve. Post-threshold, the RMS error decreases in steeper fashion than the post-threshold accuracy, due to the steep decrease in random phase error. Returning to the comparision of pre- and post-process averaging, Fig. 7.11 clearly shows that while post-process averaging may improve phase precision, it also makes overall error worse and is therefore not as useful as pre-process averaging. In addition, post-process averaging requires a larger number of readout cycles, which

results in the loss of integration time; albeit, this was not explicitly factored into the simulations.

7.2.4 Choosing Lookup Table Dimensions

Fig. 7.13 is plot of the CDF of systematic error as a function of lookup table dimensions and SNR. For higher LUT resolutions as the SNR increases, the accuracy CDF asymptotically approaches a fundamental limit; although, for particularly low resolutions, the presence of moderate levels of noise can actually improve systematic error through dithering. Comparing Figs. 7.13a and 7.13, LUT resolution is clearly a limiting factor at high SNR. Even for the worst simulated SNR level, $40:1$, the processed data results in a significant improvement in accuracy over the unprocessed reference data. An interesting

Fig. 7.14 shows precision to be impacted in a very different manner by SNR and LUT resolution in contrast to accuracy: there is no significant impact due to LUT resolution, the only impacts are due to subtle changes in how the LUT is addressed due to quantisation. As SNR is increased, precision continues to increase across the entire CDF curve – however, certain regions of the CDF curve improve less than others.

One particular situation where high noise levels occur is when τ_0 is very low – in other words, the component returns have nearly cancelled out when measured at the lowest measurement frequency. From Section 6.2.2, given a value of χ at a frequency ratio of $2:1$ we can write the equation for finding the component returns as

$$\mathrm{H} \;=\; \hat{R}^{-1}(\chi)\eta_0. \tag{7.15}$$

If we ignore the possibility that $\hat{R}^{-1}(\chi)$ might have some sort of noise cancelling influence, we can infer a lower bound on the SNR of the recovered component returns as being the SNR of η_0, the low frequency complex domain range measurement. Returning to Fig. 7.14, we can see that if this was the only noise source of importance (as would be the case for a 1×1 LUT), then the phase precision before and after processing would be identical. Therefore, in order to produce the given precision CDFs it is necessary for the SNR of χ, propagated through the inverse, $\hat{R}^{-1}(\chi)$, to result in significant noise magnification. Since we do not have a mathematical expression for the inverse, we terminate our analysis at this point.

As a general statement, we can say that accuracy depends on a combination of LUT resolution and SNR, and precision solely on SNR. This is demonstrated by

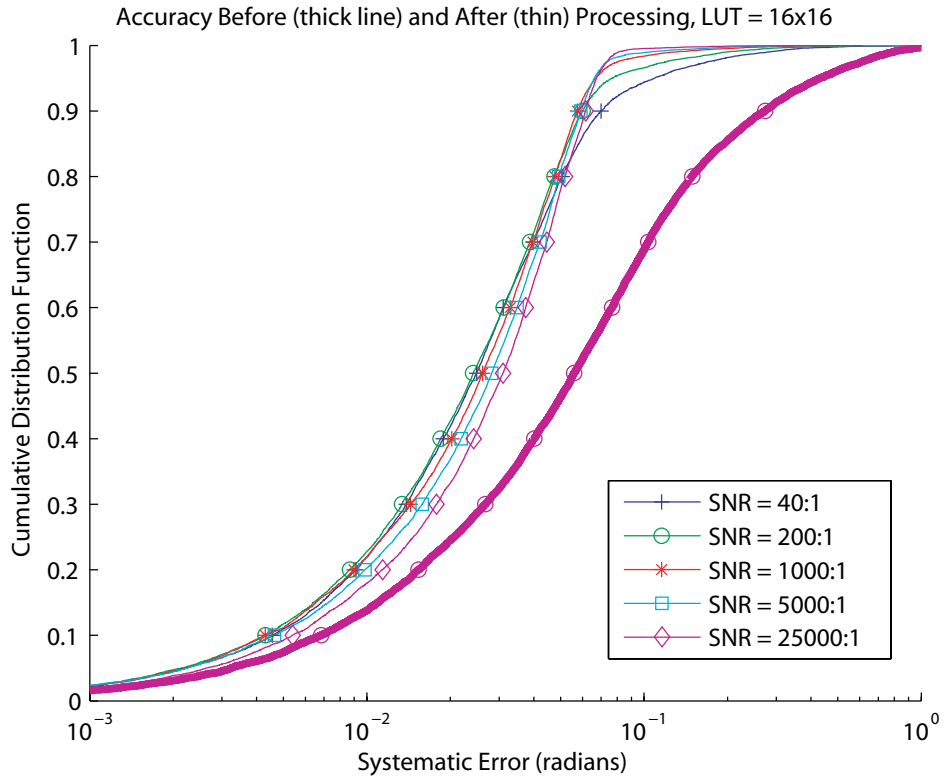

(a) Accuracy at 16×16

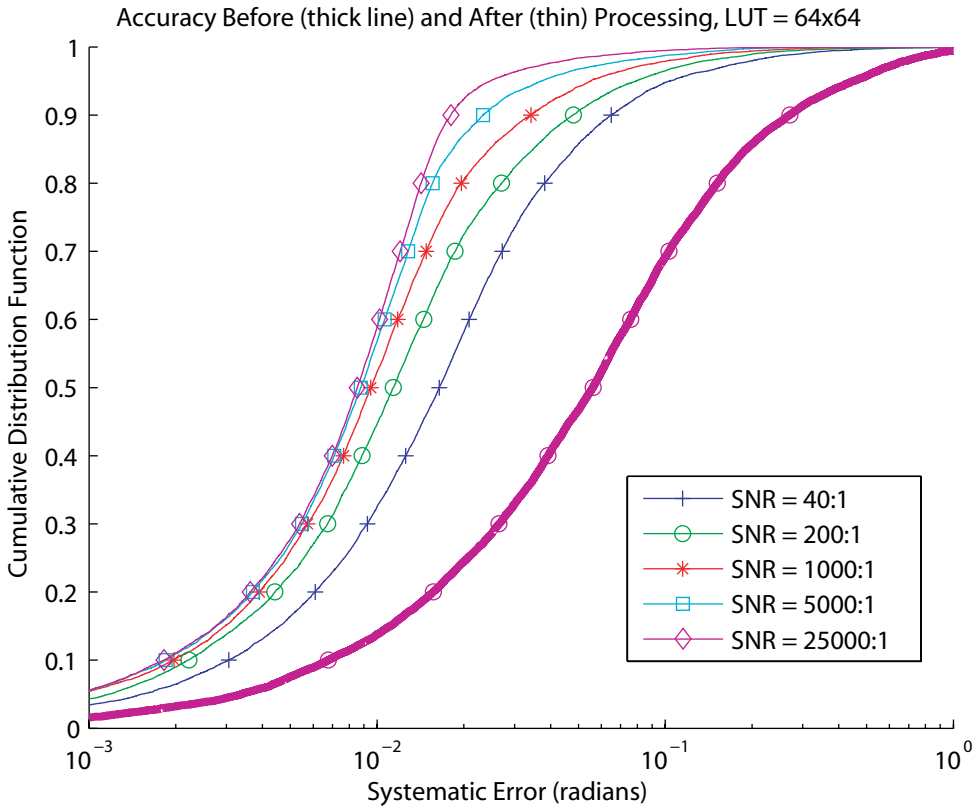

(b) Accuracy at 64×64

Figure 7.13: Accuracy versus SNR for different LUT resolutions using a translated polar layout. The thick line is a reference curve of unprocessed data.

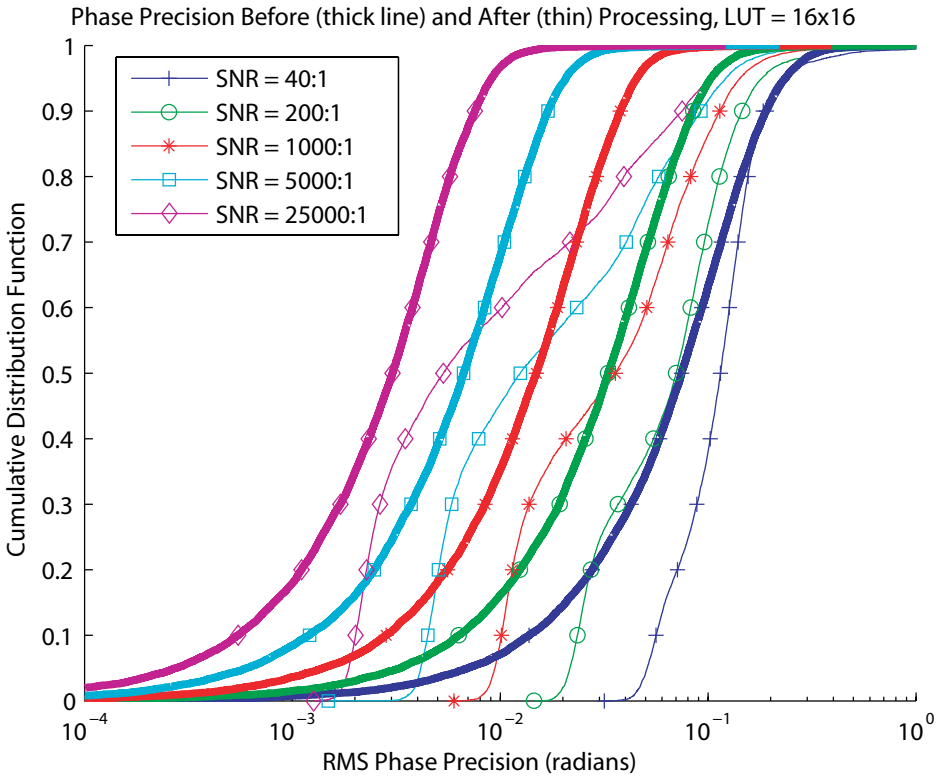

(a) Precision at 16×16

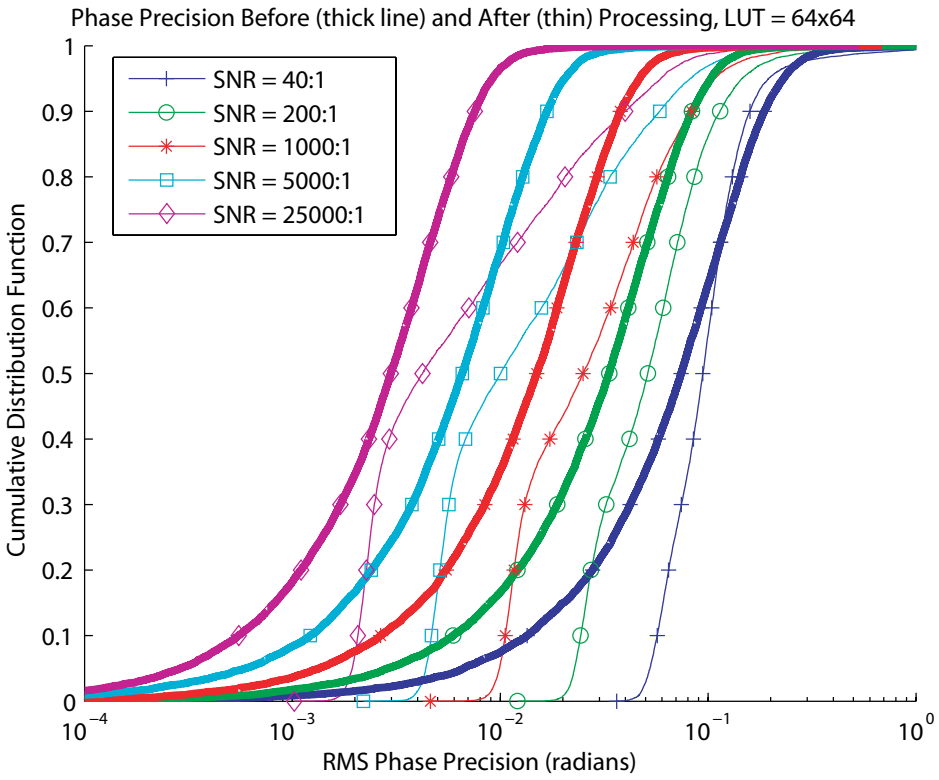

(b) Precision at 64×64

Figure 7.14: Precision versus SNR for different LUT resolutions using a translated polar layout. Thick lines represent unprocessed data.

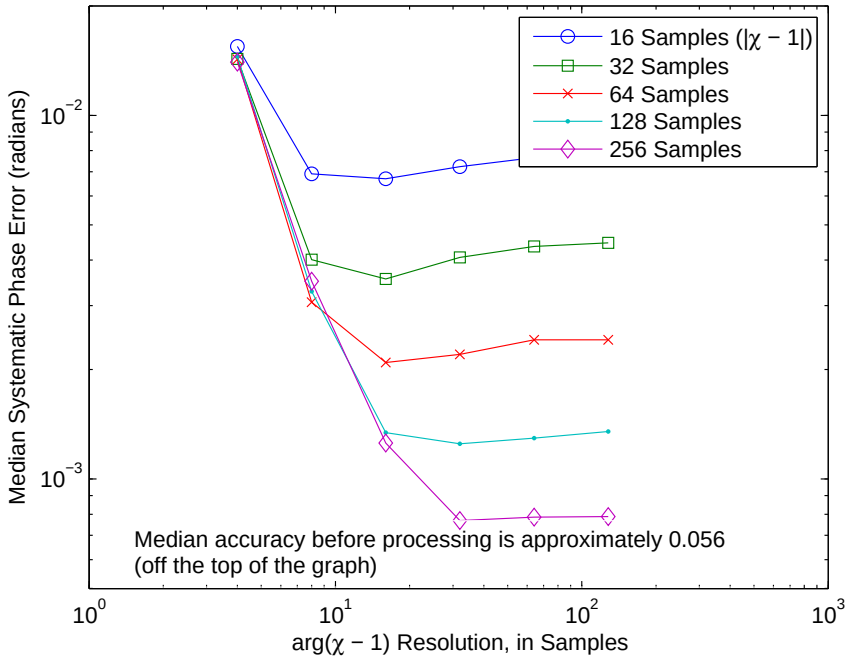

(a) Systematic Error (Accuracy)

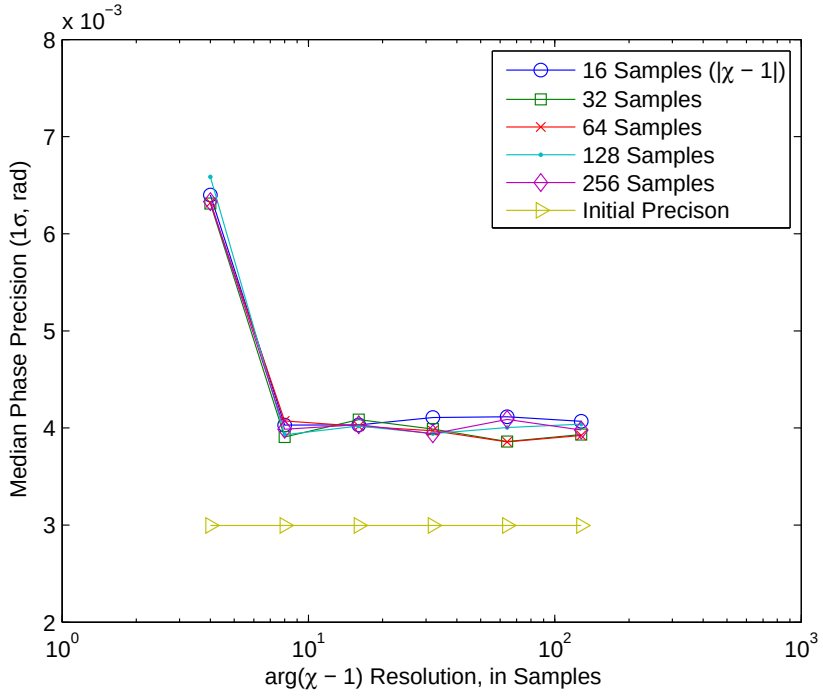

(b) Precision

Figure 7.15: Median precision and systematic phase error versus lookup table dimensions at a SNR of 25000 : 1.

plotting the median systematic phase error and median phase precision versus LUT resolution in Fig. 7.15 – median values being employed in order to avert the skewing of results by the long tail of the error distributions.

Fig. 7.15a shows that an improvement in accuracy requires an increase in both phase and amplitude resolution around $\chi = 1$. The optimal dimensional sampling ratio appears to be somewhere between $4:1$ and $8:1$ in terms of $\arg(\chi - 1)$ samples to $|\chi - 1|$ samples. There are a few unexpected values, such as at $(8, 64)$ which has better precision than $(8, 128)$ – this value is definitely correct and is almost certainly due to the confluence of quantisation and the specific mapping of $|\chi - 1|$ from $[0, \infty]$ to $[0, 2]$. While we have not endeavoured to develop an optimal projection, it is a candidate for further research.

The lowest resolution samples on the x-axis for the precision measurement in Fig. 7.15b appear to be erroneous, the ostensible genesis of the high systematic error being some sort of quantisation related error – credible given the assumption of four samples in the phase dimension around $\chi = 1$. In the majority of cases, the algorithm results in a 33% increase in median phase precision.

7.2.5 Third Components

While this method was originally intended to address the case where there is either one or two component returns within a pixel, it is nevertheless possible that there be more than two component returns. An important question to resolve is whether the algorithm continues to improve the range information in the case of additional interference. In order to evaluate this property, consider the model of Section 7.1, extended by adding an additional component return. That is,

$$a_2 \sim \mathcal{U}(0, b_{\mathrm{M}}) \tag{7.16}$$

where it is assumed that $b_{\mathrm{M}} = 0.1$. Whereas previously the relative intensity of the component returns was fully constrained by a single variable b, we now introduce two symbols b_1 and b_2, which represent the relative intensity of the second bright component return and third brightest component return relative to the brightest ie.

$$b_1 = \frac{\sup(\{a_0, a_1, a_2\} \backslash \gamma)}{\gamma} \tag{7.17}$$

$$b_2 = \frac{\inf(\{a_0, a_1, a_2\})}{\gamma} \tag{7.18}$$

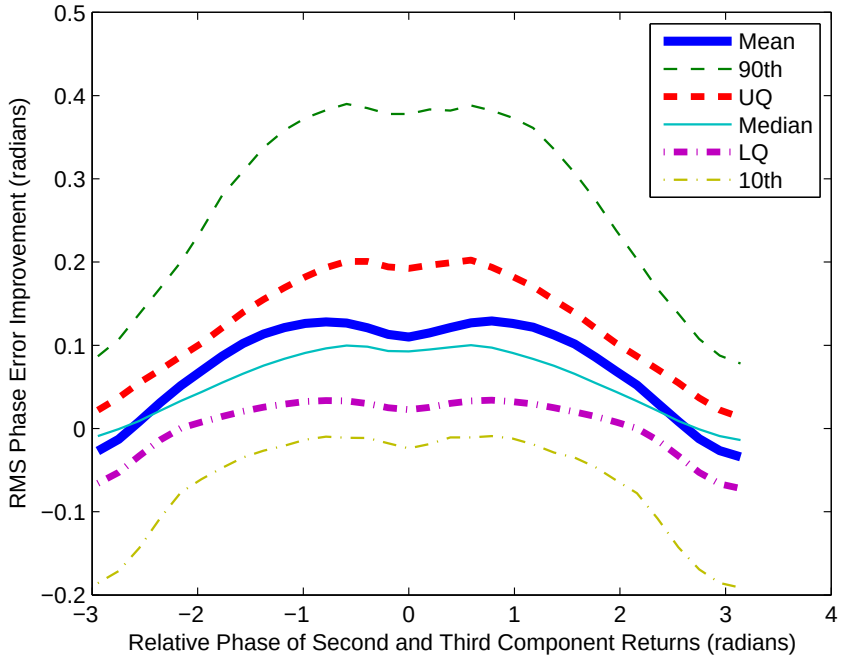

(a) Overall Error Improvement vs. Relative Phase

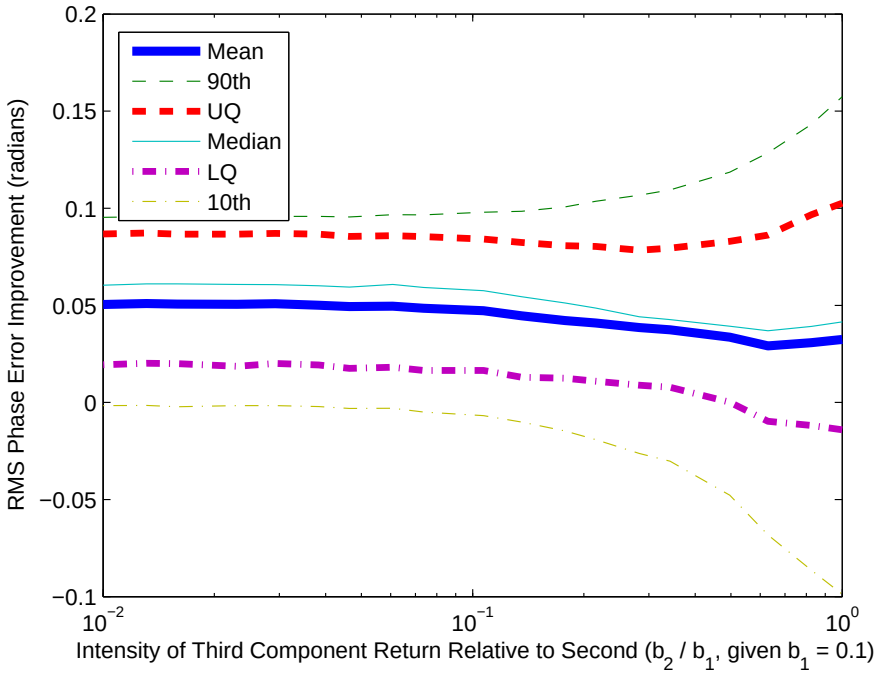

(b) Overall Error Improvement vs. Relative Intensity

Figure 7.16: Improvement in overall RMS error of the primary component return in the case of three component returns, by relative phase of second and third components and intensity of third component relative to second (b_2/b_1) at an SNR of 25000 : 1.

where $\gamma = \sup(\{a_0, a_1, a_0\})$. As before, assume a uniform prior for the phase of the additional component return over $[0, 2\pi]$.

Fig. 7.16a gives the improvement in RMS phase error as a function of the relative phase of the second and third brightest component returns. The case where the relative phase is zero corresponds to a degeneracy: the two component return case, which provides a useful reference point. As relative phase diverges from zero, initially the processed estimates improve and then become worse than the two return case at around $\pm\frac{3}{2}$. The worst possible case occurs when the two component returns are separated by π radians – in other words they are completely out-of-phase at a relative frequency of one and completely in-phase at a relative frequency of two. This potentially corresponds to a significant perturbation of the characteristic measurement, thus an erroneous restoration. Also note that for a very small region near a relative phase of π, the net impact of processing the range data is to increase error in both the mean and median RMS phase error improvement cases.

Fig. 7.16b is the improvement in RMS error improvement as a function of the intensity of the third component relative to the second component. Note that as the relative intensity increases, mean RMS error improvement decreases and the RMS error 10th and 90th percentiles diverge. In other words, there are greater extremes in terms of improvement and worsening of error, depending on the distributions of the relative phases.

Overall, the inversion in the 2 : 1 case can be considered to be partially resistant to the influence of third components in that the processing still results in an overall improvement in expected phase error.

7.2.6 The Problem with Multiple Solutions

Situations where there are multiple possible valid solutions to a particular set of measurements are particularly problematic. However, there are some physical and probabilistic constraints than can potentially be utilised to assist in the solution of the problem. Firstly, the marginal probabilities for the different solutions are not necessarily equal. Making the assumption of the distribution from Section 7.1, at a frequency ratio of 3 : 2 one particular solution has a 50% marginal probability, while the other two have 25% percent probabilites each. Additionally, there are spatial constraints; real measurements of real scenes do not consist of matricies of uncorelated data. One of the most promising approaches to determination of the correct solution is the application of Markov Random Fields. By applying spatial smoothness constraints and knowledge of the marginal probability of each particular component return it is possible to construct an energy function to minimise; for

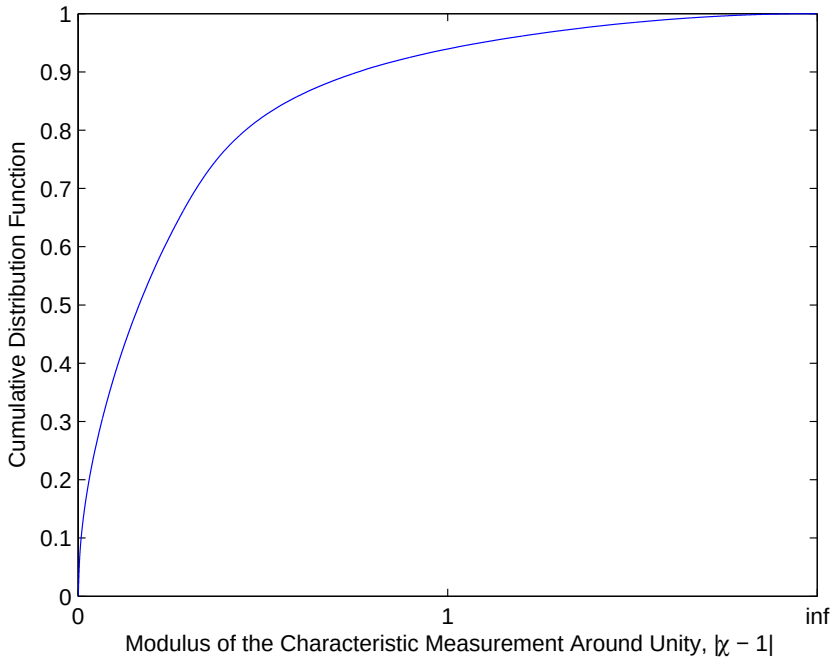

(a) Distance from $\chi = 1$

Figure 7.17: The distribution of phase peturbation and distance from $\chi = 1$ given the distribution of component returns introduced in Section 7.1 at 2 : 1.

example, as is done in Ising models (Bishop, 2006). At present we have only briefly experimented with this approach. Particular issues include computational complexity issues inherent in the solution process which make the method impractical for many real-time situations. This is deserving of further research.

Closely related to the problem of determining which solution is correct is determination of the number of components within a pixel. Like the multiple solution problem it is possible to deal with this on a single pixel basis, but there is potential for better results using spatial algorithms. From Eqn. 6.58 we know that the single component return cases are located at the r_0th roots of unity. We also know that the distance from these roots can be used as a metric of how mixed a point is; this is demonstrated in Section 6.4.6, setting minimum and maximum bounds on normalised total integrated intensity for a frequency ratio of 2 : 1. Values of normalised total integrated intensity above one indicate the partial cancellation of component returns. As the value of $|\chi - 1|$ increases, the minimum bound on normalised total integrated intensity increases, ergo the pixels become progressively more and more mixed/perturbed. We can generalise this to any non-zero value of r_0 by taking the distance from the closest root of unity. Fig. 7.17 shows that this drops off quite quickly. If we have knowledge of the signal to noise ratio in an image it is possible

to calculate the likelihood of the measurement being a single component return – implying either analysis of a sequence of images or knowledge of the total integrated intensity and ambient offset. Given that there is no plausible assumption we can make regarding the distribution of one versus two components it is not possible to calculate an estimate of the posterior probability. Section 7.2.8 gives an example scene where we have used the Mahalanobis distance from unity in the $2:1$ frequency ratio case to provide a metric of how likely a pixel is likely to be a single return. This is important, because near the roots of unity there can be a noise magnification effect leading to a significant loss of precision. One way to mitigate this is to assume that the point is a single component return and then use the measurement at the second frequency to denoise and unwrap the measurement at the first frequency instead of to separate out the component returns. For example, in the single component return case we can produce an optimal, unwrapped estimate of the phase of the underlying complex return, $\arg(\eta_0)$, using knowledge of the SNR for the estimate given by each measurement;

$$SNR_l = \left(\frac{d\arg(\eta_l)}{d\arg(\xi_l)}\right)^2 \cdot \frac{|\xi_l|^2}{\sigma_l^2}, \tag{7.19}$$

where σ_l^2 is the noise variance of the measurement ξ_l. Hence,

$$\arg(\eta_0) = \frac{\sum_{l=0}^{N-1} SNR_l \frac{\arg(\eta_l)+2\pi m_l}{r_l}}{\sum_{l=0}^{N-1} SNR_l} \tag{7.20}$$

$$= \frac{\left(\frac{r_0|\eta_0|}{\sigma_0}\right)^2 \cdot \frac{\arg(\eta_0)+2\pi m_0}{r_0} + \left(\frac{r_1|\eta_1|}{\sigma_1}\right)^2 \cdot \frac{\arg(\eta_1)+2\pi m_1}{r_1}}{\left(\frac{r_0|\eta_0|}{\sigma_0}\right)^2 + \left(\frac{r_1|\eta_1|}{\sigma_1}\right)^2}, \tag{7.21}$$

where $m_0, m_1 \in \mathbb{Z}$ are phase unwrapping constants which are chosen so as to minimise

$$\left|\frac{\arg(\eta_0)+2\pi m_0}{r_1} - \frac{\arg(\eta_1)+2\pi m_1}{r_0}\right|. \tag{7.22}$$

Eqn. 7.21 is a first order approximation that is not suitable for very low SNR, where the coupling between errors in phase and amplitude is significant. Another alternative is to simply alter the lookup table near $\chi = 1$ so as to achieve soft-thresholding second component removal, although this is not noise optimal.

Given a scene which is primarily composed of single component returns or near single component returns, summing the Mahalanobis score over all pixels of the image provides a cost function that can be used for automatic calibration of the relative

phase and intensity characteristics of the different frequencies. This is useful because there is generally an attenuation in the modulus of complicated domain range measurements as modulation frequency increases, there are also varying phase shifts. Predicting these can be a complex endeavour due to non-linearities in the electronic circuits and temperature induced phase/amplitude drift. We have frequently used this technique with the image intensifier system.

7.2.7 Overall Performance Comparison

We now compare the overall performance of the lookup table algorithm, optimisation and attenuation ratio methods. Thus far we have not addressed the performance of attenuation ratio based methods; that is because the two complex domain range measurement method has significantly better noise properties. In order to demonstrate this we have generated attenuation ratios from complex domain range measurements with our standard noise assumptions, but given noiseless total integrated intensity information. This results in attenuation ratios with lower noise levels than would otherwise hold, but avoids making more complicated assumptions which could obfuscate the results.

Fig. 7.18 gives the results of simulations at signal to noise ratios of $25000:1$ and $200:1$. There are two variations of the attenuation ratio method, listed as 'Closed Form (CF)' and 'CF Bounded τ' – the latter being a variation that adds an additional processing step before the attenuation ratio inversion algorithm limiting τ_0 to a maximum of 0.999999. Detailed analysis of the performance of the attenuation ratio method indicated that the performance was being seriously degraded by values where $\tau_0 \geq 1$, corresponding to a physically impossible situation. Fig. 7.18b at the lower SNR particularly demonstrates this degradition with over 10% of values with greater than one radian RMS phase error. After limiting the value of τ_0, fewer than 10% of values have worse than 0.1 radian RMS phase error. Even with this limit, the attenuation ratio has relatively poor performance compared to the lookup table method (notated as '$x \times y$ χ LUT'). If a 256×256 sample lookup table were plotted, it would be indistinguishable from the optimisation curve.

7.2.8 Real-data Results

In this section we demonstrate the operation of the lookup table algorithm. All measurements are taken using a frequency ratio of $2:1$ and processed using a translated polar LUT with a resolution of 4096×512.

Fig. 7.19 presents two scenes; the first captured using the Mesa Swissranger at

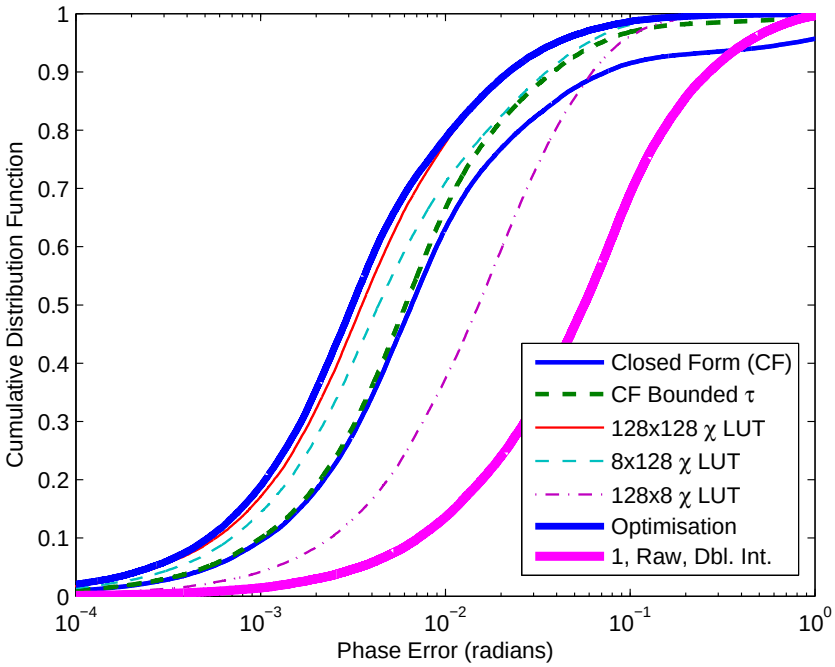

(a) SNR = 25000 : 1

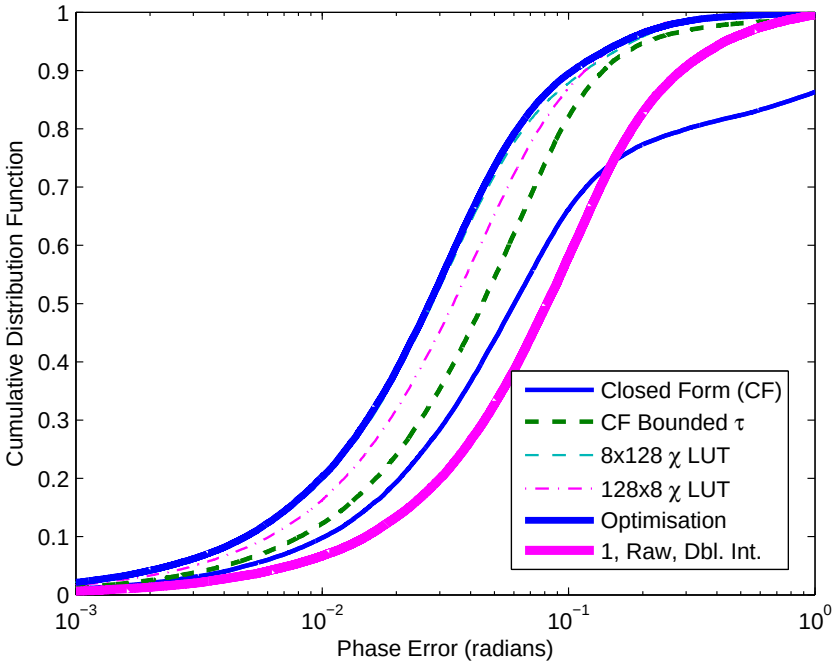

(b) SNR = 200 : 1

Figure 7.18: Overall comparison of phase error in recovered primary component return for different methods. Closed form refers to closed form attenuation ratio methods. For these graphs no noise was present in the total integrated intensity measurements used to calculate the attenuation ratios – definitively proving that attenuation ratio methods are inferior to techniques also incorporating phase information.

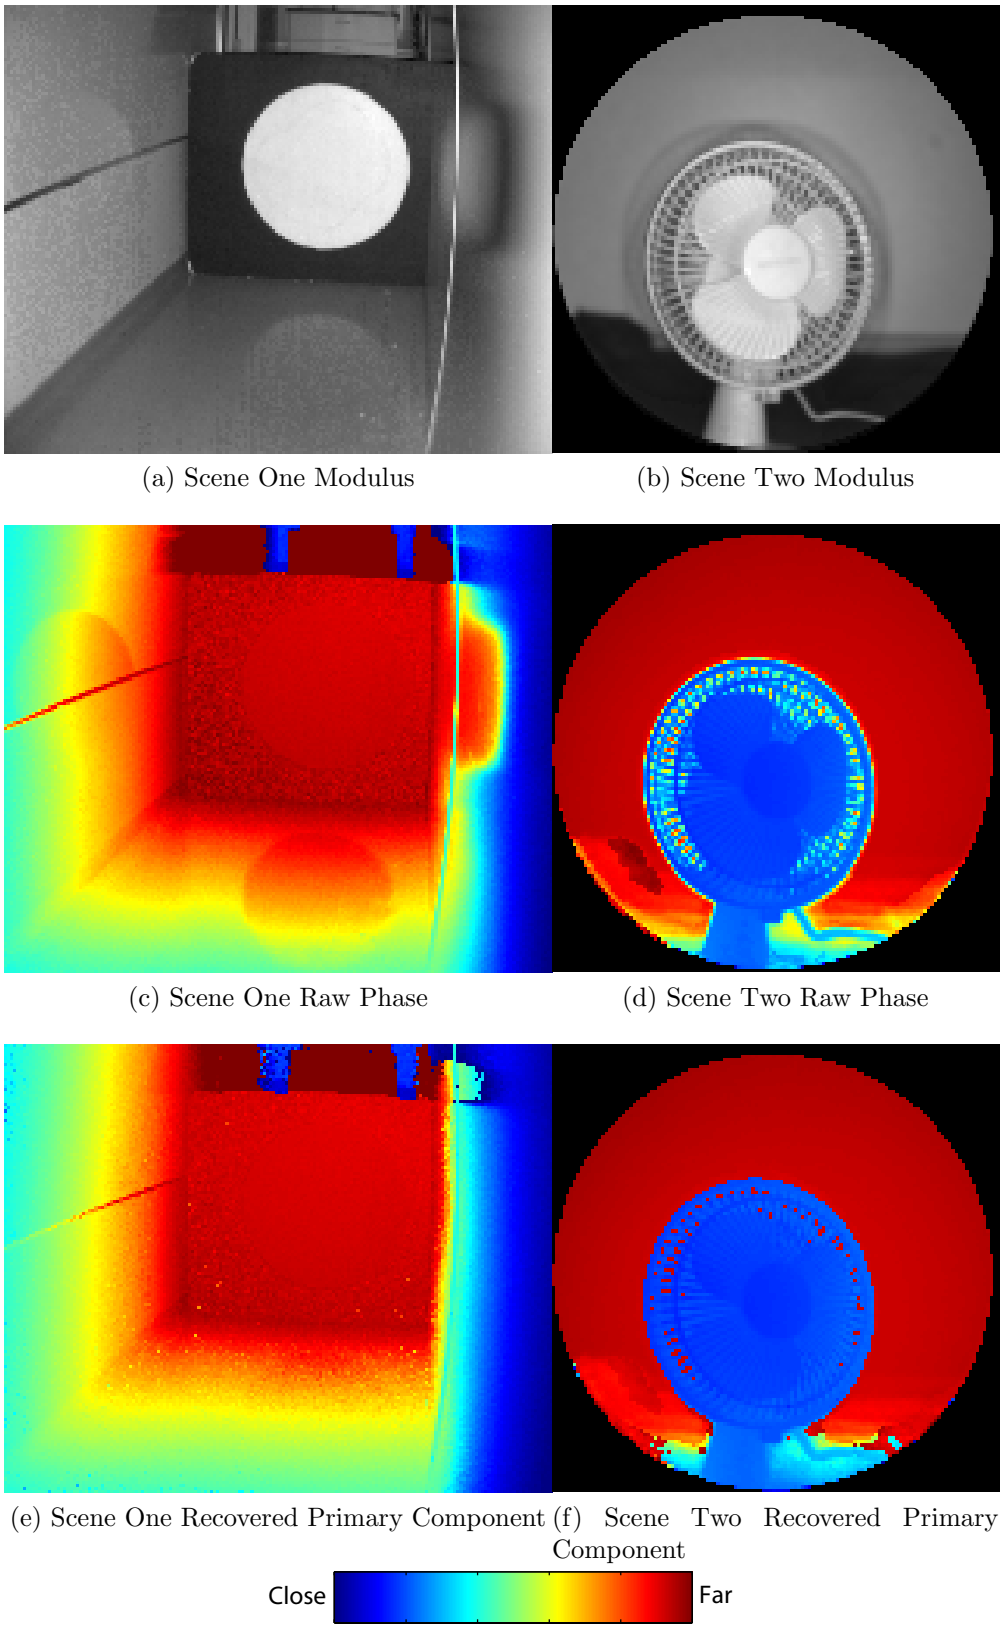

(a) Scene One Modulus

(b) Scene Two Modulus

(c) Scene One Raw Phase

(d) Scene Two Raw Phase

(e) Scene One Recovered Primary Component

(f) Scene Two Recovered Primary Component

Figure 7.19: Two example restored scenes; the first of multipath reflections taken using a Mesa Swissranger at 60/30MHz, the second mixed pixels in an image of a complicated object taken using the image intensifier system at 80/40MHz.

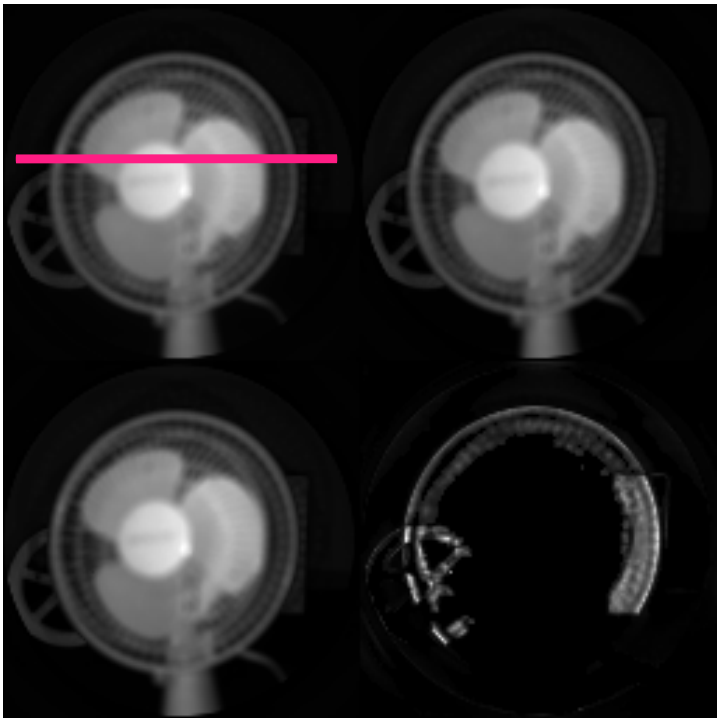

(a) Modulus

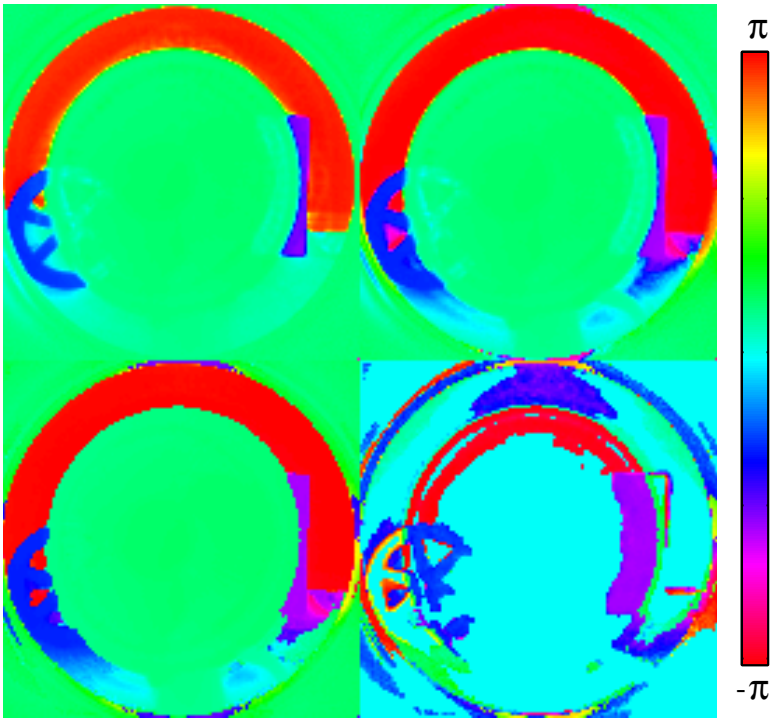

(b) Phase

Figure 7.20: Scene three: Clockwise from top left: original scene, scene with isoplanatic deconvolution based scattering reduction, second component return, first component return. Note that low probability second components have been thresholded out and that the modulus of the second component has been scaled differently from the other images. The thick line indicates the location of the slice plotted in Fig. 7.21

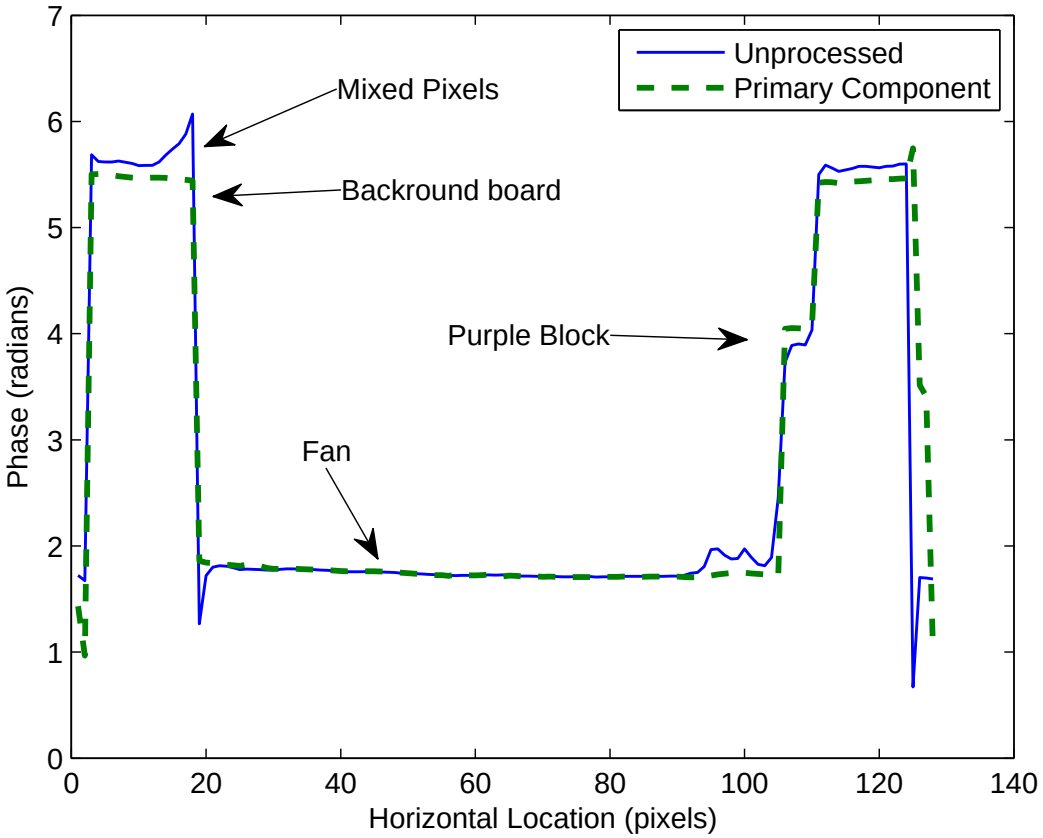

Figure 7.21: A slice through scene three, from Fig. 7.20

60/30MHz, the second using the image intensifier system at 80/40MHz. Scene one is an image of a target located in a reflective corridor. The high albedo target has produced reflections in both walls and the floor – Fig. 7.19c shows that these reflections have resulted in significant multipath perturbation of the range measurements. Passing the raw data through the algorithm gives the recovered primary component return phase plotted in Fig. 7.19e, where the majority of the multipath error has been clearly removed. However, an increase in noise is clearly visible in regions that do not have multipath problems.

The phase and modulus of scene two is shown in Fig. 7.19b and 7.19d. This is a mixed pixel problem, primarily caused by the partial masking of the wall by the fan grill. Fig. 7.19f gives the phase of the recovered primary component. The fan looks somewhat uglier, but no longer contains erroneous intermediate values – in a spatial sense processed versus unprocessed data can be compared to aliased versus anti-aliased algorithms for drawing lines in a bitmapped image.

Scene three shows a more complicated variation of the restoration algorithm. Fig. 7.20 shows a slightly defocussed scene containing three different objects on a

table with a board behind imaged using the image intensifier system. In Fig. 7.20b, the blue object is a wheel, the green object a fan and the purple object a box; both the wheel and box are partially obscured by the fan. One of the limitation of the restoration algorithm is that it is primarily designed to handle two components; the image intensifier system, in particular, has a problem with scattering within its primary optics due to the high reflectivity of the image intensifier. This results in light from each of the objects in the scene being scattered across the entire image. Previous papers have utilised deconvolution to partially mitigate this effect (Mure-Dubois and Hügli, 2007a; Kavli *et al.*, 2008). By sampling the scattering point spread function and applying an isoplanatic, spatial derivative regularised deconvolution in the Fourier domain we can largely remove this perturbation (the approach is detailed in appendix C). The upper left quadrant of Fig. 7.20b gives the phase before the deconvolution and the upper right quadrant gives the phase afterwards. The processing has removed some of the mixing around the outer rim of the fan and the wheel, it has also partially removed scattered light from the table; because the table is nearly completely unreflective, it is very susceptible to perturbation by scattered light. In the original image, the table appears to be located at a homogeneous range, despite being sloped away from the camera – the range being approximately a weighted mean of the ranges within the image (phase angle of the mean of the complex domain range measurements at each pixel over the entire image.)

The bottom two quadrants of Fig. 7.20b show the recovered range to the primary and secondard component returns, where thresholding has been applied to determine whether the second component return was valid or not. The primary component return shows crisp edges, but suffers from some minor artefacts over on the lower right hand side. The second component return is much more interesting. The range to the board at the back (red) has been clearly identified through the fan, as has the purple block on the right hand side, despite both only resulting in barely perceivable perturbations in the original raw measurement. On the left hand side, the blue wheel has been identified through the fan as has some locally scattered light. Around the outside are a number of low intensity second component returns of dubious value.

Fig. 7.21 shows a horizontal slice through the phase of the raw measurements at the lower frequency and the estimated phase of the primary component return. The graph shows clearly that the processing has removed fluctuation from the phase data which cannot be justified by the physical construction and layout of the scene. For example, in the raw data on the left the unprocessed phase measurements of the background board have a very suspicious curve, given that the board was completely

flat. After processing, the board appears slightly sloping, but otherwise flat. Similar effects are visible for the purple block and fan.

Scene four is a range-image captured at 40 MHz and 80 MHz using the image intensifier system; the layout is shown in Fig. 7.22. The first range-image (Figs. 7.23a and b) is taken without the tissue box, which acts as a source for intra-camera scattering induced perturbations. The letter notations were added after the fact to indicate particular regions of the image. This particular ranging system is highly sensitive to intra-camera scattering due to the highly reflective nature of the image intensifier. Table 7.2 gives details of the labelled regions of the image before and after processing, it can be seen that even in the 'unperturbed' data, the differing reflectivities of the regions has resulted in very slight phase perturbations. Nonethe-less, this range-image is close to ground-truth, and we have used it as a reference to estimate the performance of the separation and bounding algorithms.

The second range image in Fig. 7.23 was taken with the addition of the scattering object; the box is the bright object in Fig. 7.23c. While both Fig. 7.23a and c are gamma compressed in order to increase the apparent dynamic range, it is still very clear that there is a huge dynamic range within the second range-image. This large dynamic range increases the likelihood of the background stairs being perturbed by light scattered by the extremely bright box. As can be seen in Fig. 7.23d, about half the squares have been highly perturbed, with some of the squares appearing closer to the camera and some appearing farther away than in the unperturbed image. By applying the LUT separation algorithm, the phases in Fig. 7.23e and f are recovered. Due to the high dynamic range of the scene, for many squares the light scattered from the box was brighter than that originating from the square itself. It can also be seen that if there is only one clear backscattering source within a pixel, as within the bright box, then the estimated phase of the second component is extremely unpredictable. From Table 7.2 we can see that processing the most perturbed squares has achieved dramatic improvements in phase error (multiple radians), while squares that have been minimally perturbed appear to have fractionally worse error (¡ 10 milliradians).

Fig. 7.24 demonstrates several approaches to bounding parameters of the underlying components. In this particular case, we compare bounds generated from the raw captured characteristic measurement values, with calculated values of phase perturbation and relative intensity generated by forward transformation of the components determined by the LUT separation algorithm. While not an exact reference, the true values would be visually indistinguishable from the subfigures in Fig. 7.24. Because applying the separation algorithm to perfect single returns results in a

Table 7.2: Performance of multipath interference separation and bounding algorithms for labelled regions of scene four (Fig. 7.23) at a two-to-one frequency ratio.

Perturbed and Reference Image Statistics (Low-Frequency, ξ_0 where $r_0 = 1$)							
Region	A	B	C	D	E	F	G
Mean Unperturbed Modulus ($\times 10^3$)	8.87	86.6	53.1	80.6	154	11.2	4.65
Mean Unperturbed Phase (mRad)	4855	4825	4831	4878	4444	4742	4595
Mean Perturbed Phase (mRad)	1247	5003	5162	5033	4362	1390	1352
1σ Precision of Pert. Phase (mRad)	36.41	18.40	29.50	16.13	15.07	15.97	21.71
LUT-Based Multipath Interference Separation Algorithm Performance							
Mean Abs. Error Pert. vs. Ref. (mRad)	2665	178	331	156	83	3079	3015
Mean Abs. Error Proc. vs. Ref. (mRad)	121	180	338	161	84	15	14
Improvement in Mean Abs. Error (mRad)	2544	(2)	(7)	(5)	(1)	3064	3001
Estimated Perturbation Parameters (Calculated From LUT-Based Separation Results)							
Mean Calculated $\lvert\theta_\mathrm{f}\rvert$ (mRad)	121	180	338	161	84	15	14
Mean Calculated b	0.308	0.379	0.538	0.307	0.466	0.171	0.107
Mean Calculated $\lvert\theta\rvert$ (mRad)	2857	2829	2818	2756	3050	3070	3028
Bounds Calculated From the Characteristic Measurement							
Max($\lvert\theta_\mathrm{f}\rvert$) from arg($\chi$) (mRad)	183	263	449	248	113	25	25
Max($\lvert\theta_\mathrm{f}\rvert$) from $\lvert\chi - 1\rvert$ (mRad)	1016	1115	1265	1032	1196	784	633
Max($\lvert\theta_\mathrm{f}\rvert$) from arg($\chi - 1$) (mRad)	386	466	619	507	174	85	125
Max($\lvert\theta_\mathrm{f}\rvert$) from $\lvert\chi\rvert$ (mRad)	2344	2402	2546	2320	2530	2235	2178
Min(b) from arg(χ)	0.121	0.174	0.295	0.164	0.075	0.017	0.016
Min(b) from $\lvert\chi\rvert$	0.283	0.342	0.476	0.261	0.462	0.169	0.105
Min(θ) from arg(χ) (mRad)	122	175	299	165	75	17	16

slight increase in ranging error due to noise magnification, it is useful to be able to determine when a pixel is mixed or not. If the pixel contains only a single backscattering source, then there is no reason to attempt to separate out multiple components within it. One approach is to discriminate between single and multiple backscattering components using the likelihood that the characteristic measurement corresponds to a single component return; that is, the probability that the measured value of the characteristic measurement was produced by random chance, given that the true value actually corresponds to the r_0th root of unity[1]. A useful approximation is to use the SNR and the first order derivatives of the real and imaginary parts of the characteristic measurement relative to the real and imaginary parts of the raw measurements to determine the Mahalanobis from the closest root of unity: this is the value plotted in Fig. 7.24a. The Mahalanobis distance clearly identifies the box at the front of the scene as having only a single backscattering source and identifies the amount of mixing for other regions of the scene based upon noise statistics, in this case modelled using a model of variance as a linear function of the mean intensity of the correlation waveform[2]. Most notably, square G is the closest to being a single backscattering source because of a high noise level and that the scattered light is an order of magnitude brighter than the original unperturbed intensity of the square; in other words, the 'single backscattering source' is not necessarily always the one you might intuitively expect. An interesting option for future research would be to model the influence of noise on the various bounding methods so as to produce bounds with a specified statistical precision.

Figs. 7.24c and d give bounds on phase perturbation calculated from $\arg(\chi)$ and $|\chi - 1|$ respectively. These can be compared against the calculated phase perturbation in Fig. 7.24b. In this particular physical configuration, the $\arg(\chi)$ based method appears to give significantly tighter bounds, however as mentioned earlier, which bound is best can vary depending on the phase and intensity relationship between the backscattering sources within a pixel. Fig. 7.24f demonstrates a bound on relative intensity, b, which appears to be extremely tight. In general, minimum bounds on b are extremely useful for the detection of multiple returns. Table 7.2 gives the mean values calculated by different bounding methods for each of the

[1] As noted previously, when normalising a measurement at a relative frequency of r_1 by a measurement at a relative frequency of r_0, a single backscattering source results in $\chi = \sqrt[r]{1}$.

[2] This is only possible because this is a non-differential system, therefore the mean intensity of the correlation waveform includes both ambient light and the correlated active illumination. In a differential system, while the ambient light increases the noise level, it is difficult to measure because the differential process removes the offset from the raw phase steps. As a result, even with knowledge of the photon transfer curve it is difficult to assign an exact noise level to a pixel without taking statistics over a series of frames.

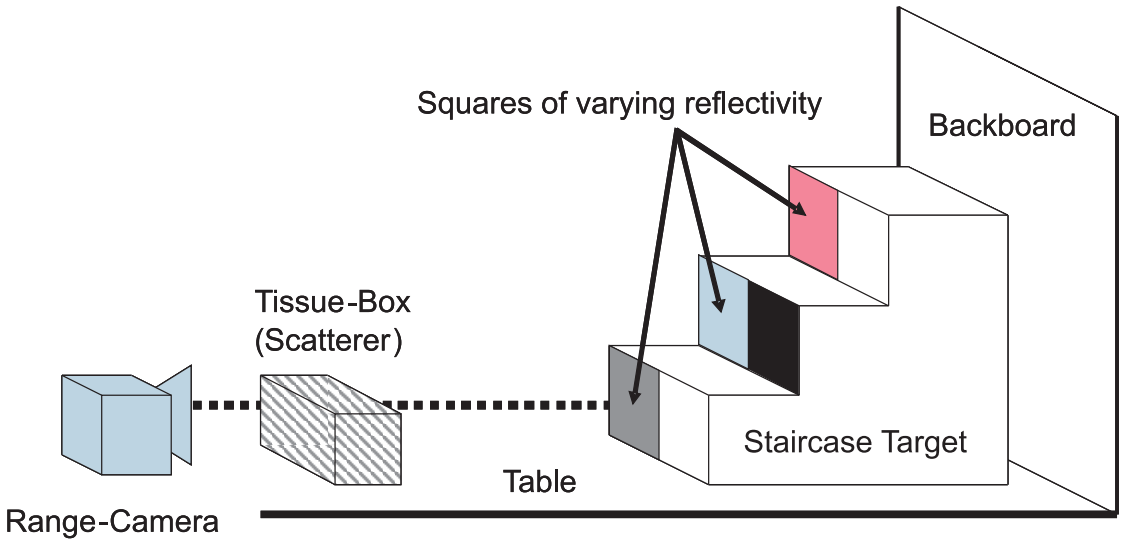

Figure 7.22: Layout of scene four

labelled squares as well as reference values.

Fig. 7.25 plots a horizontal slice through the phase data before and after processing as well as a slice through the reference image; the recovered processed data is found to match the reference data extremely closely.

Scene five was captured using the Canesta XZ-422 at modulation frequencies of 36/18MHz. Unlike the other systems, the Canesta camera uses off-axis illumination, which results in shadows in the scene. Fig. 7.26b shows the modulus of the unprocessed scene, clearly showing the deleterious impacts of this illumination approach. Figs. 7.26a and 7.26c show the phase before and after processing. On the left is a near-black sheet suspended from the ceiling; in the unprocessed data the shape of the sheet is clearly defined, albeit inaccurately, due to non-Lambertian reflectance. As a result some regions of the sheet appear darker and are more significantly perturbed by scattering within the camera optics. The processed data results in the removal of mixed pixels around the edge of the wheel, but leaves artefacts on the right side – corresponding to the region with occluded illumination. Fig. 7.26 gives the estimated relative intensity of the second component return relative to the primary. This clearly detects mixed pixels and dark objects, like the black sheet, which have been perturbed by scattered light. Fig. 7.27 gives a plot of the phase of a slice through the scene, showing separation of mixed pixels.

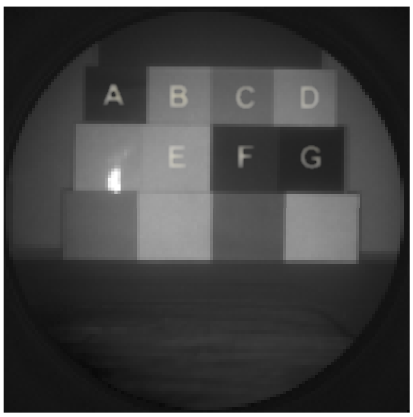

(a) Unperturbed Modulus (labels added in postprocessing, gamma compressed, $\gamma = 0.4$)

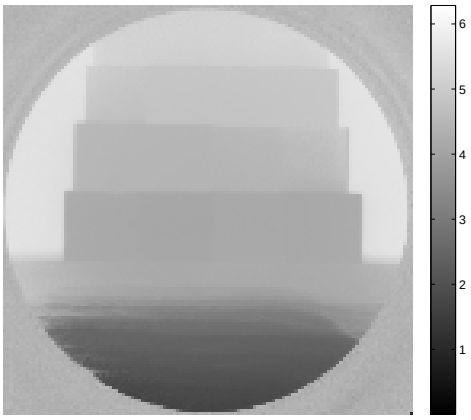

(b) Unperturbed Phase

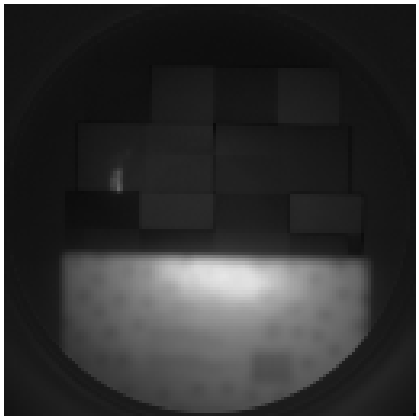

(c) Perturbed Modulus (gamma compressed, $\gamma = 0.4$)

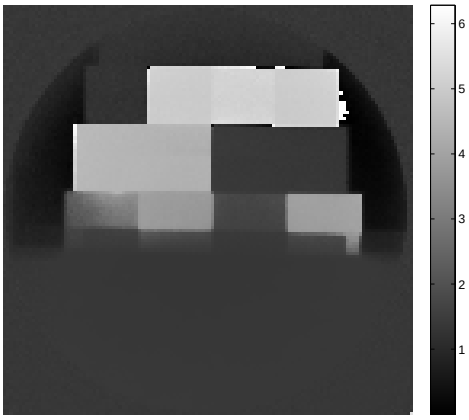

(d) Perturbed Phase

(e) Phase of Brighter Separated Component

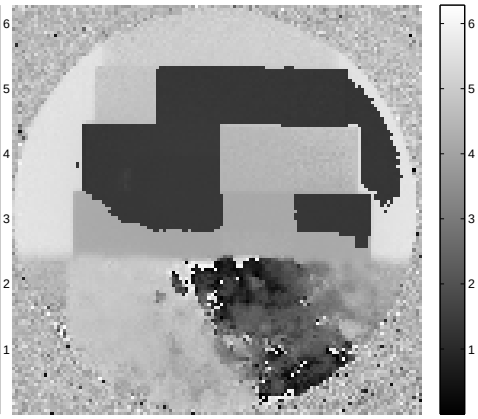

(f) Phase of Darker Separated Component

Figure 7.23: Scene four: mixed pixel separation on real data in the case of extreme intra-camera scattering against ground-truth.

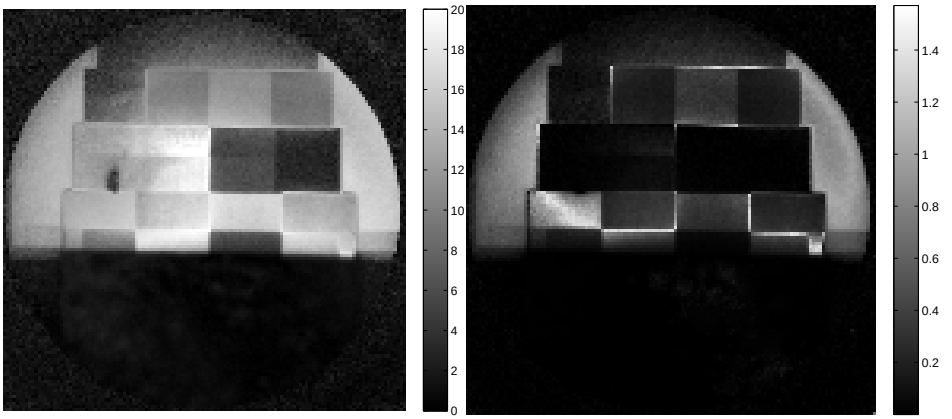

(a) Mahalanobis Distance From $\chi = 1$ (b) Calculated Phase Perturbation, $|\theta_{\mathrm{f}}|$

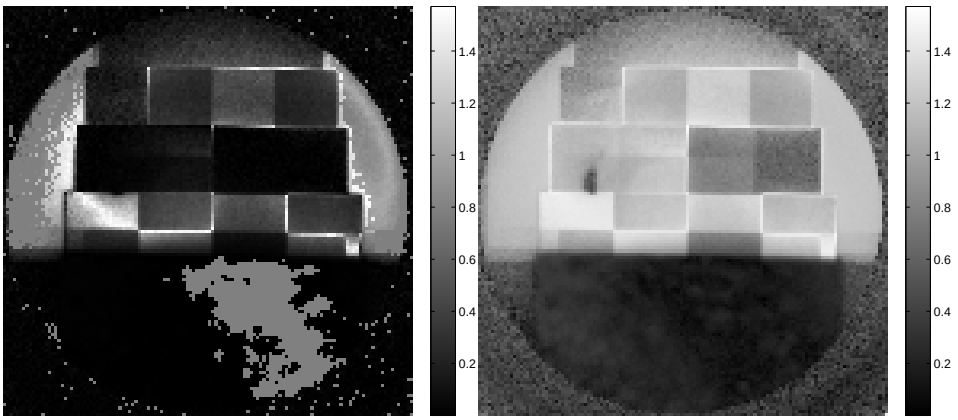

(c) $\max(|\theta_{\mathrm{f}}|)$ from $\arg(\chi)$ (d) $\max(|\theta_{\mathrm{f}}|)$ from $|\chi - 1|$

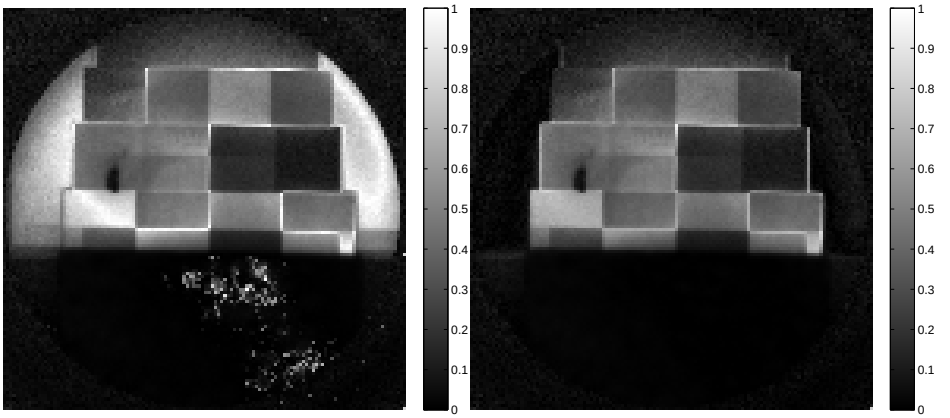

(e) Calculated Relative Intensity, b (f) $\min(b)$ from $|\chi|$

Figure 7.24: Important statistics and bounds on scene four (Fig. 7.23.)

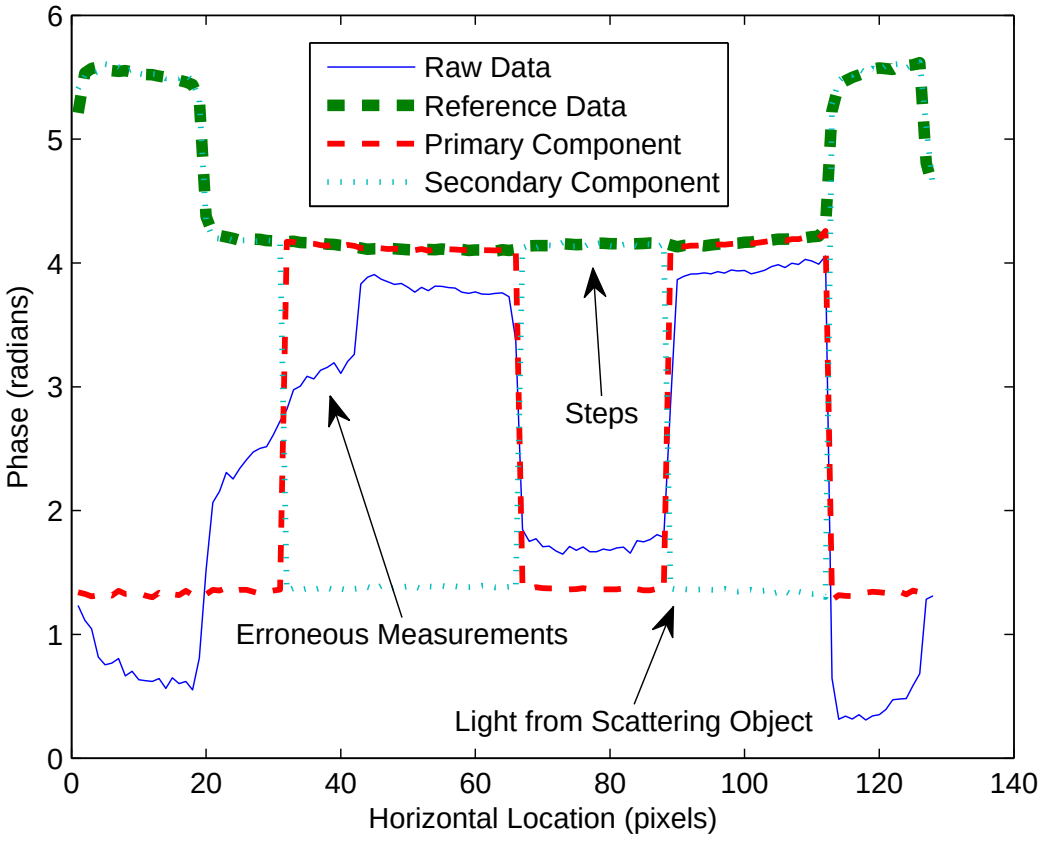

(a) Scene Four

Figure 7.25: A slice through scene four, from Fig. 7.23

7.3 Bounding Methods

In Chapter 6 we introduced a large number of bounds on relative intensity and phase perturbation among other parameters. We now present a brief overview of those bounds and an analysis of their performance and noise sensitivity using the model described in Section 7.1. While we have developed a small number of bounds which are applicable to frequency ratios other than 2 : 1, we only evaluate bounds for this particular ratio.

Applying the bounds from Table 7.3 to noiseless Monte-Carlo generated data gives the Cumulative Distribution Functions from Fig. 7.28. For each descriptive parameter, the first curve corresponds to a reference bound. Each reference bound is the least-constraining, physically possible bound and represents fundamental limitations on the parameter values. Positive slack corresponds to the amount by which a particular parameter is underconstrained, zero slack corresponds to an exact estimate of a parameter and negative slack corresponds to the amount of error

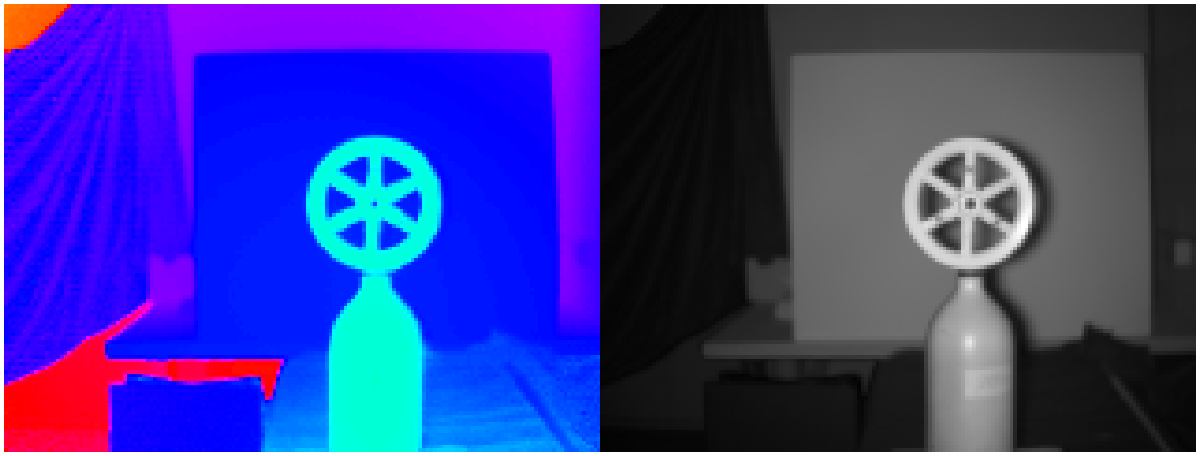

(a) Scene Five Raw Image Phase

(b) Scene Five Raw Image Modulus

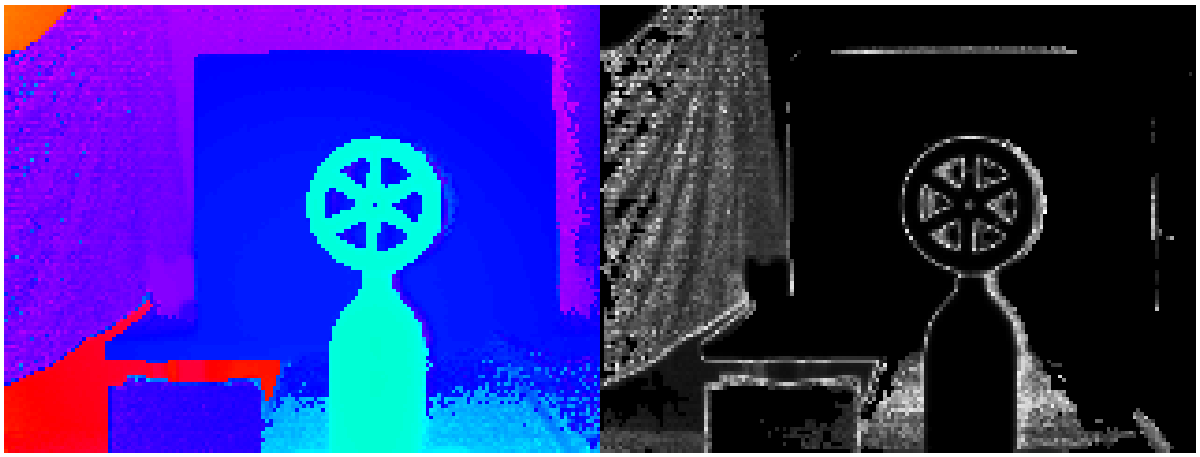

(c) Scene Five Primary Phase

(d) Scene Five Estimated Relative Intensity (b)

(e) Phase Colour Scale

Figure 7.26: Scene five before and after processing

in a bound. In order to fully understand the performance of the bounding methods we also need to understand noise performance: Table 7.4 provides insight into the impact of noise on each particular bound by providing the percentage of erroneous/overconstrained bounds as a function of the Signal-To-Noise ratio as well as the mean error.

A minimum bound on relative intensity provides an extremely useful indicator of the degree of mixing of multiple measurements. For a given $\min(b)$, it is guaranteed that at least $100\% * \min(b)/(1+\min(b))$ of the measured light is being contributed by a second backscattering source. One useful approach to identification of perturbed points is to set a threshold for $\min(b)$ and remove all points for which this threshold is exceeded. It can be seen from Fig. 7.28a that the three different bounds on $\min(b)$

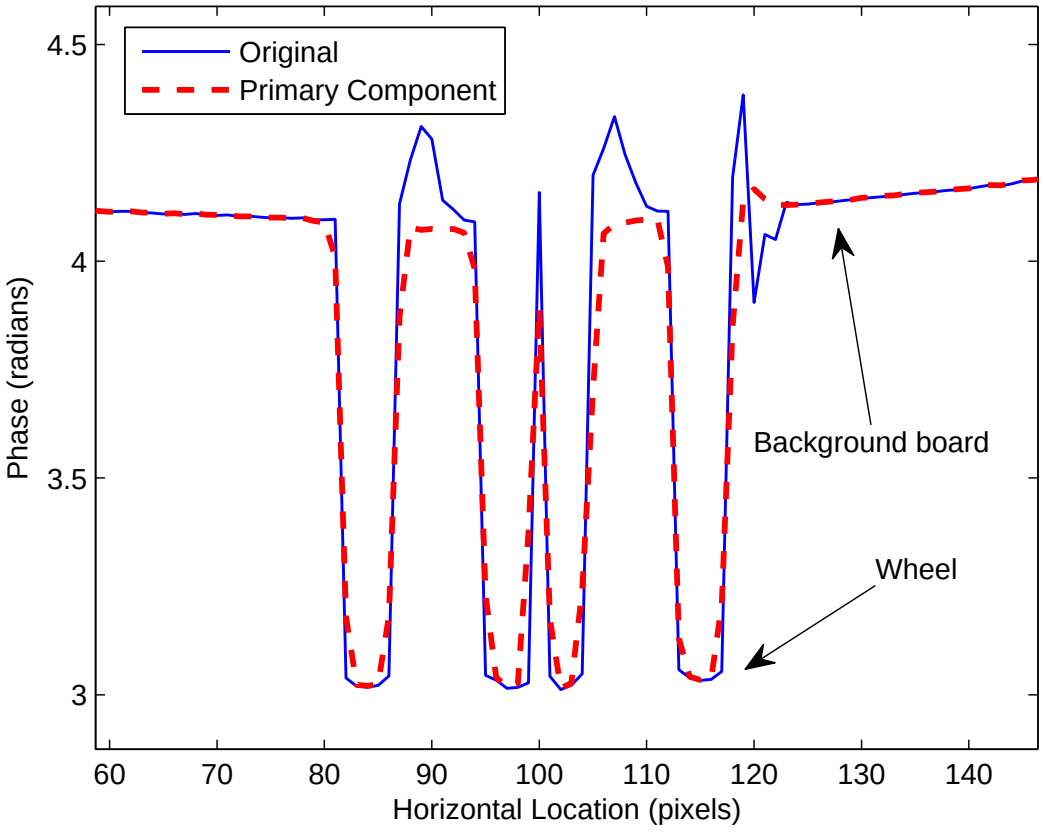

Figure 7.27: A slice through scene five, from Fig. 7.26

all result in a moderate improvement over the reference bound, this is illustrated by the CDF curves being above and to the left of the reference CDF curve. In this particular case all the bounds appear to have fairly similar performance, although the bounds do not appear to be particularly tight. From Table 7.4 it can be seen that the tightest bound in the noiseless case (from τ_0) is not necessarily the bound which is most noise sensitive (either $\arg(\chi)$ or $|\chi|$ depending on the SNR).

A minimum bound on relative phase is of less use than one on relative intensity, but still gives useful information about the spatial relationship between the component backscattering sources. Fig. 7.28b is a plot of the CDF of phase slack in bounds on $\min(\theta)$: the graph indicates that the $\arg(\chi)$ derived bounding method produces only a mediocre improvement over the reference bound, but the τ_0 bounding approach generates very tight bounds much of the time (albeit with a long tail). However, in circumstances with low or moderate SNR, the latter bound tends to produce a significant number of erroneous bound estimates. For example, at an SNR of 30 dB, 45.3% of estimated bounds are erroneous, with a mean phase error of 155 mRad. It is not unreasonable to expect that the tightest bounds tend to be

Table 7.3: Selected bounds on phase perturbation, relative phase and relative amplitude for one or two backscattering components and two measurements (freq. ratio of two-to-one).

Max. Abs. Value of Phase Perturbation of Brightest Component, $\max(\lvert\theta_f\rvert)$		
Details	Definition	Bound Summary
From τ_0	Eqn. 6.136	$\lvert\theta_f\rvert \leq \cos^{-1}(\tau_0)$
From $\arg(\chi)$	Eqn. 6.117	For $\lvert\chi\rvert \leq 1$, $\lvert\theta_f\rvert \leq \max(\pi/4, \arg(\chi)/3)$
	Eqn. 6.102	For $\lvert\chi\rvert > 1$, $\lvert\theta_f\rvert \leq \arg(\chi)/2$
From $\lvert\chi - 1\rvert$	Eqn. 6.150	$\lvert\theta_f\rvert \leq \sec^{-1}(1 + \lvert\chi - 1\rvert)$
From $\arg(\chi - 1)$	Eqn. 6.172	$\lvert\theta_f\rvert \leq \lvert\arg(\chi - 1)\rvert/2$
From $\lvert\chi\rvert$	Eqn. 6.80	$\lvert\theta_f\rvert \leq \frac{1}{2}\cos^{-1}\left(\frac{1}{4}\left(\lvert\chi\rvert^2 - \sqrt{\lvert\chi\rvert^4 + 8\lvert\chi\rvert^2}\right)\right)$
Min. Abs. Value of Relative Phase Between Components $\min(\lvert\theta\rvert)$		
From τ_0	Eqn. 6.160	$\lvert\theta\rvert \geq 2\cos^{-1}(\tau_0)$
From $\arg(\chi)$	Eqn. 6.89	$\lvert\theta\rvert \geq \lvert\arg(\chi)\rvert/3$
Min. Relative Amplitude of Darker Component, $\min(b)$		
From τ_0	Eqn. 6.135	$b \geq (1 - \tau_0)/(1 + \tau_0)$
From $\arg(\chi)$	Eqn. 6.131	$b \geq \sin(\lvert\arg(\chi)\rvert/3)$
From $\lvert\chi\rvert$	Eqn. 6.128	For $\lvert\chi\rvert \leq 1$, $b \geq (1 - \sqrt{2\lvert\chi\rvert - \lvert\chi\rvert^2})/(1 - \lvert\chi\rvert)$
	Eqn. 6.128	For $\lvert\chi\rvert > 1$, $b \geq (\lvert\chi\rvert - 1)/(\lvert\chi\rvert + 1)$

more greatly impacted by noise; the latter particular bound is a perfect example in this respect.

The remaining bounding methods we address here relate to maximum bounds on phase perturbation, $|\theta_f|$. Fig. 7.28c and d are plots of the CDFs of bound slack at frequency ratios of two-to-one and five-to-one. There are three methods that stand out as being particularly useful at two-to-one: using τ_0, $\arg(\chi)$ and $|\chi - 1|$. Of these, the $\arg(\chi)$ based method is the most noise sensitive, although the size of the bounding error is generally quite small. While the attenuation ratio based method produces the best results, it requires a different modulation frequency from that typically required for the characteristic measurement: if a capture is limited to only two different modulation frequencies it is probably better to determine the characteristic measurement rather than the attenuation ratio because it allows inversion and a greater variety of bounds, even if the bounds we have derived are not individually quite as tight. Ultimately, the best bounding approach is likely to be inversion, followed by placing noise derived bounds on the separated component sources. In this respect, there are lots of possible approaches to approximating the inverse which we have not investigated. Fig. 7.28c also includes a combined bound, which provides far better performance than any of the individual bounding approaches: the median slack is 99 mRad and the worst case is only 520 mRad. This implies that for a situation where one bound is particularly loose, there is often another much tighter

Table 7.4: Fraction of invalid bounds at a frequency ratio of two-to-one as a function of SNR

SNR	10 dB	20 dB	30 dB	40 dB	50 dB	∞
Bound on $\max(\lvert\theta_f\rvert)$	*Percentage Invalid Bounds/Mean Error of Invalid Bounds (mRad)*					
From τ_0	0.4/40	0.1/15	0.05/5	0.02/1	0.01/0	0/0
From $\arg(\chi)$	15.7/157	9.3/77	4.9/39	2.6/21	1.4/11	0/0
From $\lvert\chi - 1\rvert$	0.3/56	0.1/37	0.06/23	0.03/13	0.02/7	0/0
From $\arg(\chi - 1)$	5.2/236	2.3/148	0.9/83	0.4/45	0.01/25	0/0
From $\lvert\chi\rvert$	1.0/49	0.4/16	0.1/5	0.05/2	0.02/0	0/0
Bound on $\min(\lvert\theta\rvert)$	*Percentage Invalid Bounds/Mean Error of Invalid Bounds (mRad)*					
From τ_0	68.8/626	58.0/325	45.3/155	33.6/73	23.8/34	0/0
From $\arg(\chi)$	2.6/60	0.8/19	0.2/6	0.08/2	0.03/1	0/0
Bound on $\min(b)$	*Percentage Invalid Bounds/Mean Error of Invalid Bounds (relative intensity* $\in [0, 1]$*)*					
From τ_0	28.9/0.104	17.3/0.037	10.0/0.012	5.7/0.004	3.2/0.001	0/0
From $\arg(\chi)$	13.6/0.073	6.7/0.024	2.7/0.007	0.7/0.002	0.2/0.001	0/0
From $\lvert\chi\rvert$	13.1/0.073	6.1/0.027	3.1/0.009	1.6 / 0.003	0.9/0.001	0/0

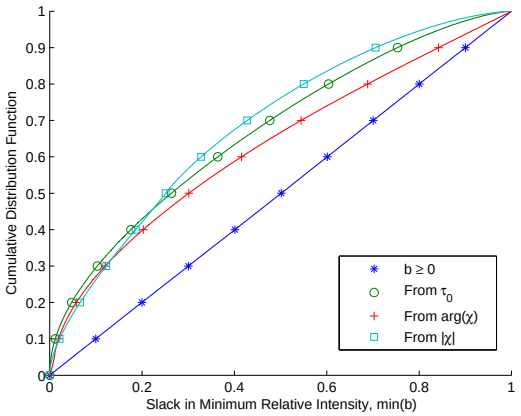

(a) Min b at 2:1

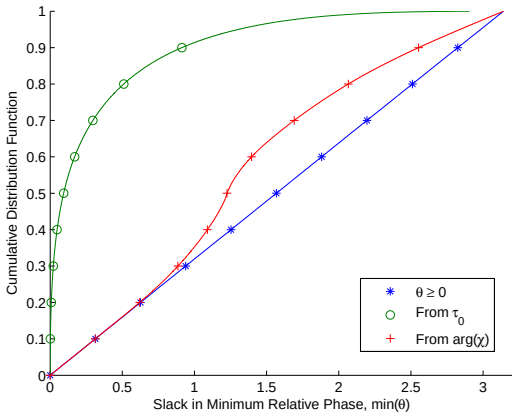

(b) Min θ at 2:1

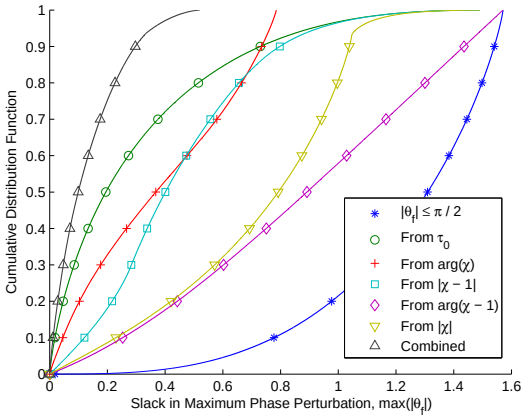

(c) Max $|\theta_f|$ at 2:1

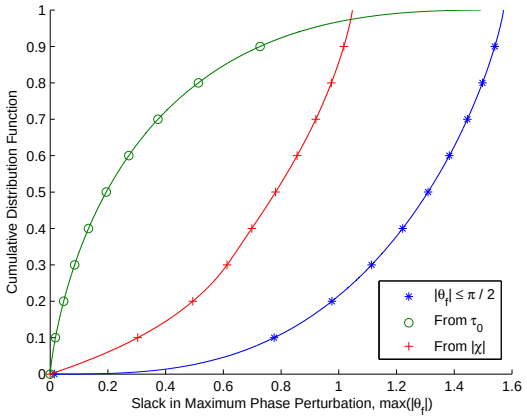

(d) Max $|\theta_f|$ at 5:1

Figure 7.28: Cumulative Distribution Functions of excess slack in bounds on parameters of the component backscattering sources at two-to-one and five-to-one by bounding method, via Monte-Carlo simulation with no added noise.

bound; this is assisted by the fact that $\arg(\chi)$ and $|\chi|$, and $\arg(\chi - 1)$ and $|\chi - 1|$ are orthogonal.

7.3.1 Summary

In this chapter we demonstrated the operation of mixed pixel bounding and separation methods using the characteristic measurement. A framework was developed for testing the lookup table separation algorithm using simulated data and then the algorithms were tested on both simulated and real data. It was determined that the ability to separate out two backscattering sources was highly dependent upon the relative phase and relative intensity of the sources, as the well as the SNR. In particular, for a given SNR if two sources are closer than a specified threshold, they

cannot be reliably resolved. We showed that the algorithm trades-off worse precision for better accuracy, although the trade-off can be altered by averaging after processing. This means that the algorithm may not be suited for all applications. Other important results include the reduction in accuracy present at certain frequency ratios, in particular when both frequencies are odd or both frequencies are even. As a general rule, frequency ratios of $n:(n-1)$ appear best, although they still result in a significant number of different solutions. Choosing between these solutions is problematic without further information or the imposition of other spatial or temporal constraints. Analysis of the noise impact of lookup table resolution showed that resolution has no measurable effect on precision, but a very significant effect on accuracy. In particular, for the translated polar table at two-to-one, it was determined that the resolution of the $|\chi - 1|$ dimension was significantly more important than the resolution of the $\arg(\chi - 1)$ dimension, by at least a factor of four. Using the lookup table method on data actually containing three component returns, rather than the two actually assumed by the restoration algorithm, was found to generally result in a decrease in the overall error present in the primary component, but with exact performance depending on the phase and intensity relationship between the second and third brightest components.

In testing on real-data, the algorithm demonstrated the separation of two components within a pixel in multiple different circumstances, including multipath interference due to intra-scene multipath and intra-camera scattering. In limited testing against ground-truth data, it was found that the separation algorithm did not appear to make values significantly worse; either values improved significantly (decreasing error by $> 90\%$), or if unable to separate out the constituents, the estimates were only slightly worse than unprocessed data ($< 5\%$). Unfortunately, while single, specific figures like these are nice (and impressive), they tend to be unrepresentative because the exact performance is dependent upon the phase and intensity relationship and the error distribution is particularly long-tailed.

(It is worth noting that at the time of post-oral examination revision, this document is significantly out-of-date and has been surpassed in many ways by later work. Necessarily, at some point a line in the sand has to be drawn, even if the evaluation in this chapter is very incomplete.)

Chapter 8

Conclusions and Further Work

The initial topic of research was defocus restoration for full-field AMCW lidar systems, however as the project progressed it became more and more obvious that solving the mixed pixel problem was at the heart of the defocus problem. Given the paucity of research into the mixed pixel problem, there have been many avenues to explore and the most fertile ground for future work is likely to be deepening the analyses of these avenues. As there is so much ground that has been covered but briefly, it is perhaps easiest to distinguish those avenues than can be considered dead ends and are not suited to further investigation. The correlation waveform shape fitting methods from Chapter 4 and the scene texture based methods from Appendix C both fall into that category. On the other hand, the majority of other material warrants further investigation, including all the material relating to defocus restoration and mixed pixel problem inversion methods using measurements at multiple different modulation frequencies.

This thesis has presented a large variety of new methods, including algorithms and approaches for mixed pixel separation and defocus restoration of full-field AMCW lidar images and new analyses of systematic error. A large number of different approaches have been presented to the mixed pixel problem and the most commercially viable algorithm – the characteristic measurement lookup table – has been subjected to detailed analysis of its noise properties and demonstrated on real-world data. The results demonstrate how the algorithm trades-off precision for accuracy, resulting in an order of magnitude improvement in ranging accuracy. As a result of this success, the implementation is now the subject of a provisional patent and possible commercialisation. A large amount of additional work remains to characterise the remaining new methods, as realistically, there was no way to fully examine these methods within the available time constraints.

On reflection, this thesis could very well have been entitled ‘Tenuously connected

investigations into AMCW lidar.' If there is one fault, it is that this thesis chases too many disparate ideas and does not analyse each one in enough detail. While investigations – such as that into the systematic perturbations from the second harmonic of the correlation waveform – are novel, they are fairly disconnected from the mixed pixel problem, which forms the primary focus of this document. I guess to a certain extent, this can be viewed as either an advantage or a disadvantage: there is a surfeit of ideas presented, but there are a lot of loose ends. In this respect, the thesis is a reflection of an, as yet, very incomplete journey.

In the remainder of this chapter, we address some of options there are for future research and discuss remaining issues with the research presented in this document.

8.1 Mixed Pixel Restoration

Chapter 6 introduced the new concept of the characteristic measurement, which models the relationship between two AMCW range measurements at different modulation frequencies. A large number of mathematical properties and bounding methods were derived from this concept for the two measurement, two component return case, although there was no success in finding a closed form inverse for the problem. Attempting to derive such an inverse remains an interesting and challenging problem as it is difficult to be certain whether such an inverse exists.

Without an inverse, it was decided to use a lookup table approximation to the inverse, derived using numerical methods. An advantage of this approach is that given sufficiently high sampling density, the mathematical properties of the LUT method are identical to those of an exact closed-form inverse, if it exists. A detailed analysis of the noise properties of the two measurement, two component return case was presented in Chapter 7 and the method was shown to operate on real data at a $2:1$ frequency ratio. A journal paper is currently being prepared based on this material.

While no closed-form inverse was found, a significant number of bounding methods were found and tested in Chapter 7. Unsurprisingly, it was found that the tightest bounding methods tended to be the most noise sensitive. Possible future research could include better analysing the noise properties of these bounds so as to put a statistical significance on each figure. One of the biggest remaining challenges is determination of the correct component return for frequency ratios other than $3:1$ and $2:1$. While there are options for heuristic methods, which are likely to be faster, the best option appears likely to be equivalent to solving some sort of Markov Random Field.

In Chapter 5 we developed two new models for measurements of mixed pixels: one using attenuation ratios and the other using Cauchy distributions to model diffuse-range returns. Both of these methods give relatively simple closed-form inverses and are suited to real-time implementation. One of the biggest remaining loose ends is a detailed analysis of the noise properties of each of these algorithms; while they are both known to work, this thesis does not present a detailed analysis of the phase, amplitude and SNR relationships comparable to that of the characteristic measurement LUT approach in Chapter 7. Completing and publishing such an analysis is one of the highest priorities for continuing research.

It was deliberately chosen to prioritise research into the characteristic measurement LUT and bounding approaches because they had the greatest potential for commercial application, ultimately leading to a provisional patent. Chapter 7 demonstrates that even assuming a noiseless measurement of total integrated intensity, the attenuation ratio method underperforms the characteristic measurement LUT approach. The Cauchy distribution method suffers from a different malaise; requiring four measurements at different modulation frequencies severely restricts the method's application to static scenes. In addition, while mathematically elegant, a Cauchy distribution is likely to be a pretty poor model for diffuse-range component returns, thus only really suited to solving for returns that are well modelled by Dirac deltas.

8.2 Defocus Restoration and Miscellanea

The choice to place the material on defocus in the appendices was made largely due to time limitations and the large amount of rework required to fully integrate the material with body of the thesis. Given that defocus was the original topic of investigation, it appears appropriate to include the material as part of a record of the research performed. Given how small a subset of the image processing community works with range-images, there is a huge untapped potential to apply computational imaging techniques developed for intensity images to range-images.

Appendix A presents a new algorithm for the spatially variant blind deconvolution of a sequence of two defocussed range images, each one taken with a different aperture stop. While the concept and algorithm is novel, unfortunately the execution was flawed. While the focal parameter determination worked relatively well, operating far better than either of the depth-from-defocus methods, the actual deconvolution was stymied due to the image intensifier being knocked out of alignment since it was last flat-fielded. While more recent CMOS sensor designs do not suf-

fer from this problem, they also have much lower resolutions – partly negating the necessity for defocus restoration in the first place. One additional reason for performance is likely to be the pillbox model for defocus, which did not fully model the PSF.

Appendix B follows on from Appendix A; whereas the latter attempted to blindly determine the in-focus distance and psf scale, Appendix B uses a fixed, known PSF that has been altered using a coded-aperture into order to make the PSF more broadband, thus more easily deconvolved. Whereas the PSF in Appendix A heavily attenuated all high spatial frequencies, the PSF in Appendix B was designed to at least partially preserve these spatial frequencies and thus make deconvolution more viable. By using the blurred range information from a defocussed range-image to estimate PSF scale, this chapter demonstrated the successful deconvolution of two range images. The algorithm was also shown to fail in another case, where very dark objects in the scene were perturbed by scattered light to the extent that the initial PSF was grossly misestimated. One particularly interesting idea for future research would be to attempt to combine the mixed pixel algorithms with the deconvolution operation, and thus partially remove the perturbing effect of scattered light and misestimated initial range. To a limited extent, this sort of approach was applied to one scene in Chapter 7, where an isoplanatic deconvolution was applied in an attempt to partially mitigate scattered light before operating the mixed pixel algorithm.

There is a particularly large gamut of possible related hardware developments related to defocus; while we investigated coded apertures further plausible developments include plenoptic range-cameras, which would be able to reconstruct range-images which are perfectly in-focus everywhere. One interesting application would be to use the range-information to speed up the determination of the scaling parameter in Plenoptic 2.0 (Lumsdaine, 2008) using the implicit range information; at present, lacking this information, huge computational resources are required to determine it. Other techniques that might be of value include wavefront coding (Dowski and Johnson, 1999), which effectively makes depth-of-field isoplanatic.

Chapter 3 developed a model of measurement formation. While based off previous work in the literature, this model introduced new ideas such as systematic perturbations from uneven phase steps and systematic errors due to the presence of a second harmonic in the correlation waveform. The chapter also developed additional mathematics describing the behaviour of aliasing. In particular, the treatment of uneven phase steps is relatively shallow and deserving of deeper analysis.

8.3 Dead Ends

Chapter 4 presented several approaches to separating out components within mixed pixels using correlation waveform harmonic content. The fundamental problem with correlation waveform shape based methods is that sampling the correlation waveform is typically an extremely inefficient method of sampling the higher spatial frequencies of the signal return model; it is far more efficient to take discrete measurements at different modulation frequencies. While it is theoretically possible to apply the shape fitting methods to data synthesised from discrete measurements at different modulation frequencies, the closed-form (and not-so-closed form) methods of Chapters 5 and 6 are defined in a manner that allows easier mathematical analysis and faster implementation. Ultimately, while Chapter 4 demonstrated that correlation waveform shape fitting was possible, it also demonstrated the challenges – such as local minima and slow execution times – that make it impractical for use in a commercial product.

The scene texture approach from Appendix C has some fairly obvious limitations that make it unviable: in particular, the requirement for scene texture. While Falie (2008) proposed using pattern projection to achieve something similar, no-one has yet to actually implement such a system. At best, the method is a novelty.

Appendix A

Blind Determination of Focal Parameters

A.1 Introduction

Full-field range imagers work on the time-of-flight (TOF) principle, acquiring range data for every pixel simultaneously. Whereas in the past imaging lidar devices have used actuated point sensors acquiring data sequentially (Blais, 2003), the use of full-field sensors (Gokturk *et al.*, 2004; Ringbeck and Hagebeuker, 2007; Dorrington *et al.*, 2007) allows much simpler mechanical design and mass production leading to a reduction in costs. However the use of full-field sensors comes at a cost; many new problem are created such as range-intensity coupling and finite depth-of-field (DOF). Mixed pixels, also known as multipath interference, occur in full-field Amplitude Modulated Continuous Wave rangers and are caused by a pixel integrating light from more than one source (Hebert and Krotkov, 1992; Adams and Probert, 1996). This can be due to light scattering within an optical system, defocus blur or unavoidably around the edges of objects. As a result the measured ranges are erroneous, often appearing to have little or no connection to the range to the component light sources. While a number of correction/separation algorithms have been proposed (Godbaz *et al.*, 2009b; Larkins *et al.*, 2009; Godbaz *et al.*, 2009a), in general mixed pixels remain a fundamentally unsolved problem.

As the resolution of sensors and the necessity for system flexibility increases, DOF issues become problematic. Firstly there are calibration issues – commercial full-field lidar sensors come with fixed focal length wide aperture lenses because if focal length were to change the calibration required to produce accurate point clouds would be invalidated. Additionally, wide apertures result in greater ranging precision due to an increase in the signal-to-noise ratio. This means that when

imaging an object which is either very close to the camera or a long way away, the range image is out of focus. In practice this results in a blurry amplitude image and the formation of a wide band of mixed pixels around the edges of objects; thus much of the recorded range data is unreliable. Secondly, even if it were possible to change focus all optical systems suffer from finite DOF – meaning that certain scenes will always suffer from some defocus. One way to mitigate this is to take multiple images at different focal lengths and combining the most in-focus regions of each (Goshtasby, 2006), instead in this paper we choose to approach this problem by deconvolving range images to extend the DOF. In order to make our algorithm as flexible as possible we determine the defocus parameters of the optical system from the captured range data – thereby allowing an operator to freely refocus the range imager and even change the lens without needing to resample the point spread function (PSF). While we do not deal with other aspects of automatic calibration, we believe that this level of flexibility is going to become more important as full-field lidar technology develops further.

A.1.1 Previous Work

Previous work on single frame blind deconvolution has primarily focussed on the isoplanatic case (Kundur and Hatzinakos, 1996); this is largely because the problem is ill-posed and it is difficult to quantify how blur changes across a single isolated image. Single image spatially variant deblurring has generally involved complex blur calibration rather than blind deconvolution (Boden *et al.*, 1996). In other applications blurring issues have been avoided by use of special hardware such as adaptive optics for telescopes (Hardy, 1998). In the case of limited DOF, high quality restoration requires implicit range information of some sort.

The use of image sequences opens up a variety of new possibilities. For example Yuan *et al.* (2007) used an isoplanatically blurred image and a noisy image of the same scene to produce an unblurred denoised image and an estimate of the PSF. Other work on atmospheric turbulence has involved either breaking anisoplanatic images into smaller isoplanatic regions or dewarping the images before deconvolution in order to make the problem tractable (Fraser *et al.*, 2003), other work has included the deblurring of range-gated lidar (MacDonald and Cain, 2006). In a similar manner, depth from defocus (DFD) methods measure the change in blurring across a pair of images at different known focal settings in order to determine range to objects in a scene. In this paper we analyse the spectral method of Pentland (1987) and the spatial S-Transform method of (Subbarao and Surya, 1994). Once the amount of blurring at each pixel across an image has been determined, the de-

convolution problem becomes much simpler. Many other DFD methods have been proposed Subbarao (1988); Rajagopalan and Chaudhuri (1997); Chaudhuri and Rajagopalan (1999): one of the most promising is the use of coded apertures to make the deconvolution problem more well-posed (Levin *et al.*, 2007).

There are several aspects which all previous DFD methods have in common: firstly they process positive intensity images, they produce range estimates from blur by a complex camera calibration process and they require high levels of scene texture/patterning in order to produce correct blur – thus range – estimates. In comparison, the method proposed in this paper uses range images in the complex domain, determines focal parameters from a combination of range and blur data, and can produce blur estimates in areas with no texture as long as there is adequate texture in other regions of the image.

A.1.2 Overview

We present a new spatially variant parametric two-stage blind deconvolution algorithm. The first stage consists of a Maximum Likelihood based method for determining the focal parameters of a range imager optical system, including the amount of blurring and the distance at which objects are in-focus, from pairs of range images taken at different aperture settings. The algorithm works by reblurring an image taken with a smaller aperture to match an image taken with a large aperture using range data. Whereas previous DFD algorithms have been able to determine the amount of blur at each pixel and thus range from intensity image pairs, since full-field lidar systems also produce range data it is possible to dynamically determine how defocus blur varies with range. Since images taken at small aperture settings are still limited by DOF, we then estimate the amount of blur at each pixel and apply spatially variant deconvolution to extend DOF.

A.1.3 Limitations of the Range Imager

The examples in the paper were generated using the University of Waikato (UoW) (Dorrington *et al.*, 2007) range imager. The UoW range imager uses the time-of-flight (TOF) principle to produce 2 dimensional matrices of range, amplitude and mean signal intensity data. A scene is illuminated with modulated laser light and the system measures the TOF induced phase shift in the illumination modulation, thus the distance to the objects in the scene. While commercial systems use modulated CMOS sensors (Gokturk *et al.*, 2004), modulated image intensifiers currently offer higher spatial resolutions. The UoW range imager has a resolution of 512x512

whereas the highest resolution modulated CMOS sensor has a resolution of 204x204 pixels (PMD PhotonICs®41k-S). In a heterodyne system, either the sensor or an image-intensifier is gain modulated with a signal at a very slightly different frequency from the illumination modulation. This results in the formation of a much lower frequency beat waveform which can be recorded over time by a CCD camera or CMOS device; this beat waveform shape is a correlation of the illumination modulation waveform with the sensor/intensifier response waveform. By taking a temporal Fourier transform of the data, it is possible to determine the phase and amplitude of the beat waveform for each pixel, which is the raw range data we concentrate on in this paper.

A.2 Model

A.2.1 Range Data

We represent the ideal, unblurred range data by a vector $\mathrm{H} = (\eta_0, \ldots, \eta_{n-1})$, where n is the number of pixels in the 2D image and $\eta_i \in \mathbb{C}$ is the the ith pixel. The range datum η_i is formed by

$$\eta_i = a_i \mathrm{e}^{4\pi j r_i f_m / c}, \tag{A.1}$$

where $a_i \in \mathbb{R}^+$ is the amplitude, $r_i \in \mathbb{R}^+$ is the range to the object, f_m is the illumination modulation frequency, c is the speed of light. For convenience, we define a function $\mathrm{H} = \pi_\Lambda(a, r)$ that produces range data from vectors of amplitude and range data respectively, which we utilise in section A.4.1.

If we consider a model of active signal returns versus range for a pixel, $f_i(r) \in \mathbb{R}^+$, we can regard a measurement at a particular modulation frequency as sampling a particular spatial frequency of the model of signal returns i.e. $\eta_i = F_i(-c/2f_m)$. Conventional processing estimates amplitude and range by

$$a_i = |\xi_i| = |\eta_i + \epsilon_i + \lambda_c|, \tag{A.2}$$

$$r_i = \tfrac{c}{4\pi f_m} \arg(\xi_i) = \tfrac{c}{4\pi f_m} \arg(\eta_i + \epsilon_i + \lambda_c), \tag{A.3}$$

where ξ_i is the recorded noisy data at pixel i, ϵ_i is noise and λ_c is light from other objects – which results in mixed pixels. Lenses, image intensifiers and reflective sensors tend to result in light scattering within full-field range imagers. This is particularly important if an image has a high dynamic range as range-intensity coupling can result. The resulting errors are one of the dirty secrets of the range-imaging community – it means that range measurements change depending on reflectivity.

A flat board can appear to be spread across multiple ranges because it is patterned. We now model the effect that limited depth of field has on range data.

A.2.2 Depth of Field

The thin lens equation is

$$\frac{1}{f_{\mathrm{L}}} = \frac{1}{d_o} + \frac{1}{d_i}, \tag{A.4}$$

where f_{L} is the focal length, d_o is the distance to the object projected onto the optical axis and d_i is the distance to the image, similarly projected. Since in most cameras the distance to the imaging plane is fixed, this means that the distance of the correctly formed image from the imaging plane varies across an image depending on the distance along the optical axis of camera to the object. While the focal length can be adjusted, it is impossible for multiple objects at different discrete ranges to all be in perfect focus with a real physical camera, although a single flat object at an angle can be imaged according to the Scheimpflug principle (Scheimpflug, 1904). The properties of defocus blur are well understood; data subject to limited depth-of-field can be regarded as being convolved with a spatially variant point spread function.

There are two commonly used theoretical models for defocus blur: the Gaussian and pillbox models. The Gaussian model is more appropriate for systems where diffraction/aberration effects dominate, particularly at small aperture sizes. The pillbox model is more representative of systems where the aperture size is larger and diffraction effects are less significant. Whether the illumination is polychromatic or monochromatic can also impact on the suitability of a particular model as the superimposed diffraction patterns at different wavelengths in white light can be approximated by a Gaussian. In general, the PSF of the UoW range imager is best approximated by a pillbox model. While most previous work has focussed on Gaussian blurs because of the nice properties in the Fourier domain, most algorithms can be applied to either sorts of blurring with the inclusion of a correction factor.

We notate the blur radius at a particular point as $\sigma_i \in \mathbb{R}^+$; we avoid the awkward concept of negative blur radii by using a symmetric PSF. The same symbol, σ, is used for both Gaussian and pillbox blur radii. If we ignore the projection onto the optical axis by assuming a sufficiently small field-of-view and instead approximate by using radial distance to the object, from geometric optics it can be determined that

$$\sigma_i \propto D \left| \frac{1}{f_{\mathrm{L}}} - \frac{1}{d_i} - \frac{1}{r_i} \right| \tag{A.5}$$

where D is the aperture diameter (Chaudhuri and Rajagopalan, 1999). This ap-

proximation allows us to avoid the lens calibrations required to correctly project onto the Z-axis. The equation can be rewritten as

$$\sigma_i \;=\; \alpha \left| 1 - \frac{\beta}{r_i} \right|, \tag{A.6}$$

where $\alpha \in \mathbb{R}^+$ is a scaling constant and $\beta \in \mathbb{R}^+$ is the distance from the camera at which objects are in focus. The tuple $\Theta = (\alpha, \beta, r)$ is considered to be the focal parameters of a system, where $r \in \{r_0, \ldots, r_{n-1}\}$ is the range to each pixel in the scene. This information completely defines our spatially variant blur model; in general we can regard the image formation process as

$$\Xi = \mathrm{H} \star_{sv} \sigma(\Theta) + \epsilon, \tag{A.7}$$

where $\Xi \in \mathbb{C}^n$ is the recorded range data and $\epsilon \in \mathbb{C}^n$ is noise. In theory, given a good estimate of Θ it ought to be possible to restore a full-field lidar image subject to limited depth of field.

In our system a complicating factor is that there are additional optics, which we cannot account for with this model. These optics couple the CCD camera to the image intensifier – since the Z-distance is fixed, this results in convolution by an additional PSF, albeit very small. For the purposes of this paper, we ignore this effect.

A.3 Depth From Defocus

A.3.1 Pentland's Method

While prior work on depth-from-focus extends back much further, the original concept of DFD can be traced back to Pentland (1987). In this work he proposed two methods of depth recovery: one using sharp edges and one using two images taken at different camera aperture settings. The second method is the ancestor of most modern depth-from-defocus algorithms; we now analyse the isoplanatic case. In the Fourier domain convolving an image, y, by a Gaussian PSF with blur radius σ is equivalent to attenuation in the spatial frequency domain, viz

$$z = y \star \sigma \Leftrightarrow Z[u, v] = Y[u, v] \mathrm{e}^{2\pi^2 (u^2 + v^2) \sigma^2}, \tag{A.8}$$

where z is the blurred image. Given two images, z_α and z_β, of the same scene taken with different focal parameters, we can measure the relative attenuation of

the spatial frequencies of the image allowing us to estimate the relative amount of blurring between the two images. For blur radii, σ_α and σ_β, the relative blurring is estimated by

$$g = \sigma_\alpha^2 - \sigma_\beta^2 = \frac{1}{4\pi^2 A_R} \sum_{(u,v)\in R_f} \frac{1}{u^2+v^2} \ln\left(\frac{|Z_\beta[u,v]|}{|Z_\alpha[u,v]|}\right), \tag{A.9}$$

where R_f is a region of the spatial frequency domain and A_R is the number of samples from the spatial frequency domain. By applying knowledge of the focal parameters of the optical system, it is possible to recover σ_α and σ_β individually. While in the most general case this involves solving a quadratic equation, in this paper we restrict ourselves to the situation in which the focal length is fixed and only the aperture size changes. For aperture diameters D_α and D_β,

$$\sigma_\alpha^2 = \frac{g}{1 - \frac{D_\beta^2}{D_\alpha^2}}, \tag{A.10}$$

$$\sigma_\beta^2 = \frac{D_\beta^2}{D_\alpha^2}\sigma_\alpha^2. \tag{A.11}$$

In the conventional case, given the blur relationships from section A.2.2, the range to objects in the scene can be determined. The implementation for this paper divides the image up into 16 by 16 pixel regions and Fourier transforms them – only a small subset of the spatial frequencies are used as most tend to be too excessively attenuated to be useful. A blur estimate is produced for each region using the method described above.

A.3.2 Subbarao's S-Transform Method

Using the S-Transform Subbarao created an alternative method that operates in the spatial domain (Subbarao and Surya, 1994). The S-Transform models an n-dimensional image as the sum of m-order polynomials. By modelling a local region of an image as a cubic polynomial, Gaussian defocus blur can be expressed as

$$z = y \star \sigma \Leftrightarrow z[x,y] + \frac{\sigma^2}{4}\nabla^2 z[x,y] = y[x,y], \tag{A.12}$$

where the Laplacian is determined by a Chebyshev polynomial based smoothed differentiation filter. The size of the differentiation filter is parametised by $N \in \mathbb{Z}^+$, giving a filter of size $2N+1$ pixels. In a similar fashion to eqn. A.9, we can estimate

g by

$$g^2 = 4\frac{\sum_{(x,y)\in R_s}(z_\alpha[x,y] - z_\beta[x,y])^2}{\sum_{[x,y]\in R_s}(\nabla^2 z_\alpha[x,y] + \nabla^2 z_\beta[x,y])^2} \tag{A.13}$$

where R_s defines a spatial domain, rather than a spatial frequency domain. These blur estimates are then denoised via a Parzen windowed histogram over a large spatial region. Although blur estimates can be computed for each individual pixel, they are not statistically independent estimates; for this paper we have produced images at the same resolution as Pentland's algorithm.

A.3.3 Algorithm Response

In order to test the response of the two DFD algorithms we generated some noiseless test data. These parent data were generated by filtering Gaussian white noise with a Gaussian filter of a particular specified width, thus simulating texture at different spatial scales in the ideal unblurred scene. These parent data were then blurred by another blur function, either pillbox or Gaussian, with a specified width, which we call the differential blur. The parent data and the differentially blurred data were then passed to the algorithms as the unblurred and blurred data respectively and the algorithm response recorded. The results are shown in figs. A.1a and A.1b. The graph key refers to the Gaussian blur radius for the parent data and the type of differential blur model.

Since texture scale is proportional to blur radius, fig. A.1a shows that the Pentland algorithm is extremely sensitive to the scale of texture. There is a systematic bias caused by attenuation of the higher spatial frequencies that pulls the estimate of blur towards zero as the accuracy of the logarithmic term is limited by finite precision and noise. A simple example of this effect is if a spatial frequency does not exist (or only noise is present) in the less blurred image, if this spatial frequency is included in the spectral region being analysed then the blur estimate is erroneously pulled towards zero. More advanced methods that dynamically determine the spectral region are possible, but are not dealt with here. Due to the extreme nature of this effect, Pentland's algorithm appears to require the addition of an arbitrary scaling coefficient for comparison to Subbarao's algorithm.

Fig. A.1b shows that the Subbarao Algorithm has a linear response in certain cases, and a non-linear response in others; much of this is due to the spatial scale of the texture. For this experiment, $N = 4$, meaning that the Laplacian was calculated using 9 pixel derivative filters. If the size of the derivative filter does not match the spatial scale of the texture, the algorithm appears to give erroneous results. Parent data which have a much larger initial blur results in a much more linear response

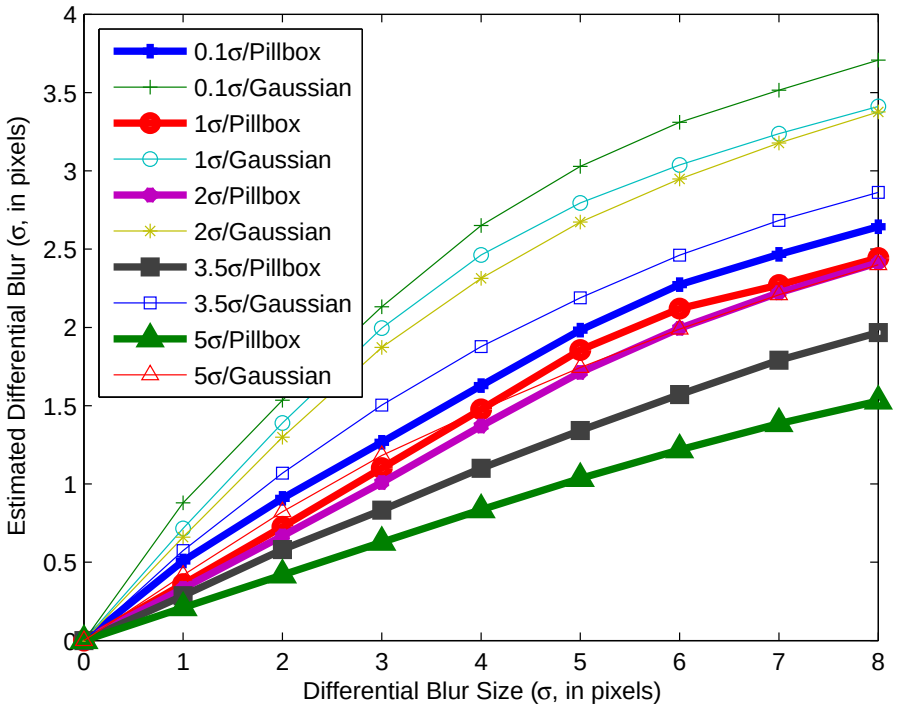

(a) Pentland's frequency domain algorithm response

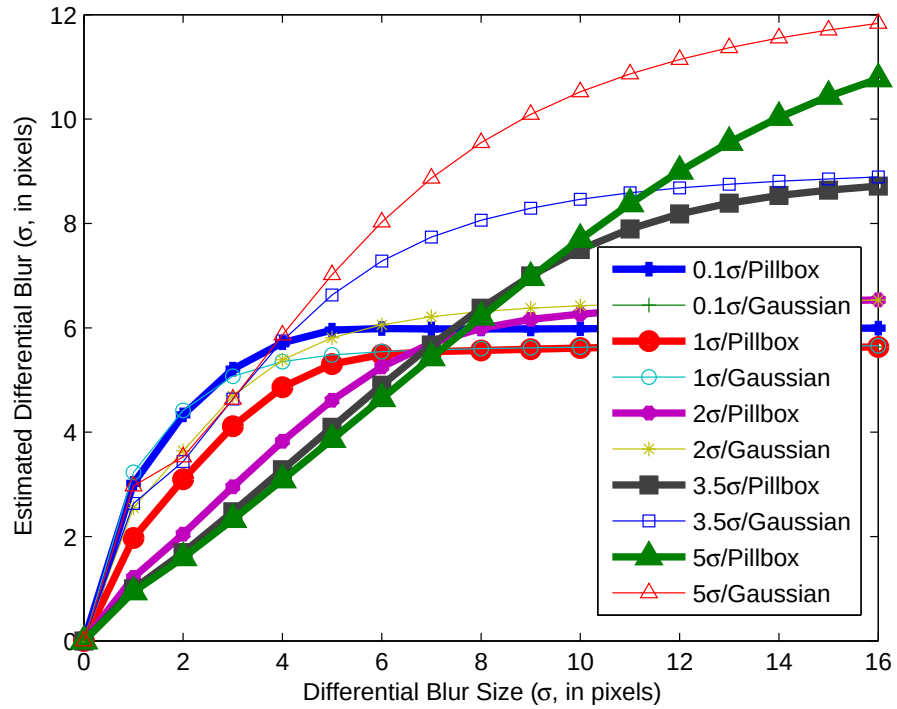

(b) Subbarao's spatial domain algorithm response($N = 4$)

Figure A.1: Depth from defocus algorithm response versus parent data blurring and differential blur type. Key = parent data blurring/differential blur type

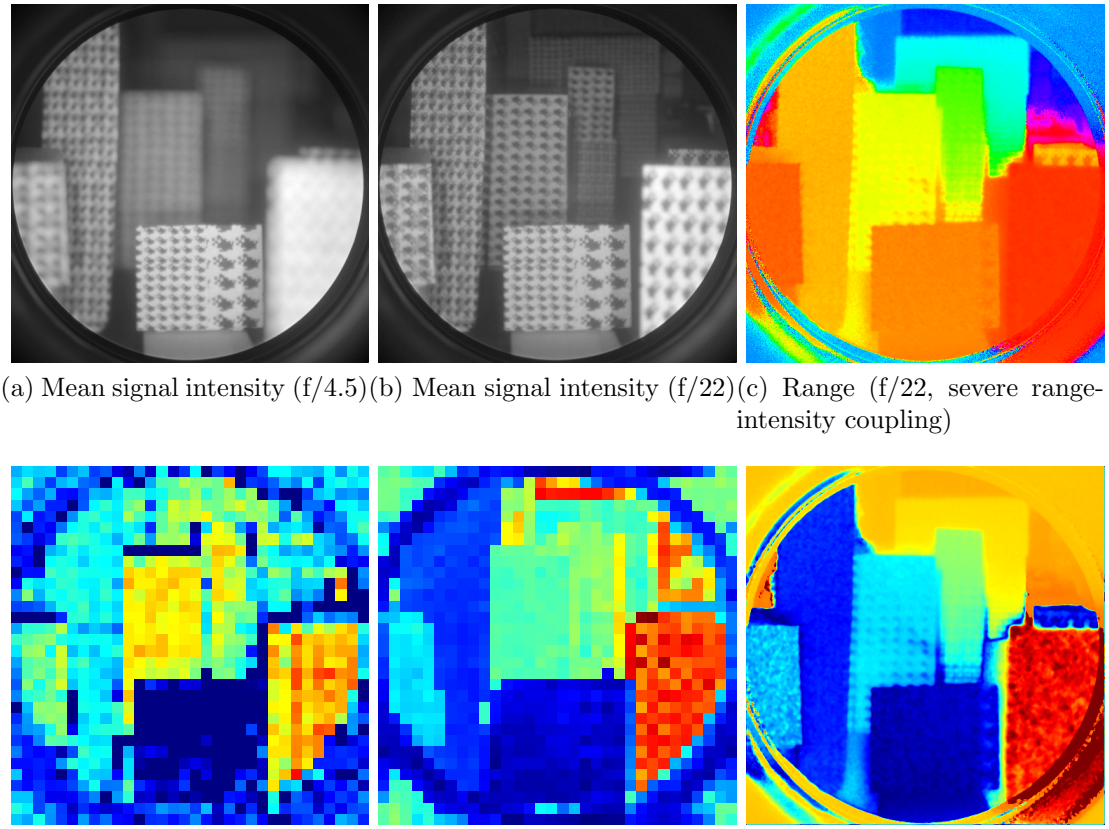

(a) Mean signal intensity (f/4.5) (b) Mean signal intensity (f/22) (c) Range (f/22, severe range-intensity coupling)

(d) Pentland's method (arbitrary scaling) (e) Subbarao's method ($N = 4$) (f) Differential blur from range and sampled PSF model (severe range-intensity coupling)

Figure A.2: Scene A depth from defocus results. Mean signal intensity images use $\gamma = 0.1$.

gradient – which appears to be close to $\frac{6}{8}$ in the pillbox case and $\frac{8}{6}$ in the Gaussian case.

A.3.4 Depth From Defocus Results

Three range images were taken of different scenes, each with a different focal length. The complex domain range images were then analysed using the two DFD algorithms and compared to blur estimates created using range data and the known focal parameters. The only difference between processing intensity images and processing complex domain images being that the absolute value of the Laplacian and relative error are taken in the Subbarao algorithm. The results are shown in figs. A.2, A.3 and A.4.

Scene A has an extremely high dynamic range; the mean signal intensity images have been highly gamma compressed using $\gamma = 0.1$. This high dynamic range results in a problem known as range-intensity couplingGodbaz *et al.* (2009b), where

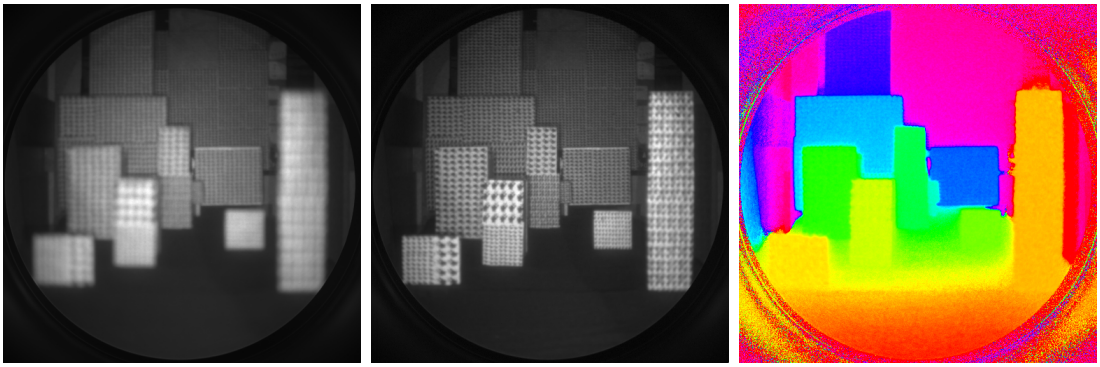

(a) Mean signal intensity (f/4.5) (b) Mean signal intensity (f/22) (c) Range (f/22)

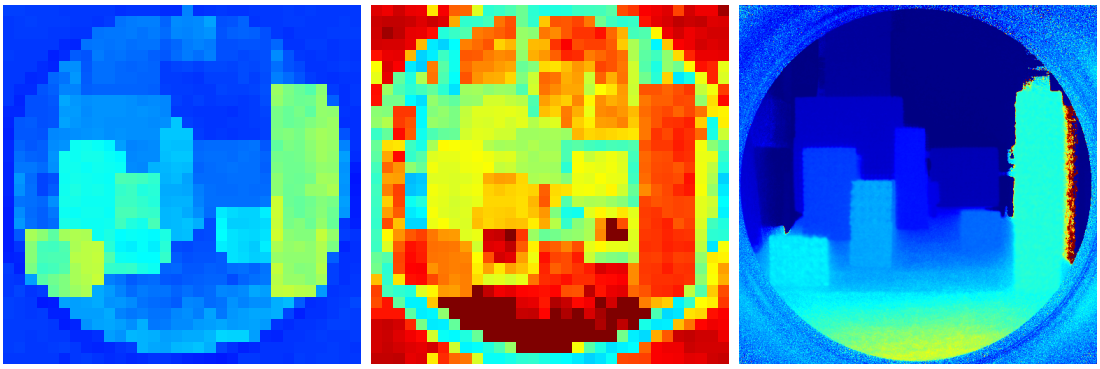

(d) Subbarao's method ($N = 1$) (e) Subbarao's method ($N = 4$) (f) Differential blur from range and sampled PSF model

Figure A.3: Scene B depth from defocus results. Mean signal intensity images use $\gamma = 0.5$.

light scattered within the range imager results in ranges that change depending on the reflectivity of the target object. Fig. A.2c shows range information, where red represents objects closer to the camera and blue/purple objects farther away. Figs. A.2a and A.2b show how decreasing the aperture size increases the depth of field, however even at the very small f/22 the depth of field is still finite and some objects appear quite blurry. Passing the range image of the scene to the DFD algorithms gives figs. A.2d and A.2e. Both algorithms struggle to accurately estimate the blur scale for objects in the background, however Subbarao's produces visibly better output. Pentland's algorithm produced artefacts at several object edges, suffered more from noise and while retaining the correct relative blur ordering across objects, did not maintain the same ratios. Subbaro's algorithm however performed creditably compared to the estimated actual blurring in fig. A.2f (blue represents less blurring, red more blurring). The blur estimates are slightly corrupted due to range-intensity coupling.

Scene B has a much smaller dynamic range than scene A; as a result the range

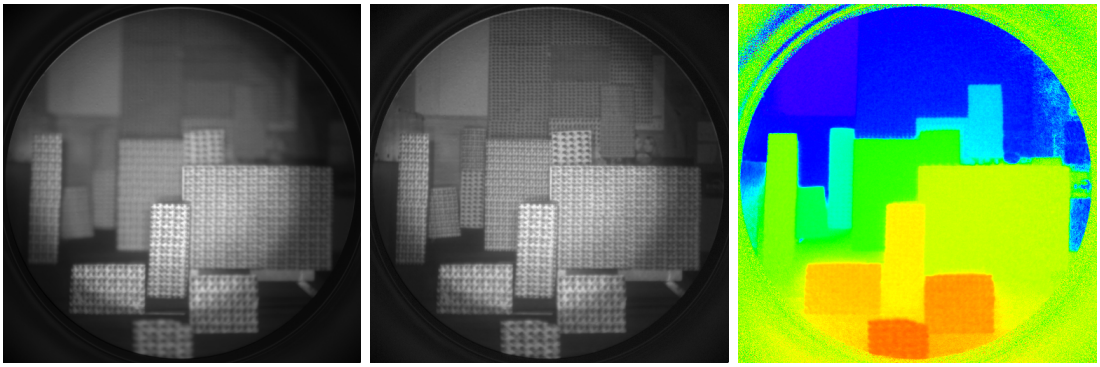

(a) Mean signal intensity (f/4.5) (b) Mean signal intensity (f/22) (c) Range (f/22)

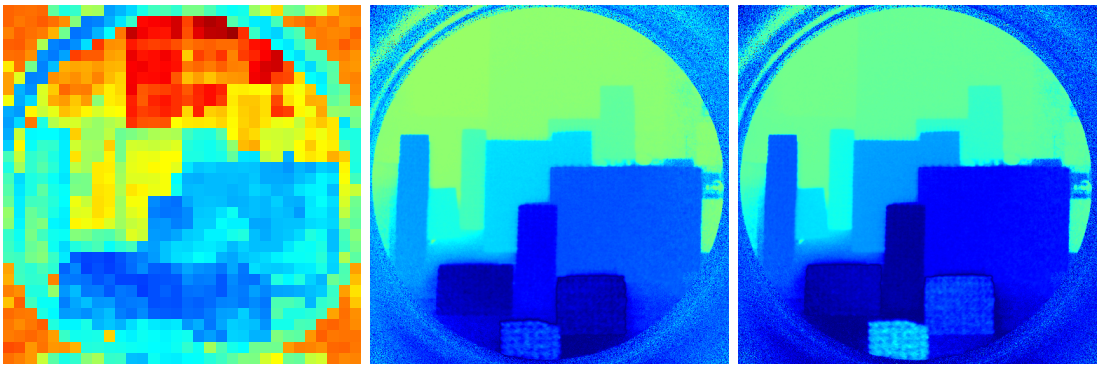

(d) Subbarao's method ($N = 4$) (e) ML focal parameter determination (f) Differential blur from range and sampled PSF model

Figure A.4: Scene C depth from defocus results. Mean signal intensity images use $\gamma = 0.5$.

data (fig. A.3c) suffers much less range-intensity coupling. However there are still minor erroneous ranges to the right/top of the rightmost box in the image. Figs. A.3d and A.3e show how the value of N affects the quality of range data returned by Subbarao's algorithm. We are currently unsure whether there is a good automated way to estimate the optimal setting. Notably, the textureless table region does not produce accurate blur estimates.

In scene C no value of N was found to produce correct blur estimates from Subbarao's algorithm. Fig. A.4e shows a blur estimate produced by the Maximum Likelihood method discussed in the next section.

A.4 Determination of Focal Parameters

A.4.1 Theory

Photon shot noise is Poisson distributed; in most imaging situations the mean intensity, μ_i, is sufficiently high that the Poisson distribution can be accurately modelled

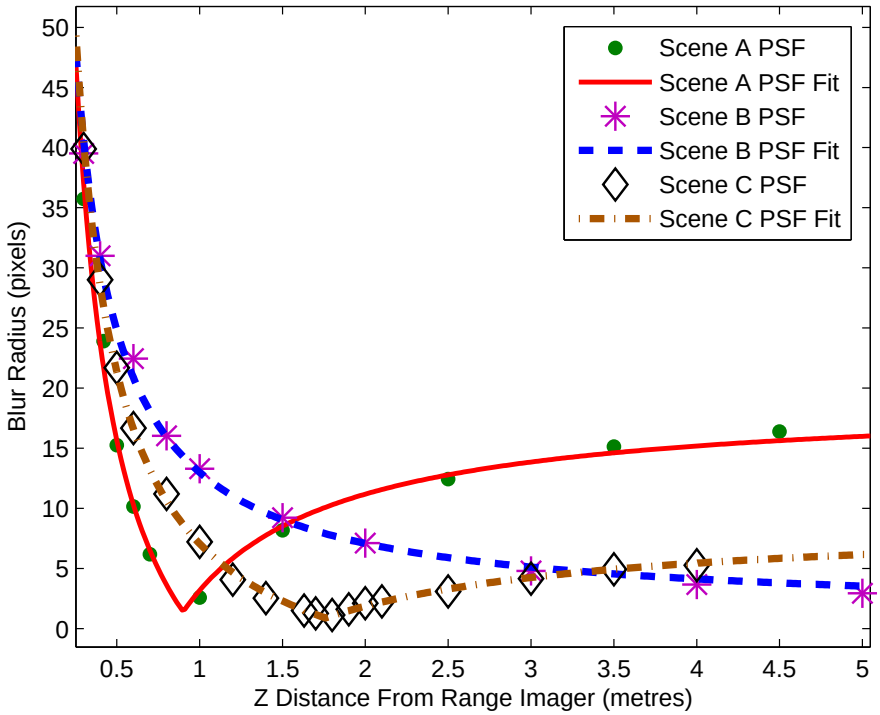

(a) Ground truth PSF measurements (f/4.5)

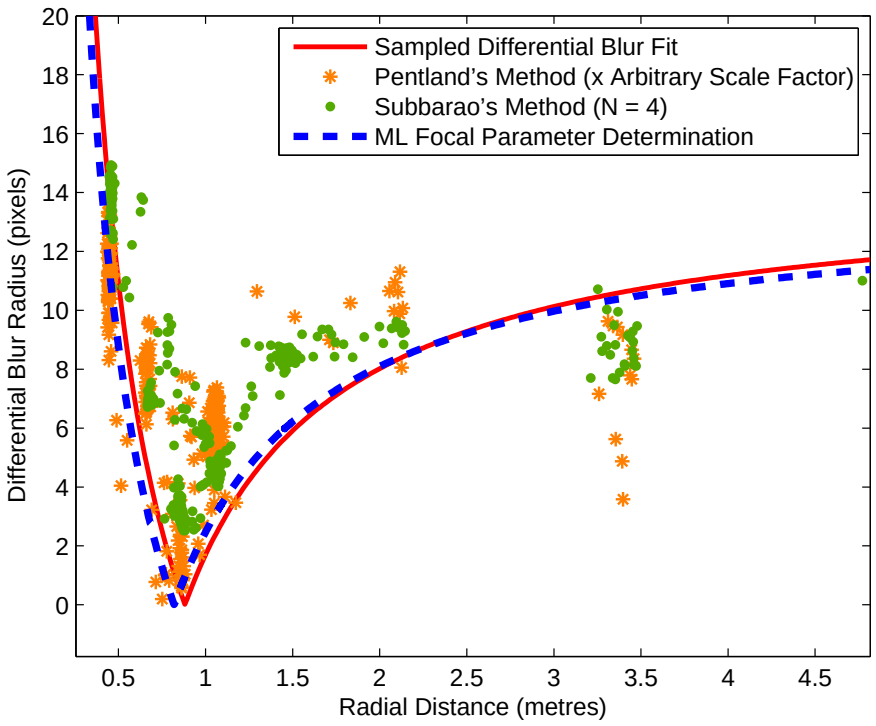

(b) Scene A

Figure A.5: ML focal parameter determination compared to depth from defocus methods and ground truth. Sampled PSF at f/22 not shown.

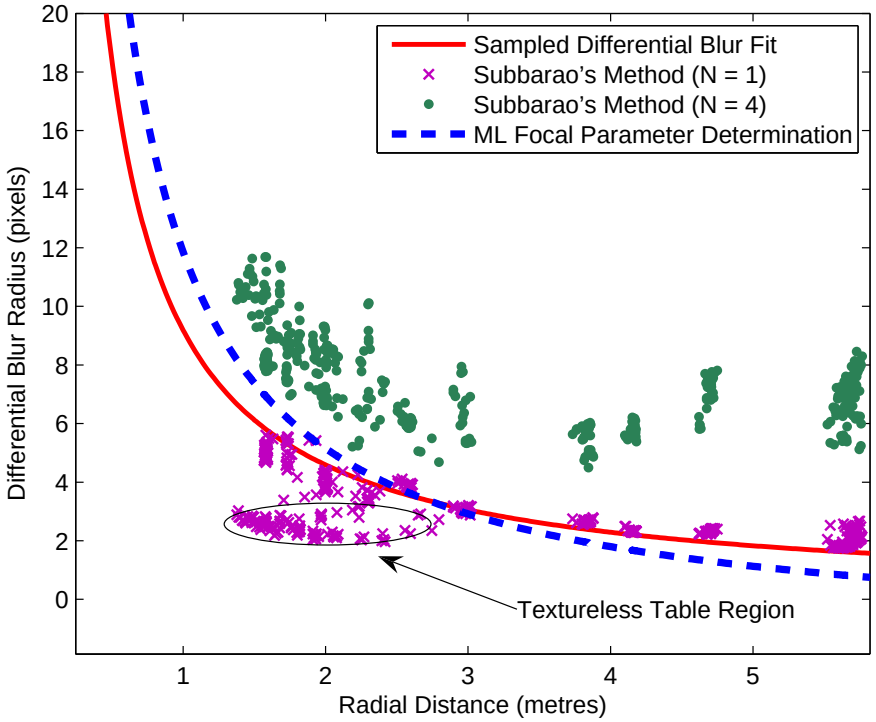

(a) Scene B

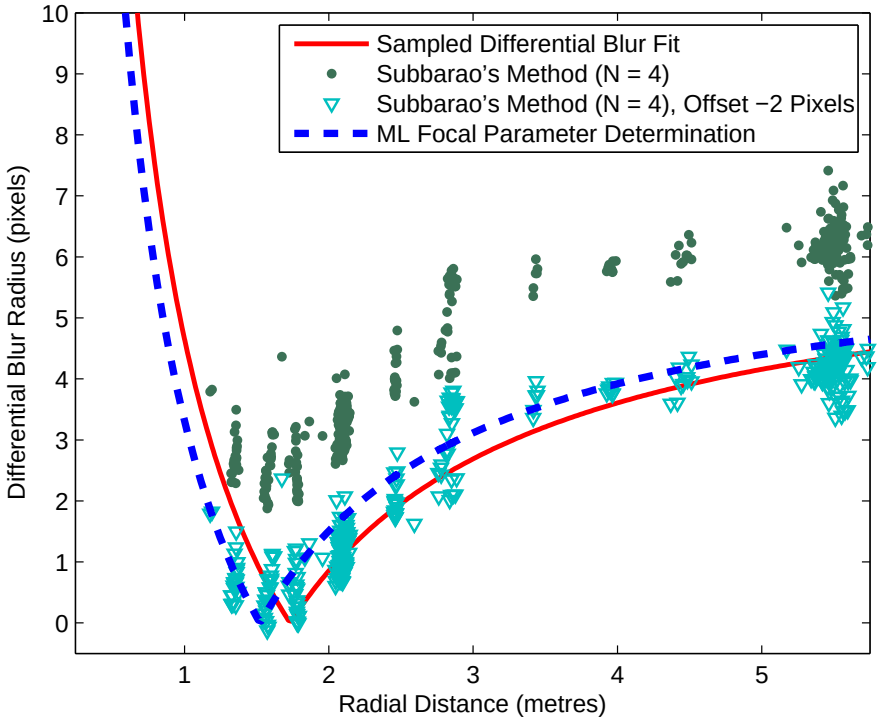

(b) Scene C

Figure A.6: ML focal parameter determination compared to depth from defocus methods and ground truth. Sampled PSF at f/22 not shown.

by a Gaussian distribution where $\sigma_i^2 = \mu_i$. In AMCW lidar noise variance is not necessarily proportional to amplitude ($\sigma_i^2 \not\propto |\eta_i|$) due to contributions from ambient lighting and partial cancellation at mixed pixels. Instead the mean signal intensity at a pixel, χ_i, takes into account all light sources within a pixel

$$\chi_i = \phi_i + \int f_i(r)dr, \tag{A.14}$$

where ϕ_i is the contribution from ambient light at pixel i. Given that the University of Waikato range-imager is a non-differential system, the mean signal intensity is equivalent to averaging the correlation waveform over all phase steps. This allows us to model the noise at a pixel, ϵ_i, as having a zero centred circularly symmetric complex normal distribution, $\epsilon_i \sim \mathcal{N}_{\mathbb{C}}(0, \chi_i)$, resulting in the absolute error at a pixel, $|\epsilon_i|$, being Rayleigh distributed. In general, mean signal intensity images look very similar to amplitude images unless there is significant ambient light or partial cancellation at mixed pixels (all data for this paper were generated in a darkened room).

If we model z_α and z_β as a convolution of ideal unblurred range data, $\hat{y} = \pi_\Lambda(a_\emptyset, r_\emptyset)$, we can express the log-likelihood of particular focal parameters and unblurred data as

$$\begin{aligned} & \mathcal{L}(\alpha, \beta, r_\emptyset, a_\emptyset | z_\alpha, z_\beta, \Sigma_\alpha, \Sigma_\beta) \\ = \; & const - \frac{1}{2}(z_\alpha - h(\alpha, \beta, r_\emptyset)\pi_\Lambda(a_\emptyset, r_\emptyset))^* \Sigma_\alpha^{-1}(z_\alpha - h(\alpha, \beta, r_\emptyset)\pi_\Lambda(a_\emptyset, r_\emptyset)) \\ - \; & \frac{1}{2}(z_\beta - h(\alpha\frac{D_\beta}{D_\alpha}, \beta, r_\emptyset)\pi_\Lambda(a_\emptyset, r_\emptyset))^* \Sigma_\beta^{-1}(z_\beta - h(\alpha\frac{D_\beta}{D_\alpha}, \beta, r_\emptyset)\pi_\Lambda(a_\emptyset, r_\emptyset)) \end{aligned} \tag{A.15}$$

where Σ_α and Σ_β are diagonal matricies consisting of the values from χ_α and χ_β respectively, and $h\colon (\mathbb{R}^+, \mathbb{R}^+, \mathbb{R}^{+n}) \to \mathbb{R}^{n \times n}$ is a function that returns the linear transformation matrix corresponding to a spatially variant convolution such that $h(\alpha, \beta, r)z = z \star_{sv} \sigma(\alpha, \beta, r)$.

A single step approach might attempt to maximise this, most probably with the addition of some sort of regularisation. Instead we assume that the least blurred range estimates r_β are correct and approximate the relationship between z_α and z_β as a spatially variant convolution parametised in the same fashion as before, giving

$$\begin{aligned} & \mathcal{L}(\alpha, \beta | z_\alpha, z_\beta, \Sigma_\alpha, \Sigma_\beta) \\ = \; & const - \frac{1}{2}\ln(\det(\Sigma_\alpha + h(\tau\alpha, \beta, r_\beta)\Sigma_\beta h^*(\tau\alpha, \beta, r_\beta))) \\ - \; & \frac{1}{2}(z_\alpha - h(\tau\alpha, \beta, r_\beta)z_\beta)^*(\Sigma_\alpha + h(\tau\alpha, \beta, r_\beta)\Sigma_\beta h^*(\tau\alpha, \beta, r_\beta))^{-1}(z_\alpha - h(\tau\alpha, \beta, r_\beta)z_\beta), \end{aligned} \tag{A.16}$$

where τ is a constant conversion factor. In an ideal case τ would be confirmed

experimentally, however for the purposes of this paper we have assumed that $\tau = \frac{D_\alpha - D_\beta}{D_\alpha}$. By approximating the error covariance by $2\Sigma_\alpha$ we can further simplify the problem, giving a constant, diagonal covariance matrix. An estimate of the remaining focal parameters (α, β) is found by numerically maximising

$$\mathcal{L}(\alpha, \beta | z_\alpha, z_\beta, \Sigma_\alpha) = const - \frac{1}{2}(z_\alpha - h(\tau\alpha, \beta, r_\beta) z_\beta)^* (2\Sigma_\alpha)^{-1} (z_\alpha - h(\tau\alpha, \beta, r_\beta) z_\beta). \tag{A.17}$$

Given an estimate of the focal parameters, a spatially variant deconvolution operation can be performed as a second step.

A.4.2 Focal Parameter Determination Results

Results from application of the ML Focal Parameter Determination algorithm to the three scenes are shown in figs. A.5 and A.6. Fig. A.5a shows some ground truth measurements of the PSF blur radius as a function of Z distance from the range imager; measurements at f/22 and f/4.5 were combined in order to estimate the differential blur which is used as a reference in each of the following graphs. Fig. A.5b shows how the DFD results compare to the actual differential blur and the results of ML focal parameter determination. In this case, both DFD methods suffer from noise issues – although the range information is corrupted slightly by range-intensity coupling, each discrete object should still have a consistent blur level, which they do not. Both methods do appear to identify roughly the correct in-focus range. The ML estimate of focal parameters very closely matches the sampled differential blurring. Fig. A.6a contains the results from Scene B. The graph shows how N affects the quality of blur estimates from Subbarao's algorithm; in the $N = 1$ case the blur estimates are very close to the actual differential blurring everywhere except for the featureless table region. Scene C results in an interesting problem with the Subbarao algorithm; fig. A.6b shows that the blur estimates appear to be offset by 2 pixels from the correct values. The ML focal parameter determination algorithm, while not producing inordinately false results, has misestimated the in-focus distance by ~ 30 cm.

A.5 Deconvolution of Lidar Images

A.5.1 Theory

Deconvolution was performed by a spatial derivative regularised, variant Landweber algorithm (Landweber, 1951). The Landweber algorithm is an iterative Maximum

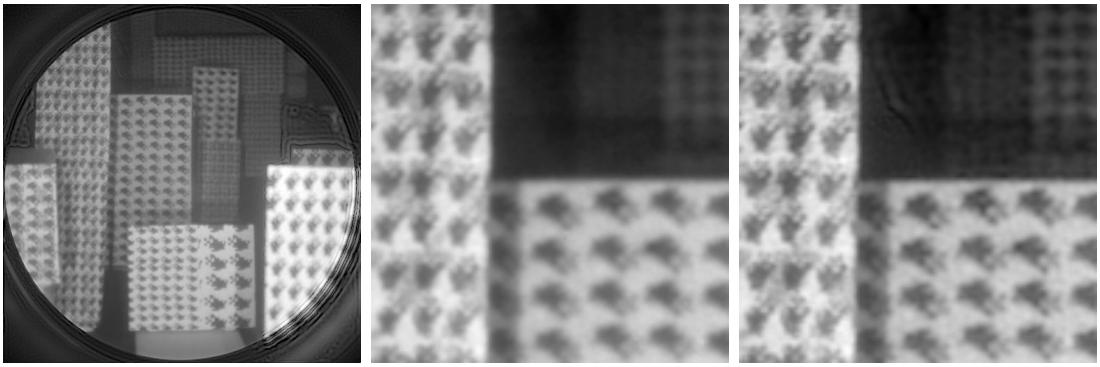

(a) Deconvolved (b) Subregion (f/22) (c) Deconvolved subregion

Figure A.7: Scene A deconvolved mean signal intensity image

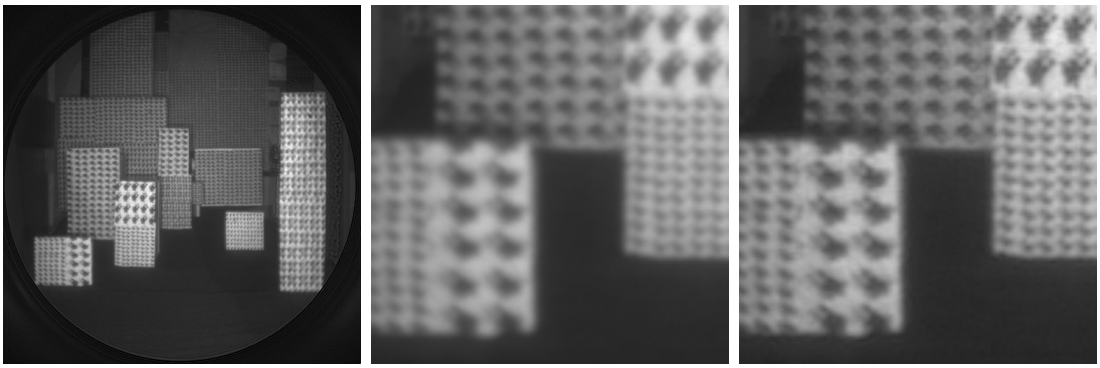

(a) Deconvolved (b) Subregion (f/22) (c) Deconvolved subregion

Figure A.8: Scene B deconvolved mean signal intensity image

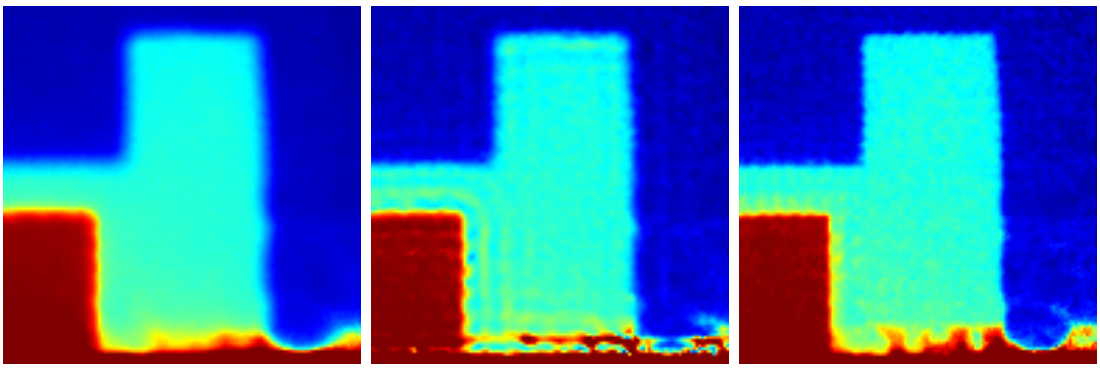

(a) Phase, limited DOF (f/4.5) (b) Deconvolved phase (from f/4.5) (c) Reference phase (f/22)

Figure A.9: Scene C phase of deconvolved complex domain range image subregion

Likelihood algorithm in the case of additive Gaussian noise, although it can be applied to other noise distributions. Each iteration is defined by

$$\hat{y}_i = \hat{y}_{i-1} + \zeta(h^T(\alpha, \beta, r_\beta)(z_\beta - h(\alpha, \beta, r_\beta)\hat{y}_{i-1}) - \lambda L\hat{y}_{i-1}), \tag{A.18}$$

where $\hat{y}_i$ is the ith estimate of the deconvolved data, ζ is gain, L is a Laplacian filter and λ is a regularisation parameter. We implement the spatially variant convolution by generating pillbox blur models at blur radii intervals of one-third of a pixel and modelling intermediate blur radii by linear interpolation. This leads to efficient computation, although linear interpolation is more suited to smooth functions. The transpose of the spatially variant convolution matrix, $h^T(\alpha, \beta, r_\emptyset)$ corresponds to a spatially variant correlation operation which can be implemented in a similar manner to the spatially variant convolution. A good overview of the implementation of spatially variant convolutions is provided by Nagy and O'Leary (1997). Unfortunately, the finite size of the image intensifier crops the images resulting in edge effects in deconvolved images.

A.5.2 Results

Figs. A.7 and A.8 show the resultant deconvolved mean signal intensity images from scenes A and B respectively. Using mean signal intensity images instead of amplitude allows us to avoid any confounding effects of phase in the complex domain. There is an improvement in depth of field, but with a difficult trade-off between recovering detail and noise amplification. The high level of gamma compression in fig. A.7a shows significant artefacts in regions where the blur scale has been misidentified due to mixed pixels. For example, there is a region on the far right above the very bright defocussed object where severe blur misestimation due to mixed pixels has resulted in irregular wave patterns.

Fig. A.9 shows the results from attempting to deconvolve a complex domain range image. An image taken at f/4.5 (fig. A.9a) is highly blurred due to limited depth of field – thus forming a large band of mixed pixels around the edge of each object, each mixed pixel being an erroneous range measurement. Deconvolution reduces the size of the regions containing mixed pixels (fig. A.9b), however it also results in ringing – contributed to by the highly bandlimited nature of large pillbox PSFs (Gibbs oscillations). A comparison image at f/22 is given in fig. A.9c. While there is some PSF misestimation, similar ringing occurs if simulated isoplanatic data with a hard discontinuity is generated and deconvolved with the correct PSF. Designing a good deconvolution algorithm involves determining how to synthesise these missing frequencies, without which the image cannot be properly restored.

These oscillations are very problematic; because the image is in the complex domain, it contains both amplitude and range information. The range information is best regularised to ensure sharper edges but less detail, yet at the same time we need to retain (and improve) the highly textured amplitude information. This suggests that it may be necessary to develop a more advanced regularisation technique: for example, one that combines Harr wavelet based regularisation in the phase/range domain and spatial derivative regularisation in the amplitude domain. We leave this to future work.

Apart from the regularisation trade-off, the biggest problem with the current algorithm is mixed pixels. Range-intensity coupling due to scattered light introduces nonexistent discontinuities in blur estimates for patterned surfaces and in dark regions can result in completely erroneous blur estimates due to wraparound effects. Edge induced mixed pixels cause similar problems, although are more predictable – resulting in artefacts around the edges of objects. Fundamentally, there is no simple solution to this problem because each individual pixel can contain multiple component signal returns, each with a different range and thus subject to a different amount of blur.

A.6 Summary

We have presented a new method for parametric spatially variant blind deconvolution of full-field lidar image pairs taken at different aperture settings. The method allows estimation of focal parameters without prior calibration of the optical setup and produces blur estimates which have better spatial resolution and less noise than previous depth from defocus blur measurement algorithms under most circumstances. The estimated focal parameters were applied to deconvolution of mean signal intensity images improving depth of field. Application to complex domain images of multiple objects requires the development of more appropriate regularisation methods due to the trade-off between different levels of amplitude and range texture in the same image. There remain issues with correct PSF determination, especially around the edges of objects and in regions affected by range-intensity coupling.

Appendix B

Extending Depth-Of-Field Using a Coded-Aperture

B.1 Introduction

Full-field amplitude modulated continuous wave (AMCW) lidar systems utilise the time-of-flight (TOF) principle to generate two dimensional matrices of intensity and radial range values using active scene illumination. Whereas point and line scanner based systems require expensive mechanical systems to sequentially capture a point cloud, full-field systems capture an entire image simultaneously and near-instantly opening up a variety of applications including games, medical imaging, security and engineering quality control.

However, despite their advantages, full-field AMCW systems introduce new challenges such as systematic errors due to multipath interference and limited depth-of-field (DOF). In full-field AMCW lidar systems limited DOF results in both erroneous range and intensity values around the edges of objects as well as a loss of detail. While most previous computational photography work has addressed the DOF problem for intensity images using techniques such as coded apertures (Levin *et al.*, 2007) and plenoptic cameras (Ng *et al.*, 2005), previous systems have relied on implicit range information. Since AMCW lidar systems produce explicit range information, albeit limited by the DOF, there is inherently more information available to assist in restoration.

Prior depth-from-defocus (DFD) techniques (Pentland, 1987; Chaudhuri and Rajagopalan, 1999) utilise the known range variant properties of the PSF to determine distance, however typically require more than one image of a scene. More modern methods have used coded apertures to make the blurring less of a low-pass filter and enable high quality restoration while requiring only a single image Levin *et al.*

(2007). Related work has changed the nature of motion blur using coded fluttered shutter patterns (Raskar *et al.*, 2006). Traditional plenoptic cameras allow refocussing without any explicit range information (Ng *et al.*, 2005) however sacrifice spatial resolution. Alternative methods like Lumsdaine and Georgiev's 'Plenoptic 2.0' (Lumsdaine, 2008), which offer a substantial increase in resolution, require the determination of a range dependent magnification parameter in order to produce an artefact free image. Other techniques for defocus invariance include wavefront coding (Dowski and Johnson, 1999) and merging multiple images at different focal settings (TODO). Deconvolution techniques have been previously applied to full-field lidar images for the purposes of light scattering reduction (Mure-Dubois and Hügli, 2007a; Kavli *et al.*, 2008). Another work (Godbaz *et al.*, 2010) blindly determined the focal parameters of a full-field lidar system and utilised them to improve DOF.

In this paper we briefly demonstrate the advantages of our coded-aperture design over a circular aperture for extending DOF and then show the deconvolution of real defocussed range-images captured using a coded-aperture full-field AMCW lidar system.

B.2 Theory

B.2.1 AMCW Lidar

AMCW lidar systems work by illuminating a target scene with modulated light and then sampling the correlation of the reflected light with a reference signal at the same or a slightly different frequency. The TOF results in a range variant phase shift in the returned illumination – this phase shift is typically measured by mixing the returned light with a reference signal using either a modulated CCD or CMOS sensor (Oggier *et al.*, 2004) or modulated image intensifier (Dorrington *et al.*, 2007). By changing the phase relationship between the illumination and reference signals it is possible to measure the range and active intensity of an object. This information is typically encoded as a complex domain value for each pixel. For a single pixel an ideal AMCW lidar measurement can be written as

$$\eta = a\mathrm{e}^{4\pi j d/\lambda} \tag{B.1}$$

where $\eta \in \mathbb{C}$ is a complex domain range measurement, a is the active intensity, d is the distance from the camera and λ is the illumination modulation wavelength.

In practice, AMCW lidar measurements are subject to systematic errors, partic-

ularly due to the impact of multipath interference. Multipath interference, of which mixed pixels are a type, is caused when a single pixel integrates light from sources at more than one range causing an erroneous range measurement. The erroneous value being the sum of the complex domain range measurements of each component return. This can result in range-intensity coupling – where the measured range is a function of intensity. When a range image is subject to limited DOF, blurring of the edges of objects results in the formation of large bands of mixed pixels containing erroneous values. One of the aims of this paper is to demonstrate that these erroneous values can be restored. Methods have been developed to mitigate (Larkins *et al.*, 2009) or find the component returns within mixed pixels (Godbaz *et al.*, 2009a, 2008), however the output from these algorithms is difficult to incorporate into a simple deconvolution model. Since each component within a mixed pixel is at a different range from the camera, each has a different PSF. For this paper we model each pixel as being at a single discrete range, which while non-ideal, retains the simplicity of a single two dimensional image array.

At the moment full-field lidar image processing research is limited by the unavailability of off-the-shelf high resolution systems and the black-box nature of many commercial devices. The custom range-imager utilised for this paper is built around an image intensifier and has an effective resolution of around 200,000 pixels – many times that of any commercially available device. However, this comes at the cost of an increase in complexity due to the additional optics required to couple the image intensifier to the CCD and an increase in scattered light.

B.2.2 Image Formation

From geometric optics, the defocus PSF is a scaled image of the aperture shape given by

$$r_p = \alpha \left(1 - \frac{\beta}{d}\right) \tag{B.2}$$

where r_p is the radius of the PSF, d is the distance from the first principal plane to the object, β is the distance from the first principal plane to the point on the optical axis at which objects are most in-focus and α is a scaling constant (Godbaz *et al.*, 2010). In the Fourier domain, convolution by a PSF corresponds to elementwise multiplication of the spatial frequencies of the image with the spatial frequencies of the PSF

$$g = f \star h \Leftrightarrow G[u,v] = F[u,v]H[u,v] \tag{B.3}$$

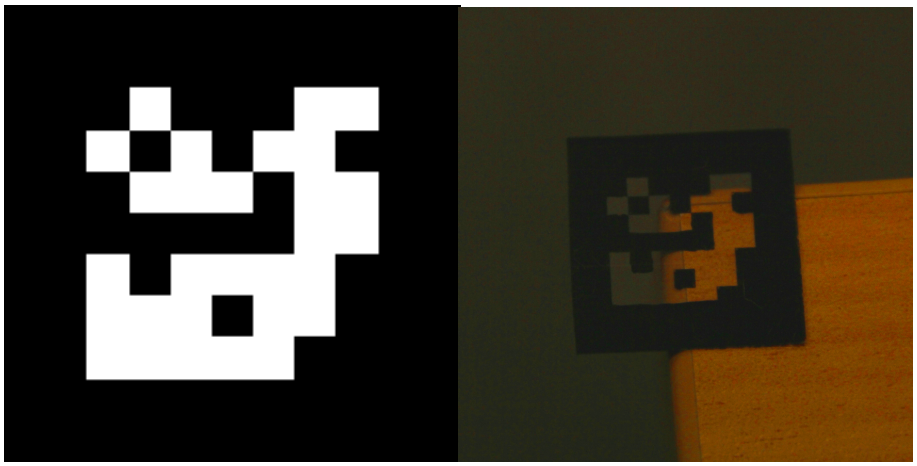

(a) Pattern (b) Example OHT Aperture

Figure B.1: The binary coded aperture pattern and an example OHT based coded aperture.

where f is the original image, g is the blurred image and h is the PSF. Any spatial frequencies missing from the PSF are lost, making high quality image restoration difficult. However by changing the nature of the PSF it is possible to whiten the power spectrum of the PSF and improve the quality of restored images. Because there is explicit range information, it is possible to aim for as broadband a PSF as possible without the constraints imposed by extraction of range information.

The image formation process for an AMCW range-imager is the same as for a standard camera with the exception that any reflections before the image intensifier result in an increased TOF and thus a phase shift in the range measurements; fully modelling this requires the utilisation of a complex domain PSF.

B.2.3 Restoration Method

We use a spatially variant Landweber (Landweber, 1951) deconvolution method using a Gaussian spatial derivative prior and a weighting mask to remove boundary effects due to the image intensifier. By writing the spatially variant convolution as a matrix transformation, $f \star_{sv} h = Tf$, each iteration becomes

$$\hat{f}_{n+1} = \hat{f}_n + \gamma(T^*((g - T\hat{f}_n) \otimes m) - \lambda \hat{f}_n \star l) \tag{B.4}$$

where $\hat{f}_n$ is the estimate of the unblurred image at the nth iteration, * is the Hermitian transpose of a matrix, $\otimes$ is elementwise multiplication, γ is a gain term, m corresponds to data weights, λ is the regularisation parameter and l is a Laplacian kernel. The initial estimate is the captured blurred range-image. Additional blank,

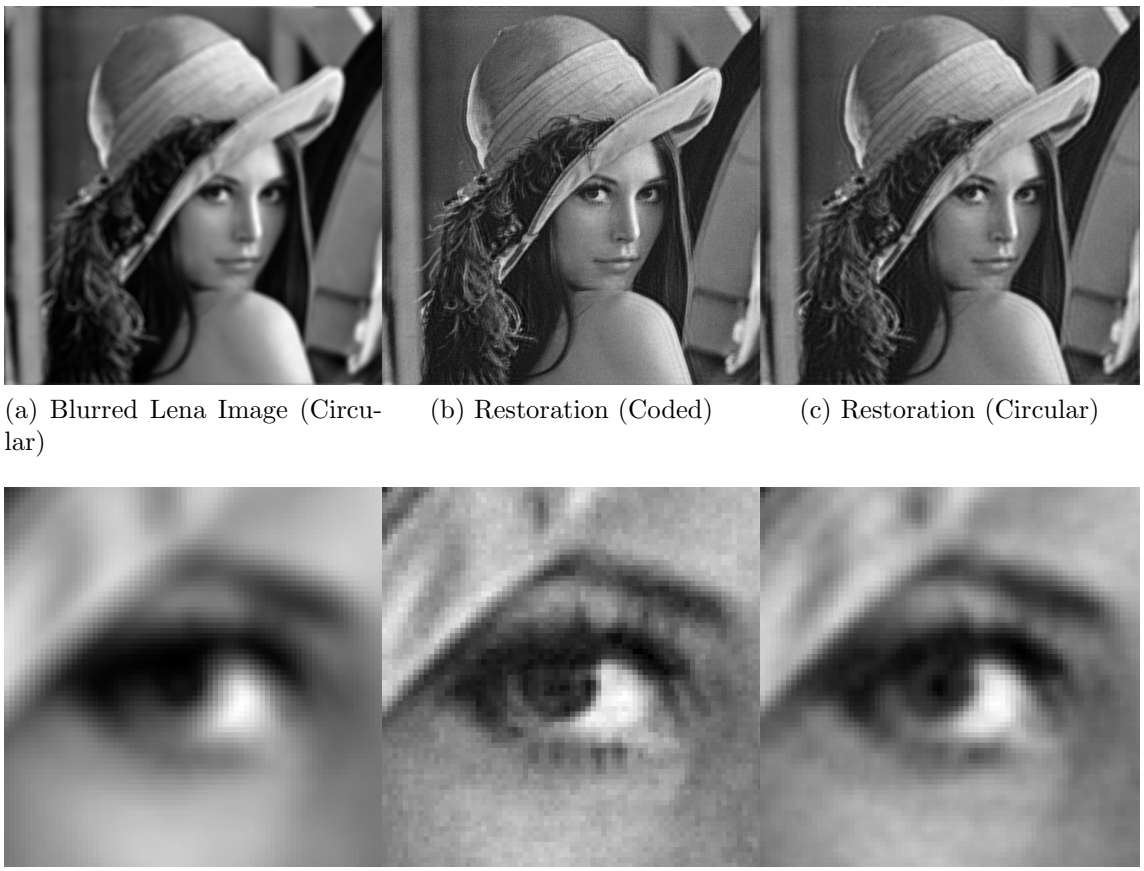

(a) Blurred Lena Image (Circular) (b) Restoration (Coded) (c) Restoration (Circular)

(d) Blurred Subregion (Circular) (e) Restored Subregion (Coded) (f) Restored Subregion (Circular)

Figure B.2: The impact of aperture choice on deconvolution restoration quality of an intensity image in the known, isoplanatic PSF case. Simulated at a SNR of 1000 : 1, $\lambda = 0.015$ with 50 Landweber iterations. (Original Lena Image: Copyright Playboy Enterprises, 1972).

zero weighted boundaries are added to each image, increasing the image size to 768x768 to mitigate wraparound effects from the use of circular convolutions.

Before each iteration the PSF is dynamically determined for each pixel using radial distance calculated from the phase angle of value in $\hat{f}_n$. In general, distance along the optical axis can be approximated without calibration by the radial distance. A threshold was set for each restoration, usually 10 iterations, at which point the PSF was no-longer dynamically updated to prevent noise amplification. This method typically works quite well in regions with edge induced mixed pixels as the values tend to converge to a sharper edge, but in regions subject to severe range-intensity coupling due to scattered light the algorithm can fail.

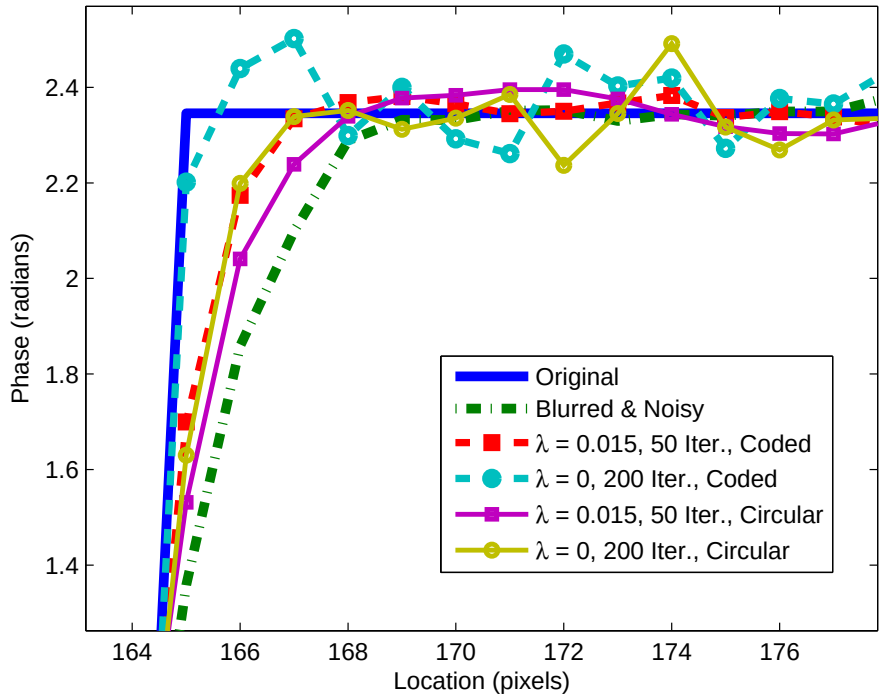

(a) Hard Discontinuity

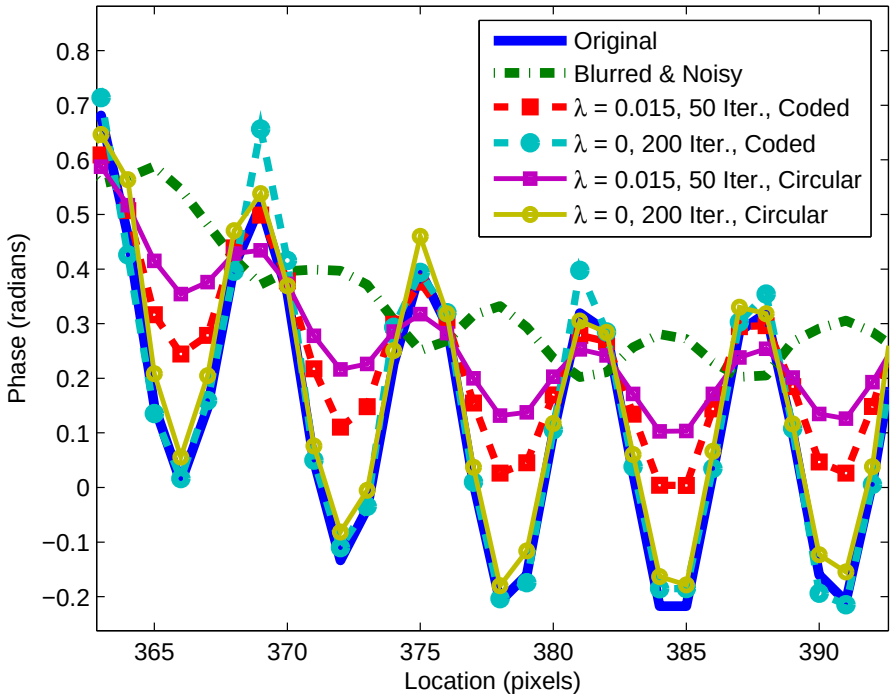

(b) Textured Object

Figure B.3: Slices through a simulated pure phase image pre- and post-deconvolution using a SNR of 1000 : 1. For a given regularisation constant the coded aperture generally results in better restoration quality than a circular aperture – the behaviour for the phase of a complex number is similar to that in the case of an intensity image, but with a slightly greater sensitivity to ringing.

B.3 Methodology

B.3.1 The Coded Aperture

The coded aperture utilised for this paper is a 7×7 random noise pattern that was printed onto an overhead projector transparency (OHT) as shown in fig. B.1. Due to the limited contrast provided by the printing process, the aperture pattern was augmented using marker pen – this resulted in slight unevenness, but had no other impact due to empirical sampling. Advantages of this method of aperture construction include low cost and that any pattern can be produced without physical constraints such as the connectivity required for a physical cut-out pattern. The biggest disadvantage is that depending on the type and quality of the OHT material, the aperture may contribute to light scattering and reflection within the ranger. Some previous approaches include cut out patterns (Levin *et al.*, 2007) and LCD screens (Marcia *et al.*, 2010).

In order to compare our coded aperture design to a similarly sized circular aperture we simulated blurred and noisy intensity and phase images. Fig. B.2 shows how the coded aperture improves the performance of deconvolution for an intensity image. For the Lena image at a SNR of $1000:1$ there is a 24% decrease in RMS error in the restored image. Fig. B.3 shows how the coded aperture affects the restoration of phase content in a pure phase image – that is a simulated range image where every pixel has a modulus of one, thus isolating the impact on phase information. The blurred phase information for the textured object counterintuitively appears to peak where there are troughs in the unblurred image due to the black regions in the centre of the aperture pattern. Despite designing the aperture for a white spectral response, limited Gibbs' phenomenon still occurs at hard discontinuities.

B.3.2 The Optical Path of the Range-Imager

Fig. B.4 shows the optical configuration of the ranger system. The scene is illuminated by modulated laser light and imaged by a Nikkor 50mm f/1.8D lens where the aperture diaphragm blades are replaced with an OHT coded aperture. The primary optics image the scene onto the mirror-like surface of the image intensifier photocathode, which is one of the main sources of multipath in range-images taken by this system. A phosphor screen displays the correlation of the returned scene illumination with the image intensifier modulation signal. This results in a temporally varying correlation waveform, where phase corresponds to object range. The phosphor screen is focussed onto a CCD image sensor using additional coupling

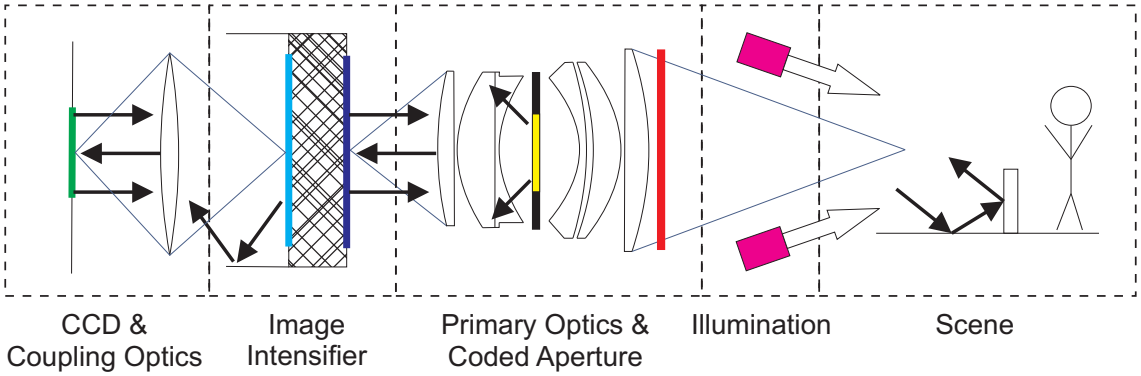

Figure B.4: The optical configuration of the range-imager. Key: modulated lasers (magenta), narrowband filter (red), coded aperture (yellow), image intensifier photocathode (blue), phosphor screen (cyan), CCD image sensor (green). Black arrows represent sources of multipath.

optics. A range-image is produced by calculating the bin of the temporal discrete Fourier transform corresponding to the correlation waveform fundamental frequency for each pixel.

B.3.3 The Point Spread Function

The empirical point spread function of our system is formed as the convolution of the fixed point spread function of the image intensifier and CCD coupling optics with the range-variant point spread function of the primary optics.

Previous papers have sampled the PSF of a full-field AMCW lidar system – both for the purpose of extending DOF (Godbaz *et al.*, 2010) and for the purpose of mitigating multipath due to scattering in the range-imager optics (Mure-Dubois and Hügli, 2007a; Kavli *et al.*, 2008). While (Kavli *et al.*, 2008) utilised retroreflective dots, we utilise a fibre-optic based point source because it offers better performance while remaining subpixel in size. Attempting to measure both the defocus PSF and scattering effects at the same time is very difficult due to the extreme dynamic range required. In particular, temperature stability is extraordinarily important because even a slight change in bias can result in a massive redistribution of intensity from the defocus component of the PSF to the scattering component.

Fig. B.5a shows how the PSF changes over range. Allowing for the image intensifier and coupling optics, the PSF scales in the manner predicted by eqn. B.2. However the PSF samples close to the ranger are much more blurred than the PSFs of similar radius at a large distance – possible causes include optical abberations and scattering from the coded aperture. There is a slight pincushion effect on the PSF shape due to radial distortion from the component lenses.

(a) PSF Range Variance (Intensity)

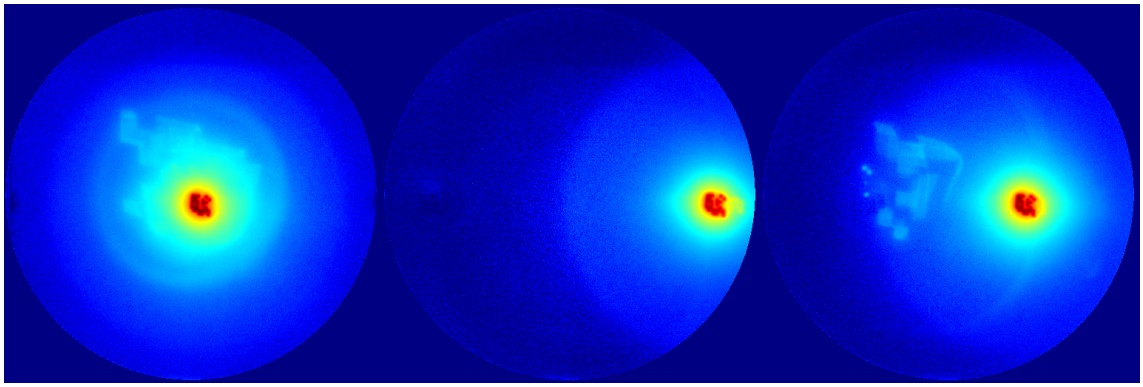

(b) PSF Spatial Variance (Log Intensity)

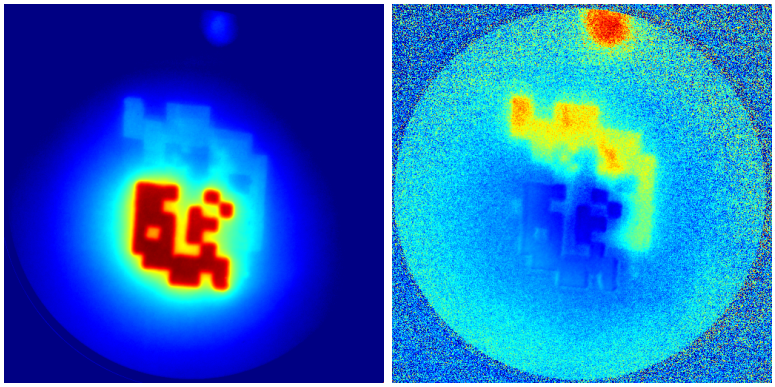

(c) Complex PSF – Log Intensity (left), Phase (right)

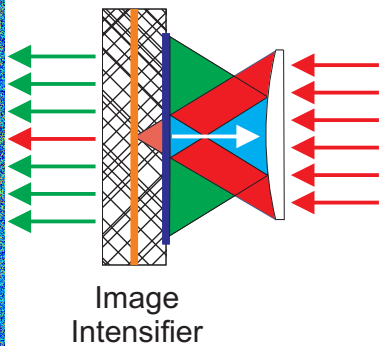

(d) PSF Formation Model

Figure B.5: Spatial and range variation in the coded-aperture range-imager PSF. In addition, the complex domain PSF is shown for the highly defocussed case – showing subtle phase shifts in the scattered light. In log-intensity images, red represents high intensity and blue low. In phase images, red represents greater distance and blue less distance. From these data we can determine the formation process for the most prominent scattering. In fig. B.5d, the initial aperture image (red) is reflected off the image intensifier and back to the final lens in the primary optics (cyan). Despite the low reflectivity of the lens, a significant amount of light is reflected back towards the image intensifier (green). The focal plane (orange) moves as the range to the point source changes, thus changing whether the primary PSF is inverted and the size of both the primary and reflected PSFs. The reflected PSF always has the same orientation.

The PSF also changes spatially; fig. B.5b shows the log intensity of the PSF in order to highlight subtle scattering effects. Most notably, there is an inverted image of the coded aperture present in the left-most image, which distorts and disappears as the point source is moved to the right side of the image – there is also a soft halo and some specular 'dots' (right-most image). Because of the spatial complexity of the PSF, we only utilise centred PSF samples, otherwise the large number of PSF samples would greatly increase the computational complexity of the restoration.

Calculating the phase of extremely dark scattered light is very difficult, so barring inordinately long exposure times or image intensifier burn-in due to oversaturation it is only possible to image the complex domain PSF with extreme defocus. High levels of defocus allow the intensity of the scattered light to be increased while keeping the maximum image intensity to a safe region for the image intensifier. Thus while we still model scattering, we cannot plausibly model the slight phase shifts inherent in the scattering PSF across the entire PSF gamut. Fig. B.5c shows the complex domain PSF for an extremely defocussed point source – note the low SNR for the darkest regions. Since the point source is within a few centimetres of the optics, the path length difference for light travelling through different sections of the aperture is visible – the path length varies by almost a centimetre within the primary/defocus PSF (blue/cyan). The reflections in the background have a much greater path length; the inverted aperture shape (yellow) has a path length at least 6cm longer than the primary PSF and the reflection at the top (red) has a path length at least 7.5cm longer. From this information, we can determine the formation process for the inverted image – this is given in fig. B.5d. We are unaware of any previous measurements of the complex domain PSF of a full-field AMCW lidar system.

B.4 Results and Discussion

Three different scenes were imaged of increasing spatial complexity: two boxes at varying distance from the ranger (fig. B.6), a garden gnome and several patterned boards (fig. B.7) and a chess set (fig. B.9). Due to the optical configuration, ground truth was unavailable. Slices through the first two scenes are shown in fig. B.8.

Scene one is an extremely simple scene containing two boxes printed with a test pattern. Fig. B.6a shows the initial blurred modulus, which using the blurred range information from fig. B.6b is restored to the point where most of the text can be read – a substantial improvement in DOF. Fig. B.8a shows how the phase is recovered during the deconvolution process – this graph shows a horizontal slice through the

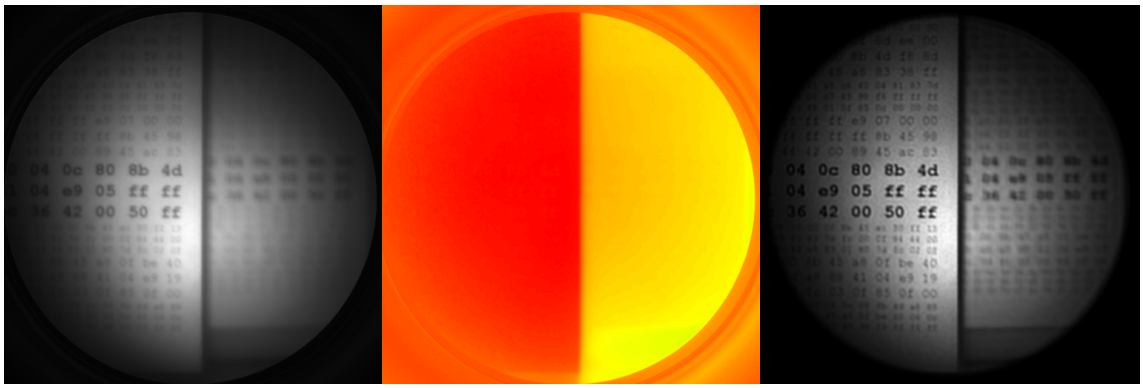

(a) Initial Modulus (b) Initial Phase (c) Deconvolved Modulus

Figure B.6: Scene One, pre- and post-deconvolution. For phase images, hue reperesents phase – the colour palette is cyclic in the same manner that phase is. The restoration of the hard phase discontinuity is shown in fig. B.8a.

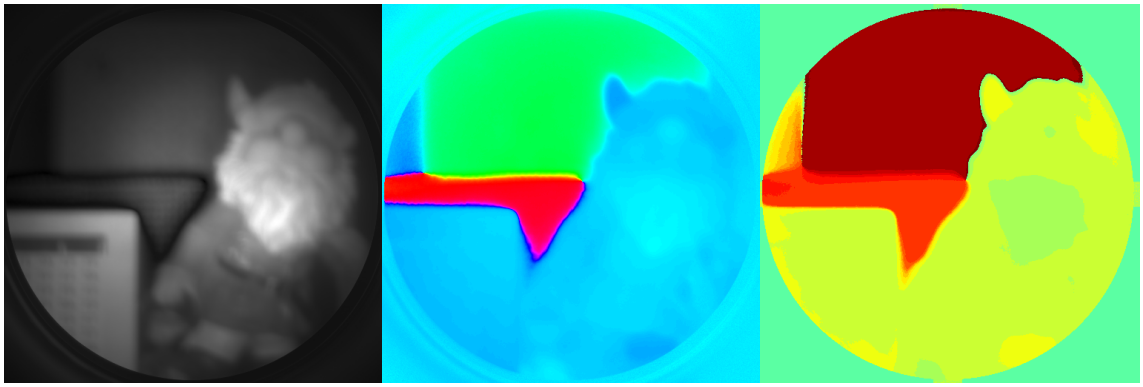

(a) Initial Modulus (b) Initial Phase (c) Initial PSF Number

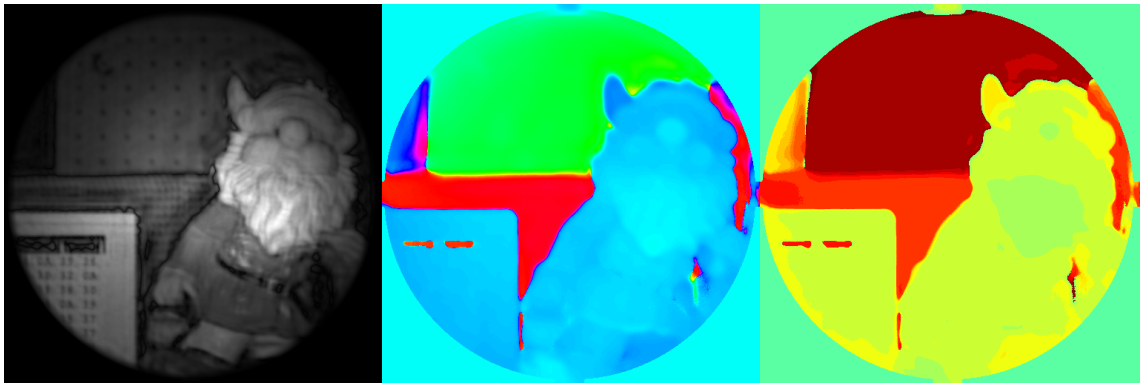

(d) Deconvolved Image Modulus (200 Iter.) (e) Deconvolved Image Phase (10 Iter.) (f) Final PSF Number (10 Iter.)

Figure B.7: Scene two, pre- and post-deconvolution. For phase images, hue reperesents phase – the colour palette is cyclic in the same manner that phase is. While a large number of iterations increases modulus resolution substantially, it tends to introduce unnecessary ringing into phase information.

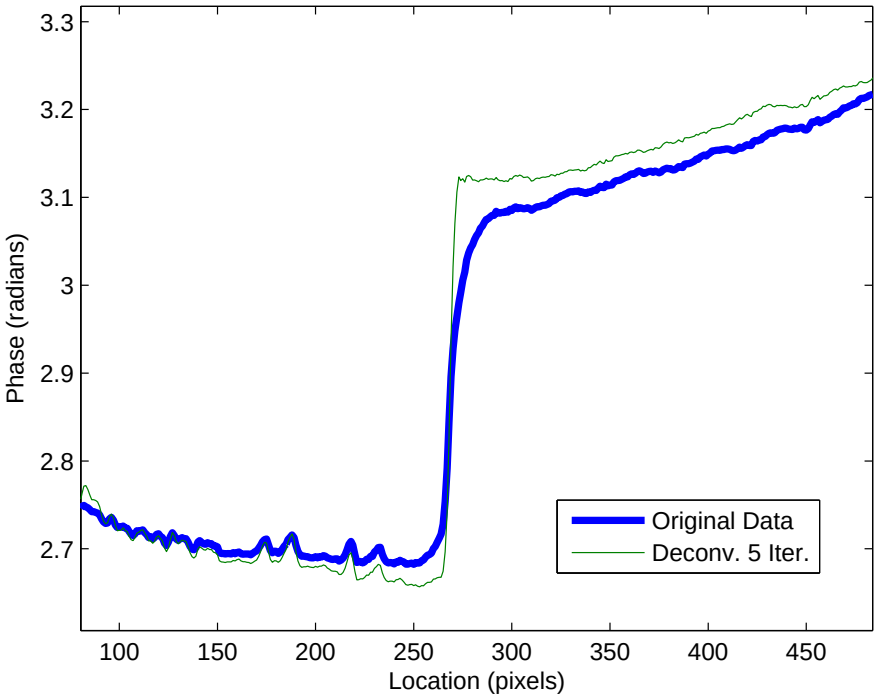

(a) Scene One

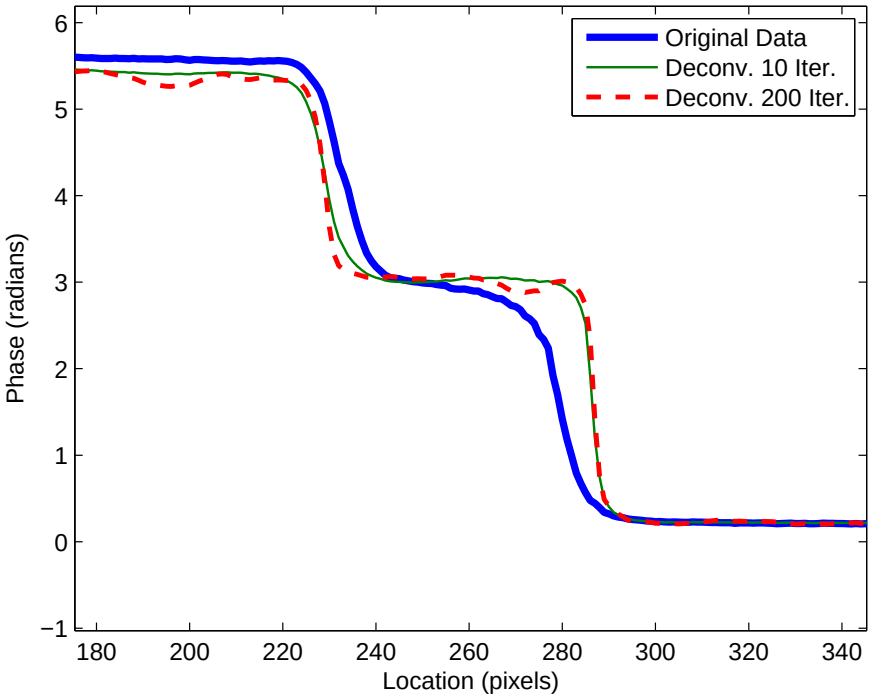

(b) Scene Two

Figure B.8: The phase of slices through scenes 1 and 2, before and after deconvolution. From a phase perspective, 200 iterations provides few benefits over 10 iterations. Fig. B.8a shows range-intensity coupling before and after restoration.

scene in the middle. The deconvolution process results in a substantial sharpening of the boundary between the two boxes as well as a significant shift in the range of the right hand box due to the partial removal of some scattered light. However there remains range-intensity coupling post-deconvolution most probably due to incomplete modelling of the spatial variance of scattering. It is extremely common in real images for range measurements to be shifted by 2-3cm due to scattered light.

Scene two is a more complicated scene. Due to the larger dynamic range, the modulus images of both scenes two and three use gamma compression of $\gamma = 0.5$. In this scene there is much more significant blurring and light scattering. Fig. B.7c shows the initial PSF used for each pixel, by the 10th iteration the PSF has changed in regions such as between the garden gnome and front-most board (fig. B.7f). In the final deconvolved range-image the modulus (fig. B.7d) and phase (fig. B.7e) components have substantially improved sharpness, although there are some notable artefacts. Most noticeable is the erroneous range value given for the black tape holding the test pattern onto the front board – the red range value is roughly equivalent to phase shifting the correct range value by π radians and this may indicate excessive compensation for scattering. There are ringing effects around the edges of objects such as the head of the gnome and the pattern. Like many real-life range-images, scene two contains a small region at the top left which is outside the range ambiguity interval – ie. due to the modulo 2π nature of phase, this region has been deconvolved by an incorrect PSF. This is unavoidable for real-world scenes unless range precision is sacrificed by using a particularly low modulation frequency or a phase unwrapping method utilised.

Unlike normal intensity images, complex domain range-images have some complicated behaviour around edges. In typical scenes the edges of objects are mixed pixels, however these tend to be heavily attenuated by the deconvolution process, resulting in dark bands at the boundaries of objects. A different type of dark band is seen in defocussed images where the objects have sufficiently different phases as to result in partial cancellation – these bands can be seen around the edges of the chess pieces in fig. B.9a. While a smoothness constraint may limit the impact of noise on the restoration, it also has a tendency to intensify dark bands between objects at significantly different ranges. If the aim of a restoration is to produce an in-focus pure intensity type image, then it may be more appropriate to deconvolve the mean signal intensity, which is essentially the total amount of light detected by the ranger (the mean of the correlation waveform over all phase steps). Albeit, most commercial ranger-imagers use a differential measurement process that removes this information from the raw measurements.

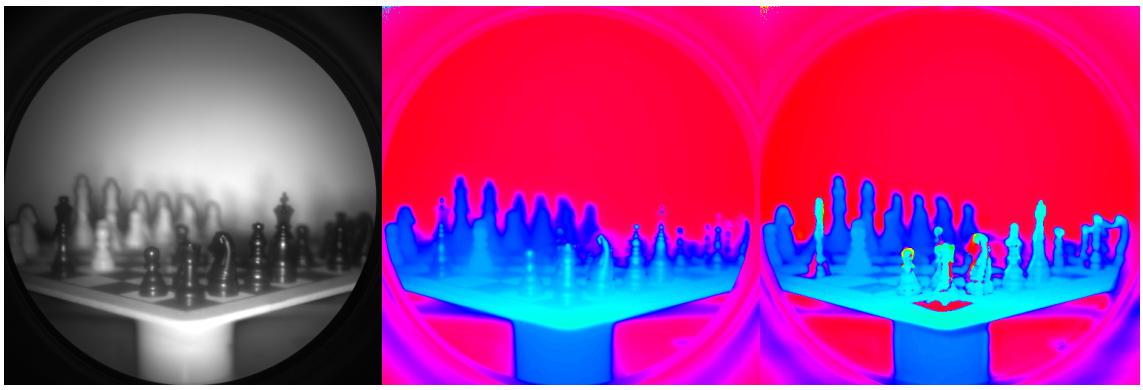

(a) Initial Modulus (b) Initial Phase (c) Failed Deconvolution (10 Iter.)

Figure B.9: Scene three, pre- and post-deconvolution. For phase images, hue reperesents phase – the colour palette is cyclic in the same manner that phase is. This scene suffers from severe multipath contamination, as shown by the range-intensity coupling for the black chess pieces and squares. A combination of multipath and high phase complexity results in an unsuccessful deconvolution.

Scene three demonstrates the current limitations of the restoration algorithm. The extreme range-intensity coupling is demonstrated by the black chess pieces. Regions such as the knight's head, which is near black, are perturbed by light scattered from the board in the background resulting in PSF misestimation in addition to having very complicated range content. Since none of the image is saturated, the regions with specular reflections have the most accurate range measurements, which are visibly different from adjacent areas. This is compounded by the fact that each component at a different range within a pixel has a different PSF. Successful restoration of this type of scene awaits a more advanced restoration algorithm that takes into account the range of possible components within each pixel rather than making a naïve assumption that each sample is of an unperturbed single component return.

B.5 Conclusion

In this paper we have designed a broadband coded-aperture for coding defocus so as to allow depth-of-field to be extended through deconvolution. We have demonstrated that the coded aperture design results in an improvement in restoration performance over a circular aperture and incorporated the coded aperture design into a real full-field AMCW lidar system. The range variation of the defocus and scattering PSFs was sampled and reflection off the image intensifier was isolated

as a significant contributor to scattered light. A naïve, proof-of-concept restoration algorithm was demonstrated to substantially improve the quality of some, but not all range-images captured using this new system – difficulties including misestimation of the restoration PSF due to multipath and the naïve assumption of a single component return.

Appendix C

Removal of Perturbations Due to Scattered Light Using Scene Texture

C.1 Introduction

Full-field amplitude modulated continuous wave (AMCW) lidar systems use the time-of-flight (TOF) principle to measure the range to a large number of points in a scene simultaneously. Whereas a normal camera captures an image where each pixel has an RGB colour value, a range image camera produces 2D matrices of radial distance values. These measurements can be converted to 3D Cartesian coordinates and used in applications like process control and robot navigation. A major unsolved problem in AMCW imaging is the mixed pixel problem, occurring when a pixel integrates range measurements from two or more different objects within a single pixel, causing erroneous range values. In order for full-field AMCW lidar to be a high quality replacement for traditional techniques such as line scanning triangulation systems and photogrammetry, mitigation techniques need to be developed.

C.1.1 Previous Work

Mixed pixels are most common around the edges of objects. Previous methods have included median filtering and normal angle thresholding techniques to enable removal of mixed points (Hebert and Krotkov, 1992; Adams, 1993; Vandapel *et al.*, 2004; Tang *et al.*, 2007). A more recent method developed by Larkins (Larkins *et al.*, 2009) uses clustering and parametric surface fitting to project mixed pixel back onto their parent surfaces; others use beat waveform harmonics to identify

the component returns (Godbaz *et al.*, 2008, 2009a). More insideous, however, are mixed pixels caused by scattered light. Fig C.1a and C.1b show an initial scene that suffers from only limited light scattering. If a much brighter object is moved into the scene (see figs. C.1c and C.1d), not only does the bright object now appear in the scene, but the ranges to the darker objects in the scene are changed due to scattered light. In particular, objects at the same distance from the camera with different brightnesses appear at different ranges – this is known as range-intensity coupling. Deconvolution has been applied to this problem (Mure-Dubois and Hügli, 2007a), but this is unable to handle light scattered from outside the field-of-view (FOV) and requires complex calibration and processing to fully compensate for the spatially variant light scattering.

C.1.2 Overview of New Mitigation Algorithm

We present a new method for mitigating range-intensity coupling due to scattered light in a full-field AMCW lidar system. The method does not require prior calibration of the ranging system and can migitate scattered light that is uncorrelated with the scene. The raw range data values corresponding to a particular object can be modelled as a linear function of reflectivity over small regions. By segmenting out each object and determining this linear function, it is possible to identify the true range to an object irrespective of any scattered light. By combining the linear functions corresponding to regions of several objects at different ranges, it is possible to create local estimates of light scattering. We show that subtracting these local estimates produces a substantial improvement in image quality.

The authors recommend printing in colour or viewing on a computer screen to ensure the clarity of the images in this paper.

C.1.3 AMCW Range-Imagers

AMCW range-imagers illuminate a scene with modulated light. Light takes longer to travel to objects farther away from the camera; this results in a phase shift in the modulated signal that is proportional to range. The range-imager measures this phase shift by correlating the illumination signal with a reference signal at the same or similar frequency. In the case of the University of Waikato Range-Imager (Dorrington *et al.*, 2007) the light is gain mixed with the reference signal using a modulated image intensifier. By using high frequency modulation signals that differ only very slightly, a technique known as heterodyning, the high frequency phase measurement problem is reduced to a low frequency phase measurement problem.

For an illumination modulation frequency, f_1, and reference signal frequency, f_0, the output beat waveform is at a frequency of $|f_1 - f_0|$. The phase offset of the beat waveform is the same as the phase offset of the returned light, thus proportional to the range of an object from the camera. The phase of the beat waveform is measured by Fourier transforming a sequence of images over time by an off the shelf charge coupled device (CCD) camera which is optically coupled to the rear of the image intensifier. While as few as three frames are sufficient to allow range to be measured, a larger number of frames prevents aliasing of harmonics above the Nyquist frequency and reduces the impact of noise.

The results presented in this paper were generated at frame rates of either 16 or 32 frames-per-second, two phase cycles per acquisiton, 0.5 Hz beat frequency. The laser modulation frequency was 40 MHz, which results in a range ambiguity distance of 3.75 metres. Having sets of optics both before and after the image intensifier makes the system particularly susceptable to scattered light; in particular, most of the light scattering in this system occurs post-intensifier. The image intensifier uses the photoelectric effect to generate photoelectrons which are then used to illuminate a phosphor screen. Because the light from this phosphor screen is not collimated, it reflects off the walls of the optical cavity. Depending on scene layout anywhere between 2% (typical) and 90% (extreme) of the total light measured at any one pixel is scattered from somewhere else in the scene. This is illustrated well by fig. C.2, which shows that in addition to specific specular reflections within the lens system, there is a diffuse scattered background.

C.1.4 Range Image Model

We henceforth distinguish between component returns and actual range measurements. A model of signal return intensity versus range, $f(r)$, can be considered to be a sparse spike train, composed of N individual component returns

$$f(r) = \sum_{i=0}^{N-1} a_i \delta(r - r_i), \tag{C.1}$$

where δ is the Dirac delta function, and a_i and r_i, are the amplitude and range of the ith component return, respectively. A well calibrated AMCW capture is equivalent to sampling $F(u)$, the Fourier transform of f, at a spatial frequency of $-c/2f_1$, where c is the speed of light. Thus a range measurement ξ is formed by

$$\xi = F\left(-\frac{c}{2f_1}\right) = \sum_{i=0}^{N-1} \eta_i = \sum_{i=0}^{N-1} a_i \mathrm{e}^{\theta_i}, \tag{C.2}$$

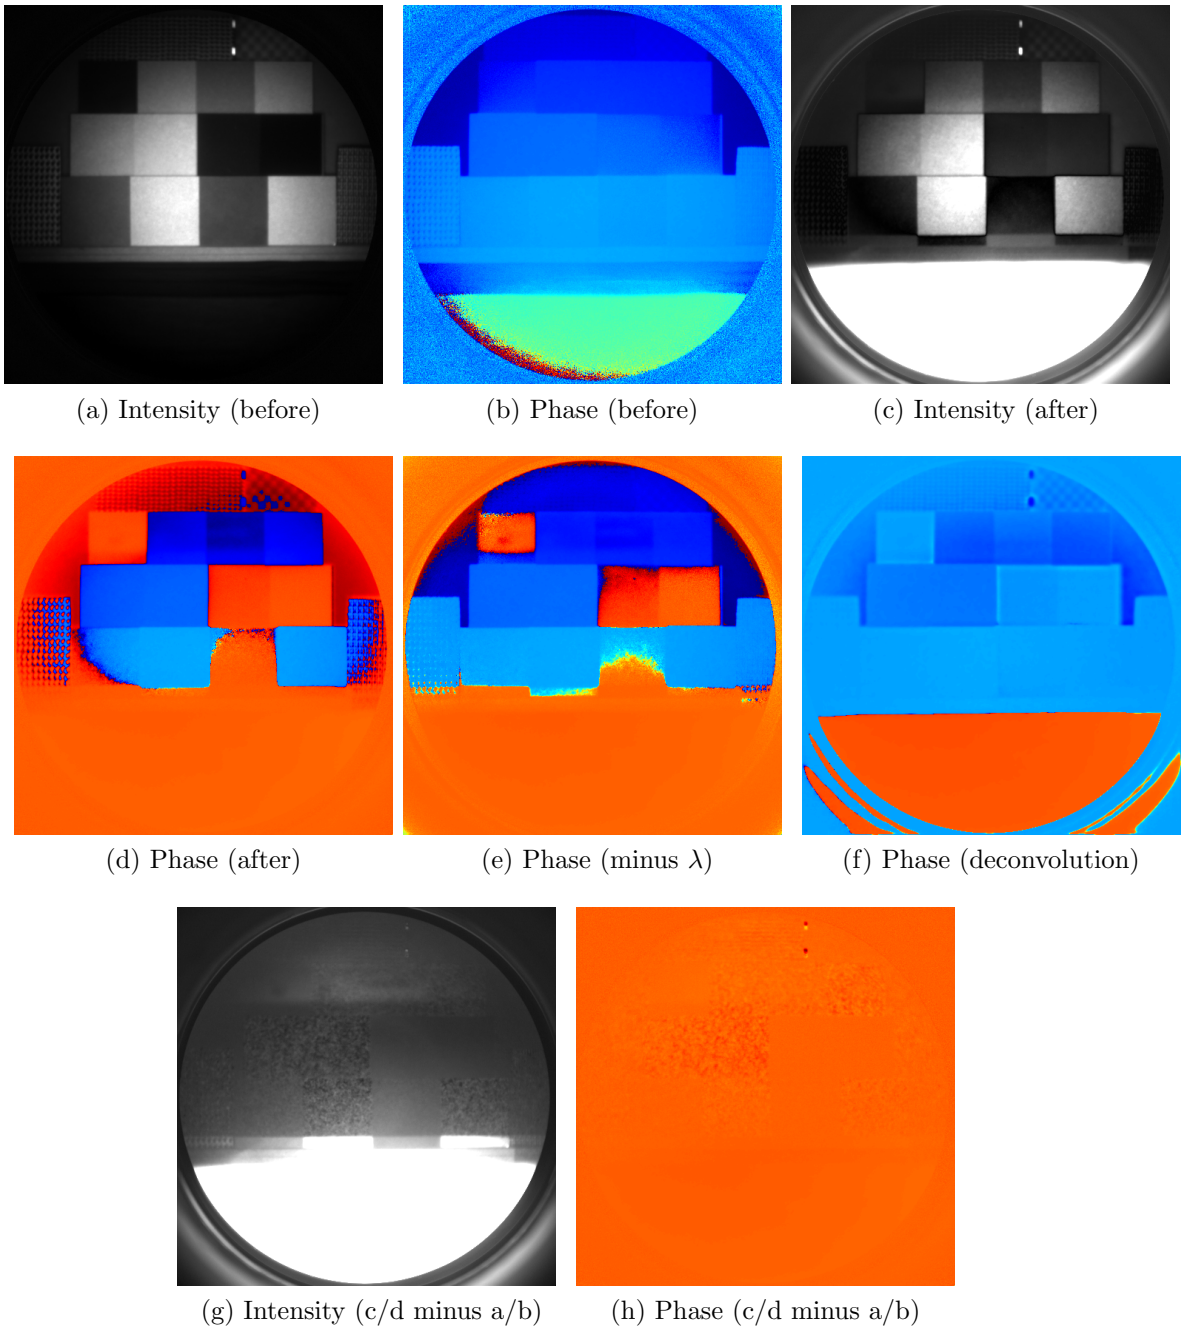

(a) Intensity (before) (b) Phase (before) (c) Intensity (after)

(d) Phase (after) (e) Phase (minus λ) (f) Phase (deconvolution)

(g) Intensity (c/d minus a/b) (h) Phase (c/d minus a/b)

Figure C.1: Range images of a real scene before and after the addition of a bright box which scatters light across the entire image and changes the range to objects in the scene. Dark red corresponds to a phase of 0 and dark blue to a phase of 2π – however mixed pixels result in erroneous ranges. C.1e and C.1f show the results of two different restoration methods, while C.1g and C.1h show the intensity and phase of the scattered light – determined by subtracting the first range image from the second.

where a single component return is represented as a phasor $\eta_i \in \mathbb{C}$. Each component return is composed of an active intensity a_i and a phase $\theta_i = 4\pi r_i f_1/c$. If we make the assumption that in the ideal, unblurred, unscattered image each pixel only contains one component return, then we can model the light scattering process as a spatially variant convolution. In actuality this is a naive assumption because light can be scattered onto the scene from bright objects outside the FOV, and mixed pixels occur frequently around the edges of objects where each pixel integrates light from two or more objects. For the ideal image, $\mathrm{H} \in \mathbb{C}^m$, where m is the number of pixels, the recorded data are formed by

$$\Xi = \mathrm{H} \star_{sv} h + \epsilon, \tag{C.3}$$

where h is the point spread function (PSF) and ϵ is noise.

C.1.5 Range-Intensity Coupling

If we have a large object in the scene that has a textured/patterned surface and also have significant light scattering, then we can model the measured phase of the object as a function of the surface brightness,

$$\theta = \arg(\xi(\rho)) = \arg(\rho\vec{\eta} + \lambda), = \arg(\eta + \lambda) \tag{C.4}$$

where $\rho \in \mathbb{R}^+$ is the brightness, $\lambda \in \mathbb{C}$ is the scattered light, which we assume to be near constant in the region of the object, and $\vec{\eta} \in \mathbb{C}$ is a unit phasor corresponding to the underlying range.

There are two approaches to determining η, one is to assume that λ is solely a function of η and apply deconvolution techniques to determine H directly; the other is to determine λ separately and then subtract it from ξ.

C.1.6 Weighted Least Squares Linear Fitting with Windowing

The new algorithm detailed in this paper is based around least squares linear fitting across an entire image in the Fourier domain. For a model $\alpha + \beta x$, this works by finding the values of α and β such that

$$(\alpha_i, \beta_i) = \arg\min_{(\alpha_i,\beta_i)}(y - \alpha_i - \beta_i x)^2 \otimes \omega \otimes \zeta[i] \tag{C.5}$$

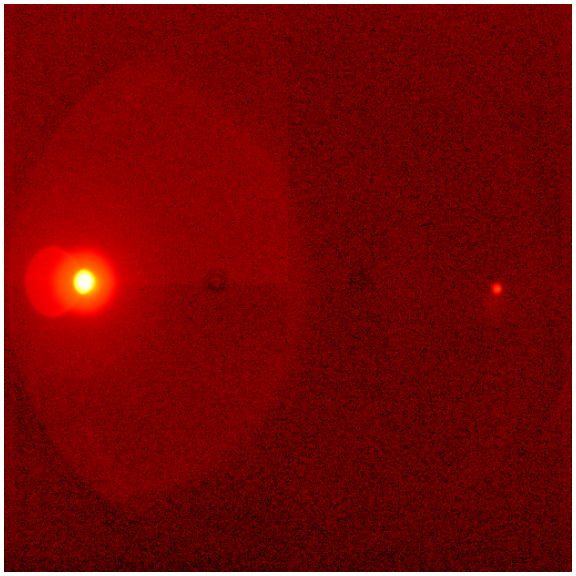

(a) Left

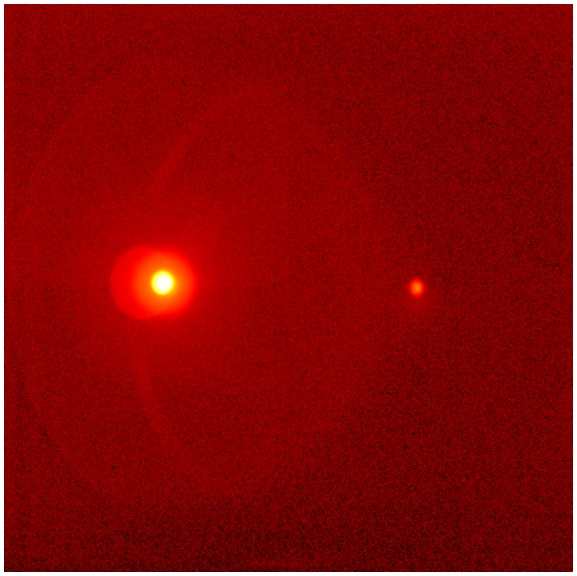

(b) Left/Centre

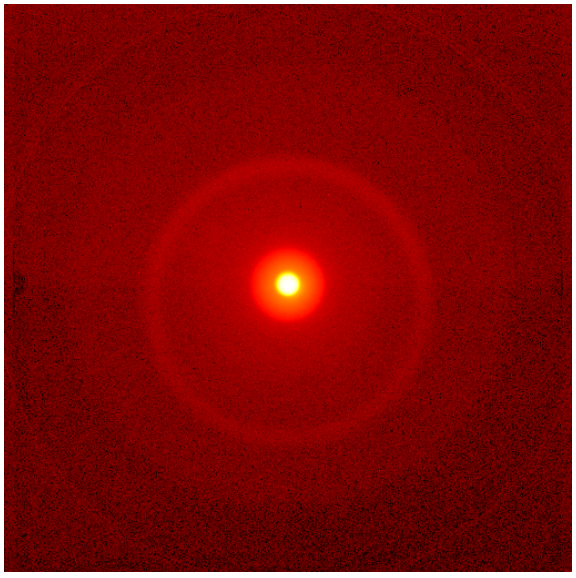

(c) Centre

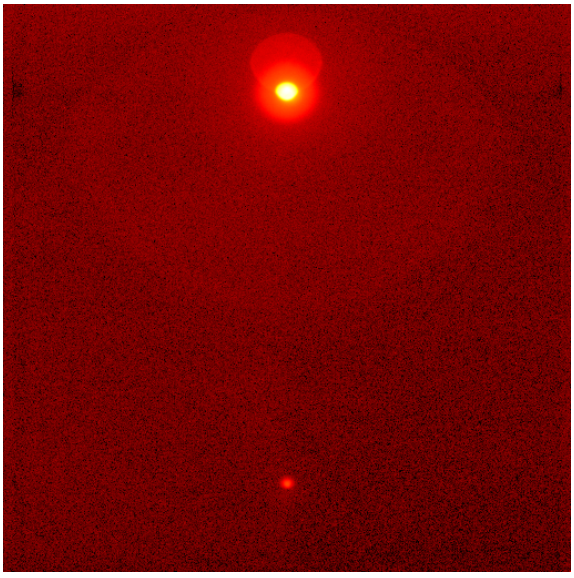

(d) Top

Figure C.2: PSF sampled at different pixel locations. The PSF is highly anisoplanatic, varying in a rotationally symmetric manner. Due to the high dynamic range, the logarithm of the intensity is plotted here, black/red is dark and white bright. There is a reflection that 'follows' the light source around, always appearing on the opposite side of the principal point.

for each pixel i, where x and y are vectors corresponding to the independent and dependent variables, ω is data weighing and ζ is a windowing function. In closed form using convolutions the solution is

$$\Phi(z) = (z \otimes \omega) \star \zeta, \tag{C.6}$$

$$\Phi_\emptyset = \omega \star \zeta, \tag{C.7}$$

$$\Upsilon_\alpha(x, y, \omega, \zeta) = \frac{\Phi(x^2) \otimes \Phi(y) - \Phi(x \otimes y) \otimes \Phi(x)}{\Phi_\emptyset \otimes \Phi(x^2) - \Phi(x) \otimes \Phi(x)}, \tag{C.8}$$

$$\Upsilon_\beta(x, y, \omega, \zeta) = \frac{\Phi_\emptyset \otimes \Phi(x \otimes y) - \Phi(y) \otimes \Phi(x)}{\Phi_\emptyset \otimes \Phi(x^2) - \Phi(x) \otimes \Phi(x)}, \tag{C.9}$$

where $\Phi(z)$ is a weighted convolution of a data element with the windowing function and $\Phi_\emptyset$ is a convolution of the weights with the windowing function. Least squares fitting is notated by two functions: $\Upsilon_\alpha(x, y, \omega, \zeta)$ which returns the value of α and $\Upsilon_\beta(x, y, \omega, \zeta)$ returns the value of β.

C.1.7 Segmentation of Range-Images

Range images can be segmented by identifying contiguous regions larger than a certain size, where all pixels are connected to each other via at least one path where the Δ phase between adjacent connected pixels is less than a certain threshold. In this paper we use a deterministic region-growing method, chosen primarily for ease of implementation. The method assigns each pixel an initial unique region code and then progressively merges regions together. At each iteration, the connected neighbours of each pixel are scanned and the current pixel is set to the minimal region code. Over time any region changes propagate over the image, eventually converging. Any regions of fewer than 100 pixels are discarded.

C.2 Algorithms

C.2.1 Phase From Texture

We can rewrite eqn. C.4 as

$$\Im(\xi) = \tan(\theta + \pi N)(\rho - \Re(\lambda)) + \Im(\lambda), \tag{C.10}$$

$$= \alpha + \beta \Re(\xi), \tag{C.11}$$

where $N \in \mathbb{Z}$ is initially unknown. If we assume λ to be constant over a region of the scene, then we can fit a linear model to $\Im(\xi)$ versus $\Re(\xi)$ using a least squares ap-

proach. In reality, the scattering PSF has a significant local component (see fig. C.2), but an assumption of large scale smoothness is necessary to allow estimation of λ. After segmenting the range image into separate contiguous surfaces as described in section C.1.7, we apply eqn. C.8 and C.9 to find

$$\alpha = \sum_{i=0}^{m-1} \Upsilon_\alpha(\Re(\Xi), \Im(\Xi), \omega \otimes \zeta[i], \zeta) \otimes \zeta[i], \tag{C.12}$$

$$\beta = \sum_{i=0}^{m-1} \Upsilon_\beta(\Re(\Xi), \Im(\Xi), \omega \otimes \zeta[i], \zeta) \otimes \zeta[i], \tag{C.13}$$

where m is the number of range-image segments, $\zeta[i]$ is a mask corresponding to the ith range-image segment, ζ is a Gaussian windowing function with $\sigma = 4$ pixels and ω are data weights. In more advanced experiments the data weights could be calculated based on known noise statistics, but for the purposes of this paper, we assign values of 1 to regions within the image intensifier field of view, and 0 to regions outside.

The mean signal intensity for a pixel, is given by

$$\chi_i = \int f_i(r) dr, \tag{C.14}$$

which is the sum of total integrated intensity and any background light. We can use information about the mean signal intensity to disambiguate the possible values of θ (remove N from eqn. C.10). This ambiguity occurs because it is impossible to determine from a line's slope which direction it points in. Since the highest integrated intensity values tend to be correlated with greatest possible values of ρ, we can use the covariance of χ with $\Re(\Xi)$ to disambiguate θ for each pixel in the image by

$$\kappa = \mathrm{cov}(\chi, \Re(\Xi), \zeta), \tag{C.15}$$

$$\theta = \arctan(\beta \otimes \kappa, \kappa), \tag{C.16}$$

where $\mathrm{cov}(x, y, \zeta)$ is the covariance operator over the data vectors x and y with the same windowing function (ζ) as above. The estimates of θ returned by this method are henceforth referred to as range/phase from texture.

C.2.2 Estimating Light Scattering From Object Texture

In order to find λ, we combine samples of α and β at multiple different ranges. For example, in fig. C.1, there are arbitrarily placed patterned regions at the left, top and right of the two scenes. Although relatively dark in the intensity images, the patterned regions are clearly visible in C.1d, where the scattering has resulted in range-intensity coupling. We now utilise these patterned regions to identify the nature of the scattered light. In fig. C.3 $\Im(\Xi)$ is plotted versus $\Re(\Xi)$ for two regions at different ranges. Since regions at different ranges result in linearly independent relationships, λ is the phasor corresponding to the intersection of the fits. In this case the intersection of the two fits is at $\lambda = 1084 + 3091j$. The patterned regions are widely spaced, which is adequate in cases where the scattering is largely constant across the entire image, but this does not always hold due to the local contributions of the scattering PSF. Fig. C.3 illustrates the limitations of estimates based on widely spaced samples because $\lambda = 1084 + 3091j$ implies that $\rho < 0$ for some data points, which is physically impossible. In order to avoid the limitations of global estimates, we attempt to generate local estimates of λ over a windowed region. Under the assumption of sufficient patterned sufaces across the entire field of view, then the problem of finding λ can be reduced to a least squares linear fit, viz

$$\lambda = \Upsilon_\beta(\beta, \alpha, \omega, \zeta_l) + j\Upsilon_\alpha(\beta, \alpha, \omega, \zeta_l), \tag{C.17}$$

where ζ_l is a Gaussian windowing function, $\sigma \geq 150$ pixels for our experiments, depending on the amount of useful information in the scene. An estimate of the unscattered data, H, can now be calculated from Ξ by subtracting λ.

C.2.3 Deconvolution

An alternative approach to removing scattered light is to apply deconvolution techniques. The PSF for our ranger system is shown in fig. C.2 – note that the PSF is highly spatially variant, albeit the variance is circularly symmetric in behaviour. There are very noticible lens reflections and halos. In order to achieve a high signal-to-noise ratio (SNR), the camera was deliberately defocussed – the PSF is not only pixel location variant, but also based on range and focal parameters. Attempting to deconvolve a range-image directly using the defocussed PSF resulted in severe ringing effects. A PSF model was generated by least squares fitting the sum of three concentric, centred weighted pillbox PSFs and one Gaussian to the centre PSF sample. This required operator intervention in order to find good solutions. Subsequently the central pillbox and Gaussian were replaced by a point source to re-

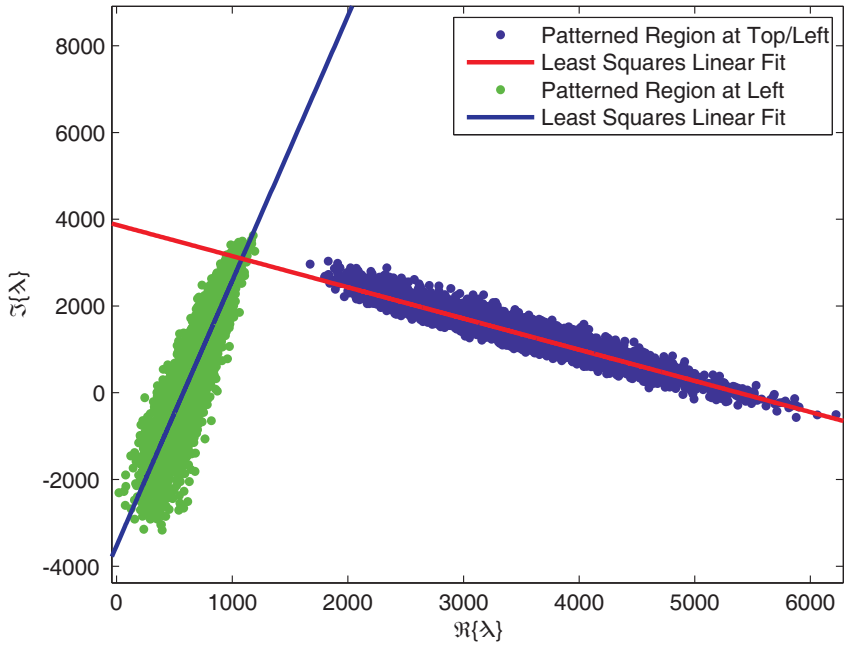

Figure C.3: The linear relationship between $\Re(\Xi)$ and $\Im(\Xi)$. Combining linear fits from two textured regions at different ranges allows λ to be estimated. In this case, the estimate of λ is based on inconsistant data – the intersection of the two fits requires $\rho < 0$ for some data points, which is physically impossible.

move the defocus and reduce ringing. Rather than the method used by Mure-Dubois and Hügli (2007a), we implemented a spatial derivative regularised, isoplanatic deconvolution similar to that used in the positive real domain by Levin *et al.* (2007), that is

$$\mathrm{H} = \arg\min_{x}(||\Xi - x \star h||_2^2 + \tau||lx||_2^2), \tag{C.18}$$

where l is a Laplacian filter, τ a regularisation constant and H the recovered image. This method makes the assumption that the spatial derivatives are Gaussian distributed. The solution was implemented as direct inversion of the normal equations using 2D Fourier transforms. This requires accurate calibration of the PSF, and cannot handle light scattered from outside the image.

C.3 Methodology and Results

C.3.1 Light Scattering Behaviour in a Typical Scene

Fig. C.1 illustrates the effect of light scattering on a typical scene. Figs. C.1a and C.1b are a test scene consisting of a staircase structure with square patches

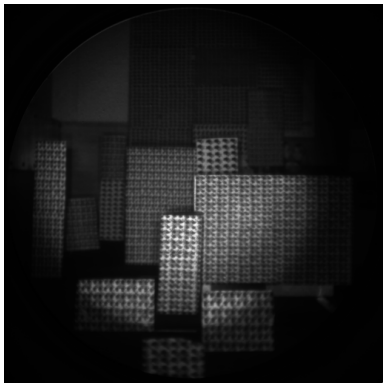

(a) Intensity (Before)

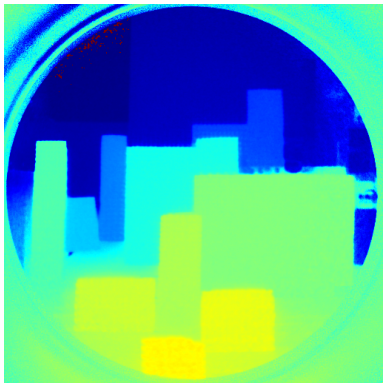

(b) Phase (Before)

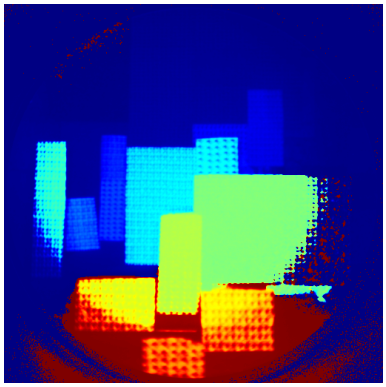

(c) Phase (Simulated Scattering)

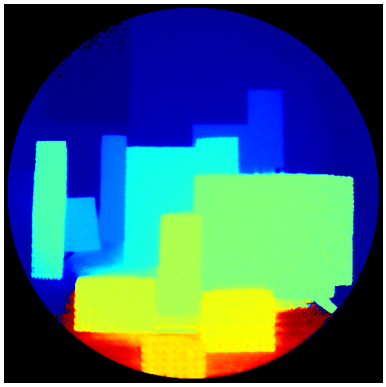

(d) Phase (Minus λ, Segmented)

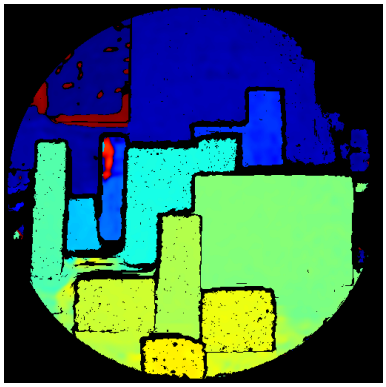

(e) Phase (From Texture, Segmented)

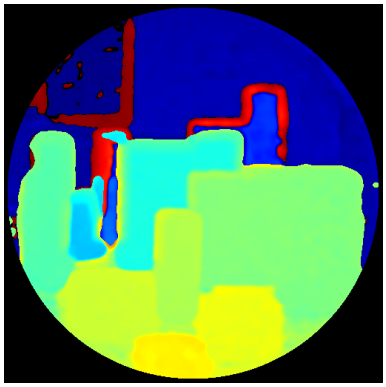

(f) Phase (From Texture, Unsegmented)

Figure C.4: The textured scene restoration process. C.4a and C.4b are the intensity and phase of the original highly textured scene. Red represented objects close to the camera and blue objects far away. Adding simulated scattered light gives C.4c, in which the patterned surfaces clearly show range-intensity coupling. Estimating λ using the segmented method and subtracting it from C.4c gives C.4d. This is a substantial improvement, but darker regions are restored less well. Applying the phase from texture algorithm in the segmented and unsegmented versions results in C.4e and C.4f respectively.

of different reflectivities. Note that in C.1b there is a very slight range-intensity coupling. Figs. C.1c and C.1d show the scene after the addition of a very bright box in the foreground. The greatest changes in range occurred in regions where the reflectivity of the staircase was very low, such as the middle step on the far right. Using the textured regions on the left, and top/left a global estimate for λ was calculated (as described in section C.2.2). These textured regions are quite widely spaced, meaning that only large scale scattering can be resolved. Subtraction of λ gives fig. C.1e, which shows significantly reduced range-intensity coupling, although the image has not been fully restored. This can be attributed to misestimation of λ. Fig. C.1f shows the results of applying deconvolution to the image. The use of

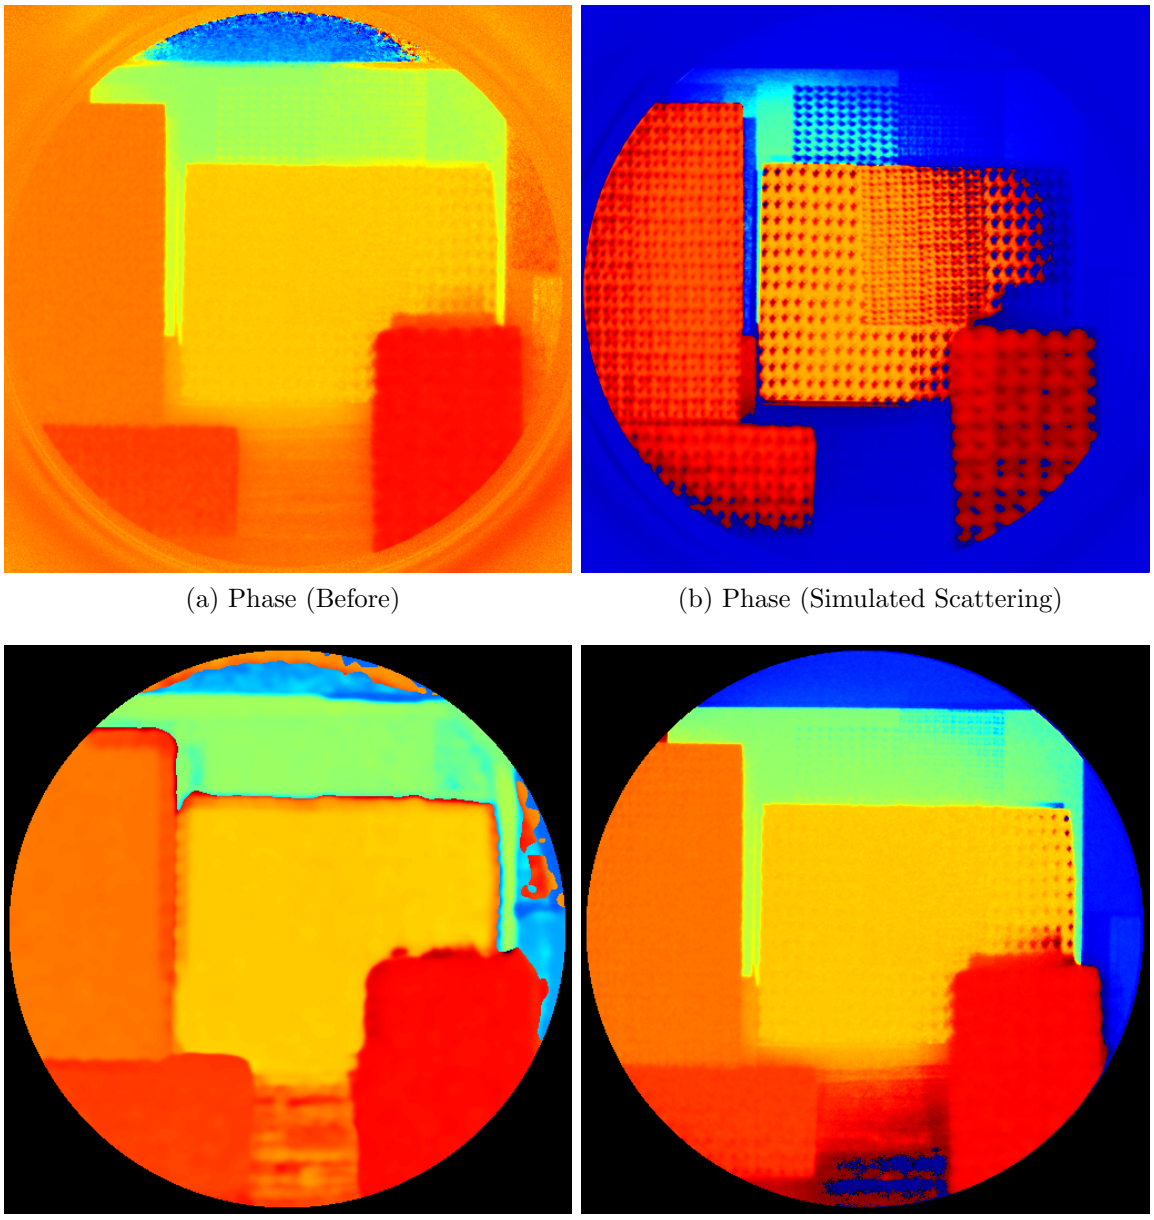

(a) Phase (Before)

(b) Phase (Simulated Scattering)

(c) Phase (From Texture)

(d) Phase (Minus λ)

Figure C.5: Restoration of a textured scene. In this particular case, we have not segmented the image before processing it. This makes the algorithm a lot faster, but means that the phase from texture information is blurred around the edges of objects.

Fourier convolution and assumption of isoplanaticity has resulted in underestimated phase across the top of the image. The results are quite different to those produced by subtracting λ. By subtracting the original, largely unscattered image from the scattered image it is possible to estimate the amount of light scattered within the system. This is shown in figs. C.1g and C.1h. The brightest regions of the background have much higher noise levels than the darker regions because the variance of a Poisson distribution is linearly proportional to intensity. Even when the intensity has been reduced by subtraction the noise still persists. The region at the top of the image exactly opposite the scattering object is slightly brighter than the rest of the scattering, which may be contributed to by the reflection seen in fig. C.2 that is always on the opposite side of the principal point from the illumination source.

C.3.2 Results of New Algorithms

Two different patterned scenes were imaged. In order to provide rough ground truth, these scenes were set up in such a manner as to minimise the amount of naturally scattered light. Simulated scattered light (uncorrelated with the scene) was then added to produce corrupted range data as input to the algorithm for testing. However, since these were based on real data, there is a small amount of natural scattering variance across the image.

Fig. C.4 was processed using the range-image segmentation algorithm from section C.1.7. Because the segmentation relies on phase differences between adjacent pixels, the corrupted range data results in the scene being incorrectly segmented. In order to accurately segment the scene a rough estimate of λ is required – this is provided by running the scene through the algorithm once without any segmentation. An initial estimate of λ was then calculated and subtracted from Ξ in order to produce a correct range image segmentation model. The results of the segmented phase from texture algorithm are shown in fig. C.4e, this shows that the phase has been correctly recovered everywhere except for a small box on the left. The resultant phase measurements are far more useful for most practical applications than the initial fig. C.4c phases. Estimating λ via the segmented algorithm and subtracting from Ξ gives fig. C.4d, which is correct in all the patterned regions, however small parts of the table, which are unpatterned, have retained the perturbed phases from fig. C.4c.

Fig. C.5 was processed in a similar fashion, but without any image segmentation. Fig. C.5b shows severe range-intensity coupling, which is substantially reduced in the output image fig. C.5d – however, the improvement is not as good as fig. C.4d. This is not entirely due to the lack of segmentation, but is contributed to by the scene

layout and much smaller number of objects. As with fig. C.4f, the unsegmented phase from texture image (fig. C.5c) contains erroneous values around edges of objects due to the size of the Gaussian window function ζ. Reducing σ could reduce these edge effects, but at the expense of increasing the impact of noise on the final results.

C.4 Summary of Results

We have presented a new method for mitigating range-intensity coupling in full-field AMCW lidar images. The method is robust to light scattered from outside the FOV and does not require complex hardware calibration like previous deconvolution based methods. Phase information can also be directly extracted from patterned surfaces and the intensity and range of the scattered light estimated independently of the output range measurements. The method results in a substantial improvement in range data quality and can be implemented in either a fast unsegmented or slower, higher quality variant. Future development work includes analysis to determine optimal scene layout to facilitate accurate removal of scattered light and alteration of this method to enable its application to the object edge mixed pixel problem.

References

Adams, M. D. Amplitude modulated optical range data analysis in mobile robotics. In: *Proceedings of 1993 IEEE International Conference on Robotics and Automation, Vol. 2*, pp. 8–13 (1993).

Adams, M. D. and P. J. Probert. The interpretation of phase and intensity data for AMCW light detection sensors for reliable ranging. *Int. J. Rob. Res.*, **15(5)**, pp. 441–458 (1996).

Bamji, C. Method and system for lossless dealiasing in time-of-flight (TOF) systems. US Patent No. 7791715B1 (2010).

Besl, P. J. Active, optical range imaging sensors. *Mach. Vision Appl.*, **1(2)**, pp. 127–152 (1988).

Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA (2006). ISBN 0387310738.

Blais, F. Review of 20 years of range sensor development. In: *Videometrics VII*, edited by J. S. Walton, volume 5013, pp. 62–76 (2003).

Boden, A. F., D. C. Redding, R. J. Hanisch, and J. Mo. Massively parallel spatially-variant maximum likelihood image restoration. In: *Astronomical Data Analysis Software and Systems V*, edited by J. Barnes, volume 101, pp. 131–134 (1996).

Büttgen, B., F. Lustenberger, and P. Seitz. Demodulation pixel based on static drift fields. *IEEE Transactions on Electron Devices*, **53**, pp. 2741–2747 (2006).

Büttgen, B., T. Oggier, M. Lehmann, R. Kaufmann, S. Neukom, M. Richter, M. Schweizer, D. Beyeler, R. Cook, C. Gimkiewicz, C. Urban, P. Metzler, P. Seitz, and F. Lustenberger. High-speed and high-sensitive demodulation pixel for 3D imaging. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, edited by B. D. Corner, P. Li, & M. Tocheri, volume 6056 of *Society*

of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pp. 22–33 (2006).

Büttgen, B. and P. Seitz. Robust optical time-of-flight range imaging based on smart pixel structures. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, **55(6)**, pp. 1512–1525 (2008). 1549-8328.

Chang, C.-I., X.-L. Zhao, M. L. G. Althouse, and J. J. Pan. Least squares subspace projection approach to mixed pixel classification for hyperspectral images. *IEEE Transactions on Geoscience and Remote Sensing*, **36**, pp. 898–912 (1998).

Chaudhuri, S. and A. N. Rajagopalan. *Depth from defocus: a real aperture imaging approach*. Springer Verlag (1999).

Chauve, A., C. Mallet, F. Bretar, S. Durrieu, M. P. Deseilligny, and P. W. Processing full-waveform lidar data: Modelling raw signals. In: *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 2007* (2007).

Christie, S., S. L. Hill, B. Bury, J. O. Gray, and K. M. Booth. Design and development of a multi-detecting two-dimensional ranging sensor. *Measurement Science and Technology*, **6(9)**, p. 1301 (1995).

Conroy, R. M., A. A. Dorrington, A. D. Payne, R. Kunnemeyer, and M. J. Cree. A power-saving modulation technique for time-of-flight range imaging sensors. In: *SPIE Optical Metrology, Videometrics, Range Imaging, and Applications*, volume 8085. Munich, Germany (2011).

Cree, M. J., A. A. Dorrington, R. M. Conroy, A. D. Payne, and D. A. Carnegie. The Waikato range imager. In: *Image and Vision Computing New Zealand (IVCNZ'06)*, pp. 233–238. Gt. Barrier Is. New Zealand (2006).

Cree, M. J., J. P. Godbaz, R. L. Larkins, W. H. Round, L. Streeter, A. A. Dorrington, R. Knnemeyer, A. D. Payne, and D. Worsley. Computer vision and image processing at the university of waikato. In: *Image and Vision Computing New Zealand (IVCNZ'10)*. Queenstown (2010).

Dändliker, R., Y. Salvadé, and E. Zimmermann. Distance measurement by multiple-wavelength interferometry Mesure de distance par interférométrie à plusieurs longueurs d'onde. *Journal of Optics*, **29**, pp. 105–114 (1998).

Dhond, U. R. and J. K. Aggarwal. Structure from stereo – a review. *IEEE Transactions on Systems, Man and Cybernetics*, **19(6)** (1989).

Donoho, D., , D. L. Donoho, and M. Elad. Maximal sparsity representation via l 1 minimization. In: *the Proc. Nat. Aca. Sci. 100*, pp. 2197–2202 (2003).

Dorrington, A. A., M. J. Cree, D. A. Carnegie, A. D. Payne, R. M. Conroy, J. P. Godbaz, and A. P. P. Jongenelen. Video-rate or high-precision: a flexible range imaging camera. In: *SPIE 6813 - Image Processing: Machine Vision Applications*, volume 6813, p. 681307. SPIE, San Jose, CA, USA (2008).

Dorrington, A. A., M. J. Cree, A. D. Payne, R. M. Conroy, and D. A. Carnegie. Achieving sub-millimetre precision with a solid-state full-field heterodyning range imaging camera. *Meas. Sci. and Tech.*, **18(9)**, pp. 2809–2816 (2007).

Dorrington, A. A., J. P. Godbaz, M. J. Cree, A. D. Payne, and L. V. Streeter. Separating true range measurements from multi-path and scattering interference in commercial range cameras. In: *SPIE 7864 - Three-Dimensional Imaging, Interaction, and Measurement*, edited by J. A. Beraldin, G. S. Cheok, M. B. McCarthy, U. Neuschaefer-Rube, A. M. Baskurt, I. E. McDowall, and M. Dolinsky, p. 786404. San Francisco Airport, California, USA (2011).

Dorrington, A. A., A. D. Payne, and M. J. Cree. An evaluation of time-of-flight range cameras for close range metrology applications. In: *ISPRS Commission V Mid-Term Symposium 'Close Range Image Measurement Techniques'*, edited by J. P. Mills, D. M. Barber, P. Miller, and I. Newton, volume XXXVIII, Part 5, pp. 201–206. Newcastle upon Tyne, United Kingdom (2010).

Dowski, E. R. and G. E. Johnson. Wavefront coding: a modern method of achieving high-performance and/or low-cost imaging systems. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 3779, pp. 137–145 (1999).

Droeschel, D., D. Holz, and S. Behnke. Multi-frequency phase unwrapping for time-of-flight cameras. In: *Proceedings of IEEE/RSJ International Conference on Intellegent Robots and Systems (IROS)* (2010).

Eloranta, E. W. Practical model for the calculation of multiply scattered lidar returns. *Appl. Opt.*, **37(12)**, pp. 2464–2472 (1998).

Falie, D. 3D image correction for time of flight (ToF) cameras. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7156 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (2008).

Falie, D. Improvements of the 3d images captured with time-of-flight cameras. Technical report, Politehnica University of Bucharest (2009).

Falie, D. and V. Buzuloiu. Noise characteristics of 3D time-of-flight cameras. In: *Proceedings of 2007 International Symposium on Signals, Circuits and Systems (ISSCS 2007)* (2007).

Falie, D. and V. Buzuloiu. Distance errors correction for the times of flight (tof) cameras. In: *Proceedings of the IEEE International Workshop on Imaging Systems and Techniques, IST 2008* (2008a).

Falie, D. and V. Buzuloiu. Wide range time of flight camera for outdoor surveillance. In: *Proceedings of Microwaves, Radar and Remote Sensing Symposium, MRRS 2008* (2008b).

Foix, S., G. Alenya, and C. Torras. Lock-in time-of-flight (tof) cameras: A survey. *Sensors Journal, IEEE*, **PP(99)**, p. 1 (2011).

Frank, M., M. Plaue, H. Rapp, U. Köthe, B. Jähne, and F. A. Hamprecht. Theoretical and experimental error analysis of continuous-wave time-of-flight range cameras. *Optical Engineering*, **48(1)** (2009).

Fraser, D., A. Lambert, M. R. S. Jahromi, M. Tahtali, and D. Clyde. Anisoplanatic image restoration at adfa. Proc. of Seventh International Conference on Digital Image Computing: Techniques and Applications. Sydney, Australia (2003).

Frey, B. J., R. Koetter, and P. N. Very loopy belief propagation for unwrapping phase images. In: *Proceedings of Neural Information Processing Systems 14, NIPS 01* (2001).

Fuchs, S. Multipath interference compensation in time-of-flight camera images. In: *Proceedings of the 2010 20th International Conference on Pattern Recognition*, ICPR '10, pp. 3583–3586. IEEE Computer Society, Washington, DC, USA (2010). ISBN 978-0-7695-4109-9.

Fuchs, S. and G. Hirzinger. Extrinsic and depth calibration of tof cameras. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2008* (2008).

Ghiglia, D. C. and L. A. Romero. Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods. *J. Opt. Soc. Am. A*, **11(1)**, pp. 107–117 (1994).

Godbaz, J., M. Cree, and A. Dorrington. Extending amcw lidar depth-of-field using a coded aperture. In: *ACCV 2010*, edited by R. Kimmel, R. Klette, and A. Sugimoto, volume 6495 of *Lecture Notes in Computer Science*, pp. 397–409. Springer Berlin / Heidelberg (2011).

Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Mixed pixel return separation for a full-field ranger. In: *Image and Vision Computing New Zealand 2008 (IVCNZ'08)*, pp. 1–6 (2008).

Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Multiple return separation for a full-field ranger via continuous waveform modelling. In: *SPIE 7251 - Image Processing: Machine Vision Applications II*, volume 7251, p. 72510T. SPIE, San Jose, CA, USA (2009a).

Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Undue influence: Mitigating range-intensity coupling in amcw 'flash' lidar using scene texture. In: *Image and Vision Computing New Zealand 2009 (IVCNZ09)*, pp. 304–309 (2009b).

Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Blind deconvolution of depth-of-field limited full-field lidar data by determination of focal parameters. In: *SPIE 7533 - Computational Imaging VIII*, volume 7533, p. 75330B. SPIE, San Jose, California, USA (2010).

Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Closed-form inverses for the mixed pixel/multipath interference problem in AMCW lidar. In: *Computational Imaging X, Proceedings of the SPIE, Vol. 8296* (2012a).

Godbaz, J. P., M. J. Cree, and A. A. Dorrington. Understanding and ameliorating non-linear phase and amplitude responses in AMCW lidar. *Remote Sensing*, **4**, pp. 21–42 (2012b).

Godbaz, J. P., M. J. Cree, A. A. Dorrington, and R. Kunnemeyer. Defocus restoration for a full-field heterodyne ranger via multiple return separation. In: *Image and Vision Computing New Zealand 2007 (IVCNZ'07)*, pp. 52–57. University of Waikato, Hamilton, New Zealand (2007).

Godbaz, J. P., M. J. Cree, A. A. Dorrington, and A. D. Payne. A fast maximum likelihood method for improving amcw lidar precision using waveform shape. In: *IEEE Sensors 2009*, pp. 735–738 (2009c).

Gokturk, S., H. Yalcin, and C. Bamji. A time-of-flight depth sensor – system description, issues and solutions. *Computer Vision and Pattern Recognition Workshop, 2004*, pp. 35–35 (2004).

Goldstein, R. M., H. A. Zebker, and C. L. Werner. Satellite radar interferometry - Two-dimensional phase unwrapping. *Radio Science*, **23**, pp. 713–720 (1988).

Goshtasby, A. A. Fusion of multifocus images to maximize image information. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6229 (2006).

Guomundsson, S. A., H. Aanaes, and R. Larsen. Environmental effects on measurement uncertainties of time-of-flight cameras. In: *Proceedings of International Symposium on Signals, Circuits and Systems 2007, ISSCS 2007* (2007).

Hardy, J. W. *Adaptive Optics for Astronomical Telescopes* (1998).

Hebert, M. and E. Krotkov. 3D measurements from imaging laser radars: How good are they? *IVC*, **10**, pp. 170–178 (1992).

Hegde, G. P. M. and C. Ye. Swissranger sr-3000 range images enhancement by a singular value decomposition filter. In: *Proceedings of the International Conference on Information and Automation, ICIA 2008* (2008).

Hofton, M. A., J. B. Minster, and J. B. Blair. Decomposition of laser altimeter waveforms. *IEEE Transactions on Geoscience and Remote Sensing*, **38**, pp. 1989–1996 (2000).

Högbom, J. A. Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines. *Astronomy and Astrophysics Supplement*, **15**, p. 417 (1974).

Hsu, S., S. Acharya, A. Rafii, and R. New. Performance of a time-of-flight range camera for intelligent vehicle safety applications. In: *Advanced Microsystems for Automotive Applications 2006*, pp. 205–219 (2006).

Hussmann, S. and T. Edeler. Performance improvement of a 3d-tof pmd camera using a pseudo 4-phase shift algorithm. In: *Proceedings of IEEE Instrumentation and Measurement Technology Conference, 2009*, pp. 542–546 (2009).

Iddan, G. J. and G. Yahav. 3d imaging in the studio (and elsewhere...). In: *Proceedings of Three-Dimensional Image Capture and Applications IV* (2001).

Jamtsho, S. and D. Lichti. Modeling scattering distortion of 3d range camera. In: *Proceedings of the ISPRS Commission V Mid-Term Symposium on Close Range Image Measurement Techniques* (2010).

Janesick, J. R. Ccd transfer method: standard for absolute performance of ccds and digital ccd camera systems. In: *Proc. SPIE Vol. 3019, p. 70-102, Solid State Sensor Arrays: Development and Applications, Morley M. Blouke; Ed.*, edited by M. M. Blouke, volume 3019, pp. 70–102 (1997).

Jankiraman, M. *Design of Multifrequency CW Radars*. Scitech (2007).

Jongenelen, A. P. P., D. A. Carnegie, A. D. Payne, and A. A. Dorrington. Development and characterisation of an easily configurable range imaging system. In: *Image and Vision Computing New Zealand 2009 (IVCNZ09)*, pp. 79–84 (2009).

Jongenelen, A. P. P., D. A. Carnegie, A. D. Payne, and A. A. Dorrington. Maximizing precision over extended unambiguous range for tof range imaging systems. In: *IEEE Instrumentation and Measurement Technology Conference (I2MTC '10)*, pp. 1575–1580 (2010).

Kaaresen, K. F. Deconvolution of sparse spike trains by iterated window maximisation. *IEEE Trans. Signal Proc.*, **45(5)**, pp. 1173–1183 (1997).

Kahlmann, T. and H. Ingensand. Calibration and improvements of the high-resolution range-imaging camera swissranger. In: *Videometrics VIII*, volume 5665, pp. 144–155. SPIE, San Jose, CA, USA (2005).

Kahlmann, T. and H. Ingensand. Calibration and development for increased accuracy of 3d range imaging cameras. *Journal of Applied Geodesy*, **2(1)**, pp. 1–11 (2008). Doi: 10.1515/JAG.2008.001 1862-9016 doi: 10.1515/JAG.2008.001.

Kahlmann, T., F. Remondino, and H. Ingensand. Calibration for increased accuracy of the range imaging camera swissranger (2006).

Kang, B., S.-J. Kim, S. Lee, K. Lee, J. D. K. Kim, and C.-Y. Kim. Harmonic distortion free distance estimation in ToF camera. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 7864 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (2011).

Karel, W. Integrated range camera calibration using image sequences from hand-held operation. In: *The International Achives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. Vol. XXXVII. Part B5.* (2008).

Karel, W., S. Ghuffar, and N. Pfeifer. Quantifying the distortion of distance observations caused by scattering in time-of-flight range cameras. In: *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII, Part 5 Commision V Symposium* (2010).

Kavli, T., T. Kirkhus, J. T. Thielemann, and B. Jagielski. Modelling and compensating measurement errors caused by scattering in time-of-flight cameras. In: *Two- and Three-Dimensional Methods for Inspection and Metrology VI*, edited by P. S. Huang, T. Yoshizawa, and K. G. Harding, volume 7066, p. 706604. SPIE (2008).

Keller, M. and A. Kolb. Real-time simulation of time-of-flight sensors. *Simulation Modelling Practice and Theory*, **17(5)**, pp. 967 – 978 (2009).

Klett, J. D. Lidar inversion with variable backscatter/extinction ratios. *Appl. Opt.*, **24(11)**, pp. 1638–1643 (1985).

Kolb, A., E. Barth, R. Koch, and R. Larsen. Time-of-Flight Sensors in Computer Graphics. In: *Eurographics 2009 - State of the Art Reports*, edited by M. Pauly and G. Greiner, pp. 119–134. Eurographics Association, Eurographics, CH-1288 Aire-la-Ville (2009).

Kormylo, J. and J. Mendel. Maximum likelihood detection and estimation of bernoulli - gaussian processes. *Information Theory, IEEE Transactions on*, **28(3)**, pp. 482–488 (1982). 0018-9448.

Kundur, D. and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, **13(3)**, pp. 43–64 (1996).

Kweon, I. S., R. Hoffman, and E. Krotkov. Experimental characterization of the perceptron laser rangefinder. Technical report, Robotics Institute, Carnegie Mellon University (1991).

Landweber, L. An iteration formula for Fredholm integral equations of the first kind. *Am. J. of Math.*, **73(3)**, pp. 615–624 (1951).

Lange, R. *3D Time-Of-Flight Distance Measurement with Custom Solid-State Image Sensors in CMOS/CCD-Technology*. Ph.D. thesis, University of Siegen (2000).

Larkins, R. L., M. J. Cree, A. A. Dorrington, and J. P. Godbaz. Surface projection for mixed pixel restoration. In: *Image and Vision Computing New Zealand 2009 (IVCNZ09)*, pp. 431–436 (2009).

Levin, A., R. Fergus, F. Durand, and W. T. Freeman. Image and depth from a conventional camera with a coded aperture. In: *Proc. SIGGRAPH 07* (2007).

Levy, S. and P. K. Fullagar. Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution. *Geophysics*, **46(9)**, pp. 1235–1243 (1981).

Lichti, D. D. Self-calibration of a 3d range camera. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. Vol. XXXVII. Part B5* (2008).

Lindner, M. and A. Kolb. Lateral and depth calibration of pmd-distance sensors. In: *Advances in Visual Computing*, volume 4292 of *Lecture Notes in Computer Science*, pp. 524–533. Springer Berlin / Heidelberg (2006).

Lindner, M. and A. Kolb. Calibration of the intensity-related distance error of the PMD TOF-camera. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6764 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (2007).

Lindner, M. and A. Kolb. Compensation of Motion Artifacts for Time-of-Flight Cameras. *Lecture Notes in Computer Science*, **5742**, p. 16 (2009).

Lindner, M., A. Kolb, and T. Ringbeck. New insights into the calibration of tof-sensors. *Computer Vision and Pattern Recognition Workshop*, **0**, pp. 1–5 (2008).

Lindner, M., I. Schiller, A. Kolb, and R. Koch. Time-of-flight sensor calibration for accurate range sensing. *Comput. Vis. Image Underst.*, **114**, pp. 1318–1328 (2010).

Lottner, O., A. Sluiter, K. Hartmann, and W. Weihs. Movement artefacts in range images of time-of-flight cameras. In: *Proceedings of the IEEE International Symposium on Signals, Circuits and Systems, ISSCS 2007* (2007).

Lottner, O., W. Weihs, and K. Hartmann. Systematic non-linearity for multiple distributed illumination units for time-of-flight (pmd) cameras. In: *Proceedings of the 12th WSEAS international conference on Systems*, pp. 752–756. World Scientific and Engineering Academy and Society (WSEAS), Stevens Point, Wisconsin, USA (2008). ISBN 978-960-6766-83-1.

Luan, X. *Experimental Investigation of Photonic Mixer Device and Development of TOF 3D Ranging Systems Based on PMD Technology.* Ph.D. thesis, University of Siegen (2001).

Lucy, L. B. An iterative technique for the rectification of observed distributions. *Astronomical Journal*, **79**, p. 745 (1974).

Lumsdaine, T., A. Georgiev. Full resolution lightfield rendering. Technical report, Adobe Systems Inc. (2008).

MacDonald, A. and S. C. Cain. Parameterized blind deconvolution of laser radar imagery using an anisoplanatic optical transfer function. *Optical Engineering*, **45(11)** (2006). 116001.

MacKinnon, D., V. Aitken, and F. Blais. Review of measurement quality metrics for range imaging. *Journal of Electronic Imaging*, **17(3)**, p. 033003 (2008).

Marcia, R. F., Z. T. Harmany, and R. M. Willett. Compressive coded apertures for high-resolution imaging. In: *Optics, Photonics, and Digital Technologies for Multimedia Applications*, volume 7723, p. 772304. SPIE (2010).

Mure-Dubois, J. and H. Hügli. Optimized scattering compensation for time-of-flight camera. In: *Two- and Three-Dimensional Methods for Inspection and Metrology V*, volume 6762. SPIE, Boston, MA, USA (2007a).

Mure-Dubois, J. and H. Hügli. Real-time scattering compensation for time-of-flight camera. In: *Proceedings of the ICVS Workshop on Camera Calibration Methods for Computer Vision Systems, CCMVS2007* (2007b).

Nagy, J. G. and D. P. O'Leary. Fast iterative image restoration with a spatially-varying psf. Technical report, University of Maryland at College Park (1997). Tech Report 270642.

Ng, R., M. Levoy, M. Bredif, G. Duval, M. Horowitz, and P. Hanrahan. Light field photography with a hand-held plenoptic camera. Technical Report CTSR 2005-02, April 2005, Stanford University (2005).

Oggier, T., B. Büttgen, F. Lustenberger, G. Becker, R. B., and H. A. Swissranger sr3000 and first experiences based on miniaturized 3d-tof cameras. Technical report, Mesa Imaging (2005).

Oggier, T., M. Lehmann, R. Kaufmann, M. Schweizer, M. Richter, P. Metzler, G. Lang, F. Lustenberger, and N. Blanc. An all-solid-state optical range camera for 3d real-time imaging with sub-centimeter depth resolution (swissranger). In: *Optical Design and Engineering*, volume 5249, pp. 534–545. SPIE, St. Etienne, France (2004).

Oprisescu, S., D. Falie, M. Ciuc, and V. Buzuloiu. Measurements with tof cameras and their necessary corrections. In: *Proceedings of International Symposium on Signals, Circuits and Systems 2007, ISSCS 2007* (2007).

Payne, A. D. *Development of a Full-Field Time-of-Flight Range Imaging System*. Ph.D. thesis, The University of Waikato (2008).

Payne, A. D., D. A. Carnegie, A. A. Dorrington, and M. J. Cree. A synchronised direct digital synthesiser. In: *First International Conference on Sensing Technology (ICST)*, pp. 174–179. Palmerston North, New Zealand (2005).

Payne, A. D., D. A. Carnegie, A. A. Dorrington, and M. J. Cree. Full field image ranger hardware. In: *Third IEEE International Workshop on Electronic Design, Test and Applications (DELTA'06)*, pp. 263–268. Kuala Lumpur, Malaysia (2006a).

Payne, A. D., A. A. Dorrington, and M. J. Cree. Illumination waveform optimization for time-of-flight range imaging cameras. In: *SPIE Optical Metrology, Videometrics, Range Imaging, and Applications*, volume 8085. Munich, Germany (2011).

Payne, A. D., A. A. Dorrington, M. J. Cree, and D. A. Carnegie. Image intensifier characterisation. In: *Image and Vision Computing New Zealand 2006 (IVCNZ'06)*, pp. 487–492. Great Barrier Island, New Zealand (2006b).

Payne, A. D., A. A. Dorrington, M. J. Cree, and D. A. Carnegie. Characterizing an image intensifier in a full-field range imaging system. *Sensors Journal, IEEE*, **8(11)**, pp. 1763–1770 (2008a). 1530-437X.

Payne, A. D., A. A. Dorrington, M. J. Cree, and D. A. Carnegie. Improved linearity using harmonic error rejection in a full-field range imaging system. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6805 (2008b).

Payne, A. D., A. A. Dorrington, M. J. Cree, and D. A. Carnegie. Improved measurement linearity and precision for amcw time-of-flight range imaging cameras. *Applied Optics*, **49(23)**, pp. 4392–4403 (2010).

Pentland, A. P. A new sense for depth of field. *IEEE Trans. Pattern Anal. Mach. Intell.*, **9(4)**, pp. 523–531 (1987).

Plaue, M. Analysis of the pmd imaging system. Technical report, Interdisciplinary Center for Scientific Computing, University of Heidelberg (2006).

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, third edition (2007). ISBN 0521880688, 9780521880688.

Radmer, J., P. M. Fust, H. Schmidt, and J. Krger. Incident light related distance error study and calibration of the pmd-range imaging camera. In: *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 2008* (2008).

Rajagopalan, A. N. and S. Chaudhuri. Space-variant approaches to recovery of depth from defocused images. *Comput. Vis. Image Underst.*, **68(3)**, pp. 309–329 (1997).

Rapp, H., M. Frank, F. A. Hamprecht, and B. Jahne. A theoretical and experimental investigation of the systematic errors and statistical uncertainties of time-of-flight cameras. *Int. J. Intell. Syst. Technol. Appl.*, **5**, pp. 402–413 (2008).

Raskar, R., R. Raskar, A. Agrawal, A. Agrawal, J. Tumblin, and J. Tumblin. Coded exposure photography: motion deblurring using fluttered shutter. *ACM Trans. Graph*, **25**, pp. 795–804 (2006).

Richardson, W. H. Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America (1917-1983)*, **62**, pp. 55–59 (1972).

Ringbeck, T. and B. Hagebeuker. A 3d time of flight camera for object detection. Proc. of Optical 3-D Measurement Techniques. Zurich (2007).

Robbins, S., B. Schroeder, B. Murawski, N. Heckman, and J. Leung. Photogrammetric calibration of the swissranger 3d range imaging sensor. In: *Optical Sensors 2008*, volume 7003, pp. 700320–10. SPIE, Strasbourg, France (2008).

Ruderman, D. L. and W. Bialek. Statistics of natural images: Scaling in the woods. *Physical Review Letters*, **73**, pp. 814–817 (1994).

Salvi, J., J. Pages, and J. Batlle. Pattern codification strategies in structured light systems. *Pattern Recognition*, **37**, pp. 827–849 (2004).

Scheimpflug, T. Improved method and apparatus for the systematic alteration or distortion of plane pictures and images by means of lenses and mirrors for photography. GB Patent No. 1196 (1904).

Schilling, B. W., D. N. Barr, G. C. Templeton, L. J. Mizerka, and C. W. Trussell. Multiple-return laser radar for three-dimensional imaging through obscurations. *Applied Optics*, **41**, pp. 2791–2799 (2002).

Schmidt, M. and B. Jähne. A Physical Model of Time-of-Flight 3D Imaging Systems, Including Suppression of Ambient Light. *Lecture Notes in Computer Science*, **5742**, pp. 1–+ (2009).

Schmidt, M., K. Zimmermann, and B. Jahne. High frame rate for 3d time-of-flight cameras by dynamic sensor calibration. In: *Proceedings of IEEE International Conference on Computational Photography, 2011* (2011).

Schwarte, R., Z. Zhang, H. Kraft, T. Moeller, J. Frey, M. Grothof, B. Buxbaum, and T. Ringbeck. Oep for extremely fast multi-channel analog and digital oe-signal processing. In: *OPTO 2004* (2004).

Simpson, M. L., M.-D. Cheng, T. Q. Dam, K. E. Lenox, J. R. Price, J. M. Storey, E. A. Wachter, and W. G. Fisher. Intensity-modulated, stepped frequency cw lidar for distributed aerosol and hard target measurements. *Applied Optics*, **44**, pp. 7210–7217 (2005).

Soumekh, M. *Synthetic Aperture Radar Signal Processing with MATLAB Algorithms*. Wiley (1999).

Stefan May, K. P. and H. Surmann. 3D Cameras: 3D Computer Vision of Wide Scope. In: *Vision Systems: Applications*, pp. 181–202 (2007).

Steiger, O., J. Felder, and S. Weiss. Calibration of time-of-flight range imaging cameras. In: *Proceedings of 15th IEEE International Conference on Image Processing, 2008*, pp. 1968–1971 (2008).

Stilla, U. and B. Jutzi. Waveform analysis for small-footprint pulsed laser systems. In: *Topographic Laser Ranging and Scanning. Principles and Processing*, edited by J. Shan, chapter 7, pp. 215–234. CRC Press (2009).

Subbarao, M. Parallel depth recovery by changing camera parameters. In: *Proc. Second International Conference on Computer Vision* (1988).

Subbarao, M. and G. Surya. Depth from defocus: a spatial domain approach. *Int. J. Comput. Vision*, **13(3)**, pp. 271–294 (1994).

Tadikamalla, P. R. Random sampling from the exponential power distribution. *Journal of the American Statistical Association*, **75(371)**, pp. 683–686 (1980).

Tang, P., B. Akinci, and D. Huber. Quantification of edge loss of laser scanned data at spatial discontinuities (2009).

Tang, P., D. Huber, and B. Akinci. A comparative analysis of depth-discontinuity and mixed-pixel detection algorithms. pp. 29–38. Los Alamitos, CA, USA (2007).

Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, **58**, pp. 267–288 (1994).

Tobler, W. R. A classification of map projections. *Annals of the Association of American Geographers*, **52(2)** (1962).

Tuley, J., N. Vandapel, and M. Hebert. Analysis and removal of artifacts in 3-d ladar data. Technical report, Carnegie Mellon University Robotics Institute (2004).

Tuley, J., N. Vandapel, and M. Hebert. Analysis and removal of artifacts in 3-D ladar data. In: *IEEE Int. Conf. on Robotics and Automation*, pp. 2203 – 2210 (2005).

Van Der Schaaf, A. and J. H. Van Hateren. Modelling the power spectra of natural images: Statistics and information. *Vision Res.*, **36(17)**, pp. 2759–2770 (1996).

Vandapel, N., D. F. Huber, A. Kapuria, and M. Hebert. Natural terrain classification using 3-d ladar data. In: *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 5, pp. 5117–5122 Vol.5 (2004).

Wagner, W., A. Ullrich, T. Melzer, C. Briese, and K. Kraus. From single-pulse to full-waveform airborne laser scanners: Potential and practical challenges. In: *International Society for Photogrammetry and Remote Sensing XXth Congress* (2004).

Wainwright, M. J. and E. P. Simoncelli. Scale mixtures of gaussians and the statistics of natural images. In: *in Adv. Neural Information Processing Systems*, pp. 855–861. MIT Press (2000).

Watanabe, M., S. K. Nayar, and M. N. Noguchi. Real-time computation of depth from defocus. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, edited by M. R. Descour, K. G. Harding, & D. J. Svetkoff, volume 2599 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pp. 14–25 (1996).

Whyte, R. Z., A. D. Payne, A. A. Dorrington, and M. J. Cree. Multiple range imaging camera operation with minimal performance impact. In: *SPIE 7538*

- *Image Processing: Machine Vision Applications III*, volume 7538, p. 75380I. SPIE, San Jose, California, USA (2010).

Xu, Z., H.-G. Heinol, R. Schwarte, and O. Loffeld. Enhanced multiprobing recovering algorithm based on color mixed nonlinear modulation and its application in a 3d vision system. **2588(1)**, pp. 200–207 (1995).

Xu, Z., R. Schwarte, H. Heinol, B. Buxhaum, and T. Ringbeck. Smart pixel – photonic mixer device (pmd). In: *Proceedings International Conference on Mechatronics and Machine Vision*, pp. 259–264 (1998).

Yahav, G., G. J. Iddan, and D. Mandelboum. 3d imaging camera for gaming application. In: *International Conference on Consumer Electronics, Digest of Technical Papers*, pp. 1–2 (2007).

Yuan, L., J. Sun, L. Quan, and S. H. Image deblurring with blurred/noisy image pairs. In: *Proceedings of SIGGRAPH 2007* (2007).

Zhang, R., P.-S. Tsai, J. E. Cryer, and M. Shah. Shape from shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**, pp. 690–706 (1999).