

DEVELOPMENT AND APPLICATION OF A
GRATING INTERFEROMETER AT THE CORNELL
HIGH ENERGY SYNCHROTRON SOURCE

A Dissertation

Presented to the Faculty of the Graduate School
of Cornell University

in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by

Robin Michelle Baur

January 2013

© 2013 Robin Michelle Baur
ALL RIGHTS RESERVED

DEVELOPMENT AND APPLICATION OF A GRATING INTERFEROMETER
AT THE CORNELL HIGH ENERGY SYNCHROTRON SOURCE

Robin Michelle Baur, Ph.D.

Cornell University 2013

Conventional x-ray absorption imaging is a powerful and commonly-used technique in a wide variety of fields, but it has limited applicability for softer, lighter materials and for specimens whose internal structure has only small differences in density. In these situations, phase-contrast imaging steps into the breach. Phase contrast is based upon refraction rather than absorption, and can be as much as 1000 times more sensitive than conventional x-ray imaging in the right situations.

This work describes the commissioning and application of a phase-sensitive x-ray grating interferometer that exploits the Talbot self-imaging effect of periodic objects illuminated by monochromatic light. Design considerations and tolerances are described in detail.

Additionally, the development and characterization of a new detector for high-spatial-resolution, high-energy, high-speed imaging is presented. This detector uses a novel fiber-optic scintillator in place of conventional settled-phosphor or film scintillators, and is shown to have significantly improved resolution over an existing lens-coupled detector used early in these experiments.

Finally, the interferometer is used to characterize the properties of the illuminating x-ray source itself and to image fossil specimens of insects embedded in amber, which pose a particular challenge to absorption-contrast imaging. Possibilities for future upgrades and imaging projects are discussed.

BIOGRAPHICAL SKETCH

Robin M Baur was born in Frankfurt, Germany and raised in Lawrence, KS and Northfield, MN. She obtained her B.S. in mathematics from Harvey Mudd College in 2006, graduating with high distinction and with honors in mathematics and the humanities. After the horrifying realization that mathematicians rarely get to play with toys, she jumped ship for physics. Following a long quest for the elusive project that was both mathematically elegant and experimentally accessible, she found x-ray physics and has never looked back.

To Eli, for everything
(this is the wonder that's keeping the stars apart)
and to the memory of Howard Jackson,
who never knew he set my feet on the road.

ACKNOWLEDGEMENTS

First and foremost, I'm endlessly grateful to Sol Gruner and Darren Dale, for taking a chance on me and starting me on this crazy ride more than three years ago. Without their unflagging support, gentle teasing, good-natured pestering, and general faith in me, I'd have run off to join the circus a long time ago.

Aaron Lyndaker came up with the design for the interferometer's flexure stage, saved my bacon at the last minute on several occasions, and (perhaps most importantly) can always be relied on to deliver my daily dose of snark. Jim Savino probably doesn't remember telling me I thought like an engineer, but the compliment still makes me smile. Nate Ellis taught me everything I know about not cutting my fingers off in the mill, and reminded me how to make fun of myself when things went wrong. Mark Lory-Moran and Vince Kotmel helped me keep my chin up during the roughest semester of my grad school life. Thinking back on my days as the lecture demo tech is always good for a grin. ("But why doesn't it start on *fire* anymore?!")

The interferometer itself would never have materialized without Christian David and his student Simon Rutishauser at the Paul Scherrer Institute in Switzerland. In addition to providing the gratings, Christian has been unfailingly gracious, helpful, and supportive. Simon, who manufactured the gratings, also calmly reassured me after several panicked emails in the early days of the project.

The Fairchild detector has a big team behind it. From Fairchild Imaging, Brian Rodricks spent many hours patiently explaining things; Duane Petrovitch and Dave Reaves helped me work out kinks in the software; and Steve Mims provided valuable troubleshooting advice. In the Gruner lab, Mark Tate spent many, many hours helping me assemble and test the system, put up with end-

less questions with grace and good humor, and generally provided the force that got the thing off the ground in the first place. Marty Novak turned the ugliest design sketch I have ever made into the sleek cryostat housing the detector has today. Richard Mead and Collimated Holes went above and beyond the call of duty in making the custom scintillators.

Without Bill Crepet and Dave Grimaldi, I'd have had no specimens to scan. I still maintain that nothing is more terrifying than realizing that someone has shipped you 100-million-year-old specimens in the mail—Dave's braver than I am! Dave's students Phil Barden and Isabelle Veal put in some long hours with me at the APS, and their expertise with amber was invaluable. Also at the APS, Wah-Keat Lee, Kamel Fezzaa, and especially Alex Deriy (all from sector 32-ID-C) were patient, knowledgeable, and surprisingly accommodating of the people who wanted to come in and tear up their beamline to fit in an interferometer. We would have a much lower supply of beautiful data if not for their help.

The Cornell aikido club and the Shinto Muso-ryu koryu kept me sane and made sure I saw the world outside the lab from time to time. Thanks to Larry sensei, Yuki sensei, T.J. Hinrichs, Ian Schachner, and especially Mark Reichert—we'll get to ganseki otoshi one of these days. Finally, a big thanks to my cheer-leading squad: my wonderful husband Eli and my mom and dad, who always believed I could do it even when I didn't. Here's hoping I made Grandpa proud.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
List of Tables	xi
List of Figures	xii
List of Abbreviations	xvii
List of Commonly Used Symbols	xviii
1 Introduction	1
1.1 Basics of x-ray interactions with matter	1
1.2 Overview of the Talbot effect	3
1.3 Theory of the Talbot effect	5
1.4 Exploiting the Talbot effect: the Talbot interferometer	12
1.4.1 Producing interference fringes: the phase grating	13
1.4.2 Detecting interference fringes: the analyzer grating	19
1.4.3 Modulating interference fringes: the specimen of interest	20
1.4.4 Data collection: phase stepping mode	23
1.4.5 Interpreting collected data	25
1.4.6 Data collection: moiré (single-shot) mode	28
1.5 Overview of the work	29
2 Design and operation	32
2.1 The CHESS Talbot interferometer	32
2.1.1 Gratings	32
2.1.2 Positioners	34
2.1.3 Detectors	39
2.2 Designing for diverging beams	41
2.3 Design considerations: tolerances	43
2.3.1 Gratings: Groove depth	44
2.3.2 Gratings: Duty cycle	45
2.3.3 Gratings: Sidewall angle	47
2.3.4 Gratings: Period	48
2.3.5 Source: Radius of curvature	51
2.3.6 Positioning: Defocusing distance	53
2.3.7 Positioning: Relative grating angle θ_z	55
2.3.8 Positioning: Rotation θ_y	56
2.3.9 Positioning: Rotation θ_x	57
2.3.10 Positioning: Vibrations along the y axis	63
2.3.11 Positioning: Monochromator vibrations	65
2.3.12 Source: Monochromaticity	67
2.3.13 Source: Transverse coherence length	68

2.3.14	Source: projected source size	69
2.3.15	Summary of tolerances	69
2.4	Use of the CHESSE Talbot interferometer	70
2.4.1	Alignment	70
2.4.2	Scan parameters	71
2.4.3	Exposure times and Poisson statistics	72
2.4.4	Monochromator detuning	72
3	Source and coherence characterization	73
3.1	Finite sources and partial coherence	74
3.2	Partial coherence and the Talbot effect	75
3.3	Source characterization	78
3.3.1	Size	79
3.3.2	Radius of curvature	84
3.4	Case study: Horizontal-vertical accelerator coupling	89
3.5	Case study: Miscut crystal effects on apparent source parameters	90
4	Data processing	95
4.1	Data processing methods	95
4.1.1	Fast Fourier Transform (FFT)	95
4.1.2	Levenburg-Marquardt algorithm (LMA)	97
4.1.3	Principal components analysis (PCA)	98
4.1.4	Generalized phase shifting algorithms (GPSAs)	105
4.1.5	Single-shot processing	112
4.2	Processing artifacts: gory details	114
4.2.1	Spectral leakage: origins in the FFT	115
4.2.2	Down the rabbit hole: least-squares fitting with LMA	116
4.2.3	Bayesian methods: generalities	124
4.2.4	Bayesian methods: specifics	126
4.2.5	Bayesian methods: results on mocked-up data	128
4.2.6	Goertzel's algorithm	129
4.2.7	Next steps: PCA and GPSAs	130
4.2.8	Gaussian filtering	130
4.2.9	Fourier filtering	131
4.3	Methods comparison	138
5	Fairchild sCMOS detector	140
5.1	Design	141
5.1.1	CMOS sensor	141
5.1.2	Coupling	142
5.1.3	Mechanicals	143
5.2	Scintillating fiber optic plate	145
5.2.1	Stopping power	146
5.2.2	Turn-on lag and afterglow	147

5.3	Characterization	148
5.3.1	Conversion gain	148
5.3.2	Dark current, read noise, and dynamic range	150
5.3.3	Dark current temperature dependence	150
5.3.4	Linearity	154
5.3.5	Sensitivity	156
5.3.6	Resolution	158
5.3.7	Detective quantum efficiency	160
5.3.8	Radiation hardness	162
5.3.9	Heating effects	165
5.4	Detector usage	167
5.4.1	Flatfield corrections	167
5.4.2	Imaging example	167
6	Fossil specimens	172
6.1	Imaging amber: sensitivity and minimum feature size	172
6.1.1	Sensitivity limitations and minimum feature thickness: the simple way	172
6.1.2	Sensitivity, resolution, and noise: the more complex case	174
6.1.3	Phase sensitivity versus absorption sensitivity	181
6.1.4	Relative sensitivity of each system	182
6.2	Working with amber	185
6.2.1	Handling and mounting	185
6.2.2	Discoloration	186
6.2.3	Deposit-to-deposit variations	187
6.3	Specimens and results	188
6.3.1	Amber inclusions at CHESS	188
6.3.2	Fossil flowers	197
6.3.3	Amber inclusions at the Advanced Photon Source	206
7	Miscellany and Future Work	220
7.1	TITAN data-processing software	220
7.2	Interferometer variants	221
7.2.1	Three-grating (Talbot-Lau) interferometers	221
7.2.2	Highly divergent sources	224
7.2.3	Two-dimensional interferometers	227
7.3	Ongoing and future work	228
7.3.1	Streamlining and beamline integration	228
7.3.2	Undulator testing	228
7.3.3	Other thermal bump/radius of curvature estimates	230
7.3.4	Dragonfly flight	232
7.3.5	Computed tomography and laminography	233
7.4	Conclusion	234

A	TITAN codebase and other snippets	236
A.1	TITAN Background	236
A.1.1	Dependencies	236
A.1.2	Obtaining dependencies	237
A.1.3	Installation	239
A.2	TITAN code	239
A.2.1	Frontend	240
A.2.2	Data analysis routines	255
A.2.3	Scripts	259

LIST OF TABLES

3.1	Calculated and relative radii of curvature for symmetric and miscut crystals	94
4.1	Correlation between N and most reliable p_2	123

LIST OF FIGURES

1.1	The Talbot carpet for an absorption grating with opening fraction 0.2	4
1.2	Free-space propagation of a wavefront from plane P_0 to plane P_1	5
1.3	The Talbot interferometer in schematic	12
1.4	Phasor plot illustrating how to find a maximally efficient beam splitter configuration.	17
1.5	Detecting interference fringes using an analyzer grating	21
1.6	Lateral deflection of interference fringes caused by a refracting specimen.	22
1.7	Illustration of phase-stepping data collection.	24
1.8	Three modes of contrast extracted from a single phase-stepping scan.	27
1.9	Moiré fringes generated by tilting the analyzer grating	29
1.10	Fourier transform of moiré fringes modulated by a refracting specimen	30
2.1	Schematic of the structure of our gold analyzer grating.	34
2.2	Piezoelectric-driven flexure stage for phase grating positioning.	35
2.3	Original 60mm rail design of the interferometer.	38
2.4	Current motorized design of the interferometer	39
2.5	Consequential physical degrees of freedom in the interferometer	43
2.6	Degrees of freedom in grating manufacture	44
2.7	Fringe height variation caused by incorrect grating depth	45
2.8	Decrease in visibility at the first imaging distance $z_T/16$ caused by incorrect grating depth	46
2.9	Fringe profile distortion caused by incorrect duty cycle	47
2.10	Effect of small grating manufacturing errors on fringe profiles.	48
2.11	Fringe profile distortion caused by nonzero sidewall angle	49
2.12	Moiré pattern of aligned but period-mismatched fringes.	50
2.13	Moiré fringe period resulting from radius of curvature mismatch	52
2.14	New transmission profile of a grating when the beam is not at normal incidence. Case 1: small angle deviation	58
2.15	New transmission profile of a grating when the beam is not at normal incidence. Case 2: large angle deviation	61
2.16	Effective shear caused by a missteered beam	66
3.1	Projection of a finite source through the phase grating plane onto the imaging plane.	75
3.2	Visibility of Talbot interference fringes given nonzero projected source size	79
3.3	Visibility and estimated source size for two-beam running and for single-beam running	81

3.4	Predicted visibility and actual measured visibility at the A2 station.	83
3.5	Tilted moire fringes produced by mismatched fringe periods . . .	84
3.6	Diagram indicating how to obtain the total beam radius of curvature given source size and divergence values	87
3.7	Moiré fringe angle as a function of grating spacing	88
3.8	Beam geometries for symmetric-cut and miscut monochromator crystals	91
3.9	Visibility curves for symmetric crystals and for 3.2° and 5° miscuts	92
3.10	Visibility curves for symmetric crystals and for 5° miscut in beam focusing orientation	93
4.1	Illustration of the parameters a_0 , a_1 , and φ relative to the overall sinusoids	96
4.2	An example illustrating the need to zero-center PCA data	100
4.3	Simulated lineouts of two recovered PCA patterns: specimen and reference	104
4.4	Illustration of the single-shot processing method	113
4.5	Absorption, dark-field, and differential phase reconstructions showing reconstruction fringe artifacts	114
4.6	FFT comparison showing pure first harmonic in the correctly-sampled case and spectral leakage in the incorrectly-sampled case.	116
4.7	Recovered sinusoid parameters for a two-sinusoid mock dataset fit with a single-sinusoid model function	118
4.8	Comparison of single-sine and modified two-sine fitting functions	119
4.9	Variation in recovered period for original and modified mock datasets fit with single-sinusoid model function	120
4.10	Variations in recovered parameter values for mock datasets fit with single-sinusoid fitting function	121
4.11	Variations in recovered a_0 values when fit period is constrained to true period	123
4.12	Variations in recovered a_1 values when fit period is constrained to true period	123
4.13	Two likelihood estimators for estimating the true period of a collection of mock datasets	129
4.14	Reducing artifacts using Gaussian filtering.	132
4.15	Gaussian filtering applied to an unsuitable dataset.	133
4.16	Scale insect (<i>Hemiptera</i> sp.) in amber. Absorption contrast image exhibiting fringe artifacts.	134
4.17	Fourier transform of Fig. 4.16, illustrating the strong first harmonic peaks due to the fringe artifacts.	135
4.18	Fourier transform of Fig. 4.16 with the first harmonic peaks wiped out.	136
4.19	Inverse Fourier transform of Fig. 4.18, showing much reduced fringe artifacts.	137

4.20	Composite comparison dataset showing FFT processing, LMA curve-fitting, PCA, and GPSA	139
5.1	Fairchild sCMOS detector with its housing removed	144
5.2	Fairchild sCMOS detector in place on the beamline	145
5.3	Lag and afterglow of the scintillating fiber optic plate	148
5.4	Log of the theoretical intrinsic semiconductor carrier density for temperatures between 200 K and 300 K.	152
5.5	Dark current of the Fairchild sCMOS detector	153
5.6	Counts per pixel in the beam spot, normalized by the incident flux, at 10 keV	155
5.7	Sensitivity of the Fairchild detector from 8 keV to 60 keV	157
5.8	Modulation Transfer Function of the Fairchild sCMOS detector	160
5.9	Gaussian fit to the line spread function	161
5.10	Detective quantum efficiency curves for the Fairchild detector at various energies, illustrating dose and energy dependence	163
5.11	DQE curves for the Finger Lakes detector showing the characteristic inverted parabolic shape	164
5.12	Mean counts per pixel per frame in a 50-frame image sequence before and after temperature equilibration with the camera framing in continuous mode; the pre-equilibration curve shows rise behavior consistent with undesirable chip heating, while the post-equilibration curve is closer to constant	166
5.13	A (100 × 100)-pixel region of an image before flatfield correction, showing characteristic features of the fiber optics	168
5.14	A (100 × 100)-pixel region after flatfield correction	169
5.15	Comparison of Finger Lakes and Fairchild images of the midsection of a desiccated yellowjacket	170
5.16	Insets from Fig. 5.15 showing significantly improved resolution from the Fairchild over the Finger Lakes	171
6.1	Egg sac (<i>Araneae</i> sp.) in Baltic amber, 42 mya. Absorption contrast.	189
6.2	Egg sac (<i>Araneae</i> sp.) in Baltic amber, 42 mya. Differential phase contrast.	190
6.3	Parasitoid wasp in Dominican amber, 20 mya. Absorption contrast.	192
6.4	Parasitoid wasp in Dominican amber, 20 mya. Differential phase contrast	193
6.5	Stingless bee (<i>Hymenoptera</i> sp.) in Dominican amber, 20 mya. Absorption contrast.	194
6.6	Stingless bee (<i>Hymenoptera</i> sp.) in Dominican amber, 20 mya. Differential phase contrast.	195
6.7	Bee (<i>Melikertes</i> sp.) in Baltic amber, 42 mya. Absorption contrast.	196
6.8	Bee (<i>Melikertes</i> sp.) in Baltic amber, 42 mya. Differential phase contrast.	196

6.9	Specimen 88, with highly mineralized interior. Absorption contrast.	198
6.10	Specimen 88, with highly mineralized interior. Differential phase contrast.	199
6.11	Specimen 269. Absorption contrast.	200
6.12	Specimen 269. Differential phase contrast.	201
6.13	Specimen 849. Absorption contrast.	202
6.14	Specimen 849. Differential phase contrast.	203
6.15	Specimen 1574. Absorption contrast.	204
6.16	Specimen 1574. Differential phase contrast.	205
6.17	Experimental setup at the Advanced Photon Source	207
6.18	Camera setup at the Advanced Photon Source	207
6.19	Drosophilid fly (<i>Electrophortica</i> sp.) in Baltic amber (56–34 mya) .	208
6.20	Two ants locked in combat in Dominican amber, ~25 mya	209
6.21	Scale insect in amber from the Charente region of France, ~100 mya. Absorption contrast.	210
6.22	Scale insect in amber from the Charente region of France, ~100 mya. Differential phase contrast.	211
6.23	Scale insect in Baltic amber	212
6.24	Scale insect (<i>Putoidae</i> sp.) in Baltic amber. Absorption contrast. .	213
6.25	Scale insect (<i>Putoidae</i> sp.) in Baltic amber. Differential phase contrast.	214
6.26	Drosophilid male in Dominican amber. Absorption contrast. . . .	215
6.27	Drosophilid male in Dominican amber. Differential phase contrast.	216
6.28	Drosophilid fly in Dominican amber	217
6.29	Ant (<i>Brownimecia clavata</i>) in New Jersey amber. Absorption contrast.	218
6.30	Ant (<i>Brownimecia clavata</i>) in New Jersey amber. Differential phase contrast.	219
7.1	Graphical user interface of TITAN data processing software . . .	221
7.2	Fringe visibilities in three-grating mode at the first and third Talbot distances	224
7.3	Comparison of recovered images using data collected at the first and third Talbot distances with a source grating	225
7.4	Geometrical optics view of the interferometer showing greater sensitivity at longer working distances.	226
7.5	Single-shot processed image of the beam profile, showing a bright ring due to the heat bump on the first monochromator crystal.	230
7.6	Diagram showing the relationship between heat bump radius of curvature, beam footprint, and the height of the heat bump. . . .	232
7.7	Specimen orientation for laminography	234

A.1 The TITAN file hierarchy. 240

LIST OF ABBREVIATIONS

APS	Advanced Photon Source
CHESS	Cornell High Energy Synchrotron Source
EMA	Extramural absorber
FOV	Field of view
GPSA	Generalized Phase Shifting Algorithm
LMA	Levenberg-Marquardt Algorithm
LSF	Line spread function
mya	Million years ago
PCA	Principal Components Analysis
PSI	Phase Shifting Interferometry
ROC	Radius of curvature
sp.	species
TEC	Thermo-electric cooler
TITAN	Talbot Interferometry Tools, Analysis, and Numerics

LIST OF COMMONLY USED SYMBOLS

a_0	Absorption channel (mean of intensity sinusoid)
a_1	Dark-field channel (amplitude of intensity sinusoid)
β_λ	Imaginary part of refractive index (wavelength-dependent)
δ_λ	Refractive index decrement (wavelength-dependent)
$\delta(x)$	Dirac delta functional
d	General distance between phase and analyzer gratings
d_m	Specific Talbot distances $mz_T/16$
L	Distance from source point to phase grating
m	Talbot imaging order <i>or</i> diffraction order
μ	Mutual intensity function <i>or</i> linear absorption coefficient
p_0	Source grating period
p_1	Phase grating period
p_2	Analyzer grating period
p_{moire}	Moiré fringe period
R	Radius of curvature
z_T	Talbot distance
φ	Differential phase contrast (phase of intensity sinusoid)
Φ	Phase profile of a specimen

CHAPTER 1

INTRODUCTION

Consider a specimen of ancient amber just extracted from a dig site. Does it have an insect or a flower embedded in it? It is often difficult to tell using non-destructive techniques: the amber may be opaque, turbid with dirt or air bubbles, or so full of flaws and cracks that it is difficult to see inside.

The conventional way of examining such an amber specimen is to section the specimen with a diamond saw, incrementally shaving off slices and polishing the exterior until any inclusions are revealed. But this is a destructive technique that may damage fragile specimens; it would be much more desirable to use a non-destructive technique such as x-ray imaging.

But if amber specimens, even specimens known to have inclusions, are examined using conventional “doctor’s office”-style radiography, there is often nothing to see on the resulting radiograph—just a wash of grey. We will establish why this is the case and describe a special x-ray instrument, the Talbot or grating interferometer, that can provide a window into amber specimens where conventional absorption radiography fails.

1.1 Basics of x-ray interactions with matter

We can write the refractive index of a material as a complex number

$$n_\lambda = \alpha_\lambda + i\beta_\lambda, \tag{1.1}$$

where α_λ governs classical refraction within the material and β_λ governs attenuation. The refractive index is generally dependent on the wavelength of the

incident light. For x-rays, which interact only very weakly with most materials, the refractive index differs only slightly from 1, and we write

$$n_\lambda = (1 - \delta_\lambda) + i\beta_\lambda. \quad (1.2)$$

(Although $n_\lambda < 1$, this does not violate the restriction that energy and information can be transmitted no faster than the speed of light in a vacuum. Only the phase velocity of the wave, which carries no information, travels faster than c , while the group velocity remains slower than c .)

We use the refractive index to define the transmission function of an object by

$$T(x, y) = \exp\left(ik \int n_\lambda(x, y, z) dz\right), \quad (1.3)$$

where k is the wavenumber $2\pi/\lambda$ and the z axis is the axis along which x-rays are propagating. If we substitute eqn. 1.2 into eqn. 1.3, we can split the transmission function into two parts:

$$\begin{aligned} T(x, y) &= \exp\left(ik \int \delta_\lambda(x, y, z) dz\right) \cdot \exp\left(ik \int i\beta_\lambda(x, y, z) dz\right) \\ &= \exp\left(ik \int \delta_\lambda(x, y, z) dz\right) \cdot \exp\left(-k \int \beta_\lambda(x, y, z) dz\right). \end{aligned} \quad (1.4)$$

The first exponential is a pure phase term, while the second exponential is a purely real term that decays as $\int \beta_\lambda dz$ increases; now we can see the connections to refraction and attenuation.

Conventional “doctor’s office” radiography is only sensitive to attenuation. The attenuation component β_λ increases linearly with the density of a material, but for x-ray interactions is typically on the order of 10^{-9} . We can see why only fairly dense structures, such as bone, show up on conventional radiographs, while softer tissues show up only faintly. Similarly, because β_λ is

density-dependent, we can now understand why embedded inclusions in amber tend not to show up in conventional radiographs; the density of the organic amber matrix is too similar to the density of the organic insect or flower embedded inside.

The refractive index decrement δ_λ , while still tiny, is quite large compared to β_λ —generally on the order of 10^{-6} compared to β_λ at 10^{-9} . The decrement δ_λ also increases linearly with density, so if we could design an instrument that was sensitive to variations in δ_λ , we could be as much as 1000 times more sensitive to small variations in density compared to conventional radiography. There are several such instruments; in the rest of the chapter we will describe a specific one based on an interference effect produced by x-ray diffraction gratings. This interference is an optical phenomenon known as the Talbot self-imaging effect.

1.2 Overview of the Talbot effect

If a periodic object such as a diffraction grating is illuminated with coherent monochromatic light, the wavefield at special distances downstream of the periodic object will form a reconstruction of the object's transmission function—the Talbot self-image. Moreover, at some fractions of these special distances, the wavefield will reconstruct the object's transmission function, but with a fraction of the original period; this is the fractional Talbot effect. This remarkable consequence of Fresnel diffraction is named after Henry Fox Talbot, who first observed the effect for visible light in 1836, using light focused through a telescope as a coherent source [51]. If one plots the intensity of the wavefield at all distances downstream of the periodic object, the result is a striking fractal




Figure 1.1: The Talbot self-imaging effect arising from near-field Fresnel diffraction. Beam enters from the top; self-images are halfway down and at the bottom.

structure called the Talbot carpet, as shown in Fig. 1.1.




Figure 1.2: Free-space propagation of a wavefront from plane P_0 to plane P_1 .

1.3 Theory of the Talbot effect

Suppose we have a monochromatic wavefield $u_0(x_0, y_0)$ in a plane P_0 and we wish to know what the wavefront will look like after it has propagated through free space to a new point (x_1, y_1) in the plane P_1 , as shown in Fig 1.2. This free-space propagation can be modeled by the Rayleigh-Sommerfeld diffraction integral,

$$u_1(x_1, y_1) = \frac{1}{i\lambda} \int dx_0 dy_0 u_0(x_0, y_0) \frac{e^{ikr}}{r} \cos(\vec{n}, \vec{r}), \quad (1.5)$$

where \vec{n} is the unit normal to the plane P_0 , \vec{r} is the vector connecting (x_0, y_0) in plane P_0 to (x_1, y_1) in plane P_1 , λ is the wavelength, and $\cos(\vec{n}, \vec{r})$ is the cosine of the angle between \vec{n} and \vec{r} [15, 29].

The Rayleigh-Sommerfeld integral is a general expression for free-space propagation of arbitrary wavefronts under essentially arbitrary conditions. The Talbot effect does not occur under arbitrary propagation conditions; we will therefore make several assumptions about the nature of the free-space propagation, restricting the type of free-space propagation we are interested in, and we will see that the Talbot effect emerges in situations for which these assumptions are valid.

To simplify the diffraction integral, we will first make the *paraxial approximation*, which is the assumption that all rays that make up the wavefront u_0 make only a small angle with respect to the optical axis of the system. (The Talbot effect can still arise for highly curved wavefronts, but the mathematical derivation becomes much more problematic.) The paraxial approximation allows us to write $\cos(\vec{n}, \vec{r}) \approx 1$, an approximation that is valid to within 5% as long as the angle between \vec{n} and \vec{r} is less than 18° [15]. Now the diffraction integral can be written

$$u_1(x_1, y_1) = \frac{1}{i\lambda} \int dx_0 dy_0 u_0(x_0, y_0) \frac{e^{ikr}}{r}. \quad (1.6)$$

Now we wish to simplify the exponential. If the planes P_0 and P_1 are a distance z apart, we can write the distance r as

$$\begin{aligned} r &= \sqrt{z^2 + (x_1 - x_0)^2 + (y_1 - y_0)^2} \\ &= z \sqrt{1 + \frac{(x_1 - x_0)^2 + (y_1 - y_0)^2}{z^2}}. \end{aligned} \quad (1.7)$$

Recall that the Taylor series of $\sqrt{1+t}$ for $t \approx 0$ is

$$\sqrt{1+t} = 1 + \frac{t}{2} - \frac{t^2}{8} + \dots \quad (1.8)$$

If z is relatively large compared to $(x_1 - x_0)$ and $(y_1 - y_0)$, which it must be under the paraxial approximation, we can use this Taylor series expansion to write

$$r = z \cdot \left(1 + \frac{(x_1 - x_0)^2 + (y_1 - y_0)^2}{2z^2} - \frac{[(x_1 - x_0)^2 + (y_1 - y_0)^2]^2}{8z^4} + \dots \right) \quad (1.9)$$

To simplify this Taylor expansion, we make the *Fresnel approximation*, which is that z is sufficiently large that we can discard the second non-unity term, but not so large that we can discard the first non-unity term as well (which would be the limit of far-field or Fraunhofer diffraction). Unlike the paraxial approximation, which can be relaxed without destroying the Talbot effect, the Fresnel approximation is crucial to the emergence of the effect; in the far-field (Fraunhofer) regime, the self-imaging effect is washed out.

Now we can write the diffraction integral as

$$u_1(x_1, y_1) = \frac{1}{i\lambda} \int dx_0 dy_0 u_0(x_0, y_0) \frac{\exp(ikz + ik[(x_1 - x_0)^2 + (y_1 - y_0)^2]/2z)}{z + \frac{(x_1 - x_0)^2 + (y_1 - y_0)^2}{2z}}. \quad (1.10)$$

Again, under the paraxial approximation, $(x_1 - x_0)^2 + (y_1 - y_0)^2 \ll z$, so we write the denominator simply as z . We note that the exponential is much more sensitive than the denominator to small changes in $[(x_1 - x_0)^2 + (y_1 - y_0)^2]/2z$, so we do not eliminate this term in the exponential. We can then write

$$u_1(x_1, y_1) = \frac{1}{i\lambda} \int dx_0 dy_0 u_0(x_0, y_0) \frac{\exp(ikz) \cdot \exp(ik/2z \cdot [(x_1 - x_0)^2 + (y_1 - y_0)^2])}{z}. \quad (1.11)$$

We pull the e^{ikz}/z outside the integral and write

$$u_1(x_1, y_1) = \frac{e^{ikz}}{i\lambda z} \int dx_0 dy_0 u_0(x_0, y_0) \exp\left(\frac{ik}{2z} \cdot [(x_1 - x_0)^2 + (y_1 - y_0)^2]\right). \quad (1.12)$$

This is one of the many ways of representing the Fresnel diffraction integral.

Now that we have made this simplification to the original Rayleigh-Sommerfeld integral, we can show that the Talbot effect emerges naturally from Fresnel propagation of a periodic wavefront. To simplify matters, we will work in one dimension only, reducing eqn. 1.12 to the equivalent expression

$$u(x, z) = \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{+\infty} dx' u(x', z=0) \cdot \exp\left(\frac{ik}{2z}(x' - x)^2\right). \quad (1.13)$$

We are only interested in the phase variations within the plane of the wavefield, rather than the overall phase offset between the planes P_0 and P_1 , so we will discard the overall phase factor e^{ikz} in the rest of our analysis.

Suppose we begin with a uniform wavefield $u_0(x) = 1$, and suppose we have an object with transmission function $T(x)$ that is periodic with period a , so that $T(x + a) = T(x)$. We now make the *projection approximation*, which is that all objects can be treated as thin layers and no scattering occurs inside them, in which case the wavefield directly behind the object can be modeled as the product of the object transmission function with the incident wavefield:

$$u(x, z = 0) = u_0(x) \cdot T(x) = 1 \cdot T(x). \quad (1.14)$$

Objects are considered to be thick when refraction within the object significantly perturbs the x-ray paths from the angle along which they would otherwise have traveled in the absence of the object. X-ray refraction is so weak that most elementally-light objects are thin in this sense; generally only scatterers within the object can significantly perturb the x-ray beam.

Assuming the projection approximation is valid, we can plug this into the convolution in eqn. 1.13 to obtain

$$u(x, z) = \frac{1}{\sqrt{i\lambda z}} \int_{-\infty}^{+\infty} dx' T(x') \cdot \exp\left(\frac{ik}{2z}(x' - x)^2\right). \quad (1.15)$$

We now seek to express the infinite integral in eqn. 1.15 in terms of a more manageable finite integral. Following [18], we will make the change of variables

$$x' = x + \eta + 2Ka, \quad (1.16)$$

where η is a continuous variable in $[0, 2a]$, and $K \in \mathbb{Z}$, where \mathbb{Z} is the set of all integers, both positive and negative. This yields a finite integral over a sum in

K :

$$u(x, z) = \frac{1}{\sqrt{i\lambda z}} \int_0^{2a} d\eta T(x + \eta + 2Ka) \sum_{K \in \mathbb{Z}} \exp\left(\frac{ik}{2z}(\eta + 2Ka)^2\right). \quad (1.17)$$

Because $T(x)$ has period a , and K is an integer, $T(x + \eta + 2Ka) = T(x + \eta)$. We can also expand the exponential as

$$\exp\left(\frac{ik}{2z}(\eta + 2Ka)^2\right) = \exp\left(\frac{i\pi}{\lambda z}\eta^2\right) \exp\left(\frac{i\pi}{\lambda z}4Ka\eta\right) \exp\left(\frac{i\pi}{\lambda z}4K^2a^2\right), \quad (1.18)$$

where we have written the wavenumber k as $2\pi/\lambda$. Consider the last term in this expansion and rewrite it slightly as

$$\exp\left(i \cdot 2\pi \cdot K^2 \cdot \frac{2a^2}{\lambda z}\right). \quad (1.19)$$

The shape of the fraction leads us to define a special distance $z_T = 2a^2/\lambda$, called the *Talbot distance*. Notice that in the special case that $z = z_T/n$ for some integer n , we can write 1.19 as

$$\exp\left(i \cdot 2\pi \cdot K^2 \cdot \frac{z_T}{z}\right) = \exp\left(i \cdot 2\pi \cdot K^2 \cdot n\right), \quad (1.20)$$

which is simply 1 since K and n are integers.

If we now consider the more general case in which $z = (n'/n) \cdot z_T$, eqn. 1.19 can still be reduced to unity by updating our change of variables: instead of eqn. 1.16, we write

$$x' = x + \eta + 2Kn'a, \quad (1.21)$$

where η is now a continuous variable over $[0, 2n'a]$ rather than $[0, 2a]$. We can then rewrite eqn. 1.17 as

$$u(x, z) = \frac{1}{\sqrt{i\lambda z}} \int_0^{2n'a} d\eta T(x + \eta) \exp\left(\frac{i\pi}{\lambda z}\eta^2\right) \sum_{K \in \mathbb{Z}} \exp\left(\frac{i2\pi}{\lambda z}2Kn'a\eta\right). \quad (1.22)$$

Now consider the sum separately. Substitute $z = (n'/n)z_T$, and recall that $z_T = 2a^2/\lambda$, so that we have

$$\sum_{K \in \mathbb{Z}} \exp\left(i \cdot 2\pi \cdot K \cdot \frac{2n'a\eta}{\lambda z}\right) = \sum_{K \in \mathbb{Z}} \exp\left(i \cdot 2\pi \cdot K \cdot \frac{z_T n' \eta}{za}\right) = \exp\left(i \cdot 2\pi \cdot K \cdot \frac{n\eta}{a}\right) \quad (1.23)$$

Recall that η and a are in general not integers. But the sum

$$\sum_{K \in \mathbb{Z}} \exp(i \cdot 2\pi \cdot Kx) \quad (1.24)$$

is only nonzero if $x \in \mathbb{Z}$, in which case it is infinite. So there is only a nonzero contribution if $\eta n/a = p$ for $p \in \mathbb{Z}$. So write $\eta = pa/n$. We can then take advantage of the identity

$$\sum_{K \in \mathbb{Z}} \exp(i \cdot 2\pi \cdot K\eta/x) = x \sum_{p \in \mathbb{Z}} \delta(\eta - px), \quad (1.25)$$

(where here δ is the Dirac delta functional), which arises from the fact that the left hand side is the Fourier series of the right hand side, to write eqn. 1.23 as

$$\sum_{K \in \mathbb{Z}} \exp\left(i \cdot 2\pi \cdot \frac{Kn\eta}{a}\right) = \frac{a}{n} \sum_{p \in \mathbb{Z}} \delta(\eta - pa/n). \quad (1.26)$$

Now return to eqn. 1.22, substituting in this sum of delta functions, to get

$$u(x, z) = \frac{a/n}{\sqrt{i\lambda z}} \int_0^{2n'a} d\eta T(x + \eta) \exp\left(\frac{i\pi}{\lambda z} \eta^2\right) \sum_{p \in \mathbb{Z}} \delta(\eta - pa/n). \quad (1.27)$$

Since η is restricted to the interval $[0, 2n'a]$, the only valid values of p are $0, 1, 2, \dots, 2nn' - 1$. (The case $p = 2nn'$ is equivalent to the case $p = 0$, by the periodicity of $T(x)$.) Take the integral over these delta functions, also substituting $z = (n'/n)z_T = (n'/n) \cdot 2a^2/\lambda$ and $\eta = pa/n$, to obtain the sum

$$\begin{aligned} u(x, z) &= \frac{a/n}{\sqrt{i\lambda \cdot n'/n \cdot 2a^2/\lambda}} \sum_{p=0}^{2nn'-1} T(x + pa/n) \exp\left(\frac{i\pi}{\lambda \cdot n'/n \cdot 2a^2/\lambda} \cdot \left(\frac{pa}{n}\right)^2\right) \\ &= \frac{1}{\sqrt{2inn'}} \sum_{p=0}^{2nn'-1} T(x + pa/n) \exp\left(\frac{i\pi p^2}{2nn'}\right). \end{aligned} \quad (1.28)$$

Although in principle this only allows us to calculate the wavefield for rational multiples of z_T , in practice we can always find a rational number arbitrarily close to any irrational number we choose, so effectively this allows us to compute the wavefield at any distance.

We now demonstrate the Talbot self-imaging effect that manifests at positions z that are integer multiples of z_T , by choosing $n = n' = 1$ in eqn. 1.28. Then we have

$$\begin{aligned}
u(x, z_T) &= \frac{1}{\sqrt{2i}} \sum_{p=0}^1 T(x + pa) \exp\left(\frac{i\pi p^2}{2}\right) \\
&= \frac{1}{\sqrt{2i}} [T(x) \exp(0) + T(x + a) \exp(i\pi/2)] \\
&= T(x) \cdot \left[\frac{1+i}{\sqrt{2i}} \right].
\end{aligned} \tag{1.29}$$

But $1 + i = \sqrt{2} \cdot \exp(i\pi/4)$ and $\sqrt{i} = \exp(i\pi/4)$, so we have

$$u(x, z_T) = T(x), \tag{1.30}$$

which is exactly what we wished to show: that at special distances $z_T = 2a^2/\lambda$ downstream of a periodic object, there is a self-imaging effect in which the wavefield exactly reconstructs the object's transmission function. Similarly we can show that at $z_T/2$, we recover the transmission function $T(x + a/2)$, i.e. the original function shifted by one half period.

If $T(x)$ is a pure phase transmission function, with no attenuation component, these self-images are not very useful; our detecting instruments cannot measure the amplitude of a wave, only its intensity, and the intensity of a pure phase function is uniform. However, in nearly all planes except z_T and $z_T/2$, the interference manifests in some degree as intensity modulations. For example, for $z_T/4$ we can show using the above method that if $T(x) = \exp(i\varphi(x))$, then [12]:

$$I(x) = 1 + \sin[\varphi(x) - \varphi(x + a/2)]. \tag{1.31}$$




Figure 1.3: The Talbot interferometer in schematic. 1: A partially-coherent source illuminates the system. 2: A weakly-interacting specimen modulates the wavefront. 3: The phase grating produces interference fringes. 4: The analyzer grating is period-matched to these fringes.

1.4 Exploiting the Talbot effect: the Talbot interferometer

In its basic form, the Talbot interferometer consists of an x-ray source, two diffraction gratings and a detection system, as shown in Fig. 1.3. Here we will describe the properties of the two gratings and the conceptual operation of the interferometer with a specimen present. In Chapter 3 we will describe the necessary properties of the x-ray source, and in Chapter 2 we will address desirable properties of the detector system, with a specific example given in Chapter 5.

1.4.1 Producing interference fringes: the phase grating

Consider a binary grating with duty cycle 0.5 for which the lines produce a π phase shift relative to the spaces; that is, the object transmission function is

$$T(x) = \begin{cases} 1 & : 0 \leq x < a/2 \\ e^{i\pi} & : a/2 \leq x < a \end{cases}, \quad (1.32)$$

with $T(x + a) = T(x)$.

For such a grating, it is straightforward but tedious to show, using the techniques of the previous section, that at distances $d_m = m \cdot z_T/16 = ma^2/8\lambda$ for odd m , there is a fractional Talbot reconstruction for which all the pure phase data in $T(x)$ is converted to intensity interference fringes with period $a/2$, so that

$$I(x) = \begin{cases} 0 & : 0 \leq x < a/4 \\ 1 & : a/4 \leq x < a/2 \end{cases}, \quad (1.33)$$

with $I(x + a/2) = I(x)$.

An interesting fact is that out of all possible binary gratings, pure phase or otherwise, the π phase grating is provably the most efficient beam splitter [37]. The interference fringes are formed mainly from interaction between the $m = \pm 1$ diffraction orders (to be defined in the following derivation), and a π phase grating diffracts $\sim 80\%$ of the incident light into these orders, more than any other grating. (It is an unfortunate fact that conventional notation in the field refers to both diffraction orders and Talbot imaging orders using the letter m .)

This can be demonstrated (closely following the derivation in [37]), by first considering a general binary grating with period L and thickness

$$t(x) = \begin{cases} A & : 0 \leq x < L/2 \\ B & : L/2 \leq x < L \end{cases}. \quad (1.34)$$

Write the (wavelength-dependent) refractive index as

$$n_\lambda = 1 - \delta_\lambda + i\beta_\lambda. \quad (1.35)$$

Then the transmission function of the object is

$$\begin{aligned} T(x) &= \exp\left(ik \int n_\lambda dz\right) \\ &= \exp\left(ik \int (-\delta_\lambda + i\beta_\lambda) dz\right) \\ &= \exp[-ik\delta_\lambda t(x)] \exp[-k\beta_\lambda t(x)] \\ &= \begin{cases} \exp[-kA(\beta_\lambda + i\delta_\lambda)] & : 0 \leq x < L/2 \\ \exp[-kB(\beta_\lambda + i\delta_\lambda)] & : L/2 \leq x < L \end{cases} \end{aligned} \quad (1.36)$$

The 1 from the refractive index contributes only an overall phase factor, which we discard.

Now, for reasons which will become clear shortly, we wish to convert $T(x)$ into its Fourier series representation. Given the square wave

$$S(x) = \begin{cases} 1 & : 0 \leq x < L/2 \\ -1 & : L/2 \leq x < L \end{cases}, \quad (1.37)$$

we can write the transmission function as

$$T(x) = \frac{c_1 + c_2}{2} + \frac{c_2 - c_1}{2} S(x), \quad (1.38)$$

where

$$\begin{aligned} c_1 &= \exp[-kB(\beta_\lambda + i\delta_\lambda)] \\ c_2 &= \exp[-kA(\beta_\lambda + i\delta_\lambda)] \end{aligned} \quad (1.39)$$

The Fourier series of a square wave is

$$S(x) = \frac{4}{\pi} \sum_{m=1}^{\infty} \frac{\sin[2\pi(2m-1)x/L]}{2m-1}, \quad (1.40)$$

so we can write the Fourier series of $T(x)$ as

$$T(x) = \frac{c_1 + c_2}{2} + \frac{2(c_2 - c_1)}{\pi} \sum_{m=1}^{\infty} \frac{\sin[2\pi(2m-1)x/L]}{2m-1}. \quad (1.41)$$

Without loss of generality, assume that the incident wavefront is uniform. Then the wavefront $u(x, z = 0)$ at the exit surface of the diffraction grating is just $T(x)$. Making use of the identity

$$\sin(x) = \frac{e^{ix} - e^{-ix}}{2i}, \quad (1.42)$$

we can rewrite $u(x, z = 0) = T(x)$ as

$$\begin{aligned} u(x, z = 0) &= \frac{c_1 + c_2}{2} + \frac{c_2 - c_1}{i\pi} \sum_{m=1}^{\infty} \frac{\exp[i2\pi(2m-1)x/L]}{2m-1} \\ &\quad - \frac{c_2 - c_1}{i\pi} \sum_{m=1}^{\infty} \frac{\exp[-i2\pi(2m-1)x/L]}{2m-1}. \end{aligned} \quad (1.43)$$

Now we can see why we wanted the Fourier series representation: we have succeeded in writing $T(x)$ in terms of a constant term plus sums of plane waves $e^{ik_x x}$ and $e^{-ik_x x}$ with

$$k_x = \frac{2\pi(2m-1)}{L}. \quad (1.44)$$

The index m defines the *diffraction order* of the plane wave. To propagate these waves along the z axis, we apply the free-space propagator $e^{ik_z z}$, noting that $k_z = \sqrt{k^2 - k_x^2}$. So we have

$$\begin{aligned} u(x, z \geq 0) &= \frac{c_2 - c_1}{i\pi} \sum_{m=1}^{\infty} \frac{\exp[i2\pi(2m-1)x/L]}{2m-1} \exp\left[iz \sqrt{k^2 - \frac{4\pi^2(2m-1)^2}{L^2}}\right] \\ &\quad - \frac{c_2 - c_1}{i\pi} \sum_{m=1}^{\infty} \frac{\exp[-i2\pi(2m-1)x/L]}{2m-1} \exp\left[iz \sqrt{k^2 - \frac{4\pi^2(2m-1)^2}{L^2}}\right] \\ &\quad + \frac{c_1 + c_2}{2} e^{ik_z z} \end{aligned} \quad (1.45)$$

The result is a forward-propagating wave (the last term in eqn. 1.45) and infinitely many plane waves (diffraction orders) propagating at angles

$$\theta_m = \arctan\left(\frac{k_x}{k_z}\right)$$

$$\begin{aligned}
&= \arctan \left(\frac{2\pi(2m-1)/L}{\sqrt{k^2 - \frac{4\pi^2(2m-1)^2}{L^2}}} \right) \\
&= \arctan \left(\frac{2\pi(2m-1)/L}{\sqrt{\frac{4\pi^2}{\lambda^2} - \frac{4\pi^2(2m-1)^2}{L^2}}} \right) \\
&= \arctan \left(\frac{(2m-1)\lambda}{\sqrt{L^2 - (2m-1)^2\lambda^2}} \right). \tag{1.46}
\end{aligned}$$

These waves are truly propagating, rather than evanescent, as long as the square root is real, i.e. $L > (2m-1)\lambda$.

The subscript m denotes the diffraction order of the plane wave, but we expect both positive and negative diffraction orders: positive to one side of the forward-propagating wave and negative to the other side. By examining eqn. 1.45, we can see that θ_{-m} (the negative diffraction order arising from the second sum) is the same as $-\theta_m$, so to be complete we should actually write

$$\theta_m = \arctan \left(\frac{(2|m|-1)\lambda}{\sqrt{L^2 - (2|m|-1)^2\lambda^2}} \right). \tag{1.47}$$

Now that we know the propagation behavior of the various diffraction orders, we can compare their relative intensities. We can read off the intensity χ_m of the m th diffraction order from eqn. 1.45 to get

$$\begin{aligned}
\chi_m &= \left| \frac{c_2 - c_1}{\pi(2|m|-1)} \right|^2 \\
&= \frac{|\exp[-kA(\beta_\lambda + i\delta_\lambda)] - \exp[-kB(\beta_\lambda + i\delta_\lambda)]|^2}{\pi^2(2|m|-1)^2} \tag{1.48}
\end{aligned}$$

The numerator does not depend on m , so we now know that the $m = \pm 1$ diffraction orders are the most intense (except possibly for $m = 0$, which we must treat separately) for any grating thicknesses A and B and for any grating refractive index components δ_λ and β_λ .




Figure 1.4: Phasor plot illustrating how to find a maximally efficient beam splitter configuration. Two arbitrary phasors (red), two difference-maximizing phasors (blue), and two phasors separated by the angle $\pi/2$ (green)

To have a maximally efficient beam splitter, we want the majority of the energy to be diffracted into the $m = \pm 1$ orders, since they are guaranteed to be the most intense. Since the denominator is fixed for a particular m , we must try to choose A , B , δ_λ and β_λ to maximize the numerator of χ_1 . The numbers c_1 and c_2 are currently arbitrary complex values with magnitude less than or equal to 1. By drawing their phasors in the complex plane, as illustrated in Fig. 1.4, we can see what relationship c_1 and c_2 should have in order to maximize χ_1 . To get both

c_1 and c_2 to have magnitude 1, we must have

$$kA\beta_\lambda = kB\beta_\lambda = 0. \quad (1.49)$$

Since A and B cannot both be 0 (or we would have no grating), it must be the case that $\beta_\lambda = 0$. Physically, this means that there should be no absorption in the grating.

To get the two phasors pointed in opposite directions, we must have

$$kA\delta_\lambda = kB\delta_\lambda + (2n + 1)\pi. \quad (1.50)$$

Physically, this means that there is π additional phase shift in the regions with thickness A compared to the phase shift in the regions with thickness B .

So we have just shown that a non-absorbing π phase-shifting grating is a maximally-efficient beam splitter. We can see that this diffracts $\sim 80\%$ of the incident energy into the $m = \pm 1$ orders by noting that

$$\chi_1 = \frac{|1 - (-1)|^2}{\pi^2(2 \cdot 1 - 1)} = \frac{4}{\pi^2} \approx 0.41, \quad (1.51)$$

so when we consider both $m = 1$ and $m = -1$ we account for 82% of the incident energy.

As an added bonus, an ideal π phase-shifting grating does not allow any direct beam through (i.e. the $m = 0$ order is missing). We can see this by noting that the constant term $(c_1 + c_2)/2$ is given by

$$\begin{aligned} \frac{c_1 + c_2}{2} &= \frac{\exp[-ikA\delta_\lambda] + \exp[-ikB\delta_\lambda]}{2} \\ &= \frac{\exp[-ikB\delta_\lambda + i(2n + 1)\pi] + \exp[-ikB\delta_\lambda]}{2} \\ &= \frac{\exp[-ikB\delta_\lambda](1 + \exp(i\pi))}{2} = 0. \end{aligned} \quad (1.52)$$

Some groups choose to use $\pi/2$ gratings rather than π gratings because the intensity fringes produced by a $\pi/2$ grating have the same period as the grating, in contrast to π gratings which produce intensity fringes with half the grating period. These larger fringes can be easier to detect, but the diffraction into the first order is only

$$\chi_1 = \frac{|\sqrt{2}|^2}{\pi^2} = \frac{2}{\pi^2} \approx 20.5\%, \quad (1.53)$$

as we can see from examining Fig. 1.4. Also, for $\pi/2$ gratings, the direct beam ($m = 0$ order) is not suppressed, which decreases the contrast of the interference fringes. Since we sometimes work in photon-starved modes, we choose to use the more efficiently-diffracting π phase gratings.

1.4.2 Detecting interference fringes: the analyzer grating

Once we have used a π phase grating to produce interference fringes, we must then be able to detect them. Recall that the intensity reconstruction distances for a π phase grating are $d_m = mz_T/16 = mp_1^2/8\lambda$ (m odd), where p_1 is the phase grating period. For x-ray wavelengths, which are on the order of $0.1 - 1\text{\AA}$, the period p_1 must be on the order of microns in order for the distances d_m to be manageable, on the order of millimeters to a meter or so. This means that the interference fringes, with period $p_1/2$, are themselves on the order of single digits of microns.

Most detectors do not have sufficient resolution to distinguish raw interference fringes on this scale. Therefore we introduce a second grating to assist in the detection. This second grating is a binary absorption grating whose period is matched to the period of the interference fringes, i.e. by “binary” we mean

that its transmission function is

$$T_{\text{abs}}(x) = \begin{cases} 0 & : 0 \leq x < p_1/4 \\ 1 & : p_1/4 \leq x < p_1/2 \end{cases} . \quad (1.54)$$

(For information on the problems that can arise if the grating is not perfectly binary, see [11]. The main effect is a degradation of detectable fringe visibility.)

If we fix the absorption grating in place and translate the phase grating so that the interference fringes move, we can deduce the relative locations of the interference fringes even though the detector pixels may be much larger than the fringe period, as shown in Fig. 1.5.

1.4.3 Modulating interference fringes: the specimen of interest

Now that we can produce interference fringes and have a way to detect them, we wish to know how the fringes are affected by the introduction of a specimen upstream of the π phase grating.



Although the effect is small, refraction within the specimen causes a slight angular deviation of the beam, in the same way that a straw placed in a glass of water appears to bend at the water's surface. This change in angle is given by [41]:

$$\Delta\alpha = \frac{\lambda}{2\pi} \frac{\partial\Phi}{\partial x}, \quad (1.55)$$



where Φ is the phase component of the object's transmission function,

$$\Phi(x, y) = \frac{2\pi}{\lambda} \int \delta_\lambda(x, y, z) dz. \quad (1.56)$$

This change in diffraction angle is assumed to be small enough that the beam can still be considered normally incident on the diffraction grating, but it does



(a) Without an analyzer grating, these two fringe positions (red) appear identical to the detector pixel



(b) With an analyzer grating (gold), we can distinguish between the two fringe positions based on whether the fringes are aligned with analyzer lines or spaces

Figure 1.5: Integrated intensity detected by a single detector pixel as a function of fringe position, with and without an analyzer grating. Two fringe periods per pixel shown for illustration purposes.




Figure 1.6: Lateral deflection of interference fringes caused by a refracting specimen.

produce a small lateral deflection

$$\Delta x = \Delta\alpha \cdot d_m \quad (1.57)$$

of the interference fringes, as shown in Fig. 1.6. Note that although the object is drawn as thick for the purposes of illustrating the refraction, we are still employing the thin-object approximation, so that the entry and exit planes are assumed to be a negligible distance apart. Therefore we do not need to worry about the position difference between the entry and exit points of rays passing through the object.

1.4.4 Data collection: phase stepping mode

Now that we know how the specimen locally modulates the interference fringes, how can we make use of this information? One possible data-collection mode is *phase stepping*.

Recall that we specified that the analyzer grating should be period-matched to the interference fringes. This means that if we can align the fringes perfectly with the gaps in the analyzer, all the diffracted intensity should be transmitted; similarly, if we can align the fringes perfectly with the absorbing lines in the analyzer, no intensity will be transmitted. Between these two extremes, the intensity is ideally a triangle wave, but is experimentally observed to vary roughly sinusoidally (like the first Fourier component of a triangle wave), and we write

$$I(x_g, x, y) = a_0(x, y) + a_1(x, y) \cos(2\pi x_g/p_2 + \varphi(x, y)) \quad (1.58)$$

where x_g is the relative position of the two gratings, p_2 is the period of the interference fringes and the analyzer grating, and (x, y) denotes a specific point or detector pixel. The parameters a_0 , a_1 , and φ (which is distinct from Φ) all have something to tell us about the nature of the specimen.

If there is no specimen in place, then as we step the phase grating and thereby sweep the interference fringes across the analyzer grating, all the pixels in the detector should see the same sinusoidal intensity pattern (in the ideal case of uniform illumination). If some of the fringes are locally distorted by a refracting specimen, some pixels will see an intensity sinusoid that is out of phase with the other pixels' intensities. This process is illustrated in Fig. 1.7, where the purple box shows an unmodulated fringe and the red box indicates a fringe that has been deflected by a specimen.

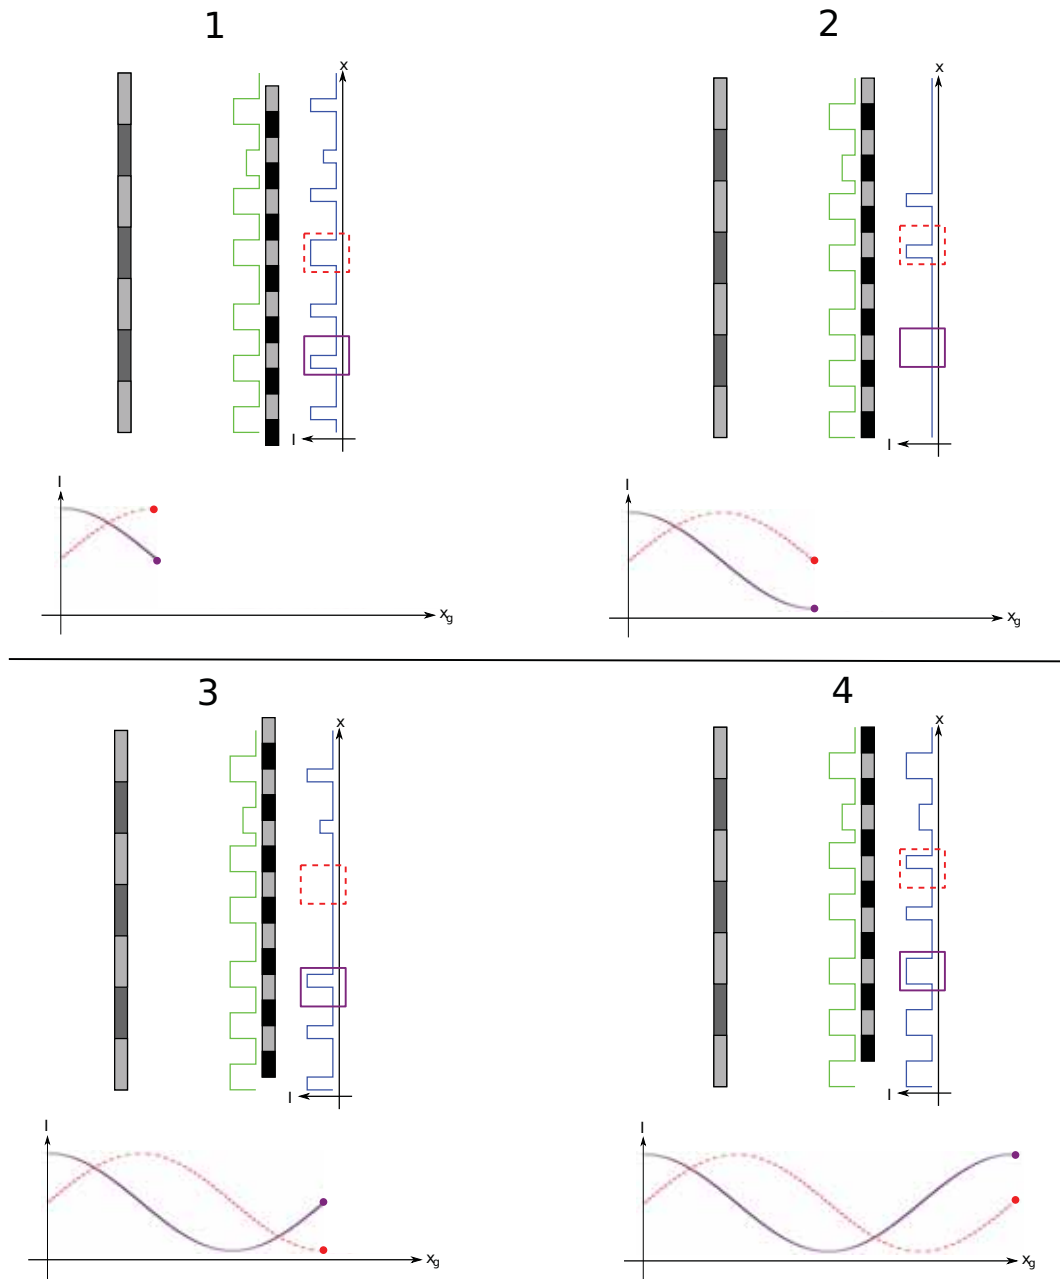


Figure 1.7: Illustration of phase-stepping data collection. As the analyzer grating is translated with respect to the phase grating, the recorded intensity in the red (dashed box) and purple (solid box) pixels differs because the fringes have been shifted by the presence of a specimen. The red (dashed) and purple (solid) curves illustrate the recorded sinusoids.

Once we have collected enough points of data, we can extract the parameters a_0 , a_1 , and φ for each pixel by various methods, described in detail in Chapter 4.

1.4.5 Interpreting collected data

Absorption contrast: a_0

For each specimen dataset $I_s(x_g, x, y)$ we also take a reference dataset $I_r(x_g, x, y)$ without the specimen present. The ratio $a_{0,s}(x, y)/a_{0,r}(x, y)$ describes how many photons successfully passed through the specimen relative to how many photons were present without the specimen in place, and therefore gives us information about the attenuation properties of the specimen (the imaginary part β_λ of the refractive index). This is the *absorption contrast*. Absorption contrast images obtained by Talbot interferometry are not intrinsically any better than conventional absorption radiographs, but we get this information “for free” as part of the data analysis process. A sample absorption radiograph is shown in Fig. 1.8(a).

Differential phase contrast: φ

Recall that when a refracting specimen is in place, the fringes are locally deflected by an amount

$$\Delta x = \Delta\alpha \cdot d_m = \frac{\lambda d_m}{2\pi} \frac{\partial\Phi}{\partial x}. \quad (1.59)$$

In radians, this corresponds to a local phase shift of the pixel intensity sinusoid of

$$\Delta\varphi = 2\pi \cdot \frac{\Delta x}{p_2} = \frac{\lambda d_m}{p_2} \frac{\partial\Phi}{\partial x}. \quad (1.60)$$

The difference $\varphi_s - \varphi_r$ corresponds to this phase shift $\Delta\varphi$ and therefore gives us information about the derivative $\partial\Phi/\partial x$ of the phase component Φ of the transmission function (the real part δ_λ of the refractive index). This is the *differential phase contrast* channel. A sample image is shown in Fig. 1.8(b).


Dark-field contrast: a_1

Visibility or contrast is defined as


$$V = \frac{I_{max} - I_{min}}{I_{max} + I_{min}}. \quad (1.61)$$

Ideally, $I_{min} = 0$ and the visibility is therefore 100%. But if there are features in the specimen that tend to scatter photons, some photons that should have been destined for the peaks of the fringes will scatter into the troughs instead, so that I_{min} is increased to a nonzero value and I_{max} is decreased correspondingly. The visibility is then degraded by an amount related to the scattering strength of the feature; especially strong scattering can wipe out the contrast completely. The contrast loss effectively serves to quantify how much we are “allowed” to violate the projection approximation (that there is no scattering within the specimen).


The contrast in one image is given by a_1/a_0 , so the ratio $(a_{1,s}/a_{0,s})/(a_{1,r}/a_{0,r})$ describes the degradation in contrast when the specimen is introduced. This signal is called the *dark-field contrast*, by analogy to dark-field microscopy, in which only the photons that are scattered from a specimen are recorded by the microscope’s condenser lens [40]. A dark-field image, showing strong black lines at sharp features, is shown in Fig. 1.8(c).



(a) Absorption contrast



(b) Differential phase contrast



(c) Dark-field contrast

Figure 1.8: Three modes of contrast extracted from a single phase-stepping scan. (The obvious artifacts in the images will be explained in Chapter 4.)

1.4.6 Data collection: moiré (single-shot) mode

Phase stepping is a sensitive technique whose resolution is typically limited only by the resolution of the detector, but it requires many images per scan in order to map out the sinusoid for each pixel. For specimens which cannot withstand much radiation dose or in cases where many hundreds of images of the specimen are desired (such as in three-dimensional computed tomography reconstructions), it is useful to have a data collection technique that requires only one snapshot of the interference fringes.

By tilting the analyzer grating a small amount with respect to the interference fringes, we can generate moiré fringes with much larger period than the original fringes. For a relative tilt angle θ and original fringe period p , the moiré fringe period is

$$p_{\text{moire}} = \frac{p}{\sin \theta}, \quad (1.62)$$

as illustrated in Fig. 1.9.

The local distortions in the original fringe pattern caused by a refractive specimen are also present in the moiré fringes, so we can use the moiré fringes as carrier fringes containing the phase information from the specimen. If we take the two-dimensional Fourier transform of the modulated interference fringes, there will be a strong harmonic corresponding to the unmodulated component of the fringes, with the modulations providing some additional local structure in Fourier space [50]. An example is shown in Fig. 1.10.

If we inverse Fourier transform only the region surrounding the harmonic, indicated by the white box in Fig. 1.10, we recover data corresponding to the phase induced by the specimen. The details of this approach are covered in




Figure 1.9: Moiré fringes generated by tilting the analyzer grating

Chapter 4.

1.5 Overview of the work

Chapter 2 will describe mechanical considerations and tolerances in the design of an interferometer and illustrate these considerations by outlining the design of an interferometer specifically for use at CHESS. This chapter will also cover the basics of interferometer operation in two-grating mode.




Figure 1.10: Fourier transform of moiré fringes modulated by a refracting specimen

The Talbot interferometer can also be used to characterize the properties of its illuminating x-ray source. Chapter 3 will describe how this is done and give examples based on CHESS operations in standard running mode and in single-beam running mode.

In chapter 4 I will discuss in more detail the data analysis methods sketched in this introduction, with special attention to the problem of eliminating artifacts remaining in the processed images in the form of “shadows” of the original fringe pattern.

In chapter 5 I will discuss the development of a new high spatial resolution x-ray detector designed for high-energy work. This detector uses a novel doped fiber optic scintillator to provide high spatial resolution without compromising stopping power at high energies.

Chapter 6 will showcase amber specimens from the American Museum of

Natural History and charcoaled fossil flower specimens from the Cornell Department of Plant Biology, describing the specific challenges of imaging amber and charcoal specimens and laying out the sensitivity limits of Talbot interferometer systems.

Finally, chapter 7 will provide a summary of the work that has been done at CHESS, including descriptions of the three major variants of the basic interferometer and comments on the undulator test run of May 2012. It will also provide a look forward at the potential future for the CHESS Talbot interferometer.

CHAPTER 2

DESIGN AND OPERATION

The CHESSTalbot interferometer was designed as a proof-of-concept low-cost interferometer. The initial outlay was less than \$5000, although we were fortunate to receive the most expensive parts (the gratings) for free as part of a collaboration.*

We will first describe the CHESSTalbot interferometer, then describe the many nuances of the tolerances and mechanical limitations that went into its design. This chapter will also address basic operation of the interferometer.

2.1 The CHESSTalbot interferometer

2.1.1 Gratings

Gratings for x-ray Talbot interferometry are most often produced by lithographic patterning of silicon wafers. For the CHESSTalbot interferometer, we use phase gratings that consist of grooves etched into the silicon to a specified depth, so that the grating lines impart a phase shift of π relative to the spaces. Since the refractive index of silicon is wavelength-dependent, the grating is tuned for a specific wavelength based on the depth of the grooves. The phase shift in a silicon line of depth h , relative to the phase shift in an equivalent depth of air, is

$$\frac{2\pi}{\lambda_{\text{air}}}h - \frac{2\pi}{\lambda_{\text{Si}}}h = \frac{2\pi}{\lambda_{\text{air}}}h - 2\pi \frac{1 - \delta_{\text{Si},\lambda_{\text{Si}}}}{\lambda_{\text{air}}}h \approx 2\pi \cdot \frac{\delta_{\text{Si},\lambda_{\text{air}}}}{\lambda_{\text{air}}}h. \quad (2.1)$$

*The gratings were a generous gift of Dr. Christian David, Paul Scherrer Institute, Switzerland.

We opted to design our interferometer for use at 15 keV ($\lambda = 0.826\text{\AA}$). At this energy, x-rays effectively penetrate even relatively thick amber specimens, while still interacting strongly enough to provide adequate sensitivity to the specimens' structure. (For details on the interferometer's sensitivity, see Chapter 6.) For this energy, the refractive index of silicon is $\delta_\lambda = 2.16 \times 10^{-6}$, so according to eqn. 2.1 we require phase gratings that are 19.1 microns deep.

The important characteristics of the analyzer grating are (1) that it be period-matched to the interference fringes, i.e. that $p_2 = p_1/2$, and (2) that it be sufficiently absorbing that no x-rays pass through the lines, i.e. that its transmission function be as close as possible to

$$T(x) = \begin{cases} 0 & : 0 \leq x < p_2/2 \\ 1 & : p_2/2 \leq x < p_2 \end{cases} . \quad (2.2)$$

To that end it is conventional to etch trenches into silicon and then fill the trenches with gold. The particular gold gratings we use for the CHESS Talbot interferometer are 30 microns deep; at 15 keV this corresponds to an x-ray transmission of only 0.01% [10]. As an aside, the analyzer grating is not a 1 micron silicon line/1 micron gold line pattern. To simplify the gold deposition, the David group etched 3-micron-wide trenches with 1-micron-wide silicon lines, then coated the whole surface in a uniform 1 micron layer of gold. The resulting structure is shown in Fig. 2.1. Note that even the "nonabsorbing" Si lines and air gaps are covered with 1 micron of gold, which absorbs 27% of all incident 15 keV x-rays. This significantly reduces the photon transfer efficiency of the analyzer grating.

The grating periods we use are $p_1 = 3.996 \mu\text{m}$ for the π phase grating and $p_2 = 2 \mu\text{m}$ for the gold analyzer grating. The slight deviation from $p_2 = p_1/2$ is




Figure 2.1: Schematic of the structure of our gold analyzer grating. Yellow is gold, grey is silicon.

intended to account for the fact that no source is perfectly parallel, but has some nonzero divergence that magnifies the fringes slightly as they propagate from the phase grating to analyzer grating (further details are in Sec. 2.2). The grating active area is $2 \text{ cm} \times 1 \text{ cm}$, well-matched to the beam size at CHESS, although gratings can be made much larger.

2.1.2 Positioners

Phase stepping

Phase-stepping data collection requires that we acquire many images over one period of the interference fringes in order to map out the intensity sinusoid. Since the interference fringes produced by the phase grating have only a $2 \mu\text{m}$ period, this means that we must be able to position the phase grating in steps of only a few hundred nanometers at most.




Figure 2.2: Piezoelectric-driven flexure stage for phase grating positioning. 1: piezo attachment point. 2: flexure arm. 3: grating attachment point. 4: flexure hinges.

To achieve this precise positioning, we use a piezoelectric-driven flexure stage designed in-house at CHESS, shown in Fig. 2.2. A piezoelectric actuator is placed in position and drives the flexure arm vertically, which moves the grating mount vertically as well, with at most a few microradians of unwanted rotational motion. The flexure hinges are placed so that the distance from the actuator to the first hinge is twice the distance between the hinges, which provides a 2:1 stroke reduction in grating motion relative to the amount driven by the actuator.

The piezoelectric actuator is a PA16-SG from Piezosystems Jena (Hopedale, MA). The full stroke of the piezo in our chosen operating mode is 12 microns, which allows for a 6 micron stroke in grating positioning after the 2:1 reduction in the stage. This is enough to scan over as much as three periods of the interference fringes, which is plenty for the phase-stepping measurement.

Piezo actuators have very fine positioning, but tend to drift over time even if the voltage applied to the actuator stack is constant; the PA16-SG has a strain gauge attached to the actuator stack which allows for negative feedback control of the actuator position (“closed-loop” operation). The end result is that the actuator position is nominally stable to within a few nanometers. We chose the PA16-SG because it was the most cost-effective actuator that could be operated in closed-loop mode to prevent drift.

Angular adjustment

In addition to linear translations for phase-stepping, we also need the ability to control the relative rotation of the two gratings around the optical axis, in order to generate or eliminate moiré fringes. For phase stepping, it is often desirable to align the grating axes as closely as possible to eliminate moiré in the field of view; in single-shot mode, the moiré fringes should be as dense as possible given the resolution of the detector.

Recall that the moiré fringe period is given by

$$p_{\text{moire}} = \frac{p_2}{\sin \theta}, \quad (2.3)$$

where θ is the angle between the phase grating and analyzer grating axes. If we want to completely eliminate moiré fringes from the field of view, we want p_{moire} to be very large; for instance, if we have a 10mm field of view, we need $p_{\text{moire}} = 20\text{mm}$ or more, which means we need to be able to align the grating axes to within

$$\theta = \arcsin\left(\frac{2\mu\text{m}}{20\text{mm}}\right) = 10^{-4}\text{rad} \approx 0.006^\circ. \quad (2.4)$$

Conversely, if we have a detector resolution of $10\mu\text{m}$ and want the densest moiré

fringes that are still resolvable, we want $p_{\text{moire}} \approx 20\mu\text{m}$, which means we need a relative grating rotation of

$$\theta = \arcsin\left(\frac{2\mu\text{m}}{20\mu\text{m}}\right) = .1\text{rad} \approx 5.7^\circ. \quad (2.5)$$

To accommodate both the gross positioning required for dense fringes and the very fine positioning required for eliminating fringes, we chose a Newport RS65 rotation stage with a micrometer for fine adjustment. Its limit of reliable positioning is roughly 0.01° , but as it turns out there are other reasons why we cannot currently eliminate all moiré fringes from the field of view (see Sec. 2.3.5), so this is sufficient for our needs.

Defocusing distance (Intergrating spacing)

Finally, we need to be able to adjust the separation between the phase and analyzer gratings, both on a small scale to locate the position of best contrast at a Talbot imaging distance d_m and on a large scale to switch from one imaging distance d_m to another. We make a distinction between the defocusing distance d , which may be any grating separation, and the imaging distances, which are specific defocusing distances corresponding to $d_m = mz_T/16$ for odd m .

Originally we designed the interferometer to be mounted on a Thorlabs 60mm rail system, with no motorized adjustment of the grating separation. A photograph of this design is shown in Fig. 2.3. In this design we found that it was too difficult to adjust the grating spacing accurately; even when lubricated with Moly Coat molybdenum disulfide powder, the stages tended to stick on the rails as we attempted to slide them, and the clearance necessary for smooth motion made it possible for the stages to tilt when the setscrews were tightened




Figure 2.3: Original 60mm rail design of the interferometer. 1: specimen mount. 2: 60mm cage system. 3: phase grating flexure stage (analyzer grating hidden in shroud).

so that the gratings were no longer normal to the beam. Also, we found that we often wanted to scan the defocusing distance and take images at many spacings, which was tedious with a non-motorized design.

To combat these problems, we switched to a design in which the flexure stage is mounted to a long-travel motorized linear stage (American Linear Manufacturers, Westbury, NY) by an aluminum bracket. This design is shown in Fig. 2.4. This design has the advantage that grating spacing scans are now easy, but the disadvantage that the gratings are decoupled, so that vibrations in one grating stage are not necessarily correlated with vibrations in the other. Also, different




Figure 2.4: Current motorized design of the interferometer

thermal expansion coefficients in the grating mounts can produce artifacts in the phase stepping data if the hutch temperature is unstable. From a stability perspective, it would be desirable to keep the grating stages coupled to reduce the number of degrees of freedom in the system. In future iterations of the design, it would be desirable to return the gratings to a rail system, but with motorized translation and smoother motion along the rails.

2.1.3 Detectors

The collection of desirable properties for a detector for Talbot imaging is the usual laundry list of desirable properties for detectors the world over: high sensitivity, low noise, high spatial resolution, large full well, and (sometimes) fast

framing. A high signal to noise ratio improves the phase sensitivity of the interferometer. So do improved Poisson statistics, so a large full well is desirable to capture as many x-ray equivalents as possible. Single-shot mode opens the door to time-resolved experiments, so Talbot interferometry setups are increasingly demanding fast-framing detectors.

We have primarily used two detectors for Talbot interferometry at CHESS: one which has reasonable signal-to-noise ratio but does not frame quickly, and one which frames quickly and has higher resolution, but does not have as much sensitivity.

Finger Lakes imager

The Finger Lakes imager is based on a backside-illuminated Kodak CCD chip, held at $-35\text{ }^{\circ}\text{C}$ to keep the dark current low, and lens-coupled to a standard scintillator [53]. For Talbot interferometry, this scintillator was typically a single-crystal $11.7\text{ }\mu\text{m}$ film of Europium-doped gadolinium gallium garnet (GGG:Eu), although other scintillators would work just as well.

Because it is lens-coupled, the Finger Lakes is very flexible. We operated it with a standard 35 mm Nikon zoom lens, stopped down to $f/4$ to improve resolution, and with a 76 mm extension tube to provide $4\times$ magnification between the scintillator plane and the chip. The Finger Lakes has 25 micron pixels on chip, so this gave a nominal pixel size of 6.25 microns. In practice, the best resolution we have achieved with the Finger Lakes was ~ 17 microns. We believe this is largely due to spreading of the optical photons in the scintillator, which appears more severe at faster lens speeds (larger acceptance cones). Stopping

the lens down to $f/4$ instead of its fastest setting, $f/1.4$, reduces the size of the acceptance cone and limits the effect of spreading in the scintillator, but cannot eliminate it entirely without sacrificing too much photon transfer efficiency.

Fairchild sCMOS detector

Seeking better resolution and high-energy stopping power than the Finger Lakes could offer, we developed a new detector: the Fairchild sCMOS detector is based on a TCAM-model prototype chip from Fairchild Imaging. This prototype has 6.5 micron pixels on a 5.5 megapixel chip, and is fiber-optically coupled by a 2:1 magnifying taper (Incom, Inc.) to a novel fiber-optic scintillator with 6.2 micron fiber pitch.

The Fairchild is fast-framing and has high spatial resolution, but its sensitivity is somewhat poor at low energies (8-20 keV). At higher energies it performs better because the fiber optic scintillator retains high stopping power without sacrificing resolution. This detector is described in excruciating detail in Chapter 5.

2.2 Designing for diverging beams

So far in our analysis we have assumed the incident x-ray beam was parallel. For realistic x-ray sources, we must account for some divergence in the beam. If we are a distance R away from a spherically-radiating source, there will be a magnification factor

$$M = \frac{R + d_m}{R} \tag{2.6}$$

between the plane of the phase grating and the plane of the analyzer grating. To correct for the magnification, we must design the phase grating with a slightly smaller period than for a parallel beam, according to

$$p_1^* = \frac{p_1}{M} = \frac{R}{R + d_m} p_1, \quad (2.7)$$

where p_1 is the design period for a parallel beam.

If we repeat the derivation in Chapter 1 but assume a spherical wavefront

$$u_0(x) = \exp\left(-\frac{i\pi x^2}{R\lambda}\right) \quad (2.8)$$

(where R is the wavefront's radius of curvature), rather than a uniform planar wavefront $u_0(x) = 1$, we can show [13] that the Talbot reconstructions now appear at distances d_m^* that satisfy:

$$\frac{1}{d_m^*} + \frac{1}{R} = \frac{1}{d_m}, \quad (2.9)$$

so that

$$d_m^* = \frac{R}{R - d_m} d_m. \quad (2.10)$$

Note that for synchrotron sources, the radius of curvature R is not the same as the distance L from the source. We must know the distance to the source, the size of the source, and the divergence of the source in order to accurately calculate R . Misunderstanding of this fact led to a slight error in the calculation of the required period of the divergence-corrected gratings for the CHESS interferometer, which is partly responsible for the data processing artifacts that we will discuss in detail in Chapter 4.




Figure 2.5: Consequential physical degrees of freedom in the interferometer

2.3 Design considerations: tolerances

With two gratings with six degrees of freedom each, and with the possible variations in the grating manufacturing process, we have a large number of variables to account for when designing an interferometer system. The consequential physical degrees of freedom are summarized by Fig. 2.5 and the grating manufacture degrees of freedom are summarized in Fig. 2.6. (As far as grating degrees of freedom, we will consider mainly the effects of errors in the phase grating manufacture; for a very thorough survey of effects of errors in analyzer grating manufacture, see [11].) The source itself often cannot be explicitly tuned, but must nevertheless satisfy some basic tolerances as well.




Figure 2.6: Degrees of freedom in grating manufacture

2.3.1 Gratings: Groove depth

In some sense errors in the groove depth are irrelevant. If the phase shift is not quite π for the design energy, we can in principle simply adjust the design energy slightly until the groove depth does give a π phase shift. In practice, we may not know exactly what the true depth of the grooves is, which makes it difficult to adjust the energy correctly. We therefore want to at least quantify the effect of a depth/energy mismatch.

The phase shift induced by grating lines of depth h relative to the grating grooves is

$$2\pi \frac{\delta_{Si,\lambda}}{\lambda} h. \quad (2.11)$$

If we intend to design a π phase grating but make a depth error Δh , the induced phase shift is instead

$$2\pi \frac{\delta_\lambda}{\lambda} (h \pm \Delta h) = 2\pi \frac{\delta_\lambda}{\lambda} h \cdot \left(1 \pm \frac{\Delta h}{h}\right) = \pi \cdot \left(1 \pm \frac{\Delta h}{h}\right). \quad (2.12)$$

Tuning the phase away from π causes a variation in the heights of the produced interference fringes, so that every other fringe is slightly brighter or less bright than expected, as shown in Fig. 2.7. At π phase shift we showed in Chapter 1




Figure 2.7: Fringe height variation caused by incorrect grating depth

that there should be no intensity in the $m = 0$ diffraction order (direct beam). We can see from the figure that when the phase shift is not π , some amount of direct beam is transmitted. As the phase gets farther from π , the ratio of fringe intensities decreases and the amount of transmitted direct beam increases, damaging the contrast as shown in Fig. 2.8.

2.3.2 Gratings: Duty cycle

The interference fringes are most regular when the grating duty cycle is exactly 0.5. When the duty cycle varies from 0.5, the fringe profile is distorted, as shown in Fig. 2.9.

Note that the two produced fringes are no longer identical. Therefore the recorded fringe profiles as we proceed through a phase-stepping scan will de-




Figure 2.8: Decrease in visibility at the first imaging distance $z_T/16$ caused by incorrect grating depth

pend on whether the detector pixel is smaller or larger than the analyzer grating pitch (1 micron line/1 micron space). If the pixel is smaller than the grating pitch, the detector can discern the difference between the two fringe heights. If the pixel is large enough to span two grating spaces, the detector cannot resolve the two different fringes and sees only an average. The two qualitatively different intensity profiles are shown in Fig. 2.10.

The gratings in use for the CHESS Talbot interferometer have a duty cycle of 0.49. If we compute the fringe profiles for the two detector pixel size cases, we find that we can expect a 4% loss of contrast due to this variation in either case.




Figure 2.9: Fringe profile distortion caused by incorrect duty cycle

2.3.3 Gratings: Sidewall angle

Ideally the gratings will have a perfectly vertical sidewall profile. However, some of the etchants used to create the grooves in the gratings etch laterally while they are etching downward, creating sloped sidewalls. The slope of the sidewall also distorts the fringe profile, as shown in Fig. 2.11.

As before, the resulting fringe profile transmitted through the analyzer grating depends on the size of the detector pixel relative to the analyzer grating pitch. The resulting fringe profiles are similar to those in Fig. 2.10. We do not know the sidewall angle of our gratings, however.




Figure 2.10: Ideal fringe profiles for a 47.5% duty cycle grating if the detector pixels are larger than the grating pitch (solid) or smaller than the grating pitch (dashed). Observe that in the latter case, there are two different peak heights.

2.3.4 Gratings: Period

If the period of the analyzer grating is not perfectly matched to the period of the interference fringes, there will be a moiré pattern even if the grating axes are perfectly aligned, as shown in Fig. 2.12. If the analyzer grating has period p_2 and the interference fringes have period $p_2 \pm \Delta p$, the period of the moiré fringes is

$$p_{\text{moire}} = \frac{p_2}{\Delta p/p_2}. \quad (2.13)$$

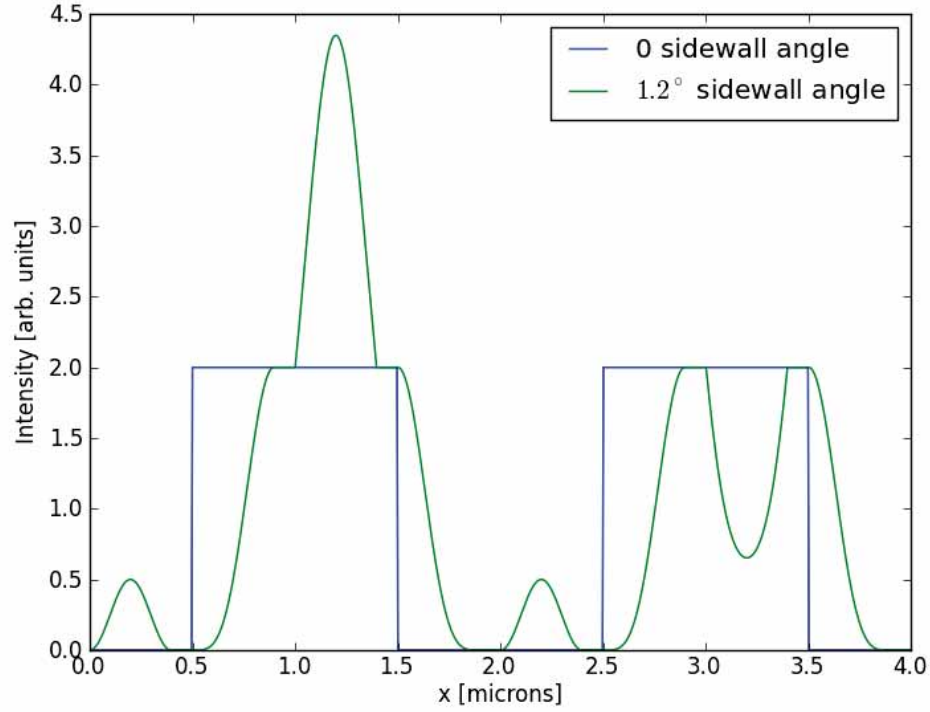


Figure 2.11: Fringe profile distortion caused by nonzero sidewall angle

Suppose that we have a field of view of height H and we wish to keep $p_{\text{moiré}} > 2H$ so that we have at most half of a moiré fringe in the field of view. Then we need

$$\frac{p_2}{\Delta p/p_2} > 2H \quad (2.14)$$

or

$$\Delta p < \frac{p_2^2}{2H}. \quad (2.15)$$

For a 10mm field of view and $p_2 = 2\mu\text{m}$, this means that we need $\Delta p < 2\text{\AA}$!

Conversely, if we have the grating axes aligned to within the limits of the rotation stage but we are still seeing moiré fringes, we can ask what period mismatch would cause such a moiré pattern. In this case we invert eqn. 2.13 to get

$$\Delta p = \frac{p_2^2}{p_{\text{moiré}}}. \quad (2.16)$$




Figure 2.12: Moiré pattern of aligned but period-mismatched fringes. The blue fringe period is 0.9 times the grey fringe period.

As an example, for some F3 datasets where we have reason to believe the gratings were aligned to within the limit of the rotation stage's positioning capability, we still see moiré fringes with a period of $1100\mu\text{m}$. In this case the required fringe period/analyzer period mismatch is

$$\Delta p = \frac{(2\mu\text{m})^2}{1100\mu\text{m}} \approx 3.5\text{nm}. \quad (2.17)$$

We initially supposed the period mismatch was due to an error in the manufacture of the phase grating. Since the interference fringes have approximately half the period of the phase grating, this would mean a manufacturing error of $\sim 7\text{nm}$. Christian David assured us that even 3nm would be an unreasonably large error in manufacturing, so we began to search for other possible sources of period mismatch.

2.3.5 Source: Radius of curvature

Recall that the fringe period at the analyzer grating is given by

$$p = Mp_1^* = \frac{R + d_m^*}{R} \frac{p_1^*}{2} = p_2. \quad (2.18)$$

So there are actually three possible causes for a mismatched fringe period: incorrect grating period p_1^* , incorrect defocusing distance d_m , and incorrect beam radius of curvature R . We will address incorrect defocusing distance in Sec. 2.3.6.

Suppose that we designed our gratings for a specific radius of curvature R , but that the radius of curvature is actually $R' < R$. Also assume that $d_m^* \approx d_m$. Then the fringes will have period

$$p = \frac{R' + d_m}{R'} \frac{p_1^*}{2} \quad (2.19)$$

when they reach the analyzer plane, so the period mismatch is

$$\begin{aligned} \Delta p &= \left(\frac{R' + d_m}{R'} - \frac{R + d_m}{R} \right) \frac{p_1^*}{2} \\ &= \left(\frac{R' + d_m}{R'} \cdot \frac{R}{R + d_m} - 1 \right) p_2 \end{aligned} \quad (2.20)$$

If we still see moiré fringes even when we believe the gratings are aligned and the imaging distance is correct, we can use the fringe period to deduce the true radius of curvature of the beam. The moiré fringe period will be

$$\begin{aligned} p_{\text{moire}} &= \frac{p_2}{\Delta p/p_2} \\ &= \frac{p_2}{\left(\frac{R' + d_m}{R'} \cdot \frac{R}{R + d_m} - 1 \right)}. \end{aligned} \quad (2.21)$$

The shape of this dependence is shown in Fig. 2.13 After some algebra, we can show that eqn. 2.21 implies

$$R' = R \cdot \frac{1}{1 + (1 + R/d_m)(p_2/p_{\text{moire}})}. \quad (2.22)$$




Figure 2.13: Moiré fringe period resulting from radius of curvature mismatch, assuming a design ROC of $R = 21.6\text{m}$

So for our example from F3 data in which we saw $p_{\text{moire}} = 1100\mu\text{m}$, this implies that the true radius of curvature is

$$R' = \frac{21.6\text{m}}{1 + (1 + 21.6\text{m}/24\text{mm})(2\mu\text{m}/1100\mu\text{m})} = 8.2\text{m}, \quad (2.23)$$

compared to an assumed radius of curvature of $R = 21.6\text{m}$. (The actual radius of curvature of the F3 source is approximately 31.1m; here we use the distance to the source, 21.6m, because that is the radius we [incorrectly] assumed when designing the gratings.)

2.3.6 Positioning: Defocusing distance

There are two reasons why we are concerned about correct defocusing distance d_m : first, an incorrect defocusing distance may disrupt the fringe period; second, it may cause a degradation of fringe contrast. We will see that in neither case is this tolerance particularly tight.

Again, the fringe period at the analyzer grating is

$$p = \frac{R + d_m}{R} \frac{p_1^*}{2} = p_2. \quad (2.24)$$

If instead we are at a defocusing distance $d_m \pm \Delta d$, we have a slightly different fringe period

$$\begin{aligned} p \pm \Delta p &= \frac{R + (d_m \pm \Delta d)}{R} \frac{p_1^*}{2} \\ &= \frac{R + d_m}{R} \frac{p_1^*}{2} \pm \frac{\Delta d}{R} \frac{p_1^*}{2} \\ &= p \pm \frac{\Delta d}{R} \frac{p_1^*}{2}. \end{aligned} \quad (2.25)$$

So the moiré fringe period caused by this defocusing is

$$p_{\text{moire}} = \frac{p_2^2}{\Delta p} = \frac{p_2^2}{(\Delta d/R)(p_1^*/2)}. \quad (2.26)$$

So if we want to restrict p_{moire} to be greater than a given value, we need

$$\Delta d < \frac{p_2^2}{(p_{\text{moire}}/R)(p_1^*/2)}. \quad (2.27)$$

For simplicity, we can assume that $p_1^*/2 \approx p_2$ in most cases. (The exceptions are situations in which the correction $(R + d_m)/R$ differs significantly from 1, i.e. when the radius of curvature is small or the defocusing distance is very large.)

Then we can write

$$\Delta d < \frac{p_2 R}{p_{\text{moire}}}. \quad (2.28)$$

For instance, if we want to keep $p_{\text{moire}} > 20\text{mm}$ at F3, we need

$$\Delta d < \frac{2\mu\text{m} \cdot 21.6\text{m}}{20\text{mm}} = 2.16\text{mm}. \quad (2.29)$$

Now consider the effect of incorrect defocusing on the fringe contrast. Fringe contrast goes sinusoidally with grating distance according to

$$V(d) = 0.5 - 0.5 \cos\left(\frac{2\pi d}{2d_1}\right), \quad (2.30)$$

modulated by an exponential envelope

$$E(d) = \exp[-(1.887\Sigma_p(d)/p_2)^2], \quad (2.31)$$

where Σ_p is the projected size of the source in the analyzer grating plane, as we will show in Chapter 3 [61]. In the following analysis, we will assume that the exponential envelope does not shift the locations of the peaks of the oscillations.

If we are defocused from a contrast peak at position $d_m = md_1$ by an amount Δd , the visibility is instead

$$\begin{aligned} V(d_m \pm \Delta d) &= 0.5 - 0.5 \cos\left(\frac{2\pi(d_m \pm \Delta d)}{2d_1}\right) \\ &= 0.5 - 0.5 \cos\left(m\pi \pm \frac{\pi\Delta d}{d_1}\right) \\ &= 0.5 \mp 0.5 \cos\left(\frac{\pi\Delta d}{d_1}\right). \end{aligned} \quad (2.32)$$

Since $V(d_m) = 1$, we see a variation ΔV given by

$$\begin{aligned} \Delta V &= V(d_m) - V(d_m \pm \Delta d) \\ &= \frac{1}{2} \left(1 \pm \cos\left(\frac{\pi\Delta d}{d_1}\right)\right). \end{aligned} \quad (2.33)$$

So to stay within a range ΔV of the maximum contrast, we need

$$\Delta d < d_1 \frac{\arccos(1 \pm 2\Delta V)}{\pi}. \quad (2.34)$$

For instance, if we wish to stay within 10% of maximum contrast at 15 keV, where $d_1 = 24\text{mm}$, we need

$$\Delta d < 24\text{mm} \frac{\arccos(1 - 0.2)}{\pi} \approx 4.9\text{mm}. \quad (2.35)$$

We can see that achieving an acceptable defocusing distance is not a challenging task.

2.3.7 Positioning: Relative grating angle θ_z

A change in the relative grating angle θ_z changes the period of the moiré fringes produced by the grating pair. How much additional rotation we can tolerate depends on what imaging mode we are using.

In phase-stepping mode, we want to keep the moiré fringes larger than the field of view (FOV). We calculated earlier that for a 10mm FOV, we needed to be able to position the gratings within 0.006° of perfect alignment. Suppose that we set the gratings to 0° (perfect alignment) within the tolerance of the stage, which is $\pm 0.01^\circ$. To introduce a single full moiré fringe into the FOV, the gratings would have to be aligned at an angle

$$\alpha = \arcsin(2\mu\text{m}/10\text{mm}) = 0.011^\circ. \quad (2.36)$$

This is comparable to the positioning uncertainty of the stage, so we are essentially guaranteed to have most of a moiré fringe in the FOV regardless of our alignment position.

That is only a question of alignment. Once the gratings are aligned and the setscrews tightened down, we do not expect the rotation stage to drift. There

may, however, be unwanted motion of the assembly itself. Aside from side-to-side flexion of the stage mounting posts, which is difficult to quantify, the most likely source of rotational misalignment is the flexure stage.

The flexure stage is designed so that the maximum amount of stage rotation over the whole stroke of its motion is approximately 1.5 microradians (0.000086°). A small rotation has its maximum effect when the gratings start out perfectly aligned; a misalignment of 1.5 microradians corresponds to moiré fringes of period

$$p_{\text{moire}} = \frac{2\mu\text{m}}{\sin 1.5\mu\text{rad}} = 1.3\text{m}, \quad (2.37)$$

so we do not need to worry about any rotational misalignment caused by the flexure stage.

2.3.8 Positioning: Rotation θ_y

For gratings where the lines are horizontally oriented, rotations θ_y change the apparent thickness of the grating. For the analyzer grating, this rotation is not problematic and may even be desirable if the depth of gold in the analyzer does not provide sufficient absorption at normal incidence. For the phase grating, a rotation θ_y has the same effect as a groove depth error. If the original groove depth is h , a rotation θ_y gives a new apparent thickness

$$h' = \frac{h}{\cos \theta_y}, \quad (2.38)$$

so the apparent groove depth error is

$$\Delta h = h' - h = h \left(\frac{1}{\cos \theta_y} - 1 \right), \quad (2.39)$$

which gives a phase shift error of

$$\frac{2\pi\delta}{\lambda}\Delta h = \pi\left(\frac{1}{\cos\theta_y} - 1\right). \quad (2.40)$$

Suppose we wish to keep the nominal visibility above 50%. This will be the case as long as the error $\Delta h/h$ is less than ~ 0.5 (see Fig. 2.8), which corresponds to a phase shift of $\pi/2$ rather than π . This occurs when $\cos\theta_y = 1/2$, so $\theta_y = 60^\circ$. (Note that fringes are still visible at $z_T/16$ when the phase shift is $\pi/2$, but they do not have uniform heights and, as noted, the visibility is poor.)

We can also use rotations θ_y to tune the energy of the interferometer somewhat. For instance, if we are working at 15 keV, the appropriate thickness is $h = 19.1\mu\text{m}$. If we wish to work at 16 keV, we require a thickness of $h' = 20.4\mu\text{m}$, which we can achieve by rotating the original grating by

$$\theta_y = \arccos(h/h') = 20.6^\circ. \quad (2.41)$$

2.3.9 Positioning: Rotation θ_x

Rotations about the x axis change the incidence angle of the beam on the grating from normal incidence to the nonzero incidence angle θ_x . There are two cases depending on the amount of rotation: either there are still regions where rays can pass entirely through the grating fins or entirely through the gaps between the fins; or alternatively, all rays must pass partially through a grating fin and partially through a gap. These two conditions are illustrated in Figs. 2.14 and 2.15. In both cases we can derive the new object transmission function in terms




Figure 2.14: Case 1: $\theta_x < \arctan(p/2h)$. The red curve shows the new apparent thickness, while the blue curve show the thickness at normal incidence.

of θ_x . Suppose the original thickness is given by

$$t(x) = \begin{cases} 0 & : x \in [0, p/2] \\ h & : x \in [p/2, p] \end{cases} \quad (2.42)$$

Case 1: $\theta_x < \arctan(p/2h)$

In this case, the tilt angle θ_x is relatively small. It is therefore still possible for some rays passing through the grating to avoid interacting with the grating fins, so in some regions the transmission function will still be 0. Likewise, some rays will still pass entirely through the grating fins, but with a longer path length than in normal incidence; therefore the maximum phase shift will be larger than the design shift of π . Between these two extremes, the imparted phase shift will vary linearly. To derive the new transmission function, we need to know (a) how big the regions are where the phase shift is minimal or maximal, and (b) the new height of the maximum.

In Fig. 2.14, the leftmost ray could be shifted anywhere in the region labeled a and still be passing entirely through the grating trench (neglecting the backing). The rightmost ray could be shifted left by a similar amount and still be passing entirely through the grating fin. As we can see from the diagram, $a = p/2 - a'$. Looking at the angle θ_x marked in Fig. 2.14, we can see that the distance a' is given by $h \tan \theta_x$. So for a region of width $a = p/2 - h \tan \theta_x$, the apparent thickness is still 0.

Similarly, by looking at the rightmost ray we can see that for another region of width a there is an apparent increase in the thickness of the grating fins given

by

$$t' = \frac{h'}{h}t = \frac{t}{\cos \theta_x}. \quad (2.43)$$

Between these two regions, the apparent thickness increases or decreases linearly between 0 and t' , as shown at the bottom of Fig. 2.14.

Case 2: $\theta_x > \arctan(p/2h)$

Case 2 is different from case 1 in two fundamental ways. First, because all rays passing through the grating must at some point intersect a grating fin, the minimal apparent thickness (again, neglecting the backing) is no longer 0. Second, because of the steepness of the incident angle, the maximum apparent thickness need not be larger than the thickness at normal incidence. Other than these points, the only difference is that the algebra is more complicated.

Between the two leftmost rays in Fig. 2.15, the apparent thickness is minimal, with a value

$$\begin{aligned} t'' &= \frac{h''}{h}t = \frac{t}{h}(H - h') \\ &= \frac{t}{h} \left(\frac{h}{\cos \theta_x} - \frac{p/2}{\sin \theta_x} \right) \\ &= \frac{t}{\cos \theta_x} - \frac{t}{\sin \theta_x} \frac{p}{2h}. \end{aligned} \quad (2.44)$$

Between the two rightmost rays in Fig. 2.15, the apparent thickness is maximal, with a value

$$t' = \frac{h'}{h}t = \frac{t'}{h} \frac{p/2}{\sin \theta_x} = \frac{t'}{\sin \theta_x} \frac{p}{2h}. \quad (2.45)$$

Note as a sanity check that this matches the previous calculation when $\theta_x = \arctan(p/2d)$, since we then have

$$t' = \frac{t}{\sin \theta_x} \cdot \tan \theta_x = \frac{t}{\cos \theta_x}. \quad (2.46)$$




Figure 2.15: Case 2: $\theta_x > \arctan(p/2h)$. The red curve shows the new apparent thickness, while the blue curve shows the thickness at normal incidence.

The width a of these minimal and maximal regions is given by

$$\frac{a}{h''} = \sin \theta_x \rightarrow a = h \tan \theta_x - p/2. \quad (2.47)$$

Again, between these two regions the apparent thickness increases or decreases linearly. The new apparent thickness function is shown at the bottom of Fig. 2.15. As a final point, note that the primary meaningful value is the *difference* between the maximal and minimal phase shifts, so the whole apparent thickness curve could be shifted down so that r'' is 0 without affecting the results of any simulation using this analysis.

Now that we have described the new apparent thickness resulting from a θ_x rotation, we can consider the effect on the fringe visibility. First we point out that if $\theta_x = \arctan(p/h)$, the apparent thickness is constant and we lose all contrast. Our phase gratings for 15 keV have period $3.996\mu\text{m}$ and depth $19\mu\text{m}$, so the absolute maximum rotation we can tolerate is $\theta_x = 11.9^\circ$. Our analyzer gratings have period $2\mu\text{m}$ and depth $30\mu\text{m}$, so the maximum tolerable rotation is $\theta_x = 3.8^\circ$.

Theoretically, as long as the transmission percentage remains approximately 0 in the gold lines, rotations of the analyzer grating should not decrease the contrast, since

$$V = \frac{I_{max} - 0}{I_{max} + 0} = 1 \quad (2.48)$$

regardless of the value of I_{max} . However, the photon transfer efficiency of the system will suffer significantly.

Rotations of the phase grating are equivalent to a combination of incorrect groove depth and nonzero sidewall angle. The distorted fringe profiles are very similar to those that only have a nonzero sidewall angle, except that the mini-

mum intensity is no longer zero.

2.3.10 Positioning: Vibrations along the y axis

When the gratings are not perfectly stable with respect to one another, we suffer loss of contrast from vibration in the plane of the grating, transverse to the grating lines. By treating the intensity fringes as classical harmonic oscillators, we can estimate how much of the integrated intensity falls in regions where we would expect to record intensity, and how much falls in regions where we would not expect to record intensity if the gratings were perfectly stable. From this calculation we can quantitatively relate the amplitude of a vibration to an associated loss of contrast. In this calculation, we assume that the detector pixel is larger than the interference fringes.

A point object oscillating about a location x_0 with amplitude Δ has (classical) one-dimensional probability density

$$\rho(x) dx = \frac{dx}{\pi \sqrt{\Delta^2 - (x - x_0)^2}} \quad (2.49)$$

and hence has a probability P_{ab} of being between two points $x = a$ and $x = b$ given by

$$P_{ab} = \frac{1}{\pi} \arcsin \left(\frac{x - x_0}{\Delta} \right) \Big|_a^b. \quad (2.50)$$

Considering the point object as a light source, the integrated intensity in a particular interval $[a, b]$ should be directly proportional to the probability P_{ab} of finding the point object in that interval.

Now consider the case of an intensity fringe with finite extent. Suppose that the fringe is a square wave which is nonzero for $x \in [0, 1]$, vibrating with

an amplitude $\Delta (< 1/2)$. We consider this oscillating fringe as being built up of oscillating point sources. Define the *central region* as the region occupied by a stationary fringe (i.e. the interval $[0, 1]$); we wish to compute the integrated intensity falling outside the central region because of the oscillation. This should be the sum of contributions from each small (point-like) part of the fringe within a distance Δ of the edge, and is therefore proportional to

$$\int_0^\Delta P_{x-\Delta,0} dx + \int_{1-\Delta}^1 P_{1,x+\Delta} dx. \quad (2.51)$$

By symmetry these terms are identical, so the proportion of intensity falling outside the central region is

$$\begin{aligned} 2 \int_0^\Delta \frac{1}{\pi} \arcsin\left(\frac{y-x}{\Delta}\right) \Big|_{y=x-\Delta}^{y=0} dx &= \int_0^\Delta \left(1 + \frac{2}{\pi} \arcsin\left(\frac{-x}{\Delta}\right)\right) dx \\ &= \Delta - \frac{2\Delta}{\pi} \int_0^{-1} \arcsin u du \\ &= \frac{2\Delta}{\pi}. \end{aligned} \quad (2.52)$$

To compute contrast, we also need to know the proportion of intensity that does fall inside the central region. Points near the edges contribute to the central region with intensity proportional to $2\Delta(1 - 1/\pi)$; we can see this either by computing a similar integral or by noting that in total these points should contribute intensity proportional to the space they take up, 2Δ . Points in the interval $[\Delta, 1 - \Delta]$ contribute to the central region with probability 1 (they are not oscillating enough to wander outside the region $[0, 1]$), so they contribute to the central region with intensity proportional to the width of this interval, $1 - 2\Delta$.

When the system is perfectly stable, the contrast is 1; when the system vibrates, the highest intensity drops to $1 - 2\Delta + 2\Delta(1 - 1/\pi) = 1 - 2\Delta/\pi$, and the lowest intensity increases to $2\Delta/\pi$, so the contrast decreases to

$$V = \frac{I_{max} - I_{min}}{I_{max} + I_{min}} = \frac{1 - 2\Delta/\pi - 2\Delta/\pi}{1 - 2\Delta/\pi + 2\Delta/\pi} = 1 - \frac{4\Delta}{\pi}. \quad (2.53)$$

We can invert this relation to find out how much we must restrict vibration (as a fraction of the fringe width, which is the same as the analyzer period) if we wish to retain a certain amount of contrast V :

$$\Delta \leq \frac{\pi}{4}(1 - V). \quad (2.54)$$

Recall that we defined Δ as a fraction of the bright portion of the fringe. For clarity, we will rewrite eqn. 2.54 in terms of a fraction $\Delta p/p_2 = \Delta/2$ of the total fringe period:

$$\frac{\Delta p}{p_2} \leq \frac{\pi}{8}(1 - V). \quad (2.55)$$

Recall that this analysis is only valid for $\Delta < 1/2$. The equivalent relation for Δp is $\Delta p/p_2 < 1/4$, or $\Delta p < 500\text{nm}$. At this level of vibration, we can expect a visibility of

$$V = 1 - \frac{8\Delta p/p_2}{\pi} = 1 - \frac{2}{\pi} \approx 36\%. \quad (2.56)$$

2.3.11 Positioning: Monochromator vibrations

The upstream crystal in the double-bounce monochromator at F3 is water-cooled, which causes the crystal to vibrate slightly as water is pumped through the cooling block. This vibration causes the beam to wander a few microradians to either side of its intended outgoing direction. This slight missteering of the beam is not enough to break our assumption that the beam is normally incident on the diffraction gratings, but it does create a small apparent shearing vibration, as shown in Fig. 2.16.

The tolerance for this apparent shear is the same as for the actual shear discussed in Sec. 2.3.10; by working backwards we can estimate the equivalent




Figure 2.16: Effective shear caused by a missteered beam

tolerance (in microradians) for the actual missteering. The angle $\Delta\alpha$ is related to the shear Δx by

$$\Delta\alpha = \frac{\Delta y}{D}. \quad (2.57)$$

We can rewrite Δy in terms of a fraction of the fringe period as $\Delta y = p_2(\Delta y/p_2)$.

Then we can make use of eqn. 2.55 to write

$$\Delta\alpha \leq \frac{p_2 \pi}{D 8}(1 - V). \quad (2.58)$$

Say we wish to maintain a minimum of 10% visibility at the first imaging distance for 15 keV x-rays. Then we require

$$\Delta\alpha \leq \frac{2\mu\text{m}}{24\text{mm}} \frac{\pi}{8}(1 - 0.1) \approx 30\mu\text{rad}. \quad (2.59)$$

If we are performing phase-stepping scans, we may prefer a more stringent vibration tolerance than the one imposed by the visibility requirements. For instance, suppose we are taking N measurements per analyzer period and we prefer that the beam should wander no more than half the step size, i.e. $\Delta y < (1/N) \cdot (p_2/2)$, so that the uncertainty regions for neighboring scan points do not overlap. Then the beam deflection from normal incidence should be no

more than

$$\Delta\alpha \leq \arctan(\Delta y/d_m) = \arctan\left(\frac{p_2}{2Nd_m}\right). \quad (2.60)$$

For $N = 5$, $p_2 = 2\mu\text{m}$ and $d_1 = 24\text{mm}$ (the first imaging distance at 15 keV), this means we require $\Delta\alpha \leq 8.3\mu\text{rad}$.

Based on measurements taken at F3, the maximum wandering appears to be $\sim 5.5\mu\text{rad}$ to either side of normal incidence, which is in the acceptable range for this measurement. However, we often take $N = 9$ measurements per period; in this case we would like to have $\Delta\alpha \leq 4.6\mu\text{rad}$.

Even though we cannot achieve the desired tolerance when taking $N = 9$ points per period, we are not prevented from taking this data. As long as we still have enough visibility to perform the measurement, this tolerance is not a hard requirement, just a desirable property. We could use the known value of the beam missteering to quantify the uncertainty in our grating positioning when we later process the data to extract the sinusoid parameters, although most of the data processing techniques do not actually account well for uncertainty in the data.

2.3.12 Source: Monochromaticity

We can view a polychromatic source as being built up of a collection of monochromatic modes. The fringe pattern in the imaging plane is a superposition of the fringe patterns of all the monochromatic modes.

We design the interferometer for a specific energy. For all of the other monochromatic modes, the design depth of the phase grating does not pro-

vide a π phase shift. The imperfect fringes arising from these monochromatic modes tend to blur the fringe profile away from the ideal profile of a purely monochromatic beam. The energy bandpass tolerance of the interferometer is nevertheless quite wide, as much as $\Delta E/E \sim 10\%$ [41].

2.3.13 Source: Transverse coherence length

The transverse coherence length of a beam is a measure of the lateral distance over which one part of the wavefront can be said to be in phase with another part.

Like most phase-contrast imaging techniques, Talbot interferometry requires some amount of transverse coherence in the beam, but its requirements are less strict than most. In essence, the constraint is that the beam must be coherent enough that waves from neighboring slits in the diffraction grating can coherently interfere in the imaging plane. So as a rough first estimate, we can say that the minimum transverse coherence length is

$$l_c \approx \frac{3}{2}p_1 \approx 6\mu\text{m}. \quad (2.61)$$

As it turns out, this estimate can be relaxed significantly if we are willing to suffer some loss of visibility. For example, the interferometer can effectively be used at F3 even though the vertical transverse coherence length is only $l_c = 2.0\mu\text{m}$ at 15 keV. We will discuss the effects of the coherence length in more detail in Chapter 3, and in Chapter 7 we will discuss a way to operate the interferometer even if the source is incoherent.

2.3.14 Source: projected source size

We will give a detailed description of the effects of projected source size in the following chapter. For now we will simply note that a large and/or nearby source, when projected through a measurement point in the specimen, casts a large “shadow” (penumbra) on the detection plane. A smaller and/or more distant source has a smaller penumbra in the detection plane. The projected source smears out features in the specimen; the amount of smearing is given by convolution of the projected source with the feature profile. We therefore cannot have a resolution better than approximately the projected source size.

At F3, given our specimen positioning, a typical projected source size is $\sim 2\text{-}3\mu\text{m}$. For the Finger Lakes detector, this is about half the pixel size and less than the resolution limit imposed by the detector itself. For the Fairchild, this is comparable to the pixel size and half to a quarter of the detector-imposed resolution limit.

2.3.15 Summary of tolerances

After evaluating all of these tolerances, we can see that the most critical restrictions are on the radius of curvature, the transverse coherence length, and the grating period. In addition, we must carefully control relative shear of the gratings along the y axis. The relative grating rotation θ_z requires precise and accurate positioning, somewhat better than we can actually hope to achieve with current equipment. Keeping the rotation angle $\theta_x = 0$ becomes increasingly important as the aspect ratio of the gratings increases; the rotation θ_y is not as critical. Monochromaticity of the source and correct defocusing distance are

relatively unimportant.

2.4 Use of the CHESS Talbot interferometer

2.4.1 Alignment

Most of the necessary alignment can be done by eye. There are not really any user-adjustable degrees of freedom in the phase grating mount. If one is using the Fairchild detector, one can ensure that the analyzer grating mount is perpendicular to the beam by pressing the mount against the detector snout. (One should not do this with the Finger Lakes detector, where the scintillator is exposed at the front of the detector snout and could be damaged.)

The two pieces of alignment that must be done carefully are calibrating the defocusing distance (intergrating spacing) and adjusting the relative grating rotation. Calibration of the defocusing distance must be done each time the interferometer is set up, and can be easily done to within 1 mm (better, with care), using a set of calipers or a machinist's scale. Alignment of the relative grating rotation requires an hour or so of iterative adjustment of the Newport's micrometer in increasingly fine steps; alignment is complete when the smallest possible adjustment causes the moiré fringes to shift almost 180° , indicating that the gratings have passed through the parallel position.

Eliminating moiré fringes entirely, even using the fine adjustment on the Newport stage, is basically impossible since the beam radius of curvature at F3 is not simply the distance to the source and therefore our gratings are not per-

fectly divergence-adjusted. Nevertheless, the fewer fringes that are in the field of view, the cleaner the resulting processed data will be. It is worth spending an hour or so getting the gratings very well aligned.

2.4.2 Scan parameters

Since we often process phase-stepping data by Fast Fourier Transform, we must choose our phase-stepping data points so that we minimize the risk of aliasing in higher harmonics. If we take at least $N = 5$ points per grating period, the lowest harmonic that can influence our data is the 9th harmonic [32]. Momose et al. call $N = 5$ the “magic number” for phase-stepping.

For a long time we also processed data by least-squares fitting. In order to improve the fits, we generally took $N = 9$ points per period. Even after it became clear that least-squares fitting could not offer a substantial enough improvement over FFT to justify the increased processing time, inertia kept us in the habit of using $N = 9$ points per period.

The angle at which the beam enters the hutch tends to drift over the course of a CESR fill; this leads to slight apparent shear of the gratings over the course of a scan, with the result that the period we scan over may be larger or smaller than the true period of the gratings. At the start of each beamtime, or anytime CESR changes its operating parameters or warms up from Tuesday maintenance, it is worth doing a few phase-stepping scans at different periods. To determine the optimal period, analyze each of these phase stepping scans using the Fourier transform method (see Chapter 4) and see which one shows the cleanest profile, then use that scan’s period for the remaining beamtime.

2.4.3 Exposure times and Poisson statistics

The phase sensitivity of the interferometer varies with the Poisson statistics for N_{ph} incident photons according to [33]

$$\Delta\varphi_{\min} \sim \frac{1}{\sqrt{N_{ph}}}, \quad (2.62)$$

a fact which we will demonstrate in Chapter 6. Assuming the sample can withstand the dose, it is worthwhile to expose for as long as is possible without saturating the detector.

2.4.4 Monochromator detuning

In a nutshell: always, *always* detune the second monochromator crystal slightly away from the intensity peak. Without detuning, higher harmonics can pass through the monochromator system. If the system were perfectly stable, this would not be a problem, but the first (upstream) monochromator crystal tends to suffer vibrations from its water cooling system. Higher harmonics then tend to flicker in and out, changing the recorded beam profile and preventing the pixel-by-pixel data from forming clean sinusoids. We experienced this problem both at CHESS and at the Advanced Photon Source and detuning fixed the problem in both cases.

At CHESS, we use the `mostab` monochromator stabilizing feedback system to keep the beam at 70% of peak. The system has a piezo that adjusts one crystal slightly so that the two crystals are not exactly parallel, which pushes the higher harmonics out of the Darwin width [1] of the second crystal and prevents their transmission.

CHAPTER 3

SOURCE AND COHERENCE CHARACTERIZATION

Up to this point in this work, we have treated our x-ray sources as pointlike and fully coherent. In dealing with real sources, we must account for the finite (nonzero) size of the source and the consequent partial coherence of the outgoing x-ray beam. We will begin with a brief overview of coherence properties of finite sources, then describe how partial coherence affects the behavior of the Talbot interferometer.

Knowing how the properties of the source affect the interferometer, we can invert the problem and use the interferometer as a tool for non-destructive beam diagnostics (in the sense that they do not require us to interfere with the electron beam producing the x-rays). A basic measurement of x-ray source size is the simplest such diagnostic, but we can also estimate the beam's local radius of curvature (ROC). Together, the source size and ROC measurements can be used to characterize the beam itself, as well as the behavior of optical elements in the system. We will describe two such efforts to illustrate the point. First, we will use source size measurements coupled with predictions from beamline-simulation software to compute the phase-space area of the beam. Second, we will describe the beam-conditioning effects of asymmetrically-cut monochromator crystals by observing the change of apparent source size as a function of miscut angle.

3.1 Finite sources and partial coherence

For thorough treatments of coherence, the interested reader is referred to [1], [37], and especially [29].

Briefly, coherence length is the distance (longitudinal or transverse) between two points in a wavefront, arising from two different points in the source, which are completely out of phase. Point sources are then necessarily fully coherent since each point in the wavefront originates from the same source point, while finite-size sources have some finite coherence length. The transverse coherence length increases with distance from the source, according to

$$l_c = L\lambda/\Sigma, \quad (3.1)$$

where L is the distance from the source, λ is the wavelength, and Σ is the (FWHM) extent of the source. (We will use the capital Σ for FWHM sizes and the lowercase σ for rms sizes.)

In order for interference fringes to be produced by the phase grating, the transverse coherence length l_c in the direction perpendicular to the grating lines must, generally speaking, be at least comparable to the distance between neighboring slits in the grating, so that waves passing through these neighboring slits can constructively interfere.

In Talbot interferometry, a more relevant figure of merit is the size of the “shadow” (penumbra) cast by a finite source on the imaging plane when projected through the phase grating plane, as shown in Fig. 3.1. When this penumbra is large relative to the interference fringes, the fringes are smeared out and visibility decreases, as we will see in the following section.




Figure 3.1: Projection of a finite source through the phase grating plane onto the imaging plane. The original source has FWHM extent Σ and the projected source has FWHM extent Σ_p .

The direct relationship between the coherence length and the penumbra (projected source size) is

$$\Sigma_p = \frac{d\lambda}{l_c}, \quad (3.2)$$

where Σ_p is the projected source size and d is the distance between the phase grating (or specimen) and the imaging plane. It is perhaps more illuminating to connect the two via the actual size of the source: the more pointlike the source is, the more coherent it is and therefore the longer the coherence length. But at the same time, the more pointlike the source is, the smaller the projected source size will be.

3.2 Partial coherence and the Talbot effect

In an ideal Talbot interferometer, the interference fringes have 100% visibility at the Talbot imaging distances $d_m = mp_1^2/8\lambda$ for m odd and 0% visibility at d_m for m even. Between these two extremes, the visibility has a complicated functional

form, which we approximate by

$$V_f(d) = 0.5 - 0.5 \cos(\pi d/d_1) \quad (3.3)$$

At any given distance d , we can approximate the interference fringes by sinusoids whose intensity is given by

$$I(x, d) = I_0(1 + V_f(d) \sin(2\pi x/p_2)). \quad (3.4)$$

If we have a finite Gaussian source of (rms) extent σ a distance L away from the phase grating, the projected image of the source in the plane of the interference fringes will also be a Gaussian, with (rms) extent $\sigma_p = (d/L) \cdot \sigma$, whose profile we can write as

$$S(x, d) = A \exp\left(-\frac{x^2}{2\sigma_p^2(d)}\right). \quad (3.5)$$

This source profile is convolved with the interference fringes to form a new fringe profile, whose visibility is decreased relative to the original profile because of the “smearing-out” effect of the nonzero source size. We can compute the visibility degradation directly by evaluating the convolution

$$\begin{aligned} I(x) * S(x) &= \int_{-\infty}^{\infty} I_0(1 + V_f(d) \sin(2\pi x_0/p_2)) \cdot A \exp\left(\frac{-(x - x_0)^2}{2\sigma_p^2(d)}\right) dx_0 \\ &= I_0 A \int_{-\infty}^{\infty} \exp\left(\frac{-(x - x_0)^2}{2\sigma_p^2(d)}\right) dx_0 \\ &\quad + I_0 A V_f(d) \int_{-\infty}^{\infty} \sin(2\pi x_0/p_2) \cdot \exp\left(\frac{-(x - x_0)^2}{2\sigma_p^2(d)}\right) dx_0 \end{aligned} \quad (3.6)$$

We can evaluate the first integral quickly using the identity

$$\int_{-\infty}^{\infty} \exp\left(\frac{-u^2}{a^2}\right) du = a \sqrt{\pi}. \quad (3.7)$$

Then use the identity

$$\sin x = \frac{e^{ix} - e^{-ix}}{2i} \quad (3.8)$$

to write the second integral as

$$I_0AV_f(d) \int_{-\infty}^{\infty} \left[\frac{\exp(2\pi ix_0/p_2) - \exp(-2\pi ix_0/p_2)}{2i} \right] \exp\left(\frac{-(x-x_0)^2}{2\sigma_p^2(d)}\right) dx_0. \quad (3.9)$$

Considering only the first term for now, we can write

$$\begin{aligned} I_0AV_f(d) \int_{-\infty}^{\infty} \left[\frac{\exp(2\pi ix_0/p_2)}{2i} \right] \exp\left(\frac{-(x-x_0)^2}{2\sigma_p^2(d)}\right) dx_0 \\ = \frac{I_0A}{2i} \int_{-\infty}^{\infty} \exp\left[\frac{-x_0^2}{2\sigma_p^2(d)} + \left(\frac{2\pi i}{p_2} + \frac{x}{\sigma_p^2(d)}\right)x_0 - \frac{x^2}{2\sigma_p^2(d)} \right] dx_0. \end{aligned} \quad (3.10)$$

To evaluate this integral, we use the less familiar identity

$$\int_{-\infty}^{\infty} \exp\left(\frac{-u^2}{a^2} + bu + c\right) du = a\sqrt{\pi} \exp\left(\frac{a^2b^2}{4} + c\right) \quad (3.11)$$

(which comes from completing the square in the exponent and then applying eqn. 3.7) to write eqn. 3.10 as

$$\begin{aligned} \frac{I_0AV_f(d)\sigma_p(d)\sqrt{2\pi}}{2i} \exp\left[\frac{2\sigma_p^2(d)(2\pi i/p_2 + x/\sigma_p^2(d))^2}{4} - \frac{x^2}{2\sigma_p^2(d)} \right] \\ = \frac{I_0AV_f(d)\sigma_p(d)\sqrt{2\pi}}{2i} \exp\left[\frac{\sigma_p^2(d)}{2} \left(\frac{2\pi i}{p_2}\right)^2 + \frac{2\pi ix}{p_2} \right]. \end{aligned} \quad (3.12)$$

If we repeat the process for the second term in eqn. 3.9, we find that the full integral is equal to

$$\frac{I_0AV_f(d)\sigma_p(d)\sqrt{2\pi}}{2i} \exp\left[\frac{\sigma_p^2(d)}{2} \left(\frac{2\pi i}{p_2}\right)^2 \right] \cdot \left[\exp\left(\frac{2\pi ix}{p_2}\right) - \exp\left(-\frac{2\pi ix}{p_2}\right) \right]. \quad (3.13)$$

We can recondense the exponentials into a sine using the identity 3.8 and combine this with the first term from eqn. 3.6 to write the new fringe intensity as

$$I(x) * S(x) = (I_0A\sigma_p(d)\sqrt{2\pi}) \left(1 + V_f(d) \exp\left(\frac{-2\pi^2\sigma_p^2(d)}{p_2^2}\right) \sin\left(\frac{2\pi x}{p_2}\right) \right). \quad (3.14)$$

The visibility of this new smeared-out fringe pattern is

$$V(d) = \frac{I_{max} - I_{min}}{I_{max} + I_{min}}$$

$$\begin{aligned}
&= \frac{1 + V_f(d) \exp(-2\pi^2 \sigma_p^2(d)/p_2^2) - (1 - V_f(d) \exp(-2\pi^2 \sigma_p^2(d)/p_2^2))}{1 + V_f(d) \exp(-2\pi^2 \sigma_p^2(d)/p_2^2) + 1 - V_f(d) \exp(-2\pi^2 \sigma_p^2(d)/p_2^2)} \\
&= V_f(d) \exp\left(\frac{-2\pi^2 \sigma_p^2(d)}{p_2^2}\right). \tag{3.15}
\end{aligned}$$

We more often deal with full-width at half-maximum (FWHM) source sizes. The FWHM projected source size $\Sigma_p(d)$ is related to $\sigma_p(d)$ by

$$\sigma_p(d) = \frac{\Sigma_p(d)}{2\sqrt{2 \ln 2}}, \tag{3.16}$$

so we can rewrite eqn. 3.15 as

$$V(d) = V_f(d) \exp\left[-\left(\frac{\pi}{2\sqrt{\ln 2}} \frac{\Sigma_p(d)}{p_2}\right)^2\right] \approx V_f(d) \exp\left(-(1.887 \Sigma_p(d)/p_2)^2\right). \tag{3.17}$$

That is, the original visibility $V_f(d)$ is modulated by a Gaussian envelope that depends on the size of the projected source relative to the size of the fringes [61].

This dependence is shown in Fig. 3.2.

3.3 Source characterization

We can predict the size and divergence of an x-ray source by knowing the properties of the electron (or positron) beam producing it. In particular, say σ_{x,e^-} and σ_{x',e^-} are the (rms) size and divergence of the electron beam, and σ_r and $\sigma_{r'}$ are the (rms) size and divergence of the radiation produced by a single electron.

Then

$$\sigma_{tot,x} = \sqrt{\sigma_{x,e^-}^2 + \sigma_r^2} \tag{3.18}$$

and

$$\sigma_{tot,x'} = \sqrt{\sigma_{x',e^-}^2 + \sigma_{r'}^2} \tag{3.19}$$

are the (rms) size and divergence of the full x-ray source [10].




Figure 3.2: Visibility $V(d)$ of Talbot interference fringes given nonzero projected source size $\Sigma_p(d)$ (FWHM)

But the size and divergence of the electron beam may not always be well-known. In this case, we would like an independent measurement of the x-ray source's properties, which we can use to make inferences about the electron source itself, as illustrated in Sec. 3.4.

3.3.1 Size

Eqn. 3.17 lets us describe the visibility of interference fringes in terms of the projected source size. By recasting the projected source size in terms of the actual source size, we can use a series of visibility measurements to extract the source size.

The phase grating of the CHSS Talbot interferometer is mounted on a long-travel translation stage so that we can position the grating at a range of positions along the optical axis. By scanning the phase grating through the translation stage's full range of motion, taking single interferograms at various distances along the scan, we can experimentally map out the visibility curve $V(d)$ for a particular source. By fitting the visibility curve to the model function

$$V(d) = V_f(d) \exp\left(-\left(1.887 \Sigma_p(d)/p_2\right)^2\right) \quad (3.20)$$

we can extract the source size Σ from the experimental data by recalling that $\Sigma_p(d) = (d/L) \cdot \Sigma$, where L is the distance from the source to the phase grating and d is the separation between the two gratings.

The size of the x-ray source can be significantly affected by changes in the electron or positron beam used to produce the x-rays. For instance, CESR normally operates in "two-beam" running mode, with beams of electrons and positrons simultaneously counterrotating in the ring, but the ring can be operated with only a single species present ("single-beam" mode). Without the complication of maintaining two noninteracting beams in the machine at once, the accelerator operators can squeeze the beam down farther, decreasing the vertical source size. (In specific terms, they can reduce both the overall emittance and the horizontal-vertical coupling; for a detailed description of these terms, see Sec. 3.4.)

Fig. 3.3 shows the effect of this tuning change: with two beams in the machine, the source size is big enough that the Gaussian envelope suppresses all but the first peak of the visibility curve, but in single-beam mode, an additional peak becomes visible, corresponding to the 3rd Talbot order (the second imaging distance). Moreover, partway through the run the accelerator operators




Figure 3.3: Visibility and estimated source size for two-beam running and for single-beam running before and after tweaking beam conditions to improve lifetimes

made an adjustment that increased source sizes in favor of improved lifetimes; that adjustment allowed beams to be refilled every three hours instead of two, but with a tradeoff of a 20% increase in source size. These curves were recorded at the F3 station in July of 2010.

It was fortunate that we performed our first source size measurements at F3. We later installed the interferometer at the A2 station, where the source size is slightly larger but is also 10 meters farther away; we expected fringe visibility at A2 comparable to what we saw at F3. In fact we saw an anomalously weak visibility. After substantial discussion, we eventually concluded that elements of the A2 beam path were the probable cause of the visibility degradation. In

particular, the A2 beam path has approximately 3 mm of unpolished beryllium windows and a 0.76 mm graphite filter, which we suspect is scattering the beam and creating an effective secondary source which is much larger and closer than the true source.

The A2 beam path also has a second graphite filter whose thickness is adjustable, just upstream of the first monochromator crystal. This filter causes significant scattering and effectively destroys the fringe visibility unless it is set to the minimum thickness (1 mm) or removed completely. With the filter completely out, the heat load on the first monochromator crystal is tremendous, causing a thermal bump that can also distort the apparent source size and divergence. A comparison of the measured and expected A2 visibility curves is shown in Fig. 3.4.

The predicted A2 source size is 1.07 mm [26]. The fit to the observed visibility estimates $\Sigma/L = 7.4 \times 10^{-5}$, which corresponds to a 2.62 mm source if $L = 35.5$ m (the nominal distance from the hutch center to the A2 source point). If, as we suspect, the 0.76 mm graphite filter is creating a secondary source, that source would be 15 m away from the hutch center; the resulting apparent source size would be 1.11 mm (FWHM).

At F3 we have found that the measured source size is often somewhat smaller than the predictions on the CHESS webpage; if we use the coupling constant derived at F3 to predict the A2 source size using the beamline-simulation software SPECTRA [52], we find an estimated vertical source size of 305 μm rms (720 μm FWHM) and an estimated source divergence of 27 μrad . From these two parameters and the knowledge that the distance from the source point to the graphite filter is 20.5 m, we find that the FWHM beam size at the graphite




Figure 3.4: Predicted visibility and actual measured visibility at the A2 station.

filter is ~ 2 mm, significantly larger than the data predicts.

It is highly unlikely that a new secondary source formed by scattering in the graphite filter should be smaller than the actual size of the beam when it reaches the filter. Therefore we conclude that not all the visibility degradation can be due to a secondary source formed by the graphite filter. There is still the possibility that the visibility degradation arises from interaction of the beam with the unpolished beryllium windows and the graphite filter, but not from the graphite filter alone.




Figure 3.5: Tilted moiré fringes produced by mismatched fringe periods. Here $p_2 = 1.1p_1$. (Image credit: [60]; notation modified to match this work.)

3.3.2 Radius of curvature

As we mentioned in Section 2.2, the gratings we use are nominally manufactured to account for the radius of curvature of the incident beam: the phase grating has a slightly smaller period than the ideal period, so that the slight magnification induced by the beam produces fringes that are correctly matched to the analyzer grating, rather than being slightly too large. If the fringes are not period-matched, tilted moiré fringes will be produced, as shown in Fig. 3.5. The

tilt angle θ_m of these moiré fringes is given by

$$\tan \theta_m = \frac{1}{2R \tan \beta} \cdot d, \quad (3.21)$$

where β is half the relative grating inclination angle, R is the wavefront radius of curvature (which, it should be noted, is not equal to the distance L between the source and the phase grating in the case of a synchrotron source), and d is the grating separation [60].

We often do not know the exact grating alignment half-angle β , but as [60] points out, it can be determined using the period of the moiré fringes, according to

$$p_m = \frac{\bar{p}}{2 \sin \beta}, \quad (3.22)$$

where p_m is the moiré fringe period and \bar{p} is the average of the analyzer period and the true fringe period. In principle, if we do not know the radius of curvature, we do not know the true fringe period either, but we can approximate \bar{p} simply by the analyzer grating period and generally be correct to within a few nanometers.

Since β is generally small enough that $\sin \beta \approx \beta$, the uncertainty in β is given by

$$\frac{\delta \beta}{\beta} = \sqrt{\frac{\delta p_m^2}{p_m^2} + \frac{\delta \bar{p}^2}{\bar{p}^2}} \quad (3.23)$$

If we approximate \bar{p} by the grating period, generally the uncertainty $\delta \bar{p}$ is $\sim 3\text{nm}$. The uncertainty in p_m comes from the method of calculation: we take the Fourier transform of a single image and find the reciprocal space location of the first harmonic peak, from which we determine the real space period p_m . Assume the uncertainty in the location of the first harmonic peak is one pixel (one bin in the Fourier transform). The resulting uncertainty in p_m will depend on the location

of the harmonic, but a typical uncertainty is $\pm 5\mu\text{m}$ for a p_m value of $150\mu\text{m}$. Combining these two uncertainties, we find that a typical fractional uncertainty in β is

$$\frac{\delta\beta}{\beta} = \sqrt{\frac{5^2}{150^2} + \frac{.003^2}{2^2}} = 0.033. \quad (3.24)$$

Given eqn. 3.21, we can measure the radius of curvature either locally or globally. From a single *square* moiré image taken at a known grating spacing, we can determine the overall average radius of curvature by taking the Fourier transform, locating the peak of the first harmonic in (k_x, k_y) reciprocal coordinates, and computing

$$\tan \theta_m = k_y/k_x. \quad (3.25)$$

For a non-square moiré image the algebra is more complicated but a similar equation applies. (This may look as though it should be incorrect by 90° because of the relationship between reciprocal and real space coordinates, but we are measuring θ_m relative to the vertical axis, which accounts for the difference.)

To determine the local radius of curvature, we subdivide the moiré image into smaller components and use the Fourier transform process on the subimages. It is necessary to have at least one full moiré fringe in each subimage in order for the first harmonic to be adequately separated from the DC component in reciprocal space, so the resolution is limited by the smallest detectable moiré fringe period, which in turn is limited by the detector resolution.

To improve the uncertainty in the measurement, we can take moiré interferograms at many different grating separations d and find $\tan \theta_m$ either locally or globally for each d , then fit the resulting data to a line which should, according to eqn. 3.21, have slope $1/(2R \tan \beta)$. (In fact, for our system we *must* take many interferograms for this measurement, since we have not ensured that the




Figure 3.6: Diagram indicating how to obtain the total beam radius of curvature given source size and divergence values from SPECTRA.

period-matched moiré fringes are vertically aligned.)

One such curve is shown in Fig. 3.7, for F3 in single-beam running mode. Note that because our setup does not guarantee vertical fringes in the case of perfect grating alignment, the line will not pass through the origin. We compute the fringe period to be $349 \mu\text{m}$, corresponding to $\beta = 0.164^\circ$. The fitted slope of the line is 0.00428 ± 10^{-5} , giving a calculated ROC of $41\text{m} \pm 1\text{m}$. We wish to compare this value to the value predicted by SPECTRA.

Given a file describing the parameters of a particular storage ring and beam-line, SPECTRA can predict the (rms) source size and divergence. Using these values, we can project back from the (finite-size) source to the location of a point source with the same divergence; see Fig. 3.6. For single-beam running, SPECTRA predicts an rms source size of $125.5 \mu\text{m}$ and an rms divergence of $20.11 \mu\text{rad}$. The distance from the source point to the center of the hutch is 21.6 m . The additional distance from the finite size source to the location of the theoretical point source is 6.2 m , so the total predicted radius of curvature is 27.8 m .




Figure 3.7: Moiré fringe angle as a function of grating spacing. The slope is a predictor of beam ROC, here estimated to be 41m.

Curiously, a similar plot for F3 in two-beam running mode predicts an ROC of 11 m, which is too *small* according to SPECTRA, which predicts an ROC of 31.1 m in two-beam mode, but which is closer to the 8.2 m value we computed in Chapter 2 that would account for the remaining moiré fringes at perfect grating alignment.

In both cases the discrepancy is much larger than is reasonably attributable to uncertainty in β or in the slope of the line. If the two-beam case were presented alone, we would hypothesize that a heat bump had disrupted the radius of curvature of the incident beam, but if that were the case, we should expect to see a similar effect in the single-beam case. Indeed, in the single-beam case, an ROC-decreasing effect from a heat bump should be even more severe, since the

beam has higher power density in the single-beam case than in the two-beam case. (The machine was running at the same current in both cases, but the beam had a smaller footprint in the single-beam case.)

We must assume that there is some other systematic error in the experimental setup that we have not yet been able to pinpoint. It is suggestive that the residuals of the data with respect to the fit seem to be periodic in the grating spacing with a period that seems to be equal to $d_1 = z_T/16$, but the implications are unclear.

3.4 Case study: Horizontal-vertical accelerator coupling

The *emittance* ε of a beam describes the beam's extent in phase space, i.e. if a source has transverse size σ_x and σ_y and divergence $\sigma_{x'}$ and $\sigma_{y'}$, the emittance is

$$\varepsilon = \varepsilon_x \varepsilon_y = (\sigma_x \sigma_{x'}) \cdot (\sigma_y \sigma_{y'}) \quad (3.26)$$

The accelerator operators tune the machine to a particular design horizontal emittance ε_x . The design vertical emittance is 0, but there is always some coupling between the horizontal and vertical tuning parameters that leads to a nonzero vertical emittance ε_y . We can compute the vertical emittance based on measurements of the source size and radius of curvature. Knowing the design horizontal emittance, we can then compute the coupling percentage. Typically the coupling is desired to be <1%, but the accelerator scientists often increase the coupling in order to improve beam lifetimes.

The maximum visibility observed with the Talbot interferometer at F3 during the March 2010 CHESS run was 45%. After accounting for vibrational insta-

bility in the monochromator, we predict the visibility would be 54% in the perfectly stable case. Backing out from the visibility formula 3.17, with $V_f(d) = 1$, we find a source size of 780 microns FWHM, or 330 microns rms.

We then turn to SPECTRA to determine the vertical emittance, since we are not sufficiently certain of the radius of curvature measurements to use them in this calculation. Using Twiss parameters [4] provided by accelerator scientist Mike Forster, and assuming a design horizontal emittance of 130 nmRad, we find that SPECTRA predicts $\sigma_y = 330$ microns when the horizontal-vertical coupling is 3.7%. This is in good agreement with measurements taken at E-line in 2005, which found a vertical emittance of 4.9 nmRad, corresponding to 3.8% coupling.

3.5 Case study: Miscut crystal effects on apparent source parameters

In a standard symmetrically-cut monochromator crystal, the lattice planes are parallel to the crystal surface. But a crystal can also be “miscut” so that the lattice planes and the surface are at a slight angle to one another. In this case the miscut crystal can be used to focus the beam (at the cost of increased divergence) or to expand the beam, as shown in Fig. 3.8. It is reasonable to expect that this beam compression or expansion reflects a change in the apparent source size, and that this change should be apparent in the fringe visibility curve.

CHESS has obtained a collection of silicon monochromator crystals with 3.2°, 5°, and 6.2° miscut angles. In March of 2012, we installed these crystals se-




Figure 3.8: Beam geometries for symmetric-cut (top) and miscut (bottom) monochromator crystals. The miscut can also be reversed to compress the beam rather than expanding it.

quentially in the A2 monochromator box and recorded visibility curves for each crystal; in the case of the 5° miscut, we tested the crystal both in beam expansion configuration and in beam focusing configuration.

The overall visibility at A2 is not as high as at F3, but there is enough visibility to see that the apparent source size is increased by the 3.2° miscut crystal over the “bare” source size from the symmetrically-cut crystal, as illustrated in Fig. 3.9. Oddly, the 5° miscut has approximately the same effect on apparent source size as the 3.2° miscut; we would expect a larger apparent source corresponding to the larger miscut. Bizarrely, the 6.2° miscut showed $> 80\%$ visibility.




Figure 3.9: Visibility curves for symmetric crystals and for 3.2° and 5° miscuts. Visibility is decreased by approximately the same amount for both miscuts.

This was so inexplicably anomalous that we limit the following analysis to the symmetric cut and the 3.2° and 5° asymmetric cuts.

In the beam focusing configuration of the 5° miscut crystal, we expect to see a smaller apparent source size. The visibility curve bears this out, as we can see in Fig. 3.10: the second peak of the visibility is suppressed by the Gaussian envelope for the symmetrically-cut crystals, but is present in the data for the 5° miscut.

Since the phase space area of the beam is fixed, it is also reasonable to suppose that an increase (or decrease) in apparent source size should lead to a decrease (or increase) in apparent source divergence. Using the same series of




Figure 3.10: Visibility curves for symmetric crystals and for 5° miscut in beam focusing orientation. The second peak of the visibility is not suppressed for the focused beam, indicating decreased apparent source size.

images from which the visibility is calculated, we can compute the apparent radius of curvature for each miscut crystal configuration (except the 6.2° miscut, where visibility was too anomalous to draw any conclusions).

The computed radii of curvature (ROCs) and their relative values, normalized to the radius of curvature for the symmetrically-cut crystals, are shown in Table 3.1. The calculated ROCs are quite large; the ROC predicted by the SPECTRA-estimated size and divergence of the source should be closer to 47 m for symmetric crystals. (As in the single- and two-beam ROC measurements at F3, the discrepancy is too large to attribute to errors in p_m or β . The source of the discrepancy remains unclear.) Nevertheless, the normalized ROCs show

Miscut	ROC	Normalized ROC
Symmetric	104 m	1
3.2°	184 m	1.8
5°	347 m	3.3
5° (backwards)	70 m	.7

Table 3.1: Calculated and relative radii of curvature for symmetric and miscut crystals

the expected trend: increasing miscut increases the relative ROC, corresponding to a less divergent beam, while reversing the miscut into focusing geometry decreases the relative ROC, corresponding to a more divergent beam.

CHAPTER 4

DATA PROCESSING

Recall that in the phase-stepping data, each pixel's intensity can be approximated by the sinusoid

$$I(x_g) = a_0 + a_1 \cos(2\pi x_g/p_2 + \varphi). \quad (4.1)$$

We need ways to extract the parameters a_0 , a_1 , and φ (illustrated graphically in Fig. 4.1) in order to obtain absorption, dark-field, and differential phase-contrast images of our specimens.

There are numerous methods to extract some or all of these parameters, most of which are ideal in theory but problematic in execution. We will first describe several of these methods briefly, then illustrate recurring artifacts arising in the processed data and elaborate on a long series of attempts to manage or eliminate these artifacts.

4.1 Data processing methods

4.1.1 Fast Fourier Transform (FFT)

The Fast Fourier Transform (FFT) is by far the most common data processing method in grating interferometry. In this method, the FFT is applied separately to each pixel's sinusoidal data. The absorption parameter a_0 is extracted from the DC (zero-frequency) term and the dark-field and DPC parameters a_1 and φ are extracted from the magnitude and phase of the first harmonic peak.

The FFT has the advantage of being very fast, taking only a few seconds on




Figure 4.1: Illustration of the parameters a_0 , a_1 , and φ relative to the overall sinusoids. Black line/triangles: a reference dataset with no specimen in place. Blue line/circles: a dataset modulated by a specimen. Phase φ is indicated by vertical dashed lines.

an Intel i7 2.2 GHz processor even for 5.5 megapixel images with 18 phase steps. The main disadvantage of the FFT is that the data must be sampled with very high precision: the phase stepping sampling range must be an exact multiple of the true period of the data in order to prevent “spectral leakage” of the Fourier data out of its correct frequency bin into neighboring bins. Spectral leakage leads to artifacts in the recovered parameters in the form of shadows of the original moiré fringes; we will describe such artifacts in more detail in Sec. 4.2.

4.1.2 Levenburg-Marquardt algorithm (LMA)

The Levenburg-Marquardt algorithm (LMA) is an iterative nonlinear least-squares fitting routine [30]. Like all curve fitting, LMA fitting requires a model function with variable parameters. We used either

$$f_1(x) = a_0 + a_1 \cos(2\pi x/p + \varphi) \quad (4.2)$$

or

$$f_2(x) = a_0 + a_1 \cos(2\pi x/p + \varphi) + a_2 \cos(4\pi x/p + \varphi) \quad (4.3)$$

in which all the parameters including the period p were allowed to vary during the fitting process. The second cosine term in f_2 accounts for the fact that moiré fringes are not true sinusoids but rather truncated triangle waves. We could make successively better triangle wave approximations by adding higher-order cosine terms to the model function, at the cost of increased fitting time for each additional term.

LMA fitting is also prone to artifacts, but they tend to be slightly less severe than those resulting from FFT processing. The main disadvantage of LMA fitting is that it is extremely slow; a 1 megapixel dataset takes approximately half an hour to process on a machine with an Intel i7 2.2 GHz processor and 8 GB of RAM (with no parallel processing implemented). Least-squares fitting is also known to be ill-conditioned [56] when fitting sinusoids. In one sense, to say a problem is *ill-conditioned* is to say that small changes in the inputs may sometimes lead to large and/or discontinuous changes in the outputs, so one reason sinusoid-fitting is ill-conditioned is the presence of degenerate minima at $\varphi + 2\pi k$ for $k \in \mathbb{Z}$. The slowness of fit that results from the ill-conditioning is sufficiently problematic that, although we used LMA fitting routinely for many

months, we would not now recommend its use, and we include it here mainly for completeness.

4.1.3 Principal components analysis (PCA)

Principal components analysis (PCA) is a technique borrowed from the field of optical phase-shifting interferometry (PSI). Researchers in PSI take similar phase-stepping curves, sometimes with mirrors that may be positioned incorrectly by their piezoelectric transducers during the measurement [20]. Therefore the exact position of the mirror at each phase step may not be known. PCA separates the specimen phase from the mirror (or, for us, grating) displacement phase and thereby makes it possible to recover the specimen phase without knowing the displacements precisely.

PCA in layman's terms

Given a cluster of data points, PCA is an attempt to determine a set of vectors (the principal components) that describe the data as well as possible. By this we mean that the set of vectors should contain no redundant information, i.e. they should be orthogonal; and collectively they should account for as much of the variance in the data as possible. Specifically, the first vector should account for as much of the variance as it can; the second should account for as much of the remaining variance as possible once the variance described by the first vector is removed; and so on.

Given this set of vectors, PCA can also be used to reduce the data into a

lower-dimensional space as well as possible, by projecting onto the most descriptive possible set of vectors (as specified by PCA) with the appropriate dimensionality. This is the technique we will use on our data: we begin with a set of N images and reduce it to the best possible 2-dimensional dataset, from which we extract phase information.

The PCA algorithm

PCA consists of three basic tasks:

1. Find the covariance matrix C of the data.
2. Diagonalize the covariance matrix as $D = ACA^T$.
3. Apply the diagonalization transform matrix A to the original data.

Finding the covariance matrix

Finding the covariance matrix is an easy step: if we have N images of size $N_x \times N_y$, we rewrite the i th image into a single column vector x_i with length $N_x \times N_y$ and stack all the images together into a matrix

$$X = [x_1, x_2, \dots, x_N]^T. \quad (4.4)$$

We need to mean-subtract (zero-center) the data so that the first principal component is meaningful; otherwise the vector will tend to point toward the center of the data cluster and not actually correctly account for the variance, as shown in Fig. 4.2. We do this by subtracting a vector m_x that corresponds to the mean of each column of X , i.e. the mean of each pixel's data. Then the covariance matrix




Figure 4.2: An example illustrating the need to zero-center PCA data. (a) The first principal component points only to the mean of the data. (b) The first principal component accurately accounts for a large quantity of the variance.

C is

$$C = (X - m_x)(X - m_x)^T. \quad (4.5)$$

Our data is N -dimensional (one dimension for each image, since we have written our images out in one-dimensional vectors instead of in 2D), so the covariance matrix is $N \times N$ and describes the covariance of each image with respect to each other image.

Diagonalizing the covariance matrix

Because C is real and symmetric, the spectral theorem applies [56] and we can diagonalize C and be assured that the eigenvectors will all be orthogonal. We can then write

$$D = ACA^T, \quad (4.6)$$

where D is the diagonal matrix of eigenvalues of C and A is the matrix of eigenvectors of C . The orthogonality of the eigenvectors is what makes them viable principal components.

Transforming the original data

By applying the matrix A to the original data $(X - m_X)$ we can project the data onto the principal component vectors. So our final principal components are

$$y = A(X - m_X). \quad (4.7)$$

Applying PCA to interferometry data

We assume as usual that the pixel data is well-represented by the model function

$$f(x_n) = a_0 + a_1 \cos(2\pi x_n/p_2 + \varphi), \quad (4.8)$$

where $2\pi x_n/p_2$ are the phase steps and φ is the specimen phase, as usual. In order to apply PCA to the data, we rewrite the model function as

$$f(x_n) - a_0 = a_1 \cos(2\pi x_n/p_2) \cos \varphi - a_1 \sin(2\pi x_n/p_2) \sin \varphi = A_n \cos \varphi + B_n \sin \varphi. \quad (4.9)$$

Now for each pixel we have zero-centered data in the variables $\cos \varphi$ and $\sin \varphi$, assuming a_0 is well-approximated by the mean of the pixel's data. (This is not

as simple as it sounds; if the range spanned by the sampled points is not an exact multiple of p_2 , the mean of the data will not exactly correspond to the mean a_0 of the pure sinusoid. The effect is that PCA is not immune to the fringe artifacts that also plague the FFT, although they do tend to be less severe.) We will use the data for all the pixels together to construct the covariance matrix C , as described in Sec. 4.1.3.

If we can project all the data down onto the two variables $\cos \varphi$ and $\sin \varphi$, we can recover the phase data for all the pixels at once by taking

$$\varphi = \arctan\left(\frac{P_{\sin \varphi}}{P_{\cos \varphi}}\right), \quad (4.10)$$

where $P_{\sin \varphi}$ is the principal component projection corresponding to $\sin \varphi$ (likewise for $\cos \varphi$). (It is an unfortunate confluence of terminology that both the principal component vectors and the projections of the data onto them are called “principal components.”) Therefore we will seek a relationship between $\cos \varphi$ and $\sin \varphi$ and the principal components. If they are themselves the principal components, projecting onto them leads to the best possible 2D representation of the data; and furthermore that in the event they are not the principal components, we can still recover the phase data up to an overall global phase shift. To show that $\cos \varphi$ and $\sin \varphi$ are candidates to be principal components, we must first show that they are orthogonal, which is equivalent to showing they are uncorrelated.

Two signals are approximately uncorrelated if their covariance is significantly less than the variance of either signal separately. For our signals, this means

$$\sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \cos \varphi(x, y) \cdot \sin \varphi(x, y) = \frac{1}{2} \sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \sin^2 \varphi(x, y) \ll \sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \cos^2 \varphi(x, y) \quad (4.11)$$

(where (N_x, N_y) is the size of our image) [57, 63], or more strictly,

$$\sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \cos \varphi(x, y) \cdot \sin \varphi(x, y) \approx 0. \quad (4.12)$$

The signals $\cos \varphi$ and $\sin \varphi$ are only genuinely orthogonal when sampled over an exact multiple of 2π . In particular, if the samples do not even span one period, condition 4.12 will certainly not hold. However, as long as there are enough fringes (generally, two or more) in the field of view, this condition holds for interferometry data; the more fringes there are, the better the approximation 4.12 will be. So our variables are typically approximately uncorrelated and therefore are good candidates for principal components in two dimensions.

Moreover, since $\cos \varphi$ and $\sin \varphi$ are orthogonal, they form a basis for the 2D plane in which they lie. Any other orthogonal basis for this space can be at most different from $\cos \varphi$ and $\sin \varphi$ by a rotation over an angle θ . So if $\cos \varphi$ and $\sin \varphi$ are not the principal components, we can nevertheless write the principal components P_1 and P_2 in terms of φ and θ by

$$\begin{aligned} P_1 &= \cos \theta \cos \varphi + \sin \theta \sin \varphi \\ P_2 &= -\sin \theta \cos \varphi + \cos \theta \sin \varphi \end{aligned} \quad (4.13)$$

so the ratio of the principal components is

$$\frac{-\sin \theta \cos \varphi + \cos \theta \sin \varphi}{\cos \theta \cos \varphi + \sin \theta \sin \varphi} = \frac{\sin(\varphi - \theta)}{\cos(\varphi - \theta)} \quad (4.14)$$

which means we can recover the phase φ up to a global phase shift θ by taking P_2/P_1 and taking the arctangent. In fact, it turns out to be the case that $\theta = 0$ in all cases; for a pleasantly lucid derivation, see [58].

As a caveat, it need not be the case that $\cos \varphi$ is always P_1 and $\sin \varphi$ is always P_2 . There exist datasets for which the two are swapped, so that $\cos \varphi$ is P_2

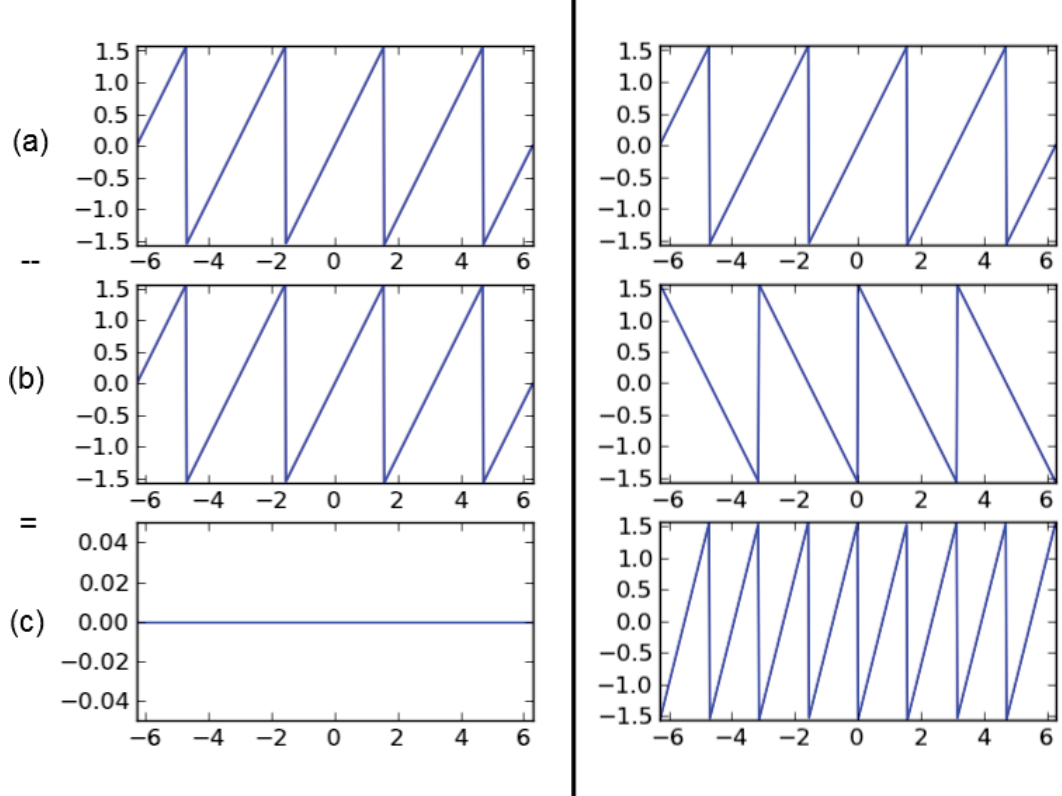


Figure 4.3: Simulated lineouts of two recovered PCA patterns: (a) specimen and (b) reference. Left column: (a) and (b) have the same principal component order, giving a uniform recovered phase profile. Right column: (a) and (b) have different principal component order, leading to striping in the recovered profile.

and $\sin \varphi$ is P_1 . In that case, $P_2/P_1 = \cot \varphi$. Note that $\arctan(\cot \varphi)$ differs from $\arctan(\tan \varphi)$ by an overall sign and by a $\pi/2$ phase shift. If a specimen dataset and a reference dataset have different principal component order (i.e. P_1 is $\cos \varphi$ for one and $\sin \varphi$ for the other), then the subtraction of the reference phase from the specimen phase will lead to stripes in the overall recovered phase, as shown in Fig. 4.3. If one encounters this behavior in PCA-processed data, the simple fix is to swap P_1 and P_2 for one (not both) of the datasets.

The PCA is approximately as fast as the FFT, taking only a second or two on an Intel i7 2.2 GHz processor with 8 GB of RAM. Moreover, the PCA can be

computed synchronously, recalculating every time a new image is collected until the user is happy with the PCA result. Also, the PCA is phase-step agnostic, requiring no particular sampling protocol.

4.1.4 Generalized phase shifting algorithms (GPSAs)

Generalized phase shifting algorithms (GPSAs) have, to the best of our knowledge, never before been applied to grating interferometry data. Papers describing GPSAs tend to be dense and based primarily in the field of PSI, which has surprisingly little contact with grating interferometry considering the similarity of the analysis methods.

GPSAs are appealing because they can be made insensitive to both higher harmonics and detuning from the true data frequency, two problems which plague other analysis techniques. Unlike PCA, they do not require mean subtraction in order to function well. Finally, applying a GPSA requires only the solution of a small linear system of equations, which allows them to challenge even the FFT for speed.

To illustrate GPSAs and show how they can be made insensitive to harmonics and detuning, we will follow the methods and (mostly) the notation of [54] sections 3-5, with attention given especially to illuminating in layman's terms the rather terse argument in the paper. We begin as in [54] with a basic GPSA which is insensitive to harmonics but not to detuning, and then show how this basic GPSA can be refined into a more useful form. What follows is a lengthy motivation for what, in the end, will turn out to be a simple prescription for analyzing grating interferometry data in the presence of detuning.

The basic algorithm

First we define our data as samples s_k of a model function with a finite number M of harmonics:

$$s_k = \sum_{m=1}^M a_m \cos[m(\varphi - \phi_k)], \quad k = 1, \dots, K, \quad (4.15)$$

where ϕ_k are the phase steps and φ is the specimen phase shift, as usual. Note that for now we are considering only a single pixel's data, but we will see that the results are generally applicable to arrays of pixel data.

Next define the sequential differences between samples

$$\Delta s_k^{k+1} = s_{k+1} - s_k \quad (4.16)$$

$$= \sum_{m=1}^M a_m \left[\cos(m\varphi) \Delta \cos(m\phi)_k^{k+1} + \sin(m\varphi) \Delta \sin(m\phi)_k^{k+1} \right], \quad (4.17)$$

where $\Delta \cos(m\phi)_k^{k+1} = \cos(m\phi_{k+1}) - \cos(m\phi_k)$ and likewise for the sine term.

If we assume we know the phase steps ϕ_k , the $\Delta \cos$ and $\Delta \sin$ terms are known scalars. We can use them to define systems of linear equations in variables A_k^p and B_k^p ($p = 1, 2, \dots, M$) as follows:

$$\sum_{k=1}^{K-1} (\Delta \sin(m\phi)_k^{k+1}) A_k^p = 0 \quad \forall m = 1, 2, \dots, M \quad (4.18)$$

$$\sum_{k=1}^{K-1} (\Delta \cos(m\phi)_k^{k+1}) A_k^p = \begin{cases} 0 & \forall m \neq p \\ 1 & m = p \end{cases} \quad (4.19)$$

$$\sum_{k=1}^{K-1} (\Delta \cos(m\phi)_k^{k+1}) B_k^p = 0 \quad \forall m = 1, 2, \dots, M \quad (4.20)$$

$$\sum_{k=1}^{K-1} (\Delta \sin(m\phi)_k^{k+1}) B_k^p = \begin{cases} 0 & \forall m \neq p \\ 1 & m = p \end{cases} \quad (4.21)$$

For each p we now have two systems each with $2M$ equations and $K - 1$ unknowns. In order to have enough data to solve the systems, we need $K - 1 = 2M$,

i.e. $K = 2M + 1$ sampled points. (Note that we are not guaranteed that the systems have a solution even if we have $K = 2M + 1$ sampled points; it is remotely possible to choose the phase steps ϕ_k in such a pathological way that the resulting matrix for the linear system is rank-deficient [56], in which case the system *may, but need not*, have a solution. In general, however, it will.)

Why have we bothered to write such a system of linear equations? Consider the ratio

$$\frac{\sum_{k=1}^{K-1} B_k^p \Delta s_k^{k+1}}{\sum_{k=1}^{K-1} A_k^p \Delta s_k^{k+1}}. \quad (4.22)$$

If we plug in the expanded form of Δs_k^{k+1} , we have

$$\begin{aligned} \sum_{k=1}^{K-1} B_k^p \Delta s_k^{k+1} &= \sum_{k=1}^{K-1} B_k^p \cdot a_m \left[\cos(m\varphi) \Delta \cos(m\phi)_k^{k+1} + \sin(m\varphi) \Delta \sin(m\phi)_k^{k+1} \right] \\ &= a_m \cos(m\varphi) \sum_{k=1}^{K-1} B_k^p \Delta \cos(m\phi)_k^{k+1} \\ &\quad + a_m \sin(m\varphi) \sum_{k=1}^{K-1} B_k^p \Delta \sin(m\phi)_k^{k+1} \\ &= a_p \sin(p\varphi) \end{aligned} \quad (4.23)$$

because of eqns. 4.20 and 4.21. Likewise

$$\sum_{k=1}^{K-1} A_k^p \Delta s_k^{k+1} = a_p \cos(p\varphi). \quad (4.24)$$

Then the ratio 4.22 is equivalent to

$$\frac{\sum_{k=1}^{K-1} B_k^p \Delta s_k^{k+1}}{\sum_{k=1}^{K-1} A_k^p \Delta s_k^{k+1}} = \frac{a_p \sin(p\varphi)}{a_p \cos(p\varphi)} = \tan(p\varphi), \quad (4.25)$$

and therefore for $p = 1$ we can recover the phase information for our specimen by taking

$$\arctan \left(\frac{\sum_{k=1}^{K-1} B_k^1 \Delta s_k^{k+1}}{\sum_{k=1}^{K-1} A_k^1 \Delta s_k^{k+1}} \right). \quad (4.26)$$

The detuning-insensitive algorithm

Suppose we sample the data at a reference frequency f_r so that $\phi_k = (2\pi f_r)x_n$ for discrete-space values x_n . This reference frequency may be detuned from the true frequency f of the model function. Suppose further (without loss of generality) that K is odd, so that $K = 2\tau + 1$ for some $\tau \in \mathbb{Z}$. Also, let us require that the sampled points ϕ_k are symmetrically distributed around 0, so that

$$\phi_j = -\phi_{K+1-j}, \quad j = 1, \dots, \tau, \quad \phi_{\tau+1} = 0. \quad (4.27)$$

Finally, let us assume that we are only interested in $p = 1$, so that we can simplify our notation from B_k^p and A_k^p to B_k and A_k . Note that all of these suppositions are valid for our interferometry protocols.

We will define two frequency-dependent functions Am_N and Am_D (known for technical reasons as amplitude functions) as

$$Am_N(f) = \sum_{j=1}^{\tau} 2B_j \Delta \sin(f\phi/f_r)_j^{j+1} \quad (4.28)$$

$$Am_D(f) = \sum_{j=1}^{\tau} 2A_j \Delta \cos(f\phi/f_r)_j^{j+1}. \quad (4.29)$$

As an aside, the nomenclature Am_N and Am_D stands for “numerator” and “denominator” because the amplitude functions are conventionally used in a more sophisticated form of eqn. 4.22. Although we will not describe the amplitude functions further here, we retain the notation for consistency with [54].

First, we reconstruct the system of equations 4.18–4.21 when $f = f_r$; that is, when there is no detuning, this algorithm should return the same result as the original algorithm. In the new notation of the amplitude functions, we write

$$\sum_{j=1}^{\tau} (2\Delta(\sin \phi)_j^{j+1}) B_j + \sum_{j=1}^{\tau} 0 \cdot A_j = 1$$

$$\begin{aligned}
\sum_{j=1}^{\tau} 0 \cdot B_j + \sum_{j=1}^{\tau} (2\Delta(\cos \phi)_j^{j+1}) A_j &= 1 \\
\sum_{j=1}^{\tau} (2\Delta \sin(m\phi)_j^{j+1}) B_j + \sum_{j=1}^{\tau} 0 \cdot A_j &= 0 \quad m = 2, \dots, M \\
\sum_{j=1}^{\tau} 0 \cdot B_j + \sum_{j=1}^{\tau} (2\Delta \sin(m\phi)_j^{j+1}) A_j &= 0 \quad m = 2, \dots, M
\end{aligned} \tag{4.30}$$

with the 0-multipliers included to make it clear that we have a system with both B_j and A_j as unknowns in each equation.

We add another equation to render the algorithm insensitive to detuning in the first harmonic (the harmonic of interest):

$$\frac{dAm_N(f_r)}{df} = \frac{dAm_d(f_r)}{df}. \tag{4.31}$$

This means that, in the more sophisticated form of eqn. 4.22 in which the amplitude functions are used, the numerator and the denominator should be changing at the same rate near the sampling frequency. In expanded notation, this corresponds to

$$\sum_{j=1}^{\tau} ((2/f_r)\Delta\phi \cos \phi_j^{j+1}) B_j + \sum_{j=1}^{\tau} ((2/f_r)\Delta\phi \sin \phi_j^{j+1}) A_j = 0. \tag{4.32}$$

Finally, we add a similar set of equations to make the algorithm insensitive to higher harmonics in the presence of detuning:

$$\frac{dAm_N(mf_r)}{df} = \frac{dAm_D(mf_r)}{df} = 0, \tag{4.33}$$

or in expanded notation,

$$\begin{aligned}
\sum_{j=1}^{\tau} ((2/f_r)\Delta\phi \cos(m\phi)_j^{j+1}) B_j + \sum_{j=1}^{\tau} 0 \cdot A_j &= 0 \quad m = 2, \dots, M \\
\sum_{j=1}^{\tau} 0 \cdot B_j + \sum_{j=1}^{\tau} ((-2/f_r)\Delta\phi \sin(m\phi)_j^{j+1}) A_j &= 0 \quad m = 2, \dots, M
\end{aligned} \tag{4.34}$$

Now we have a total of $4M - 1$ equations: $2M$ from eqn. 4.30, 1 from eqn. 4.32, and $2M - 2$ from eqn. 4.34. There are $2\tau = K - 1$ unknowns (τ of the A_j and τ of the B_j), so to have enough data to solve the system we need $K - 1 = 4M - 1$, which would imply that K is even, but we assumed K was odd. To eliminate this problem, we remove the requirement that we be insensitive to the M th harmonic in the presence of detuning, reducing the total number of equations to $4M - 2$ and allowing K to be odd. (Note again that it is possible to choose the phase steps ϕ_k in such a way that the linear system has no solution, but this is highly unlikely for typical phase-stepping protocols.)

Applying the detuning-insensitive GPSA

To apply the detuning-insensitive data to interferometry data, we note that we nearly always sample our data over two grating periods, $4\mu\text{m}$, which corresponds to $8\mu\text{m}$ stroke in the piezo because of the 2:1 reduction given by the flexure arm. We use the middle of the piezo's total 0– $12\mu\text{m}$ stroke to minimize nonlinearities, so our sampling points (in microns) are typically

$$x_n = [2, 2.8, 3.6, \dots, 9.2, 10], \quad (4.35)$$

assuming we are sampling two periods, which we can zero-center as

$$x'_n = [-4, -3.2, \dots, 0, \dots, 3.2, 4]. \quad (4.36)$$

to work in accordance with eqn. 4.27. The reference frequency f_r of these samples is $1/4 = 0.25$.

We have $K = 19$ data points, so we can be insensitive to harmonics up to $M = (K + 1)/4 = 5$ and insensitive to detuning in the harmonics up to $M - 1 = 4$

(because, recall, we removed the last equation in the harmonic detuning set of equations so as to have a solvable set of equations). Because we know our sampling points $\phi_n = (2\pi f_r)x_n$, we know the differences $\Delta \cos(m\phi)_j^{j+1}$ for all j and m and likewise for sine.

We can therefore build up the system of equations from eqns. 4.30, 4.32, and 4.34 explicitly, with known values for the coefficients which will be consistent over all datasets as long as they are sampled with the same sample points. We can solve this system of equations to find

$$\begin{aligned} A &= [-6.069, -5.540, -4.709, -3.933, -3.549, -3.724, -4.389, -5.259, -5.958] \\ B &= [-0.457, -1.186, -1.443, -1.141, -0.435, 0.348, 0.842, 0.804, 0.218] \end{aligned} \quad (4.37)$$

Because we wrote $K = 2\tau + 1$ and we only allowed A_i and B_i to range from 1 to τ , we have only half the requisite coefficients to recover the phase. It is possible but tedious to show that the correct full set of coefficients is

$$\begin{aligned} A_{full} &= [A_1, A_2, \dots, A_\tau, -A_\tau, \dots, -A_2, -A_1] \\ B_{full} &= [B_1, B_2, \dots, B_\tau, B_\tau, \dots, B_2, B_1] \end{aligned} \quad (4.38)$$

Then in order to recover the phase information from our data sample points s_j , we simply take

$$\arctan \left(\frac{\sum_{j=1}^K B_{full} \Delta s_j^{j+1}}{\sum_{j=1}^K A_{full} \Delta s_j^{j+1}} \right). \quad (4.39)$$

This is extremely fast: well under 1 second on an Intel i7 2.2GHz processor with 8GB of RAM. As long as the same sampling points are used, these coefficients need not be recomputed; to do GPSA processing, one needs only compute the sums and take the arctangent as in eqn. 4.39.

4.1.5 Single-shot processing

In the single-shot method, only one image is collected per scan, with fringes as dense as the detector can resolve. The fringes are used as carrier fringes, whose modulations contain information about the phase structure of the specimen. The analysis protocol is as follows:

- Fourier transform the single-shot image;
- Locate the first harmonic peak and its surrounding nonzero information;
- Make a new array of zeroes the same size as the original Fourier transform;
- Transfer the data from the first harmonic peak into the center of the new array;
- Inverse Fourier transform the array containing just the first harmonic peak.

This process is illustrated in Fig. 4.4. The main advantage of the single-shot method is that it is completely immune to the fringe artifacts that plague all the other methods, since those fringes are serving entirely as carriers of the interesting data. The disadvantage is that its resolution is limited by the same fact: we can do no better than the period of the carrier fringes. In contrast, the other techniques are limited by the detector resolution, which is much smaller.

Since it requires only two single-image Fourier transforms, the single-shot method will be fast on nearly all machines.




Figure 4.4: Illustration of the single-shot processing method. (a) The original image. (b) The Fourier transform of the image. (c) The first harmonic peak of the Fourier transform, shifted to the center. (d) The inverse Fourier transform, containing phase data.




Figure 4.5: Absorption, dark-field, and differential phase reconstructions showing reconstruction fringe artifacts in the images of a charcoalified fossil flower. (Specimen: Dr William Crepet, Cornell Plant Biology)

4.2 Processing artifacts: gory details

Most if not all of our datasets exhibit fitting artifacts in the form of shadows of the original moiré fringes. Using a metric of severity that compares the visibility of the fringe artifacts to the natural visibility scale of the actual data, the artifacts tend to be most severe in the dark-field data, next most problematic in the phase data, and often mild (but rarely absent) in the absorption data. Moreover, the artifacts in the dark-field and phase data have half the period (twice the frequency) of the artifacts in the absorption data, which match the original moiré fringes. These artifacts are illustrated in Fig. 4.5.

Other groups seem to have similar problems. In correspondence with a student from Christian David's group, he commented that fringe artifacts are nearly always present to some degree, and in conversation with other grating interferometry researchers, they sometimes express surprise that we operate with any fringes in the field of view at all, rather than perfectly aligning the gratings to remove moiré fringes. The reason is that our gratings are not perfectly

divergence-matched because of a calculation error when defining their specifications, so that it is actually not possible for us to remove all fringes from the field of view. Other groups that have a longer history in the field tend to be more tightly coupled to grating manufacturing groups and can iteratively work to get the exact correct grating period.

However, there does not seem to be any published information to this effect in the grating interferometry field. The best methods for mitigating the problem turn out to be from another field entirely: optical interferometry, which obtains similar phase-stepping data.

4.2.1 Spectral leakage: origins in the FFT

As mentioned above in the description of the FFT method, the primary source of artifacts in FFT-processed data is spectral leakage from the main harmonic peak into neighboring frequency bins. This spectral leakage is due to a mismatch between the true period of the fringes and the period that is assumed when sampling; the Fourier transform algorithm assumes that the sampled data consists of an integer number of periods of a periodic function, and the period error causes the sampled data to violate this condition slightly.

This problem is illustrated in Fig. 4.6. In both cases the period of the sampled sinusoid is *assumed* to be 4, and the data is sampled at points $x = [-4, -3.2, -1.6, \dots, 3.2, 4]$. In panel (a) the true period of the sampled sinusoid is also 4, so that the sampling matches the periodicity. In panel (b), the true period of the sampled sinusoid is only 3.6. (We use an extreme example so the problem is clearly illustrated.) The clean first harmonic peak from panel (a) has




Figure 4.6: FFT comparison showing (a) pure first harmonic in the correctly-sampled case, where the sampling period matches the sinusoid period, and (b) spectral leakage in which some information from the harmonic peak has moved into neighboring frequency bins in the incorrectly-sampled case, where the sampling period deviates from the true sinusoid period by 10%.

undesirable wings in panel (b).

4.2.2 Down the rabbit hole: least-squares fitting with LMA

After discovering the artifacts resulting from FFT processing, we began using LMA least-squares fitting in an attempt to minimize or eliminate the appearance of the artifacts. Unfortunately, LMA-processed data also exhibits the same artifacts, but unlike the spectral leakage in the FFT, the source of the fringe arti-

facts in LMA-processed data is not obvious.

Although we spent significant amounts of time attempting to describe and understand the artifacts arising in LMA-processed data, we never managed to track them down to their source. Eventually we abandoned this line of inquiry and moved on to other processing methods. Nevertheless, the results of this work are included here for completeness.

Mock data: single sinusoid

Assume the data is defined by

$$y = 20000 + 10000 \cos(2\pi x/p_2 + \varphi). \quad (4.40)$$

Define a set of sampling points $x = \text{linspace}(2, 10, 19)$ [45] (to match our usual piezo positions) and collections of periods and phases $p_2 = \text{linspace}(3.5, 4.5, 20)$ and $\varphi = \text{linspace}(0, 2\pi, 40)$, respectively. From these, build up a 20×40 array of 19-point sinusoidal data functions, one for each (p_2, φ) pair. Then fit the datasets using `scipy.optimize.leastsq` [47] with two model functions,

$$f_1(x) = a_0 + a_1 \cos(2\pi x/p_2 + \varphi) \quad (4.41)$$

and

$$f_2(x) = a_0 + a_1 \cos(2\pi x/p_2 + \varphi) + a_2 \cos(\pi x/p_2 + \varphi + \pi/4), \quad (4.42)$$

using initial parameter guesses

$$[a_0, a_1, a_2, \varphi, p_2]_0 = [\text{mean}(y), (\text{max}(y) - \text{min}(y))/2, 1, 1, 4]. \quad (4.43)$$

In this case it is not possible to replicate the fringe artifacts; in all cases the recovered parameters matched the input parameters after correcting for the fact that occasionally a_1 is recovered as negative and φ was consequently off by π .




Figure 4.7: a_0 (recovered), a_1 (recovered), and $\varphi_{\text{rec}} - \varphi_{\text{input}}$ for a 38×20 collection of two-sinusoid mock data (results from first two φ values stripped). Vertical scale: φ index; horizontal scale: p_2 index.

Mock data: two sinusoids

Using the same x , p_2 , and φ as before, we created mock datasets given by

$$y = 20000 + 10000 \cos(2\pi x/p_2 + \varphi) + 1000 \cos(\pi x/p_2 + \varphi + \pi/4). \quad (4.44)$$

I then fit the data with model function $f_1(x)$. In this case, the recovered parameters *do* vary from the input parameters periodically with the input phase φ_{input} , as shown in Fig. 4.7.

I could not replicate this behavior using the model function $f_2(x)$; the two-sinusoid mock data is recovered correctly by the two-sinusoid model (up to the previously mentioned $a_{1,\text{rec}} = -a_{1,\text{input}}$, $\varphi_{\text{rec}} = \varphi_{\text{input}} + \pi$ errors).

Real data: modifying the model function

Originally we began using the model function $f_2(x)$ because the second term seems to follow logically from accounting for the structure of the analyzer grating. After an embarrassingly long time it occurred to me to look at the residuals




Figure 4.8: Comparison of single-sine fit (green) and new two-sine fit (blue). The red curve is the correction to the one-sine fit to arrive at the two-sine fit.

from a single-sine fit of my data; those residuals in fact suggested that the appropriate correction term is

$$a_2 \cos(4\pi x/p_2 + \varphi) \tag{4.45}$$

rather than

$$a_2 \cos(\pi x/p_2 + \varphi + \pi/4). \tag{4.46}$$

This suggests that the correction from the fact that moiré fringes are actually truncated triangle waves is more significant than the correction from grating structure.

Fig. 4.8 shows the one-sine fit using the model function $f_1(x)$ and a two-sine fit with the updated second term. It is clear that the two-sine fit is objectively better than the one-sine fit.




Figure 4.9: $p_{2,\text{rec}} - p_{2,\text{input}}$ for previous mock data (left) and updated mock data (right). Vertical scale: φ index; horizontal scale: p_2 index.

Mock data: double sinusoid redux

I constructed new mock data

$$y = 20000 + 10000 \cos(2\pi x/p_2 + \varphi) + 1000 \cos(4\pi x/p_2 + \varphi) \quad (4.47)$$

and a new model function

$$f_{2'}(x) = a_0 + a_1 \cos(2\pi x/p_2 + \varphi) + a_2 \cos(4\pi x/p_2 + \varphi). \quad (4.48)$$

As before, fitting the two-sinusoid mock data with the two-sinusoid model function $f_{2'}(x)$ did not reproduce the fringe artifacts.

A fit done with the single-sinusoid model function $f_1(x)$ did show less severe (but not no) artifacts in the recovered period, as shown in Fig. 4.9. Several points are worth mentioning:

- It is not clear why the character of the artifacts should have changed as it




Figure 4.10: Normalized variations in recovered parameter values for $N = 101$ sampling points over $x \in [2, 10]$.

seems to have done (compare the two panels of Fig. 4.9).

- In nearly all cases (old mock data, new mock data, all input periods, all input phases), the recovered period is smaller than the input period. Looking back at Fig. 4.7, we see that the same is true for φ , but not for a_0 and a_1 .
- Looking down the columns of Fig. 4.9, it is apparent that for some input periods, the amount of variation in $p_{2,\text{rec}}$ is smaller, taken over all input periods $\varphi \in [0, 2\pi]$, than for others.

Exploring the amount of variation in recovered parameters

In Figs. 4.7 and 4.9 we see that the recovered parameter values oscillate with the input phase φ_{input} of the data. We wanted to know under what circumstances this oscillation was minimized. For each column in the datasets exemplified by Figs. 4.7 and 4.9, we took the max - min of the recovered values and normalized to the input values. This is shown in Fig. 4.10 for $N = 101$ sampled points. Since we used $N = 101$ points to generate Fig. 4.10, the numerical values should not

be taken to be representative of our usual $N = 19$ datasets. Instead we will point out some qualitative features:

- There is always at least *some* variation in the recovered parameter values depending on the input φ , regardless of the input period p_2 .
- The period for which one parameter is recovered most consistently (although not necessarily correctly) is in general not the same as the period for which other parameters are recovered most consistently.

I had two followup questions:

- How are these inconsistent recovered values affected by the number N of sampling points?
- How well can we recover the parameters if we constrain the fit period to the correct value?

I repeated the fitting test on the same mocked-up data as before, using the single-sinusoid fitting function $f_1(x)$ on a new-style two-sinusoid dataset generated from the model function $f_2(x)$, but constraining the period p_2 to the correct value rather than allowing it to be a fitted parameter. Figs. 4.11 and 4.12 show the results of this analysis for a_0 and a_1 .

Note that in this case there *does* in fact seem to be a period for which the input a_0 and a_1 are recovered consistently, *but* it is dependent on N . Table 4.1 shows this relationship. Essentially, we only recover the parameters correctly when $2p_2$ would fall into the sequence of sampled points eventually; that is, only if we *could* have sampled the true period perfectly if we had continued sampling.




Figure 4.11: Variations in recovered a_0 values when fit period is constrained to true period, for several values of N .




Figure 4.12: Variations in recovered a_1 values when fit period is constrained to true period, for several values of N .

N	Microns per step	Period p_2 with min error	$2p_2$
11	0.8	4.4	8.8
19	0.444	4.222	8.444
41	0.2	4.1	8.2
101	0.08	4.04	8.08

Table 4.1: Correlation between N and most reliable p_2 .

Finding a name for the problem

As the preceding analysis shows, a significant part of the least-squares fitting problem is that it is difficult to determine the true period of the data. Searching through the `comp.dsp` newsgroup at length, we discovered several things:

- Least-squares fitting is known to behave badly for sinusoidal model functions.
- More general terms for this problem include “spectral estimation” (preferred by the signal processing community) and “harmonic inversion” (preferred by the NMR spectroscopy community)
- This is generally acknowledged to be a hard problem, which doesn’t actually help one make progress but does make one feel better about struggling with it.

4.2.3 Bayesian methods: generalities

The `comp.dsp` newsgroup pointed me in the direction of Bayesian likelihood estimation, which is purported to be more reliable than least-squares fitting in determining the period correctly. The general idea is as follows, based on the arguments in Chapter 3 of [7].

Construct a model function

$$f(x) = \sum_i B_i G_i(x, \{\omega\}), \quad (4.49)$$

where B_i are prefactors and G_i are functions depending on x and on a set of frequencies $\{\omega\}$. Denote a dataset $D = d_1, d_2, \dots, d_i, \dots, d_N$ sampled from a process

$y(x)$ at points x_1, \dots, x_N , not necessarily uniformly spaced, and put

$$d_i = y(x_i) = f(x_i) + e_i, \quad (4.50)$$

where e_i is noise with (not necessarily known) rms level σ . We can then write down a *likelihood estimator* L for a given choice of prefactors $\{B\}$, frequencies $\{\omega\}$, and noise level σ , based on the noise-normalized least-squares distance of the data from the model function:

$$L(\{B\}, \{\omega\}, \sigma) = \sigma^{-N} \times \exp \left[-\frac{1}{2\sigma^2} \sum_{i=1}^N (d_i - f(x_i))^2 \right]. \quad (4.51)$$

This estimator is largest when $(d_i - f(x_i))^2$ is minimized, under the Bayesian idea that the most likely collection of parameters $\{B\}$ and $\{\omega\}$ is the one for which the greatest part of the data can be attributed to signal rather than to noise (conveniently giving one plenty of rope with which to hang oneself if one's choice of model f is poor).

Rather than trying to minimize $(d_i - f(x_i))^2$ directly as in least-squares fitting, we first modify the likelihood estimator to remove irrelevant information. By converting the general model functions $\{G\}$ and their amplitudes $\{B\}$ into a new set of *orthonormal* model functions $\{H\}$ with new amplitudes $\{A\}$, [7] shows that we can rewrite L as

$$L(\{A\}, \{\omega\}, \sigma) \propto \sigma^{-N} \times \exp \left[-\frac{N}{2\sigma^2} \left[\overline{d^2} - \frac{2}{N} \sum_{j=1}^m A_j h_j + \frac{1}{N} \sum_{j=1}^m A_j^2 \right] \right], \quad (4.52)$$

where $h_j = \sum_{i=1}^N d_i H_j(x_i)$ is the projection of the data onto the j th orthonormal model function and $m = |\{H\}| = |\{A\}|$.

With the likelihood estimator written this way, it is possible to integrate away the parameters $\{A\}$ and the rms noise level σ , leaving only the frequencies $\{\omega\}$. But by doing this integration we are potentially removing information,

so we want to do it in the most naive possible way, without making any restrictive assumptions about the amplitudes or the noise. Therefore we assume a uniform prior probability for the amplitudes $\{A\}$ (that is, that all values of $\{A\}$ are equally likely), and we assume the *Jeffreys prior* $1/\sigma$ for the noise level (this is a scale-invariant prior that makes minimal assumptions about σ). Once we perform the integrations, we are left with the joint posterior probability for the frequencies $\{\omega\}$ given our data D and the priors I . For the naive priors (uniform, Jeffreys), we get the *Student t-distribution*

$$P(\{\omega\}|D, I) \propto \left[1 - \frac{m\overline{h^2}}{Nd^2}\right]^{\frac{m-N}{m}}. \quad (4.53)$$

4.2.4 Bayesian methods: specifics

The favorite basic model for grating interferometry is

$$f(x) = a_0 + a_1 \cos(2\pi x/p_2 + \varphi). \quad (4.54)$$

For a given dataset D , construct a new dataset $D' = D - \overline{D}$, so that the data have mean 0. Then note that

$$\cos(\omega x + \varphi) = \cos \varphi \cos(\omega x) - \sin \varphi \sin(\omega x), \quad (4.55)$$

so we can write our mean-0 model as

$$f(x) = B_1 \cos(\omega x) + B_2 \sin(\omega x). \quad (4.56)$$

This exact problem is covered in Sec. 6.1 of [7]. The orthonormal model functions are

$$H_1(x, \omega) = \frac{\cos(\omega x)}{\sqrt{\frac{N}{2} + \frac{\sin(N\omega)}{2\sin\omega}}} \equiv \frac{\cos(\omega x)}{\sqrt{c}} \quad (4.57)$$

and

$$H_2(x, \omega) = \frac{\sin(\omega x)}{\sqrt{\frac{N}{2} - \frac{\sin(N\omega)}{2 \sin \omega}}} \equiv \frac{\sin(\omega x)}{\sqrt{s}}. \quad (4.58)$$

Note that the projections $\sqrt{c}h_1 = \sum_i d_i \cos(\omega x_i)$ and $\sqrt{s}h_2 = \sum_i d_i \sin(\omega x_i)$ are the real and imaginary parts $R(\omega)$ and $I(\omega)$ of the DFT. If we substitute these projections into the Student t-distribution, we get

$$P(\omega|D', I) \propto \left[1 - \frac{R^2(\omega)/c + I^2(\omega)/s}{Nd^2} \right]^{\frac{2-N}{2}}. \quad (4.59)$$

There are two important points to make about this result:

- These particular orthonormal basis functions are constructed assuming that the sampling points are uniformly spaced (generally true for our datasets).
- This expression is valid for all N , even if N is small. The same is not necessarily true of the FFT, as we can observe from the spectral leakage problem.

Furthermore, to take advantage of the fact that we have many channels of data (1024×1024 for the Finger Lakes camera, 2592×2160 for the Fairchild camera), all of which should have the same frequency, we can construct a multichannel joint posterior probability by taking the product of the single-channel probabilities, to get

$$P(\omega|\{D_k\}, I) \propto \prod_k \left[1 - \frac{R_k^2(\omega)/c + I_k^2(\omega)/s}{Nd^2} \right]^{\frac{2-N}{2}} \quad (4.60)$$

Because the Bayesian likelihood function tends to be strongly peaked, and because the right-hand side is not normalized, this probability quickly becomes problematically large even on a 64-bit computer. It is useful to work instead with the log likelihood

$$\log P \propto \sum_k \frac{2-N}{2} \left[1 - \frac{R_k^2(\omega)/c + I_k^2(\omega)/s}{Nd^2} \right]. \quad (4.61)$$

4.2.5 Bayesian methods: results on mocked-up data

Construct sampling points $x = \text{linspace}(-4, 4, 19)$ and sample data

$$y = \cos(2\pi x/4) + 0.1 \cdot \text{standard_normal}(y.\text{shape}), \quad (4.62)$$

that is, a cosine with 10% noise added [46]. We can generate two likelihood estimators: $(R^2(\omega) + I^2(\omega))/(N/2)$ and $R^2(\omega)/c + I^2(\omega)/s$. The first is the *Schuster periodogram*, the continuous version of the information returned by the FFT. The second is the version of the Schuster periodogram which has been corrected for small N , called the *sufficient statistic*. These two estimators are normalized to $N\bar{d}^2$ and plotted together in Fig. 4.13.

Note that somehow both of these estimators are larger than 1 at the peak. This is not supposed to be possible according to Bessel's inequality, because the model functions are orthonormal, but we never tracked down the problem. When these normalized estimators exceed 1, the Student t-distribution becomes singular, so it is not possible to correctly compute the log likelihood. However, we will point out that the maximum of the Schuster periodogram is at $p_2 = 3.90$ and the maximum of the second estimator is at $p_2 = 3.96$, so the correction for small numbers of data points N is significant, although neither one correctly returned the true period of $p_2 = 4$.

At this point we admitted defeat at the hands of attempts to recover the correct period, and moved on to other techniques.




Figure 4.13: Two likelihood estimators (the sufficient statistic and the Schuster periodogram), with the correct frequency shown by the vertical dashed bar. The value 1 is shown by the horizontal dash-dotted line to illustrate the puzzling behavior by which the statistics become greater than 1. This is assumed to be due to a coding error.

4.2.6 Goertzel's algorithm

Goertzel's algorithm is a variant of the Discrete-Time Fourier Transform (DTFT) that picks out a particular frequency, even if that frequency would not ordinarily fall in the center of one of the DFT bins [36]. At first glance, this seems as if it would be optimal in a situation where slightly incorrect sampling has moved the correct frequency out of the center of a DFT bin. Unfortunately, Goertzel's algorithm is still a DFT at heart, and therefore suffers from the same spectral leakage problems that are seen in the FFT: although it may find the true peak, that peak will not have the correct value because some data has leaked into

neighboring frequencies.

4.2.7 Next steps: PCA and GPSAs

Eventually we broke into the phase shifting interferometry (PSI) literature and found references to PCA and GPSAs. We have not had sufficient time to explore these techniques in detail, but we have basic implementations of both that we have used to process datasets.

GPSAs are extremely fast, but we have not managed to tweak their detuning insensitivity to be as strong as it is supposedly capable of being; see Sec. 4.3 for the somewhat disappointing results. Tests on mocked-up data have been extremely precise in recovering correct phases even in the presence of detuning, so this is probably a worthwhile area for further investigation.

PCA generates datasets that are relatively clean, and is also relatively fast (a few seconds for a 1 megapixel dataset with an Intel i7 2.2 GHz processor and 8 GB of RAM). It is not clear how well it will function in the presence of significant noise, since the PCA projection onto $\cos \varphi$ and $\sin \varphi$ will not accurately capture all the data if the points lie significantly out of the $\cos \varphi - \sin \varphi$ plane due to noise.

4.2.8 Gaussian filtering


Gaussian blurring is the convolution of an image with a Gaussian kernel which has characteristic width larger than a pixel. This functions as a low-pass filter, where the filtering is stronger the wider the kernel. Since the fringe artifacts

tend to be lower-frequency than the features of interest, we can use this low-pass Gaussian filter to clean up the processed data: we blur the processed image to recover a new image which ideally contains only the fringe artifacts (but in practice contains other very low-frequency data), then subtract this blurred image from the original image to recover a dataset with less severe fringe artifacts. This process is illustrated in Fig. 4.14, where we used it not on fringe artifacts but on the irregular beam profile in an absorption scan.


Gaussian filtering is not suitable for all datasets. If the fringe artifacts are of a characteristic size similar in scale to important features in the data, the Gaussian filter can subtract unintended pieces of the dataset, as shown in Fig. 4.15.

4.2.9 Fourier filtering


Since the fringes look like they should be strongly peaked in Fourier space, one attractive option is to Fourier transform a processed image, wipe out the harmonic peak, and then inverse transform to recover the image, hopefully now fringe-free. In some cases this works very well. For instance, in Fig. 4.16 we see a scale insect whose features are obscured by fringes. If we take the Fourier transform of this image, as shown in Fig. 4.17, and wipe out the first harmonic peaks to obtain the new Fourier transform shown in Fig. 4.18, we can see that the recovered image, Fig. 4.19, shows little trace of the original fringe artifacts.



(a) The original image.




(b) The Gaussian blurred image of the irregular beam profile only.




(c) The filtered image.


Figure 4.14: Reducing artifacts using Gaussian filtering.



(a) The original image.



(b) The Gaussian blurred image of fringe artifacts, which includes specimen features.



(c) The filtered image, which has lost some information in the process of subtracting the Gaussian blurred image.

Figure 4.15: Gaussian filtering applied to an unsuitable dataset.




Figure 4.16: Scale insect (*Hemiptera* sp.) in amber. Absorption contrast image exhibiting fringe artifacts.




Figure 4.17: Fourier transform of Fig. 4.16, illustrating the strong first harmonic peaks due to the fringe artifacts.




Figure 4.18: Fourier transform of Fig. 4.16 with the first harmonic peaks wiped out.




Figure 4.19: Inverse Fourier transform of Fig. 4.18, showing much reduced fringe artifacts.

4.3 Methods comparison

Fig. 4.20 shows a composite comparison dataset with pieces processed by FFT, LMA, PCA, and GPSA.

FFT-processed and LMA-processed data show very similar fringe artifacts, underscoring the undesirability of LMA processing. The PCA data is the cleanest, presumably because it is the most tolerant of phase steps that do not exactly match the desired grating stepping positions (detuning). The GPSA result is disappointing; given the insensitivity to detuning that is built into the algorithm, one would hope that the result would be significantly cleaner, but instead it is the worst of the lot. The cause of this is unclear and is a target for further study.




Figure 4.20: Composite comparison dataset showing (a) FFT processing, (b) LMA curve-fitting, (c) PCA, and (d) GPSA.

CHAPTER 5

FAIRCHILD SCMOS DETECTOR

For early studies with the CHESS Talbot interferometer, we used the lens-coupled Finger Lakes Imager (FLI) as the detection system [53]. The lens coupling makes the FLI detector very flexible in terms of magnification and field of view, and the detector has a very low dark current ($0.30 e^-/\text{pixel}/\text{sec}$ when the detector is operated at -35°C), which is desirable in phase-contrast imaging, which depends on good signal-to-noise ratio for its sensitivity. But the best resolution we ever managed to obtain with the FLI camera was 17 microns, and the detector read-out takes nearly 30 seconds per frame, making the FLI unsuitable for many types of experiments.

Seeking a detector with faster framing and better resolution, we turned to a prototype CMOS chip from Fairchild Imaging. Historically, CMOS sensors have not been able to operate at the very low noise levels that are possible with CCD sensors [17], but the Fairchild CMOS prototype is comparatively low noise. The on-chip pixel size of the sensor is 6.5 microns (compared to the 25.4 micron pixel size of the FLI), and it has the additional advantages of a 100% pixel fill factor and electronic shuttering, as opposed to the mechanical shuttering required for CCD systems. This allows for fast, high-efficiency readout even in full-frame mode.

For detectors using a settled-phosphor scintillation screen, the detector resolution is typically limited by the thickness of the scintillator, since optical photons tend to random-walk isotropically through the phosphor grains before they escape toward the rest of the detection chain [17]. At high energies, very thin high-resolution phosphors have unacceptably low stopping power.

In single-crystal scintillating films, photons tend to travel ballistically through the film, so the resolution is determined by the acceptance cone of the coupling optics. With appropriate microscope objectives, micron resolution can be obtained with film thicknesses on the order of 10 microns [14]. This still, however, provides very limited stopping power for very hard (> 30 keV) x-rays.

Although the Talbot interferometer described in this thesis was designed for use at 15 keV, we wanted the new detector to be effective at higher energies as well. To bypass the problem of low stopping power in high-spatial-resolution phosphors, we opted not to use a traditional phosphor at all. Instead, we used a Terbium-doped fiber optic plate as the scintillator, which effectively decouples the resolution from the detection efficiency. The resolution is set by the pitch of the fiber optic, while the stopping power can be quite high for a sufficiently thick fiber optic.

5.1 Design

5.1.1 CMOS sensor

The sensor is a prototype TCAM model CMOS chip from Fairchild Imaging (Milpitas, CA; now owned by BAE Systems, Inc.), which has been described previously [44]. * The chip has 2160×2560 pixels in two halves of 1080×2560 pixels each, with an on-chip pixel pitch of 6.5 microns. In full frame, the frame rate ranges from 0.2 Hz to 100 Hz; in region-of-interest (ROI) mode the frame rate can be as high as 1760 Hz.

*The CMOS chips and camera electronics were a generous gift of Fairchild Imaging.

The framerates are set by the pixel clock rate, which can be set to certain values between 100 MHz and 300 MHz. The framing time is the amount of time required to clock out 2160×2560 pixels (or an ROI) at the chosen pixel clock rate, plus a certain amount of overhead. To frame with exposure times slower than 30 Hz, “dummy” lines are clocked out until the desired exposure time is obtained, at which point the actual recorded signal is clocked out of the chip. The 0.2 Hz lower limit is a hardware limitation of this prototype; there is a hardware register responsible for storing the number of lines to clock out, and its maximum stored value limits the exposure time to just over 6 seconds per frame, which we approximate by 0.2 Hz to be safe. Commercial cameras employing the sCMOS chip (the successor to the TCAM model) do not have this limitation.

5.1.2 Coupling

The chip was prepared by cleaning dust particles off with a single-hair brush under a laminar flow hood. Once the chip was free of contaminants, the chip was bonded with epoxy (TRABOND F114, Emerson & Cuming; Irvine, CA) to a fiber optic taper (formulation BLE-359-6, Incom Inc.; Charlton, MA)[†] using a home-built bonding jig. The bonding was done with the assistance of Dr. Mark Tate. The taper provides a 1:2 magnification from the scintillator plane to the plane of the chip, with fiber pitch starting at 3 microns on the narrow end and growing to 6 microns on the larger end, which is bonded to the CMOS imager. This fiber optic taper was coupled by removable optical grease (#Q2-3067, Dow Corning; Midland, MI) to a scintillating fiber optic plate, described further in

[†]The fiber optic taper was a generous gift of Michael Detarando at Incom.

Section 5.2.

5.1.3 Mechanicals

The printed circuit board is too large to effectively enclose, but the chip itself is housed in a custom vacuum enclosure/cryostat (built by Marty Novak, the Gruner group's machinist). The cryostat contains a copper cold finger whose temperature is controlled by a Peltier thermoelectric cooler (CP 1.4-127-045L, Melcor Inc.; Trenton, NJ). The TEC maintains the cold finger at -20°C to reduce the otherwise considerable dark current of the sensor. (The cold finger is large enough that there is a thermal gradient across its thickness, so the chip is actually held at an unknown temperature somewhat warmer than the TEC.) The thermal contact between the cold finger and the chip is improved by a layer of silver-containing thermal paste. Heat from the cold finger is removed by a recirculating water chiller with lines running through the copper block.

Fig. 5.1 shows the detector with the clamshell vacuum housing removed. The chip and fiber optic taper are in the center of the circuit board. The cold finger's water lines and the vacuum hose port can be seen at the back of the clamshell, as well as the power and controller port for the TEC. At the front of the clamshell is the detector snout, a brass KF40 stub with a central bore to allow x-rays in and a light-tight plastic window to keep stray optical illumination out. Fig. 5.2 shows the detector "in action" with the clamshell in place.




Figure 5.1: Fairchild sCMOS detector with its housing removed. The 4 CameraLink cables that take data off the board are on the left (1). The chip, with bonded fiber optic taper (2), is in the center of the circuit board. The left half of the clamshell has the detector snout (3); the right half has the power connector for the TEC, the water lines for the cold finger, and the vacuum port (4).




Figure 5.2: Fairchild sCMOS detector in place on the beamline. 1: a dragonfly specimen. 2: the grating interferometer. At right: the detector body and circuit board.

5.2 Scintillating fiber optic plate

The scintillator is a Terbium-doped fiber optic faceplate, custom-made for this detector by Collimated Holes, Inc. (CHI, Campbell, CA) While the idea of using scintillating fiber optics for x-ray detection is not new [23, 38, 49], to the best of our knowledge this is the finest-pitch such scintillator to date.

The fibers consist of a heavy-element-doped core glass (the precise formula is proprietary) with a standard high-index cladding. As the fibers are bundled, each one is individually clad in an extramural absorber (EMA), a black glass that prevents optical photons from escaping into neighboring fibers if their trajectories exceed the critical angle for the fiber in which they are generated. This

“superclad” EMA is in contrast to the more standard “statistical” EMA, in which certain fibers are replaced with black fibers in a regular pattern. Compared to statistical EMA, superclad EMA decreases crosstalk between nearby fibers significantly and therefore tends to improve the point spread function of the detector.

This resolution improvement comes at some cost; CHI’s goal is a core:center-to-center ratio of 0.8; i.e. the fiber cores have 5 micron diameter and the center-to-center spacing is 6.2 microns. This means that the fill fraction of x-ray detecting core glass on the face of the plate is only ~65%, which limits the efficiency of the scintillator.

5.2.1 Stopping power

Although the core glass formula is proprietary, we do know that the effective atomic number of the glass is $Z = 31$ [8]. The size of the fiber optic block is $12\text{ mm} \times 12\text{ mm} \times 10\text{ mm}$, with the last number being the length of the fibers. Based on data for gallium (which also has $Z = 31$) from [34], the stopping power of the scintillator should be effectively 100% up to 60 keV and >99.5% at least up to 80 keV.

This statement comes with a caution: although the data predicts 100% stopping power, we observed radiation damage on the fiber optic taper behind the scintillator after heavy dose, indicating that a nontrivial quantity of radiation was transmitted through the scintillator. We will discuss radiation damage to the scintillator in Sec. 5.3.8.

5.2.2 Turn-on lag and afterglow

All scintillators have some inherent delay between x-ray illumination and the production of optical photons (“lag”) and also some inherent tendency to continue to fluoresce after the x-ray illumination is removed (“afterglow”). To ensure uniform imaging and to reduce “ghosting” of one frame into the next frame, it is desirable to minimize both lag and afterglow.

To measure the turn-on lag and afterglow of the scintillating face plates, scintillators were mechanically coupled to an RCA 6655 photomultiplier tube (PMT) which was then shrouded against stray optical illumination. The PMT output was passed to an oscilloscope and traces were recorded as the mechanical shutter of a Cu K_α rotating anode source was switched open and closed. Comparison traces for a single-crystal Europium-doped gadolinium gallium garnet (GGG:Eu) were also recorded.

As shown in Fig. 5.3, the GGG:Eu curves show classic lag behavior, while the scintillating fiber optics closely follows the shutter open and close behavior. We also recorded the x-ray fluence during shutter transit using a beamstop pin diode; as expected, this signal varies linearly as the shutter blade cuts across the x-ray beam. When all the curves are normalized to a minimum of 0 and a maximum of 1, the fiber optic’s curve tracks the shutter blade’s curve, showing that the lag and afterglow behavior are fast enough to be dominated by shutter transit.




Figure 5.3: Lag and afterglow of the scintillating fiber optic plate relative to a GGG:Eu screen and to the mechanical transit of the rotating anode shutter. (CHI SFOP: Collimated Holes, Inc. Scintillating Fiber Optic Plate)

5.3 Characterization

5.3.1 Conversion gain

The *conversion gain* of the CMOS sensor is the number of digital counts recorded per electron produced in the sensor. Since the detector outputs its data in the form of digital counts, it is necessary to know the conversion gain in order to reconstruct the original electronic signal.

The TCAM chip has two gain modes, each with its own conversion gain:

high gain, when there is minimal signal, and low gain, when signal is significant. There is a user-settable threshold dividing the two gain modes; by default this threshold is 2000 digital counts. The output signal from the detector is 12-bit: 11 data bits and one gain bit which is 0 in high-gain mode and 1 in low-gain mode. By knowing the ratio of the conversion gains, one can reconstruct the full dynamic range of the instrument from this 12-bit data.

To measure the conversion gain, we use the photon transfer curve method described in [25]: two images are collected and one is subtracted from the other to remove pixel-by-pixel variations, pedestal, etc. leaving only the variance of the measurement. This variance, when divided by the signal in the sum of the two images, yields the conversion gain.

To perform this measurement, we need a signal that is uncorrelated from pixel to pixel; otherwise the variance will seem to be reduced because of the coupling between neighboring pixels. For the high-gain measurement, we used the dark current of the detector as the signal; for the low-gain measurement, we removed the scintillator and recorded the signal from a diffuse source of optical photons (overhead fluorescent lights with paper in front of the detector as a diffuser). Since fluorescent lights have a 120 Hz flicker, we were careful to keep the exposure times long enough to avoid sampling the flicker differently from frame to frame. Based on these measurements and the photon transfer curve method, we compute a high-gain conversion gain of 1.8 digital numbers per electron (DN/e^-) and a low-gain conversion gain of $0.049 \text{ DN}/e^-$.

5.3.2 Dark current, read noise, and dynamic range

The *dark current* is the signal that is recorded by the detector due to thermal generation of electron-hole pairs in the sensor, even in the absence of an actual optical signal. The Fairchild CMOS is nominally relatively low-noise as CMOS sensors go; nevertheless, its dark current is considerable. We measured the dark current with the cold finger held at a range of temperatures from 10 °C to -20 °C; the results are shown in Fig. 5.5. At -20 °C the dark current is 40 e^-/s . In Si depletion layers, the dark current should be halved for each 7° degree decrease in chip temperature.

5.3.3 Dark current temperature dependence

We can think of the thermal generation of electron-hole pairs in silicon as an Arrhenius process: electrons in the valence band must, through thermal fluctuation, acquire enough energy to jump the band gap into the conduction band. The law of mass action states that

$$n_c \cdot p_v = N_c \cdot P_v \cdot e^{-E_g/kT}, \quad (5.1)$$

where n_c and p_v are the electron and hole densities in the conduction and valence bands respectively, N_c and P_v are the effective densities of states for the conduction and valence bands, and E_g is the band gap.

We expect in this case that the number of electrons should be equal to the number of holes, so we can modify the law of mass action to instead read

$$n_i = \sqrt{N_c P_v} e^{-E_g/2kT}, \quad (5.2)$$

where n is the intrinsic concentration of either species. The densities of states can be written

$$N_c = 2.5 \left(\frac{m_c}{m_e} \right)^{3/2} \left(\frac{T}{300K} \right)^{3/2} \times 10^{19} / \text{cm}^3, \quad (5.3)$$

and likewise for P_v , substituting the effective mass m_v of the hole in place of the effective mass m_c in the conduction band. So the density n_i can be written

$$n_i = 2.5 \left(\frac{m_c}{m_e} \right)^{3/4} \left(\frac{m_v}{m_e} \right)^{3/4} \left(\frac{T}{300K} \right)^{3/2} e^{-E_g/2kT} \times 10^{19} / \text{cm}^3. \quad (5.4)$$

We now summarize the argument in [62]. In a CCD at or below room temperature, the dominant source of dark current is assumed to be the behavior of the carriers that remain in the depletion layer. We assume that the same is true of a CMOS sensor.

We can still think of the dark signal per pixel per second to be

$$D_{\text{dep}} = \frac{V_{\text{dep}} n_i}{2\tau}, \quad (5.5)$$

where V_{dep} is the volume of the depletion region per pixel and τ is the recombination time.

The obvious physical interpretation of this equation is that $V_{\text{dep}} \cdot n_i$ carriers are generated per unit time τ and these carriers get swept into the collection wells. The factor of 2, however, suggests that in fact the true interpretation is more subtle. In fact the density of carriers in the depletion region is much less than n_i , and the factor of 2 has no immediately obvious physical interpretation without more detailed derivation, for which see [62].

Of course, D_{dep} has a temperature dependence. We wish to show that over operationally relevant temperature ranges, the dark current typically halves




Figure 5.4: Theoretical $\ln n_i(T)$ for temperatures between 200 K and 300 K.

with every 7° decrease in temperature. The temperature dependence is primarily in $n_i(T)$, so we take $\ln n_i(T)$ and examine the slope, which should be $\sim \ln(2)/7$ if this rule of thumb holds. In doing so, we can ignore the plethora of constants in front and just take

$$\ln n_i(T) - \ln c = \frac{3}{2} \ln\left(\frac{T}{300 \text{ K}}\right) - \frac{1.12 \text{ eV}}{2 \cdot 8.5 \times 10^{-5} \text{ eV/K} \cdot T} \quad (5.6)$$

(1.12 eV is the approximate band gap of silicon.) In Fig. 5.4 we can see this function plotted for temperatures between 200 K and 300 K. Typical detector operating ranges are 30°C to -50°C (303 K to 223 K). We can see that over this range, $\ln n_i$ is indeed approximately linear, with a slope of $(-22 + 33)/100 = 0.11$. In fact $\ln(2)/7 = 0.099$, so over this temperature range it is in fact the case that the dark current halves approximately every 7 degrees.

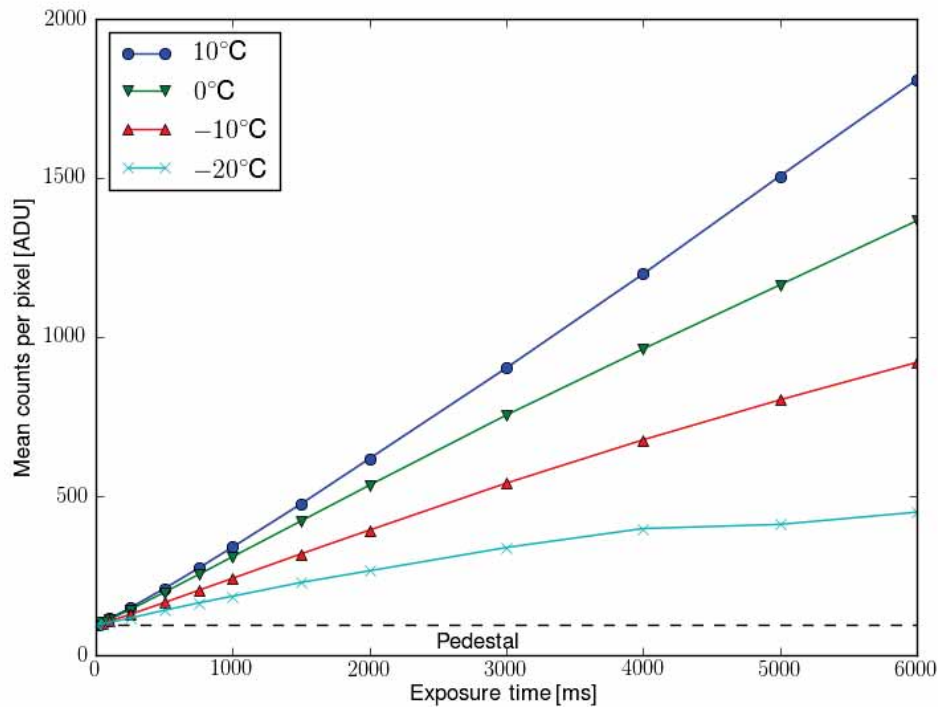


Figure 5.5: Dark current of the Fairchild sCMOS detector

Note that the curves in Fig. 5.5 do not obey this rule of thumb; to see this, examine the data points at 6000 ms exposure time. At 10 °C, we recorded ~1800 counts per pixel. At 0 °C, we recorded ~1400 counts per pixel, which is certainly not less than half of the 10 °C value. Likewise at -10 °C we recorded ~900 counts per pixel, which is not less than half of the 0 °C value.

This deviation presumably due to a significant thermal gradient across the cold finger, so that the chip is not being held at the same temperature as the TEC whose temperature we are setting. The difference between the 10 °C and the 0 °C measurements corresponds to a true temperature difference of ~4.6 °C; the difference between the 0 °C and -10 °C measurements corresponds to ~5.6 °C; and the difference between the -10 °C and -20 °C measurements corresponds to ~7 °C.

The *read noise* or *zero-dose noise* of the detector is the projection of the dark signal back to 0 integration time, i.e. the noise that will always be recorded by the detector no matter how short the exposure. Based on the data in Fig. 5.5, we compute the read noise to be 4.5 digital counts, or $2.5 e^-$. The read noise is distinct from the *pedestal*, a constant value that is intentionally added to all pixels to prevent sending a negative value to the A-D converter [5]. For the Fairchild system, the pedestal is 92 digital counts.

The *dynamic range* of the detector is the ratio of the full-well depth to the read noise, i.e. the maximum recordable value divided by the minimum recordable value. Since the well depth of the sCMOS sensor is $30 ke^-$ (nominal), the dynamic range is

$$\frac{30000}{2.5} = 12000 : 1 \quad (5.7)$$

5.3.4 Linearity

Ideally, all detector response would be perfectly linear, so that for a constant input flux, the recorded counts would be exactly proportional to the exposure time. It is normal for detectors to have some degree of nonlinearity, but hopefully the nonlinearity is quite small. For the Fairchild, the response is something of a mixed bag.

As we can see from Fig. 5.5, the dark current linearity is quite good, with the exception of a slight blip at the higher end of the -20°C line. (We speculate that this is due to a change in temperature of the chip over the course of the integrations.) The x-ray response is another story. To measure the x-ray response linearity, we set up a beam spot of known size ($2\text{ mm} \times 2\text{ mm}$), took exposures




Figure 5.6: Counts per pixel in the beam spot, normalized by the incident flux, at 10 keV. The green line is a linear fit to the data.

at a range of energies, then integrated the detected beam spot. The integrated beam spot was normalized to adjust for differences in the incident flux, then plotted and fit to a straight line. (The flux was computed using recorded ion chamber counts and the CHESS online flux calculator [42], then adjusted to account for absorption in the 1 m of air between the ion chamber and the detector.) A typical result (for 10 keV) is shown in Fig. 5.6. The rolloff at the higher end is partly due to saturation of the detector, but we can see that even if we cut off the last two or three points, there is a distinct change in the concavity of the data at around 1500ms exposure, so the linearity does not improve much when we cut off the saturated points.

Ideally we would make similar plots at a range of energies and fluxes and try

to understand the shape of the nonlinearity, so that we could set up a correction factor that could be applied to straighten out the data. This is an object for further study.

5.3.5 Sensitivity

The sensitivity of a detector is a measure of how many electrons are produced in the sensor by each x-ray photon incident on the scintillator. Since the number of optical photons produced by an x-ray conversion event is energy-dependent, the sensitivity will in general increase with energy, modulo any decrease in scintillator stopping power as the energy increases. Since the fiber optic scintillator is quite thick, we expect it to stop the vast majority of x-ray photons at the energies accessible at CHESS, and therefore we expect a roughly linear increase in sensitivity with energy.

To measure the sensitivity at low energies, we installed the detector at the CHESS F3 station. The station was tuned from 8 keV to 20 keV in 1 keV steps, and sensitivity measurements were taken at each energy. Using paired Huber beam-defining slits 1.3 m upstream of the detector, the beam was slitted down to a $2 \text{ mm} \times 2 \text{ mm}$ spot, much smaller than the detector field of view. This slitted beam was passed through a calibrated nitrogen-filled ion chamber to determine the flux in the entire beam at each energy. Since there was no helium-filled flight tube in the beam path, the flux was adjusted to account for absorption in approximately 1 m of air using transmission percentages from [10].

We varied the exposure time from 50 ms up to 5 s, taking sequences of 100 exposures at each exposure time. Because of the heating effect described in




Figure 5.7: Sensitivity of the Fairchild detector from 8 keV to 60 keV

Sec. 5.3.9, we used only the last 4 frames in each 100-image sequence for the sensitivity measurements, just to be sure that the detector had completely settled.

To measure the sensitivity at higher energies, we installed the detector at the CHESS A2 station. Again we slitted the beam down to a $2 \text{ mm} \times 2 \text{ mm}$ spot and measured the flux with an ion chamber immediately downstream of the slits. At this station we did have a He flight path in place, so we did not need to correct the flux for air absorption. At A2 we took measurements in 5 keV steps from 15 to 60 keV. Since the fluxes at A2 are so much higher than at F3, we did not need to take such long exposures. We varied the exposure time from 10 ms to 0.5 s, and even then we still saturated the detector at the longer exposure times.

In Fig. 5.7 we see the detector sensitivity from 8 keV to 60 keV, from both the F3 and the A2 data. Note that the curve does not point linearly to 0 sensitivity at 0 energy, as it should; we suspect that there is a dead layer in the scintillator causing a significant falloff in sensitivity at low energies.

For comparison, we predict the sensitivity of the Fairchild system using an 11.7 micron GGG:Eu film in place of the scintillating fiber optic would be so small that it would appear pegged at 0 if we were to plot it on the same graph.

The take-away message here is that although the system's sensitivity is not particularly good, it is significantly better than a comparable system with a standard thin scintillator thanks to the fiber optic's drastically improved stopping power. It might also be possible to improve the sensitivity of the fiber optic system by approximately a factor of two by tightly coupling a sheet of aluminized mylar to the front of the scintillator or evaporating a thin film of aluminum onto it, to reflect optical photons that are produced traveling away from the chip.

5.3.6 Resolution

We measured the resolution of the detector by the Modulation Transfer Function (MTF) method. We placed a tantalum knife edge on the front of the detector snout, angled at 3.3° with respect to the vertical pixel axis. The recommended method for performing the MTF measurement is described in ISO 12233 [9, 24], as follows:

1. Collect an image of the slanted knife edge.
2. For each line in the image, find the location of the edge.

3. Fit these locations to a straight line $f(i)$.
4. Bin the lines according to the fractional part of the edge location in each line, so that if the fitted line satisfies $f(i) - \text{int}(f(i)) < 0.25$ in row i the line is assigned to bin 1, $0.25 < f(i) - \text{int}(f(i)) < 0.5$ to bin 2, etc.
5. Combine these bins to build up a single line that forms a super-resolved image of the knife edge.
6. Differentiate the super-resolved knife edge to find the line spread function (LSF).
7. Take the Fourier transform of the LSF; its modulus is the MTF.

However, the knife edge we used was not sufficiently straight for this method to be effective; the straight line fit to the edge locations is extremely poor, causing lines to be assigned to incorrect bins and broadening the LSF unrealistically. Instead, we chose to use a simpler method, taking seven consecutive lines and summing them to reduce noise, then taking the Fourier transform of this sum. The result is shown in Fig. 5.8.

We expect the resolution of the system to be no better than the pitch of the fibers in the scintillator, 6.2 microns. In fact, because the fibers are not in perfect registry with the pixels on the chip, we expect the resolution to be somewhat worse than the optimal 6.2 micron value.

The MTF generally provides more information than a single resolution number does, but to extract a single number from the MTF, the resolution is sometimes defined as the point at which the MTF falls to some percentage of its normalized peak value, frequently 50% or 10%. Based on the straight-line fit in




Figure 5.8: Modulation Transfer Function of the Fairchild sCMOS detector. A straight line fit is included as a guide to the eye.

Fig. 5.8, the MTF50 is 27 lp/mm, or 37 microns, while the MTF10 is 57 lp/mm, or 17 microns.

Both of these values are more pessimistic than the resolution determined by taking the FWHM of the LSF, which, for the same combination of 7 consecutive lines in the knife-edge image, is calculated to be 8.2 microns, or 122 lp/mm, as shown in Fig. 5.9.

5.3.7 Detective quantum efficiency

The *detective quantum efficiency* (DQE) is a measure of how much noise is added by the detector to the noise already present in the input signal. Formally, we




Figure 5.9: Gaussian fit to the LSF, giving a FWHM of 8.2 microns

write

$$\text{DQE} = \frac{(S/N)_{out}^2}{(S/N)_{in}^2} \quad (5.8)$$

In an ideal, noiseless detector, the DQE is unity, i.e. there is no degradation of the signal-to-noise ratio as the input signal is passed through the detection chain. In practice there is degradation along every step of the chain.

The input signal-to-noise ratio $(S/N)_{in}$ is derived from the incoming x-ray photons. We assume the arrival of x-rays on the detector is a Poisson process, and hence that the $(S/N)_{in} = \sqrt{N_{ph}}$. We used the same ion chamber flux measurements as in the sensitivity data to compute N_{ph} .

To compute $(S/N)_{out}$, we need a collection of uncorrelated beam spots. We then take two sequential images of these uncorrelated spots. We compute the

sum and difference of the two images for each of the spots; the mean of the spot sums gives the output signal, while the standard deviation of the spot differences gives the output noise [53]. One method of acquiring uncorrelated spots is to use a mask with an array of well-separated pinholes; we did not have immediate access to such a mask, so we instead used the $2\text{mm} \times 2\text{mm}$ beamspots from the sensitivity measurements and considered 3×3 pixel subregions that were well-separated relative to the length of the tail of the point spread function, which should minimize correlations.

The DQE is generally dose- and energy-dependent, as we can see in Fig. 5.10, which shows DQE curves for 8, 20, 45, and 60 keV. We expect the DQE to be an inverted parabola, as shown in the DQE curves for the Finger Lakes detector, shown in Fig. 5.11. The DQE curves for the Fairchild frequently cut off before the right arm of the parabola can appear; this is because the detector is saturated at higher doses. It is clear from Fig. 5.10 that a better determination of the DQE is warranted to fully characterize the detector. However, for the purposes of this work, the results show that the detector has roughly unity DQE for hard x-rays.

5.3.8 Radiation hardness

Radiation damage in glass often takes the form of *color centers*, which are formed when ionizing radiation frees an electron or hole which is then re-localized into a vacancy defect in the glass [3]. In glass, the vacancy is typically an anionic oxygen vacancy or a cationic dopant vacancy [65]. Thermal annealing can sometimes free the electrons or holes from their vacancy traps and restore the original transparency of the glass [16].




Figure 5.10: DQE curves for the Fairchild detector at various energies, illustrating dose and energy dependence. The cutoff of the curves at higher doses is the point of detector saturation.

While the scintillator is relatively radiation-hard, it does become browned after exposure above roughly 10^7 rad (100 kGy) of dose [8]. This browning is reversible by annealing the scintillator according to the following schedule [31]:

1. Beginning at ambient temperature, increase temperature 3° per minute until 585°C .
2. Hold at 585°C for one hour.
3. Decrease temperature at $0.3^\circ/\text{min}$ until 400°C .
4. Decrease temperature at $0.6^\circ/\text{min}$ until 350°C .
5. Decrease temperature at $3^\circ/\text{min}$ back to ambient temperature.




Figure 5.11: DQE curves for the Finger Lakes detector showing the characteristic inverted parabolic shape. (Image credit: [53])

We first noticed radiation browning after the long series of sensitivity tests performed at A2. The detector had previously been used at F3 nearly continuously for a week with no ill effects. Moreover, as a caution, we noted that after the sensitivity tests the fiber optic taper was also browned behind the scintillator, meaning that the stopping power of the scintillator is not 100% at least at 60 keV and possibly at lower energies as well, despite its high effective Z . Since the taper is bonded permanently to the chip, we currently know of no way to reverse the browning of the taper other than to replace the taper and chip outright.

The scintillator core fibers have been stabilized against radiation damage by

heavy metal dopants. The EMA, to the best of our knowledge, has no such stabilizers, nor does the fiber optic taper behind the scintillator. If the EMA is standard borosilicate glass, its transmission percentage ranges from effectively 0% at 10 keV up to 65% at 60 keV [35]. This means that at most energies there is significant transmission through the 35% of the scintillator that consists of EMA, which is certainly the cause of the radiation damage to the fiber optic taper. This limits the detector's usefulness at high energies, even though the scintillator itself can be cured by annealing. This also suggests that a future version of the detector might consist of lens-coupling to the scintillating fiber optic block, with a 45° mirror in the lens-to-scintillator path to avoid radiation damage to the CMOS chip.

5.3.9 Heating effects

During initial testing of the camera, we noticed that at sufficiently long exposure times (more than 50 ms), mean counts per pixel tended to rise over the course of a sequence capture. Moreover, over longer timescales (several minutes), the mean counts per sequence overall tended to increase. Together these occurrences pointed to the heating of some component of the detector system.

After testing several components on the printed circuit board by watching the mean counts per pixel as the components were subjected to a cold spray, we concluded that the likeliest culprit was the chip itself. The problem was being compounded by the way we were collecting images: taking finite-length sequential captures and letting the camera return to its idle state (and hence cool down) between sequences.




Figure 5.12: Mean counts per pixel per frame in a 50-frame image sequence before and after temperature equilibration with the camera framing in continuous mode. The lower curve shows the undesirable heating behavior; the upper curve is more consistent.

The solution to the problem was to rewrite the image capture code so that the camera now frames constantly, but only saves the resulting image buffers when commanded. The result is that after 10-15 minutes of equilibration, the mean counts per pixel per frame becomes consistent.

5.4 Detector usage

5.4.1 Flatfield corrections

Although the scintillating fiber optic has superclad EMA, the fiber optic taper between the scintillator and the chip has statistical EMA laid out on a hexagonal grid. Moreover, the scintillator has defects where the superclad EMA material has migrated through the face plate and diffused into the fibers, leaving black blotches in the field of view. The statistical EMA and the scintillator defects are quite obvious in acquired images, as illustrated in Fig. 5.13, but their appearance can be mitigated by appropriate flatfield correction.

The protocol is quite simple: occasionally during a run, one should collect an image with nothing in the beam except the detector. Then one can divide all other collected images by this image, and the resulting corrected image will have much less obvious EMA spots, as illustrated in Fig. 5.14. This should be done at least every time the detector is set up; since the scintillator is not permanently bonded to the taper, it may shift slightly during transport and mounting, invalidating the flatfield correction for scintillator defects. The EMA in the taper, of course, does not shift.

5.4.2 Imaging example

Since the field of view is relatively small, the Fairchild detector is best suited for relatively small-scale imaging experiments. For example, it has been used for examination of the amberized specimens described in further detail in Chap-




Figure 5.13: A (100×100) -pixel region of an image before flatfield correction, showing significant features. The regular hexagonal grid pattern is the statistical EMA in the fiber optic taper.

ter 6 and for imaging of a desiccated dragonfly in preparation for potential live-dragonfly imaging.

For comparison with the Finger Lakes detector, we use a desiccated yellow-jacket wasp, which has good phase contrast and interesting internal structure. In Fig. 5.15 we can see one image from each detector of the same region of the wasp (its thorax on the left and its abdomen on the right). In Fig. 5.16 are magnifications in which we can see the resolution improvement of the Fairchild over the Finger Lakes.




Figure 5.14: The same (100×100) -pixel region after flatfield correction. The statistical EMA pattern is absent.




Figure 5.15: Finger Lakes (top) and Fairchild (bottom) images of the midsection of a desiccated yellowjacket. Boxed regions show improved resolution of Fairchild compared to Finger Lakes; see Fig. 5.16.




Figure 5.16: Boxed regions from Fig. 5.15. The Fairchild image (bottom), though much noisier, shows significantly improved resolution over the Finger Lakes image (top).

CHAPTER 6
FOSSIL SPECIMENS

6.1 Imaging amber: sensitivity and minimum feature size

There are (at least) two ways to think about resolution, sensitivity, and minimum detectable feature thicknesses. First, we will derive a simple relation between the refractive index decrement δ of the specimen and the minimum thickness T of a detectable feature, which gives a satisfyingly intuitive $T_{min} \propto 1/\Delta\delta$ dependence. Second, we will give a more complex example which takes detector parameters and measurement noise into account. In that section we will also describe the relative sensitivity of absorption and phase contrast, given some information about the detector and the system.

We have worked with amber at CHESS, carbonized flowers at CHESS, and amber at the Advanced Photon Source. For each one, we will give an estimate of the sensitivity of the system given the source parameters and the nature of the specimens, and indicate whether phase or absorption contrast is generally superior.

6.1.1 Sensitivity limitations and minimum feature thickness: the simple way

In the following analysis, we assume a discontinuous step from δ_A (amber) to δ_B (bug) for an insect embedded in amber. We seek a relationship between the

decrement difference $\delta_B - \delta_A$ and the minimum feature thickness T (in the beam propagation direction) for which this decrement difference remains detectable.

Suppose we have a discontinuous step from δ_A to δ_B . When convolved with a projected source of (FWHM) size Σ_p , the discontinuity is smeared over a distance Σ_p as well, so that the convolved decrement profile is approximately

$$\delta(x) = \delta_A + \frac{\delta_B - \delta_A}{\Sigma_p} \cdot x \quad (6.1)$$

for $0 \leq x \leq \Sigma_p$. Then the phase acquired by the beam as it travels through the sample is

$$\Phi(x) = \frac{2\pi}{\lambda} \int_B \left(\delta_A + \frac{\delta_B - \delta_A}{\Sigma_p} \cdot x \right) dz + \frac{2\pi}{\lambda} \int_A \delta_A dz, \quad (6.2)$$

where \int_B and \int_A are integrals over the thickness of the bug and the surrounding amber, respectively. Assuming the amber has been polished so that $(\partial/\partial x) \int_A \delta_A dz = 0$, the differential phase detectable by the interferometer is

$$\frac{\partial\Phi}{\partial x} = \frac{2\pi}{\lambda} \int_B \frac{\delta_B - \delta_A}{\Sigma_p} dz = \frac{2\pi}{\lambda} \cdot \frac{\delta_B - \delta_A}{2\Sigma_p} \cdot T, \quad (6.3)$$

where T is the thickness of the bug in the direction of beam propagation. The actual phase shift recorded in the detector plane is

$$\varphi = \frac{\lambda d_m}{p_2} \frac{\partial\Phi}{\partial x} = 2\pi \cdot [(\delta_B - \delta_A)T] \cdot \left[\frac{mp_2}{2\lambda\Sigma_p} \right], \quad (6.4)$$

where the first bracketed quantity contains values we cannot tune and the second contains ones we can tune. (The expression is slightly misleading; it makes it appear that it is beneficial to reduce the x-ray wavelength, when in fact δ has a λ^2 dependence which makes reducing λ ultimately detrimental). So if our system has a minimal detectable phase difference $\Delta\varphi_{\min}$, we find that the minimum detectable decrement difference $\Delta\delta_{\min}$ is related to the minimum detectable feature thickness T_{\min} by

$$\Delta\delta_{\min} \propto \frac{1}{T_{\min}}. \quad (6.5)$$

6.1.2 Sensitivity, resolution, and noise: the more complex case

The preceding derivation describes the minimal detectable decrement difference $\Delta\delta_{\min}$ in relation to the minimal detectable feature thickness T_{\min} , given a certain minimal detectable phase shift $\Delta\varphi_{\min}$. We then need a way to determine $\Delta\varphi_{\min}$.

Intuitively we can see that $\Delta\varphi_{\min}$ should depend in part on the uncertainty in the measured datapoints relative to the amplitude of the fringes; the higher the relative uncertainty, the “sloppier” a fit can be and still remain inside the uncertainty range. In the following derivation, we will quantify this relationship, and also discover other parameters that relate to $\Delta\varphi_{\min}$. The mathematics in the derivation is due to [64]; the explanation is my own. In [64] the authors assume a cone-beam (expanding) geometry; we will simplify the derivation by assuming a parallel-beam geometry, which is approximately valid for synchrotron sources. This analysis will address only random sources of uncertainty/error; a truly thorough treatment would include the possibility of systematic errors.

Setup

Suppose that we are collecting K points per phase-stepping scan and that on the k th step, the second grating is displaced by an amount kp_2/K . Then the signal we measure in one detector pixel is

$$\begin{aligned} N_{\text{meas},k}(x, y, m) &= I_k(x, y, m) + N_{\text{err},k} \\ &= (\Delta S I_0 t / R^2) \eta_k + N_{\text{err},k}, \end{aligned} \tag{6.6}$$

where m is the Talbot imaging order; N_{err} is the error in the measurement due to Poisson statistics, read noise, dark current, and so forth; η_k is the photon transfer

efficiency of the interferometer system at the position kp_2/K ; and ΔS is the size of a detector pixel, I_0 is the intensity of the source per solid angle per time, R is the distance of the detector from the source, and t is the exposure time, so that $(\Delta S/R^2)I_0t$ is the total photon number incident on the detector pixel.

Photon transfer efficiency η_k

The photon transfer efficiency η_k is determined in part by the Fourier coefficients of the analyzer grating and of the intensity fringes in the ideal case of perfect coherence, and in part by the mutual coherence function of the beam, which accounts for degradation in the ideal fringes due to partial coherence. Finally, we must also account for the moiré pattern formed by overlapping the fringes with the analyzer grating and for its distortions due to introducing a specimen into the beam.

Setting aside the details of the Fourier coefficients for a moment, the moiré pattern is given by the Fourier sum

$$\frac{I_k(x, y, m)}{(\Delta S I_0 t / R^2)} = \sum_n q_n \exp \left[2\pi i n \left(\frac{k}{K} + \frac{mp_2}{8\lambda} \varphi(x, y) \right) \right]. \quad (6.7)$$

Suppose the ideal fringes have Fourier coefficients b_n and the analyzer grating's transmission function has Fourier coefficients c_n . Then because we treat the analyzer grating as optically thin, we can get the Fourier coefficient of the wave behind the analyzer grating simply by the multiplication $b_n c_n$. Then to account for the beam's partial (rather than perfect) coherence, we multiply by the mutual coherence function $\mu_n = \mu(nmp_1)$, where again m is the Talbot order and p_1 is the period of the interference fringes. For a Gaussian beam, the mutual coherence

function in the x direction (perpendicular to the grating lines) is

$$\mu(x) = \exp(-x^2/2l_c^2), \quad (6.8)$$

where in this case the coherence length l_c is defined in terms of the rms source size σ as

$$l_c = \lambda L/(2\pi\sigma). \quad (6.9)$$

So at last we have $q_n = \mu_n b_n c_n$ and therefore

$$\eta_k = \sum_n \mu_n b_n c_n \exp \left[2\pi i n \left(\frac{k}{K} + \frac{mp_2}{8\lambda} \varphi(x, y) \right) \right]. \quad (6.10)$$

Analyzing the error term in relation to the signal

First, assume that the Poisson statistics are the dominant source of error in the system. This is true for most detectors at reasonable photon numbers, even for the Fairchild whose dark current is considerable. In this case, the error due to the photon Poisson noise has mean zero and is independent from measurement point to measurement point; that is,

$$\langle N_{\text{err},k} N_{\text{err},k'} \rangle = 0 \quad (6.11)$$

unless $k = k'$, in which case it is $I_k(x, y, m)$.

First we wish to show that the random error term N_{err} in the recorded signal does not distort the mean calculated value of the phase profile $\Psi(x, y)$ in the detection plane. In the Fourier transform method of analysis, the recovered phase profile is the angle (or argument) of the first harmonic peak of the transformed data, which is, by definition,

$$\arg \left[\sum_{k=0}^{K-1} N_{\text{meas},k}(x, y, m) \exp \left(-2\pi i \frac{k}{K} \right) \right]. \quad (6.12)$$

By splitting $N_{\text{meas},k}$ into I_k and $N_{\text{err},k}$, we can rewrite eqn. 6.12 as

$$\arg[\mathcal{I} + \mathcal{N}_{\text{err}}], \quad (6.13)$$

where \mathcal{I} and \mathcal{N}_{err} (not to be confused with N_{err}) are the Fourier transforms of $\{I_k\}$ and N_{err} , respectively. We can then rewrite this as

$$\begin{aligned} \arg[\mathcal{I} + \mathcal{N}_{\text{err}}] &= \arg[\mathcal{I} \cdot (1 + \mathcal{N}_{\text{err}}/\mathcal{I})] \\ &= \arg[\mathcal{I}] + \arg[1 + \mathcal{N}_{\text{err}}/\mathcal{I}] \end{aligned} \quad (6.14)$$

By definition, $\Psi = \arg[\mathcal{I}]$, so the remaining term must be the error in Ψ due to photon statistics:

$$\Delta\Psi_{\text{err}} = \arg[1 + \mathcal{N}_{\text{err}}/\mathcal{I}]. \quad (6.15)$$

(The source text in [64] marks both the preceding equations as approximations, but we can see no reason why they ought not to be exact.) If the signal-to-noise ratio is sufficiently high, we have that $|\mathcal{N}_{\text{err}}| \ll |\mathcal{I}|$, and therefore

$$\begin{aligned} \Delta\Psi_{\text{err}} &= \arg[1 + \mathcal{N}_{\text{err}}/\mathcal{I}] \\ &= \arctan\left(\frac{\text{Im}(\mathcal{N}_{\text{err}}/\mathcal{I})}{1 + \text{Re}(\mathcal{N}_{\text{err}}/\mathcal{I})}\right) \\ &\approx \arctan\left(\frac{\text{Im}(\mathcal{N}_{\text{err}}/\mathcal{I})}{1}\right) \\ &\approx \text{Im}(\mathcal{N}_{\text{err}}/\mathcal{I}). \end{aligned} \quad (6.16)$$

Recall that we were trying to show that N_{err} did not cause a systematic distortion when trying to recover the phase profile Ψ . This is equivalent to showing that $\langle \Delta\Psi_{\text{err}} \rangle = 0$. If we think of $\mathcal{N}_{\text{err}}/\mathcal{I}$ as a complex number e^{ix} , then $\text{Im}(\mathcal{N}_{\text{err}}/\mathcal{I})$ is $(e^{ix} - e^{-ix})/2i$, so

$$\langle \Delta\Psi_{\text{err}} \rangle = \frac{1}{2i} \left(\frac{\langle \mathcal{N}_{\text{err}} \rangle}{\mathcal{I}} - \frac{\langle \mathcal{N}_{\text{err}}^* \rangle}{\mathcal{I}^*} \right). \quad (6.17)$$

(Note that we only need to take the ensemble average over N_{err} , since \mathcal{I} has no random behavior.) But $\langle N_{\text{err}} \rangle = 0$, so $\langle \mathcal{N}_{\text{err}} \rangle = 0$ as well. Therefore $\langle \Delta\Psi_{\text{err}} \rangle = 0$ and hence there are no systematic errors in Ψ due to the measurement error N_{err} .

Next we wish to determine the magnitude of the error $\Delta\Psi_{\text{err}}$, which will tell us the minimum detectable phase distortion. To do this, we compute the mean square $\langle\Delta\Psi_{\text{err}}^2\rangle$. We first note that

$$\langle\Delta\Psi_{\text{err}}^2\rangle = -\frac{1}{4}\left[\frac{\langle\mathcal{N}_{\text{err}}^2\rangle}{I^2} - 2\frac{\langle|\mathcal{N}_{\text{err}}|^2\rangle}{|I|^2} + \frac{\langle\mathcal{N}_{\text{err}}^{*2}\rangle}{I^{*2}}\right]. \quad (6.18)$$

Now we seek to calculate $\langle\mathcal{N}_{\text{err}}^2\rangle$ and $\langle|\mathcal{N}_{\text{err}}|^2\rangle$.

$$\begin{aligned} \langle|\mathcal{N}_{\text{err}}|^2\rangle &= \langle\mathcal{N}_{\text{err}}\mathcal{N}_{\text{err}}^*\rangle \\ &= \left\langle\sum_{k=0}^{K-1} N_{\text{err},k} \exp(-2\pi ik/K) \cdot \sum_{k'=0}^{K-1} N_{\text{err},k'} \exp(2\pi ik'/K)\right\rangle \\ &= \left\langle\sum_{k=0}^{K-1} \sum_{k'=0}^{K-1} N_{\text{err},k} N_{\text{err},k'} \exp\left(-2\pi i\frac{(k-k')}{K}\right)\right\rangle. \end{aligned} \quad (6.19)$$

Then because the ensemble average of a sum is the sum of the ensemble averages, and because the exponential is a constant factor (i.e. has no random behavior), we can write

$$\langle|\mathcal{N}_{\text{err}}|^2\rangle = \sum_{k=0}^{K-1} \sum_{k'=0}^{K-1} \langle N_{\text{err},k} N_{\text{err},k'} \rangle \exp\left(-2\pi i\frac{(k-k')}{K}\right). \quad (6.20)$$

But we argued earlier that $\langle N_{\text{err},k} N_{\text{err},k'} \rangle = 0$ unless $k = k'$, so

$$\langle|\mathcal{N}_{\text{err}}|^2\rangle = \sum_{k=0}^{K-1} \langle N_{\text{err},k}^2 \rangle. \quad (6.21)$$

Since we are assuming the error is Poisson noise, we know that $\langle N_{\text{err},k}^2 \rangle = I_k$. So we plug in the definition of I_k and find

$$\begin{aligned} \langle|\mathcal{N}_{\text{err}}|^2\rangle &= \sum_{k=0}^{K-1} (\Delta S/R^2) I_0 t \sum_l q_l \exp\left[2\pi il\left(\frac{k}{K} + \frac{mp_2}{8\lambda} \varphi(x, y)\right)\right] \\ &= (\Delta S/R^2) I_0 t \sum_l q_l \exp\left(2\pi il \cdot \frac{mp_2}{8\lambda} \cdot \varphi(x, y)\right) \sum_{k=0}^{K-1} \exp\left(2\pi il\frac{k}{K}\right). \end{aligned} \quad (6.22)$$

Unless $l = l'K$ for some l' , the sum over k is 0. Furthermore, if l is $l'K$, then the sum over k is equal to K . If we make this substitution, we get

$$\langle|\mathcal{N}_{\text{err}}|^2\rangle = K(\Delta S/R^2) I_0 t \sum_{l'} q_{l'K} \exp\left(-2\pi l' K \frac{mp_2}{8\lambda} \varphi(x, y)\right). \quad (6.23)$$

We calculate $\langle \mathcal{N}_{\text{err}}^2 \rangle$ in a similar way. First we note that

$$\langle \mathcal{N}_{\text{err}}^2 \rangle = \sum_{k=0}^{K-1} \sum_{k'=0}^{K-1} \langle N_{\text{err},k} N_{\text{err},k'} \rangle \exp\left(-2\pi i \frac{(k+k')}{K}\right). \quad (6.24)$$

As before, we note that $\langle N_{\text{err},k} N_{\text{err},k'} \rangle$ is 0 unless $k = k'$, in which case it is I_k .

Therefore

$$\langle \mathcal{N}_{\text{err}}^2 \rangle = \sum_{k=0}^{K-1} I_k(x, y, m) \exp\left(-2\pi i \frac{2k}{K}\right). \quad (6.25)$$

Also as before, we substitute in the definition of I_k , but this time the sum over the final exponential is

$$\sum_{k=0}^{K-1} \exp\left(2\pi i \frac{(l-2)k}{K}\right). \quad (6.26)$$

Instead of making the change of variables $l \rightarrow l'K$, this time we need to make the change of variables $l \rightarrow l'K + 2$. Then we find

$$\langle \mathcal{N}_{\text{err}}^2 \rangle = K(\Delta S/R^2) I_0 t \sum_{l'} q_{l'K+2} \exp\left(2\pi i (l'K + 2) \frac{mp_2}{8\lambda} \varphi(x, y)\right). \quad (6.27)$$

If the higher-order terms q_n ($|n| \geq 2$) are small relative to q_0 and q_1 , then in eqn. 6.23 only the $q_{0 \cdot K} = q_0$ term survives, since the next term is $q_{1 \cdot K}$. In eqn. 6.27, the first q term is $q_{0 \cdot K+2}$, which we assume to be negligible. Therefore

$$\langle \mathcal{N}_{\text{err}}^2 \rangle = 0 \quad (6.28)$$

and

$$\langle |\mathcal{N}_{\text{err}}^2| \rangle \approx K(\Delta S/R^2) I_0 t q_0, \quad (6.29)$$

neglecting the exponential (which we can do if φ is small).

Coming back to eqn. 6.18, we can now write

$$\langle \Delta \Psi_{\text{err}}^2 \rangle \approx \frac{1}{2} \left[\frac{K(\Delta S/R^2) I_0 t q_0}{|\mathcal{I}|^2} \right]. \quad (6.30)$$

Recall that

$$\mathcal{I}(x, y, m) = \sum_{k=0}^{K-1} I_k(x, y, m) \exp\left(-2\pi i \frac{k}{K}\right), \quad (6.31)$$

so if we plug in the definition of I_k we find that

$$\mathcal{I} = (\Delta S/R^2)I_0t \sum_l q_l \exp\left(2\pi i l \frac{mp_2}{8\lambda} \varphi(x, y)\right) \sum_{k=0}^{K-1} \exp\left(2\pi i \frac{(n-1)k}{K}\right). \quad (6.32)$$

Similar to the previous equations involving such a sum, we will make the change of variables $l \rightarrow l'K + 1$. Then the Fourier coefficients are $q_{l'K+1}$. The only non-negligible Fourier coefficient is therefore q_1 , since we said q_n was negligible for $|n| \geq 2$. So if we again neglect the exponential involving φ , we can write

$$\mathcal{I} \approx K(\Delta S/R^2)I_0tq_1. \quad (6.33)$$

If we plug this into eqn. 6.30, we have

$$\begin{aligned} \langle \Delta\Psi_{\text{err}}^2 \rangle &\approx \frac{1}{2} \frac{K(\Delta S/R^2)I_0tq_0}{[K(\Delta S/R^2)I_0tq_1]^2} \\ &= \frac{1}{2} \frac{q_0}{K(\Delta S/R^2)I_0tq_1^2}. \end{aligned} \quad (6.34)$$

So at last we have the rms error due to the photon counting statistics:

$$\Delta\Psi = \sqrt{\langle \Delta\Psi_{\text{err}}^2 \rangle} = \frac{\sqrt{q_0}}{q_1} \frac{1}{\sqrt{2}} \frac{1}{\sqrt{I_{\text{total}}(\Delta S/R^2)}}, \quad (6.35)$$

where $I_{\text{total}} = KI_0t$ is the total number of incident photons per unit solid angle over the entire phase-stepping scan. If the detector has some efficiency ϵ , we can account for it by thinking of it as a coefficient of I_k in the derivation above, in which case it appears in eqn. 6.35 as a factor of $1/\sqrt{\epsilon}$.

In an unfortunate conflict of notation, what Yashiro et al. call $\Delta\varphi_{x_s}$ is what we would call $\Delta\alpha$, the change in deflection angle. They assert that

$$\Delta\alpha = \frac{\lambda}{2\pi} \frac{\Delta\Psi}{\Delta x_s}, \quad (6.36)$$

where Δx_s is the single-step translation of the gratings during phase stepping.

Looking back at eqns. 1.59 and 1.60, we find that our phase sensitivity $\Delta\varphi$ is then given by

$$\begin{aligned}
\Delta\varphi &= \frac{2\pi}{p_2} d_m \cdot \Delta\alpha \\
&= \frac{\lambda d_m \Delta\Psi}{p_2 \Delta x_s} \\
&= \frac{\lambda d_m}{p_2 \Delta x_s} \frac{1}{q_1} \frac{1}{\sqrt{2}} \frac{1}{\sqrt{I_{\text{total}}(\Delta S/R^2)}}
\end{aligned} \tag{6.37}$$

6.1.3 Phase sensitivity versus absorption sensitivity

Although the refractive index decrement δ for x-rays is 1000 times larger than the imaginary component β , that does not always mean that phase-contrast imaging is superior to absorption-contrast imaging. We wish to derive a criterion that will allow us to determine when one or the other is preferable. The following argument is due to [33].

The authors in [33] first note that the error $\tilde{\mu}$ in determining the absorption coefficient μ of a weakly-interacting specimen is given by

$$\frac{1}{\sqrt{I_{\text{total}}}} \propto \tilde{\mu} T, \tag{6.38}$$

where T is the thickness of the specimen, and then conclude that

$$\Delta\varphi = \frac{8\lambda}{2\pi m p_2} \frac{1}{\sqrt{2}} \frac{\sqrt{q_0}}{q_1} \tilde{\mu} T. \tag{6.39}$$

(Note that eqn. 6.38 can also be made to incorporate the detector efficiency in the same way as in the previous section.)

To relate $\Delta\delta$ to $\Delta\beta$, they then posit the existence of a columnar structure of radius a and refractive index $1 - \delta_c + i\beta_c$ embedded in a medium of refractive

index $1 - \delta_m + i\beta_m$. The maximum differential phase shift caused by this object occurs at its edge, and is given by

$$\varphi_{\max} = 2 \sqrt{\frac{a}{S}} \Delta\delta, \quad (6.40)$$

where S is the 1D width of the pixel (as opposed to ΔS from the preceding section, which is the 2D extent of the pixel). We will not be able to detect the object if $\Delta\varphi > \varphi_{\max}$. The maximum absorption contrast occurs at the object's center, and is given by

$$(\Delta\mu T)_{\max} = \Delta\mu a. \quad (6.41)$$

The figure of merit is therefore a comparison of the signal-to-noise ratios $\varphi_{\max}/\Delta\varphi$ and $(\Delta\mu T)_{\max}/(\tilde{\mu}T)$. When these two SNRs are of comparable value, we have

$$\Delta\delta = \frac{8}{mp_2} \frac{\sqrt{q_0}}{q_1} \frac{1}{\sqrt{2}} \sqrt{aS} \Delta\beta, \quad (6.42)$$

after substituting $\Delta\beta = (\lambda/4\pi)\Delta\mu$.

6.1.4 Relative sensitivity of each system

We will start with eqn. 6.37 to estimate the minimum detectable phase difference for each of our setups, and then use eqn. 6.42 to estimate whether phase or absorption contrast is superior.

Prefatory notes common to all calculations: The zeroth and first Fourier components of the square wave, which we need in order to do the estimation, are 1 and $4/\pi$, respectively. The wavelength at 15 keV is 0.826\AA . The grating period p_2 is 2 microns, as is the fringe period. The Talbot distance d_1 is 24mm. Instead of using photon fluxes per steradian and multiplying them by $\Delta S/R^2$, which is

the solid angle taken up by one pixel, we will use photon fluxes per mm^2 and multiply them simply by ΔS . (Recall that ΔS is the area of the detector pixel.)

Additionally, in order to employ eqn. 6.37, we need to know the typical total exposure per scan, the pixel size, the detector efficiency, and the coherence length of the source.

Finger Lakes camera at the CHESS F3 station

The Finger Lakes' effective pixel size in our typical configuration (with $4\times$ magnification) is $(6.35\mu\text{m})^2$. The coherence length of the source is $2.0\mu\text{m}$ at 15 keV. The typical photon flux at F3 after the $20\times$ attenuation of the interferometer is 10^9ph/sec/mm^2 . Total scan times were $50\text{ sec} \times 19$ scan points, or 950 seconds total. These values correspond to a minimum detectable phase shift of

$$\Delta\varphi = \frac{2.88 \times 10^{-3}}{\sqrt{\epsilon}}, \quad (6.43)$$

where ϵ is the detector efficiency.

To evaluate the relative strength of phase vs. absorption sensitivity, we return to eqn. 6.42. Here we find that

$$\Delta\delta = 12.0 \sqrt{a} \Delta\beta, \quad (6.44)$$

where a , recall, is the object's radius in microns. For objects smaller than

$$a = \frac{1}{144} \left(\frac{\Delta\delta}{\Delta\beta} \right)^2, \quad (6.45)$$

phase contrast imaging should be more effective than absorption contrast imaging. This requires that we know δ and β both for the inclusion and for its medium. In the case of insects in amber we know neither for certain, since

the density of different deposits of amber and the chemical makeup of insect exoskeletons can vary wildly.

In the case of carbonized flowers in air, we use $\delta_C = 2.03 \times 10^{-6}$ and $\beta_C = 9.96 \times 10^{-10}$ [10], although the density of fusinite (carbonized plant material) is not necessarily the same as the density of elemental carbon. For simplicity we will use $\delta_{\text{air}} = \beta_{\text{air}} = 0$. This means that as long as the features are smaller than 29mm, phase contrast imaging will be the superior technique.

Fairchild camera at the CHESS F3 station

The Fairchild's effective pixel size, after accounting for the 1:2 magnifying FO taper, is $(3.1\mu\text{m})^2$. The detector efficiency is again $\sim 0.1e^-/\text{photon}$ at 15 keV. As before, the coherence length of the source is $2.0\mu\text{m}$ at that energy and the typical flux is $3 \times 10^9 \text{ph}/\text{sec}/\text{mm}^2$. Because the bit depth of the Fairchild is less than that of the Finger Lakes, exposure times were commensurately shorter: typically 10 seconds \times 19 scan points, or 190 seconds total. These values collectively give us a minimum detectable phase shift of

$$\Delta\varphi = \frac{1.31 \times 10^{-3}}{\sqrt{\epsilon}} \quad (6.46)$$

Here we find that the relative strength of phase contrast over absorption contrast is only

$$\Delta\delta = 8.4 \sqrt{a} \Delta\beta. \quad (6.47)$$

We have not imaged the flowers with the Fairchild, but if we had, we would expect to find that phase contrast is superior as long as the features of interest are smaller than 57mm.

Sensicam QE at the Advanced Photon Source station 32-ID-C

The Sensicam QE [39] at the APS had a 10× objective, giving it an effective pixel size of $(700\text{nm})^2$. We do not know the total efficiency of the detector chain, so we will assume $\epsilon = 1$, noting that this will give us an overly optimistic estimation of $\Delta\varphi$. Finally, the coherence length of the source is $70 \text{ m} \cdot 0.826 \text{ \AA} / 15 \text{ \mu m} = 385 \text{ \mu m}$. Typical scans were 5.5 seconds long total, with photon fluxes of $\sim 5 \times 10^{12}$ photons/sec/mm², attenuated by a factor of 20 by the interferometer to 2.5×10^{11} photons/sec/mm². Taken together, this gives us a minimum detectable phase shift of

$$\Delta\varphi = \frac{1.31 \times 10^{-3}}{\sqrt{\epsilon}} \quad (6.48)$$

Finally, here we find that, according to eqn. 6.42,

$$\Delta\delta = 1.5 \sqrt{a} \Delta\beta. \quad (6.49)$$

This indicates that phase sensitivity is stronger than absorption sensitivity in every reasonably-sized specimen.

6.2 Working with amber

6.2.1 Handling and mounting

Our primary suggestion is to leave amber handling to the experts when possible. Many pieces are tiny, fragile, or both. In particular, some have cracks running across the whole piece, which are prone to fracture under even slight pressure. We cracked one piece in this way, fortunately far from the valuable inclusion, by insufficiently delicate handling.

At CHESS we originally mounted the specimens by gripping them gently in a pair of jaws intended to hold glass optics. This was generally successful, but after the cracking incident we would no longer recommend this method. At the APS, we instead used a goniometer head with an adapter made to hold SEM pucks. The intent was to mount the specimens to SEM pucks using dental wax, which is strong enough to hold the specimens but which can be removed easily without leaving residue. In actual fact we had no SEM pucks because of a miscommunication; we used coins in their place, to great success. We would recommend the dental wax and SEM puck method in the future.

6.2.2 Discoloration

At the APS we worked at sector 32-ID-C, an undulator beamline with significantly more flux than the F3 bending-magnet beamline where we did the rest of the amber research. Amber is known to darken under intense x-ray exposure, presumably due to the formation of color centers [3] in the amber matrix. Mild discoloration of amber, like similar discoloration in some glasses, can be cured by prolonged exposure to UV light. Nevertheless we were uncertain whether severe discoloration could be likewise cured.

Therefore before examining actual inclusions, we took one piece from each amber deposit and illuminated it, away from the inclusion, for one scan's exposure time, then evaluated the severity of the darkening and the ability to regain the amber's original clarity under UV exposure. In most cases, clearing the discoloration took several hours under an intense UV lamp.

On the other hand, we never observed any discoloration of the amber spec-

imens sent to CHESS, despite using one of them for a trial computed tomography effort (for which the data turned out to be unsuitable). It is possible that in part it is the intensity of the exposure, not the total dose, that causes the discoloration.

As a caveat, we never exposed any spot on each specimen for more than 6 seconds in the undulator beam (note that this is still a lot of flux, approximately 3×10^{13} photons/mm²), and yet nearly all the APS specimens exhibited some discoloration. Extreme care should be taken when contemplating computed tomography of amber specimens under intense illumination.

6.2.3 Deposit-to-deposit variations

We found that there were distinct variations in the amber's response to x-ray microscopy depending on the deposit of origin. Indian amber, for instance, tends to be pale, and all the Indian pieces we had were quite clear and amenable to light microscopy. However, many of them had no reasonable phase-contrast signal. Presumably the density of the inclusions was too similar to the density of the surrounding amber matrix. Specimens that behaved relatively well were often from Baltic and New Jersey deposits. This is not a hard-and-fast rule, only a generalization; there were some Indian specimens with good contrast and some Baltic specimens with poor contrast.

The important thing is that each deposit's amber darkened differently in response to undulator exposure. Baltic and Dominican amber darkened visibly under 5.5 seconds of exposure; New Jersey amber showed no discoloration after the same amount of time. Each deposit must be tested for discoloration before

scanning valuable specimens.

6.3 Specimens and results

6.3.1 Amber inclusions at CHESS

Our primary set of specimens was a handful of amber specimens* covering a range of insect species and originating from several different digsites, ranging in age from 12 mya (“million years ago”) to 100 mya.

All of our specimens had already been prepared by AMNH staff for inspection by light microscopy and so were quite clear. Nevertheless, there are many specimens for which light microscopy is inadequate. It is common for amber to be turbid or opaque, for instance; also, researchers may want to view the internal organs of the inclusion. Each specimen was selected by Dr. Grimaldi with an x-ray imaging goal in mind that would not be possible with light microscopy. We imaged six amber specimens at CHESS, of which four are particularly notable. We will present them in order from least successful imaging to most successful.

All of the images in this section were collected with the Finger Lakes detector with a resolution of ~17 microns. Furthermore, we have only one or two 2D projections of each specimen. All the specimens would benefit from higher resolution: for instance, the 8.2 micron resolution of the Fairchild or the 2 micron resolution we obtained at the Advanced Photon Source (see Sec. 6.3.3).

*The amber specimens were a generous loan of Dr. David Grimaldi at the American Museum of Natural History in New York City.




Figure 6.1: Egg sac (*Araneae* sp.) in Baltic amber, 42 mya. Absorption contrast.

Moreover, the specimens would also generally benefit from three-dimensional computed tomography (CT) reconstructions. We have not yet had the chance to obtain higher-resolution or CT datasets at CHESS, simply due to beamtime availability. Radiation dose is also a consideration in obtaining CT datasets; see Sec. 6.2 for some thoughts.

The specimen in Fig. 6.2 is a spider egg sac (*Araneae* sp.). The goal was to image spiderlings within the individual eggs. In many eggs there are structured patches, but it is not clear from a single two-dimensional projection whether




Figure 6.2: Egg sac (*Araneae* sp.) in Baltic amber, 42 mya. Differential phase contrast.

these are distortions of the eggs themselves, or whether they indicate the presence of spiderlings which are not well resolved. Nevertheless, the images are classic poster children for phase contrast; compare the contrast in the absorption image with that in the DPC image.

The next specimen was a parasitoid wasp that had not yet finished emerging from a large insect egg, shown in Figs. 6.3 and 6.4. The goal was to observe the parts of the wasp that remain in the egg, which are not visible in light mi-

croscopy. In the absorption contrast, we can clearly see the wasp's lower body still ensconced in the egg. In this case, however, we see one way in which phase contrast can work against us: the bubbles in the amber show up so strongly in phase contrast that they distract the eye from the wasp's features. This specimen is somewhat disappointing in terms of resolution of meaningful specimen features. It requires a loving eye even to recognize the actual orientation of the head. To clarify, the head fills the whole circle indicated in Fig. 6.4 and is oriented so that the viewer is seeing it from the top. It is not, despite how it seems on first glance, the smaller, denser subregion to the top right of the circle, which is actually one eye.

In Fig. 6.6 we have a stingless bee (*Hymenoptera* sp.) which under light microscopy is already quite striking. The goal was simply to penetrate the outer surface of the bee and view the internal organs and musculature. In the x-ray image, there is enough detail to see the segmentation in the antennae, and also to see the gaster (stomach) in the abdomen of the bee. The flight musculature in the thorax is not so distinct; it is likely that we lacked the resolution to discern the muscle fibers.

Finally, we come to the specimen that makes the strongest case for phase-contrast microscopy. In Fig. 6.8, the right-hand specimen is a well-preserved bee (*Melikertes* sp.), which under light microscopy is obscured by fine fractures in the amber and fine hairs on the bee itself. (The left-hand specimen is quite decomposed and not the focus of this imaging effort.) The goal was to penetrate these fractures and hairs to view the internal organs. In the absorption contrast the specimen is crisp, but very faint; in phase contrast the specimen pops out of the image, internal organs become visible in the abdomen, and the wing vena-




Figure 6.3: Parasitoid wasp in Dominican amber, 20 mya. Absorption contrast.




Figure 6.4: Parasitoid wasp in Dominican amber, 20 mya. Differential phase contrast. White circle indicates entire extent of head.




Figure 6.5: Stingless bee (*Hymenoptera* sp.) in Dominican amber, 20 mya. Absorption contrast.

tion shows up much more strongly. The left-hand specimen, although decomposed, is now clearly seen. This is exactly the sort of work we initially hoped to do when we originally commissioned the interferometer.




Figure 6.6: Stingless bee (*Hymenoptera* sp.) in Dominican amber, 20 mya. Differential phase contrast.




Figure 6.7: Bee (*Melikertes* sp.) in Baltic amber, 42 mya. Absorption contrast.




Figure 6.8: Bee (*Melikertes* sp.) in Baltic amber, 42 mya. Differential phase contrast.

6.3.2 Fossil flowers

The second set of specimens was a collection of tiny fossilized flowers, no more than 3-4mm in any dimension.[†] The mechanism of their fossilization was *fusainization*, in which fast pyrolysis of the plant converts the whole structure into fusain (a form of charcoal), a process that is generally (but not universally) accepted to be due to ancient wildfires. For a thorough survey of the arguments on the origin of fusain, see [48].

One might expect that this should make the flowers easier to work with than the amber specimens, since the contrast between fusain and air is stronger than the contrast between amber and its inclusions. On the other hand, since the specimens are so tiny, the interesting internal features tend to be too small to reasonably observe with grating interferometry. Also, occasionally parts of the flowers become mineralized and very dense, so that the transmitted signal is very weak and noisy in those regions, which gives the data analysis trouble, as in Fig. 6.10. Overall these specimens were very challenging for the grating interferometer, and the resulting data were not as visually pleasing as for most of the amber specimens. Still, in some cases we were able to resolve larger internal structure, as in Figs. 6.12, 6.14, and 6.16.

[†]The flower specimens were a generous loan of Dr. William Crepet, Cornell Dept. of Plant Biology.




Figure 6.9: Specimen 88, with highly mineralized interior. Absorption contrast.




Figure 6.10: Specimen 88, with highly mineralized interior. Differential phase contrast. Note that high-density regions have noisy, poor-quality phase reconstruction.




Figure 6.11: Specimen 269. Absorption contrast.




Figure 6.12: Specimen 269. Differential phase contrast.




Figure 6.13: Specimen 849. Absorption contrast.




Figure 6.14: Specimen 849. Differential phase contrast.




Figure 6.15: Specimen 1574. Absorption contrast.




Figure 6.16: Specimen 1574. Differential phase contrast.

6.3.3 Amber inclusions at the Advanced Photon Source

We also traveled to the Advanced Photon Source (APS) with students from the American Museum of Natural History to work with a large set of amber specimens (50-60 pieces) focusing largely on scale insects, ants, and flies.

The experimental setup at 32-ID-C consisted of the CHSS interferometer, a Uniblitz shutter to protect the samples when exposures were not being collected, and a Cooke Sensicam QE [39] with a 10x microscope objective, producing an $0.65\ \mu\text{m}$ effective pixel and a resolution slightly better than $2\ \mu\text{m}$. The scintillator was a single-crystal lutetium aluminum garnet (LuAG). The experimental setup and camera system are shown in Figs. 6.17 and 6.18.

With the better resolution of the APS detector system, we were able to obtain some very beautiful images. (To appreciate them fully, we recommend you view the PDF version of this work and make use of the zoom feature.) Two of the most striking are shown in Figs. 6.19 and 6.20. Fig. 6.19 shows a fruit fly (*Drosophila* sp.) in great detail, including wing venation, structure of the mouthparts, and the individual facets of the eye. Fig. 6.20 is notable for two reasons: first, the larger inclusion is a well-preserved trapjaw ant with good detail of the mandibles; second, the piece captures a moment in time in which the smaller ant has bitten and latched onto the larger ant!

We conclude with a rogue's gallery of the most beautiful images from the APS trip. Figs. 6.21 and 6.22 show the absorption and phase contrast signals for a scale insect in amber from the Charente region of France, from the Cretaceous (~100 mya). In comparing the two, note especially the wing venation, which shows up much more strongly in the phase-contrast image. Fig. 6.23 is a



Figure 6.17: APS experimental setup. From left: goniometer mount for specimens, flexure stage, analyzer grating rotation stage, detector system. Upstream to the left of the image is the Uniblitz shutter.



Figure 6.18: APS camera setup. From left: scintillator mount, microscope objective (silver), tube lens (black), camera body.




Figure 6.19: Drosophilid fly (*Electrophortica* sp.) in Baltic amber (56–34 mya).
Insets show eye facets and detail of mouthparts.




Figure 6.20: Two ants locked in combat in Dominican amber, ~25 mya. Larger ant: *Odontomachus* sp., colloquially known as the “trapjaw” ant for its large, fast-moving mandibles.



Figure 6.21: Scale insect in amber from the Charente region of France, ~100 mya.
Absorption contrast.




Figure 6.22: Scale insect in amber from the Charente region of France, ~100 mya.
Differential phase contrast.



Figure 6.23: Scale insect in Baltic amber. Left: absorption contrast. Right: differential phase contrast. Arrow: stylet within abdomen.

comparison of the absorption and phase contrast images of a tiny scale insect in Baltic amber. The antennae and some internal features are much more visible in the phase-contrast image; note also the faint visibility of a coiled stylet in the abdomen. In Figs. 6.24 and 6.25, the overall difference in visibility is extreme. Note especially the features of the segmented section of the body, as well as the tip of the upraised leg. The specimen in Figs. 6.26 and 6.27 is slightly decomposed, but the scan still reveals the structure of the mouthparts, flight musculature,




Figure 6.24: Scale insect (*Putoidae* sp.) in Baltic amber. Absorption contrast.




Figure 6.25: Scale insect (*Putoidae* sp.) in Baltic amber. Differential phase contrast.




Figure 6.26: Drosophilid male in Dominican amber. Absorption contrast.

and abdomen. Eye facets are faintly visible. In Fig. 6.28 we have yet another Drosophilid specimen. Once again the eye facets can be made out. The segmentation in the body is quite sharp, and the fine serrations on the lower legs are evident when zoomed in. Finally we include the specimen with the most remarkable difference between absorption contrast and phase contrast. The ant is nearly invisible in Fig. 6.29, but pops out in fine detail in Fig. 6.30. Especially of note are the antennae, the tips of the feet, the partially-obscured gaster, and the pitting on the upper parts of the legs.




Figure 6.27: Drosophilid male in Dominican amber. Differential phase contrast.




Figure 6.28: Drosophilid fly in Dominican amber. Top: absorption contrast. Bottom: differential phase contrast.




Figure 6.29: Ant (*Brownimecia clavata*) in New Jersey amber. Absorption contrast.




Figure 6.30: Ant (*Brownimecia clavata*) in New Jersey amber. Differential phase contrast.

CHAPTER 7

MISCELLANY AND FUTURE WORK

7.1 TITAN data-processing software

TITAN (Talbot Interferometry Tools, Analysis, and Numerics) is a software package written by me for managing and processing grating interferometry datasets. TITAN is written in Python and depends heavily on the `numpy`, `scipy`, and `matplotlib` packages, which provide advanced numerical support and plotting capabilities. For data storage it uses the High Density File (HDF) format, supported by the `h5py` package. GUI development is done in `PyQt4`.

TITAN's main data processing interface is shown in Fig. 7.1. It includes a graphical browser for the hierarchical structure of HDF5 files as well as one- and two-dimensional plotting windows for examination of data. From this screen, data can be preprocessed (background-subtracted, flat-field corrected, rescaled to counters), fit by any of the methods described in Chapter 4, and postprocessed by appropriate rescaling to a reference dataset, then brought up in the viewer for examination.

At the time of this writing, TITAN is under active development, with the eventual goal of being deployed to CHESS stations as a direct interface with the grating interferometer. The TITAN codebase will shortly be available as a GitHub repository at <http://github.com/rmbaur/titan>.

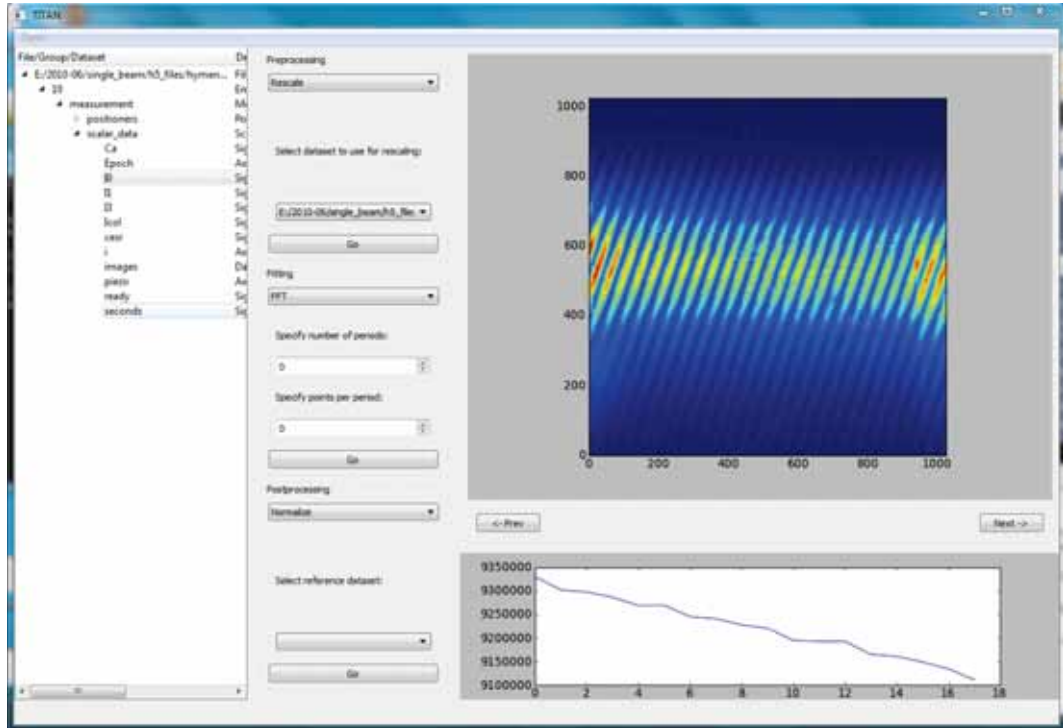


Figure 7.1: Graphical user interface of TITAN data processing software

7.2 Interferometer variants

7.2.1 Three-grating (Talbot-Lau) interferometers

Grating interferometry relies on the coherence of the source to produce visible interference fringes. In the event that the source is not sufficiently coherent, we could imagine creating a secondary source using a slit (rather than a pinhole, since we require coherence only perpendicular to the grating lines). By adjusting the width of the slit, we can adjust the apparent source size and therefore the coherence length of the secondary source.

Using a single slit wastes nearly all the incident flux, so to improve x-ray throughput we can use a large number of slits, produced by a gold line grating

(called G0 or the “source” grating). This grating is placed well upstream of the interferometer and has a period p_0 given by

$$p_0 = \frac{L}{d}p_2, \quad (7.1)$$

where L is the distance from the source grating to the phase grating, d is the distance from the phase grating to the analyzer grating, and p_2 is the period of the analyzer grating [41].

In this way we create an array of line sources which are individually coherent but mutually incoherent. Nevertheless, because of the carefully chosen period p_0 , the line sources will constructively interfere in the imaging plane. By reducing the opening fraction of the source grating, we can make the secondary sources as coherent as we like, up to the limits of photolithography and reasonable flux throughput. The interferometer can then produce fringes with close to 100% visibility.

In case it is not obvious from eqn. 7.1, the source grating is not a magic bullet for all distances. If d is changed but L is not, the secondary sources will no longer constructively interfere in the image plane and the resulting fringe visibility will be poor. Source gratings must be tuned specifically for the Talbot distance with which they will be used [61].

Three-grating interferometry at CHESS

We have successfully tested a three-grating interferometer at the F3 station. We obtained source gratings* for both $d_1 = 24$ mm and $d_3 = 72$ mm for use at $L =$

*The source gratings were manufactured by Daron Westly, Cornell Nanofabrication Facility.

1.75 m, with periods of 149 μm and 48.2 μm respectively. Both gratings had an opening fraction of 0.2.

In both cases we were able to obtain interference fringes with the grating lines vertical (which requires good horizontal coherence length), although in neither case was the visibility as good as was expected based on the source grating parameters. With an opening fraction of 0.2 and a period of 149 μm , we have 30 μm slits at $L = 1.75$ m and therefore expect a coherence length of

$$l_c = \frac{\lambda L}{w} = \frac{0.861 \text{ \AA} \cdot 1.75 \text{ m}}{30 \mu\text{m}} = 5 \mu\text{m}. \quad (7.2)$$

Likewise, for the gratings with 0.2 opening fraction and 48.2 μm period, we have 9.64 μm slits and expect a coherence length of 15.6 μm . We expect fringe visibilities of 86% in both cases.

What we instead saw were fringe visibilities of 22% at d_1 and <25% at d_3 , as illustrated in Fig. 7.2. (The very thin lines in Fig. 7.2(b) are not interference fringes, but rather shadows of the source grating itself. This could be improved by using a grating designed to be placed farther upstream to let divergence overcome the shadowing effect.) Nevertheless, we were able to obtain useful reconstructions of the data, as shown in Fig. 7.3. In the geometrical optics view of the interferometer, the longer the working distance, the larger the displacement of the deflected rays from the specimen, as shown in Fig. 7.4. Therefore we expect the interferometer's phase sensitivity to increase as the working distance is increased. We can see from Fig. 7.3 that this is in fact the case: although both subfigures (a) and (b) are shown with the same color scale, the features in the insect and the surrounding amber are much more distinct in the data taken at d_3 , even though the d_3 data was generally lower quality.




Figure 7.2: (a) 22% fringe visibility at $d_1 = 24\text{mm}$. (b) $\sim 25\%$ fringe visibility at $d_3 = 72\text{mm}$.

7.2.2 Highly divergent sources

In Chapter 2 we described the grating period modifications necessary to use the interferometer in the presence of small amounts of divergence. The same principles can be used to make the interferometer function in the presence of large divergence as well.

In the large-divergence case, if the specimen is placed well upstream of the interferometer, the expansion of the beam will magnify features in the specimen. In this way it becomes possible to detect features that are smaller than the resolution limit of the same interferometer using a plane-wave beam. Magnified imaging has been demonstrated on features as small as 900 nm with the assistance of a Fresnel zone plate to provide the magnification [6].





Figure 7.3: (a) Recovered data taken at d_1 using a source grating. (b) Recovered data taken at d_3 using a source grating.




Figure 7.4: Geometrical optics view of the interferometer showing greater sensitivity at longer working distances.

Magnifying geometry at CHESS

CHESS possesses an Oxford Instruments Nova600 microfocus x-ray tube source (90 W max power, tungsten target) which has a small enough source spot (15 microns nominal) that it can be used without a source grating. The tungsten K_α line is sufficiently high energy that we do not necessarily want to use it for grating interferometry, but the interferometer's tolerance of polychromaticity [15] means that the mean energy of the bremsstrahlung background can be used instead, by intentionally biasing the source to a voltage that is too low to excite the characteristic line.

During the summer of 2012, CHESS began construction on a hutch specifically to house this and other tube sources, so that it can be used for grating interferometry, capillary inspection, detector calibration, and so forth. When the hutch is complete, the Nova600 will be used to attempt submicron imaging of the amber specimens described in Chapter 6.

7.2.3 Two-dimensional interferometers

Since the typical grating interferometer is sensitive to phase variations only in the direction perpendicular to the grating lines, it is natural to desire a system that has sensitivity along both axes. Such a system is indeed possible and can be made by crossing two sets of 1D gratings or, to ensure that the axes are perpendicular, manufacturing the phase grating as a silicon checkerboard and the analyzer grating as a gold mesh [66].

The gratings can then be raster scanned to obtain differential phase maps and dark field images in two perpendicular directions. These two maps can then be integrated and combined to make a full phase map for the specimen [67].

Two-dimensional interferometry (not yet) at CHESS

The two-dimensional interferometer has not yet been explored at CHESS. There is in principle no reason why it could not be deployed at CHESS beamlines; however, because of attenuation in the silicon grating backing, we would need a correctly-fabricated pair of checkerboard/mesh gratings, rather than the “poor man’s 2D interferometer” consisting of crossed 1D gratings, in order to keep the exposure times reasonable. We would also need a source grating in the horizontal direction, since no current beamline at CHESS has enough horizontal coherence to operate an interferometer.

7.3 Ongoing and future work

7.3.1 Streamlining and beamline integration

The current version of the interferometer has the two gratings completely decoupled so that the phase grating can be translated for source characterization scans. This decoupling can lead to mutual vibration of the two gratings, which degrades contrast; in addition, there may be differential thermal expansion of the two grating mounting systems (although we have tried to make as many pieces as possible out of aluminum to reduce this effect), producing a slow variation in the position of the gratings which can also affect data quality.

In standard user operation, there is no need to change the grating spacing. To streamline the system, we would like to make two monolithic aluminum mounting blocks, one for $d_1 = 24$ mm and one for $d_3 = 72$ mm, on which the phase grating flexure stage and the analyzer grating rotation stage could be mounted together. The resulting system would be much more vibrationally stable and also far easier to drop into the beamline during short station changeover periods.

7.3.2 Undulator testing

In April 2012, CHESS installed and tested a new design of 1 m undulator [55], feeding A2 (with electrons) and G-line (with positrons). The grating interferometer and the Fairchild detector were installed at A2 during this testing period to attempt to measure the profile of a thermal bump on the first monochromator

crystal at a range of beam currents. The intent was to compare the interferometer measurement with a second measurement using a pinhole mask and a camera installed in the optics coffin [43]. The main benefit of using an interferometer for this measurement is that installation and removal of the equipment can be done without incursions into the coffin vacuum.

For beam currents ranging from 1 mA to 10 mA, we took two datasets at each current: one with very dense fringes and one with relatively sparse fringes. In order to measure the heat bump, we need to do a local radius of curvature measurement by segmenting the input images into smaller squares, taking the Fourier transforms of the segments, and locating the first harmonic peak. The analysis is complicated by only-partially-successful flatfield corrections, which leave traces of the hexagonal pattern of EMA fibers. This hexagonal pattern forms very strong peaks in the diffraction pattern, which can overshadow the peaks from the fringes if they are not carefully eliminated from the transformed data. Moreover, there is a tradeoff between more local measurements (smaller segments) and increased granularity of the transform data. The calculated fringe angles are more reliable for larger segments, but the large segments may miss fine changes in the fringe shape.

We have not yet worked all the kinks out of this measurement. There are fringe images in which a (presumably heat-bump-related) fringe deflection is clearly visible, but for which the local ROC measurement does not detect the deflection. This is a target for further analysis.

However, we did find success in locating the heat bump using single-shot data processing. Taking images from the dense-fringed datasets, we located and inverse-transformed the first harmonic peak in the Fourier data from those




Figure 7.5: Single-shot processed image of the beam profile, showing a bright ring due to the heat bump on the first monochromator crystal.

images. In the 10 mA data, shown in Fig. 7.5, there is a clear bright ring around the edges of the beam, which we believe to be due to the curvature imparted by the heat bump.

7.3.3 Other thermal bump/radius of curvature estimates

Recall that in Sec. 2.3.5 we estimated a radius of curvature (ROC) at F3 of 8.2 m based on the residual moiré fringes that were still visible with the gratings nominally aligned parallel to one another. We can use this to estimate the height of the heat bump.

If we have a divergent source incident on the heat bump, and if the heat bump then imparts some further curvature to the beam, we can treat the outgoing wave as a convolution of the incoming divergent profile and the heat bump's own divergence-inducing profile (assuming that the bump is not so severe as to push incoming rays out of the Darwin width of the first crystal). In such a convolution, the divergence of the outgoing beam is the sum of the divergence of the incident beam and the divergence added by the heat bump.

To determine the divergence of the outgoing beam, we imagine that there is a virtual source point somewhere upstream of the first crystal, and use our estimated beam ROC to determine that virtual source's divergence. Since the first crystal is 5 m away from the center of the hutch, that means we expect the outgoing beam to have a ROC of 3.2 m as it is just leaving the surface of the crystal. The height of the beam as it encounters and departs from the first crystal is ~ 2.1 mm, which together with the ROC allows us to find a divergence of

$$\arctan\left(\frac{2.1 \text{ mm}/2}{3.2 \text{ m}}\right) = 328 \mu\text{rad}. \quad (7.3)$$

The inherent divergence of the incident beam is $34.7 \mu\text{rad}$, so the divergence imparted by the heat bump is then $328 - 34.7 = 293.3 \mu\text{rad}$.

We also know that the footprint of the beam on the crystal is ~ 1.7 cm. Knowing this and the divergence, we can determine the heat bump ROC and then the bump height using the geometric relationships illustrated in Fig. 7.6. The ROC is given by

$$R = \frac{0.85 \text{ cm}}{\sin 293.3 \mu\text{rad}} = 29 \text{ m}. \quad (7.4)$$

Then the heat bump height is given by

$$h = R - \frac{0.85 \text{ cm}}{\tan 293.3 \mu\text{rad}} = 1.2 \mu\text{m}. \quad (7.5)$$




Figure 7.6: Diagram showing the relationship between heat bump radius of curvature, beam footprint, and the height of the heat bump.

This is in good agreement with measurements taken at A2 using a different technique [43], though in fact we should expect the A2 heat bump to be slightly more severe given the increased power density of the wiggler source over the hard-bending-magnet source at F3.

7.3.4 Dragonfly flight

The inner workings of insect flight musculature and the motion of wing joints is of great interest to researchers on flight dynamics. Synchrotron studies of fruit fly flight have been attempted [28], but fruit flies beat their wings at frequencies of approximately 200Hz [2], which can be a challenge to capture with sufficient temporal resolution without levels of irradiation that damage the subject.

By contrast, dragonflies beat their wings at frequencies of 30–40Hz [59] and their bodies are much larger, typically several millimeters, so that the muscle groups of interest are also much easier to detect. This makes them more accessible candidates for full-field x-ray microscopy. During the undulator run, we

imaged a desiccated dragonfly with the Fairchild detector to determine that the requisite muscle groups were discernible; they were. We hope in the future to use the Fairchild to image tethered dragonflies as they execute aerial maneuvers.

7.3.5 Computed tomography and laminography

Computed phase tomography has not yet been pursued at CHESS, primarily because of the exposure times required. CHESS also currently lacks reconstruction software for computed tomography datasets. There are situations in which exposure time is not a concern, such as examination with the Nova600 (rather than the use of scarce beamtime), as long as we can stay ahead of the dark current of the detector; and there exists open source software [19] for CT reconstruction. Computed tomography capability is therefore high on the list of upgrades for the CHESS grating interferometer.

Laminography is a modification of the CT algorithm for use with sheet-like specimens; the rotation axis is oriented normal to the plane of the specimen but at an angle with respect to the beam, as shown in Fig. 7.7 [22, 27]. The reconstruction algorithm is standard filtered back projection, as in CT [21]. As long as we are upgrading the system for standard CT, there is no reason not to also upgrade the system for laminography.




Figure 7.7: Specimen orientation for laminography

7.4 Conclusion

Phase contrast imaging is a valuable tool for peering inside specimens when they cannot be cut apart to view the innards and when their absorption contrast is poor. We saw in Chapter 6 how powerful this technique can be for amber and other fossils. Nevertheless, there is significant work to be done to improve the data processing protocols to minimize (and hopefully eliminate) the artifacts caused by spectral leakage. We elaborated on some of these techniques in Chapter 4, but there are still many targets for further study, particularly in the application of generalized phase shifting algorithms.

The interferometer has seen frequent use as an in-house diagnostic tool, as described in Chapter 3. Radius of curvature measurements remain problematic, with few results that are well-matched to SPECTRA predictions, but source size measurements have been a great success, and fringe visibility measurements have also identified problems with the storage ring, such as failing magnets.

After significant in-house development, CHESS is now prepared to open

the grating interferometer for outside users. Data can be processed by simple Python scripts, and a fuller-featured data processing suite (TITAN) is nearly ready for release. The new Fairchild detector (as well as a commercial detector, the Andor Neo, using the same sCMOS chip) provide high spatial resolution imaging at high frame rates. Finally, plans are in place to streamline the system for easier drop-in setup on the beamline, which will simplify user support.

The range of accessible specimens for phase contrast interferometry is huge. Tissue specimens of all types, fossil specimens such as the amber we have already worked with, live biological specimens such as dragonflies, and even integrated circuits and chips at higher energies. We anticipate a busy future for the CHESS grating interferometer.

APPENDIX A
TITAN CODEBASE AND OTHER SNIPPETS

A.1 TITAN Background

The TITAN codebase is under active development, and therefore the code presented here should not be assumed to be fully functional or necessarily up-to-date. For up-to-date code, consult <http://github.com/rmbaur/titan/>.

A.1.1 Dependencies

TITAN is dependent on several existing packages. The non-Python dependencies are

- `HDF5` 1.8 or later – a data storage format
- `libtiff` 4.0.2 or later – an image manipulation library
- `Qt` 4.8.3 or later – a Graphical User Interface toolkit
- `svn` – a version control system

The Python dependencies are

- `numpy` 1.6.1 or later – numerical data manipulation support
- `scipy` 0.10.1 or later – additional scientific data manipulation capabilities
- `matplotlib` 1.1.0 or later – a plotting and visualization package
- `h5py` 2.0.0 or later – Python bindings for HDF5

- PyQt4 4.9.4 or later – Python bindings for Qt
- `pylibtiff` – Python bindings for `libtiff`
- `praxes` 0.6.0 (not 1.1 or 1.2) – provides scientifically-meaningful structure to HDF5 files

A.1.2 Obtaining dependencies

A note on Python 2 vs. Python 3

There are substantial differences between the Python 2.x.x series and Python 3. Many vital packages are not yet ready for use with Python 3, and TITAN is not yet Python 3 compatible. In order to run TITAN, you must use a 2.x.x version of Python. If you are running Mac or Linux, you can check your Python version by typing `python --version` in a terminal. On Windows, the Python installation directory should automatically be named according to the installed Python version (typically `C:\PythonXX` where XX is the version number).

Linux

Most Linux distributions should have all of these dependencies available in their repositories with the exceptions of `pylibtiff` and `praxes`. Use the usual repository commands (e.g. `yum install libtiff` for Fedora-based distros).

The `pylibtiff` package can be obtained by `svn` checkout from Google Code at <https://code.google.com/p/pylibtiff/source/checkout>. The `praxes` package can be obtained from GitHub either as a cloned repository

or as a zipball at <https://github.com/praxes/praxes>. Both packages must then be built from source using the usual `python setup.py install` command.

Mac

Apple systems now ship with Python and some basic packages such as `numpy` preinstalled. The preference in the Python community, however, is to avoid using the system Python when possible and instead install your own, either from <http://python.org> or from a larger package system such as the Enthought Python Distribution (EPD).

Most of the major packages (`numpy`, `scipy`, etc.) have `.dmg` files available on their websites. Unless you are familiar with the `scipy` ecosystem, however, we recommend you install the EPD (either EPD Free or EPD Academic, as appropriate) from <http://www.enthought.com>. This will provide `numpy`, `scipy`, `HDF5`, `h5py`, and `matplotlib`.

For `libtiff`, `svn`, `Qt`, and `PyQt4`, we recommend the use of a Mac-based package manager, preferably `homebrew`. This provides Linux-like package management and these packages can be installed in the usual way by `brew install <package>`.

`praxes` and `pylibtiff` will still have to be built from source as for Linux.

Windows

On Windows we also recommend the use of the EPD to simplify the installation of the core dependencies. `libtiff`, `svn`, `Qt`, and `PyQt4` all have Windows installers available on their websites, and this is probably the simplest way to obtain these dependencies. Again, `praxes` and probably `pylibtiff` will need to be installed from source.

A.1.3 Installation

Once the dependencies have been installed, first obtain the `titan` package from <http://github.com/rmbaur/titan> or from this thesis. To run TITAN, install it using the usual `python setup.py install` command and then execute `titan` from a terminal.

A.2 TITAN code

The TITAN file hierarchy is shown in Fig. A.1.

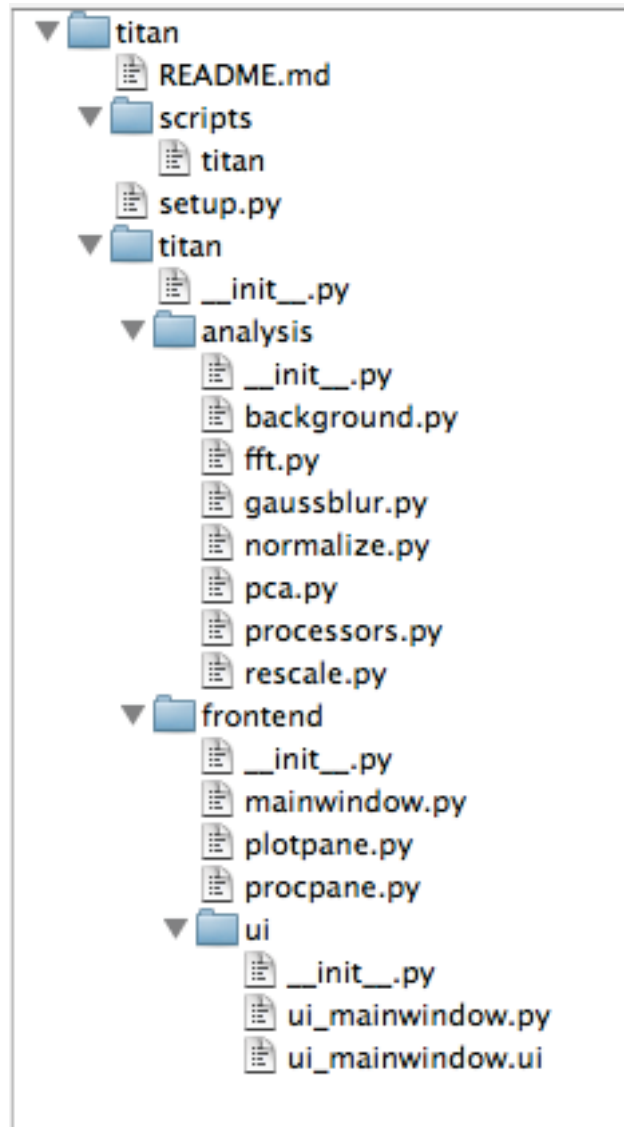


Figure A.1: The TITAN file hierarchy.

A.2.1 Frontend

The `mainwindow.py` module provides the framework for the GUI frontend of TITAN. It holds references to all the widgets that are placed in the main window and implements menu functionality for importing images and `spec` data and for opening HDF5 files.

```
1 import sys
```

```

2 import os
3
4 from PyQt4 import QtCore, QtGui
5 from libtiff import TIFF as tiff
6 import numpy as np
7
8 from praxes.frontend.phynx import FileModel, FileView
9 from praxes.io.phynx.migration.spec import convert_to_phynx
10 from praxes.io.phynx.measurement import Measurement
11 from praxes.io.phynx.dataset import Dataset
12
13 from .ui import ui_mainwindow
14 from .plotpane import DataModel, DataView
15 from .procpane import ProcPane
16
17
18 class MainWindow(ui_mainwindow.Ui_MainWindow, QtGui.QMainWindow):
19
20     proxy_changed = QtCore.pyqtSignal()
21     proc_options_changed = QtCore.pyqtSignal()
22
23     def __init__(self, parent=None):
24         super(MainWindow, self).__init__(parent)
25
26         self.setupUi(self)
27
28         self.fileModel = FileModel(self)
29         self.fileView = FileView(self.fileModel, self)
30
31         self.procPane = ProcPane()
32
33         self.dataModel = DataModel()

```

```

34     self.dataView = DataView(self.dataModel)
35     self.connect_data_model_view()
36
37     self.splitter.insertWidget(0, self.fileView)
38     self.splitter.insertWidget(1, self.procPane)
39     self.splitter.insertWidget(2, self.dataView)
40
41     self.proxy = None
42     self.fileView.clicked.connect(self.get_proxy_from_index)
43     self.proxy_changed.connect(self.update_proxy)
44
45     self.fileModel.fileAppended.connect(self.find_signals_axes)
46     self.proc_options_changed.connect(self.procPane.update_options)
47
48     def get_proxy_from_index(self, index):
49         self.proxy = self.fileModel.getProxyFromIndex(index).getNode()
50         self.proxy_changed.emit()
51
52     def connect_data_model_view(self):
53         self.dataView.buttons.prev.clicked.connect(self.dataModel.prev)
54         self.dataView.buttons.next.clicked.connect(self.dataModel.next)
55         self.dataModel.imshow_current_changed.connect(self.dataView.
56             imshow_draw)
57         self.dataModel.plot_changed.connect(self.dataView.plot_draw)
58
59     def update_proxy(self):
60         try:
61             if np.ndim(self.proxy) == 1:
62                 self.dataModel.plotdata = self.proxy
63             elif np.ndim(self.proxy) > 1:
64                 self.dataModel.imshowdata = self.proxy
65         except ValueError:

```

```

65     pass
66
67     def find_signals_axes(self):
68         signals = {}
69         axes = {}
70         files = self.fileModel.rootItem.children
71         for f in files:
72             scans = f.children
73             for s in scans:
74                 group = s.getNode('%s/measurement/scalar_data' % s.name)
75                 signals[f.file_name] = group.signals
76                 axes[f.file_name] = group.axes
77             self.procPane.signal_options = signals
78             self.procPane.axes_options = axes
79             self.proc_options_changed.emit()
80
81     @QtCore.pyqtSignature("")
82     # Slimmed down version from Praxes
83     def on_actionImportSpecFile_triggered(self, filename=None):
84         if filename is None:
85             filename = QtGui.QFileDialog.getOpenFileName(
86                 self,
87                 "Select spec file to import",
88                 os.getcwd(),
89                 "Spec files (*.dat *)")
90         )
91         if filename:
92             h5_filename = QtGui.QFileDialog.getSaveFileName(
93                 self,
94                 "Save to which file",
95                 os.getcwd(),
96                 "HDF files (*.h5 *.hdf *.hdf5)",

```

```

97         )
98     if h5_filename:
99         h5file = convert_to_phynx(str(filename), h5_filename=str(
100             h5_filename))
101         h5file.close()
102         self.fileModel.openFile(str(h5_filename))
103
104 @QtCore.pyqtSignature("")
105 def on_actionImportImages_triggered(self):
106     if not isinstance(self.proxy, Measurement):
107         confirm_import = QtGui.QMessageBox()
108         confirm_import.setText("The selected parent group is not a
109             Measurement.")
110         confirm_import.setInformativeText("Do you want to import images
111             anyway?")
112         confirm_import.setStandardButtons(QtGui.QMessageBox.Yes | QtGui.
113             QMessageBox.No)
114         confirm_import.setDefaultButton(QtGui.QMessageBox.No)
115         choice = confirm_import.exec_()
116         if choice == QtGui.QMessageBox.No:
117             return
118
119     filenames = QtGui.QFileDialog.getOpenFileNames(
120         self,
121         "Select images to import",
122         os.getcwd(),
123         "TIF images (*.tif *.tiff)"
124     )

```

```

125     for name in sorted(filenamees):
126         images.append(tiff.open(str(name), mode='r').read_image())
127
128     self.proxy.create_dataset("images", data=np.array(images))
129
130
131 @QtCore.pyqtSignature("") # Magic that prevents double signal-emits
132 def on_actionOpenHDF5File_triggered(self, filename=None):
133     if filename is None:
134         filename = QtGui.QFileDialog.getOpenFileName(
135             self,
136             "Select HDF5 file to open",
137             os.getcwd(),
138             "HDF files (*.h5 *.hdf *.hdf5)"
139         )
140     if filename:
141         self.fileModel.openFile(str(filename))

```

The main window's UI was designed in QtDesigner, a WYSIWYG package that comes with Qt. QtDesigner exports its data in the form of XML files, which are automatically converted to Python files by the `pyuic4` command (part of PyQt4) during installation. The necessary graphics objects are then rendered by Python when the program is run.

```

1 <?xml version="1.0" encoding="UTF-8"?>
2 <ui version="4.0">
3 <class>MainWindow</class>
4 <widget class="QMainWindow" name="MainWindow">
5 <property name="geometry">
6 <rect>
7 <x>0</x>

```



```
8     <y>0</y>
9     <width>1200</width>
10    <height>800</height>
11    </rect>
12  </property>
13  <property name="sizePolicy">
14    <sizepolicy hsize="Preferred" vsize="Preferred">
15      <horstretch>0</horstretch>
16      <verstretch>0</verstretch>
17    </sizepolicy>
18  </property>
19  <property name="windowTitle">
20    <string>TITAN</string>
21  </property>
22  <widget class="QWidget" name="centralwidget">
23    <property name="sizePolicy">
24      <sizepolicy hsize="Preferred" vsize="Preferred">
25        <horstretch>0</horstretch>
26        <verstretch>0</verstretch>
27      </sizepolicy>
28    </property>
29    <widget class="QSplitter" name="splitter">
30      <property name="geometry">
31        <rect>
32          <x>0</x>
33          <y>-1</y>
34          <width>1201</width>
35          <height>751</height>
36        </rect>
37      </property>
38      <property name="sizePolicy">
39        <sizepolicy hsize="Expanding" vsize="Preferred">
```

```
40     <horstretch>0</horstretch>
41     <verstretch>0</verstretch>
42 </sizepolicy>
43 </property>
44 <property name="orientation">
45     <enum>Qt::Horizontal</enum>
46 </property>
47 </widget>
48 </widget>
49 <widget class="QMenuBar" name="menubar">
50 <property name="geometry">
51 <rect>
52 <x>0</x>
53 <y>0</y>
54 <width>1200</width>
55 <height>21</height>
56 </rect>
57 </property>
58 <widget class="QMenu" name="menuOpen">
59 <property name="title">
60 <string>Open</string>
61 </property>
62 <addaction name="actionImportSpecFile"/>
63 <addaction name="actionImportImages"/>
64 <addaction name="actionOpenHDF5File"/>
65 </widget>
66 <addaction name="menuOpen"/>
67 </widget>
68 <widget class="QStatusBar" name="statusbar"/>
69 <action name="actionImportSpecFile">
70 <property name="text">
71 <string>Import spec file</string>
```

```

72 </property>
73 <property name="iconVisibleInMenu">
74 <bool>>false</bool>
75 </property>
76 </action>
77 <action name="actionOpenHDF5File">
78 <property name="text">
79 <string>Open HDF5 file</string>
80 </property>
81 <property name="iconVisibleInMenu">
82 <bool>>false</bool>
83 </property>
84 </action>
85 <action name="actionImportImages">
86 <property name="text">
87 <string>Import images</string>
88 </property>
89 </action>
90 </widget>
91 <resources/>
92 <connections/>
93 </ui>

```

The `plotpane.py` module provides plotting for 1D datasets and for collections of images. It also holds a model of the data so that it can tell which plot command, if any, is called for. The data viewer itself is passed to the `mainwindow.py` module for display.

```

1 from PyQt4 import QtCore, QtGui
2 from matplotlib.backends.backend_qt4agg import FigureCanvasQTAgg
3 from matplotlib.figure import Figure
4

```

```

5
6 class DataModel(QtCore.QObject):
7
8     imshow_current_changed = QtCore.pyqtSignal()
9     plot_changed = QtCore.pyqtSignal()
10
11 def __init__(self):
12     super(DataModel, self).__init__()
13
14     self._plotdata = None
15
16     self._imshowdata = None
17     self._current = 0
18     self._min = 0
19 try:
20     self.max = len(self._imshowdata) - 1
21 except TypeError:
22     self._max = 0
23
24 @property
25 def imshowdata(self):
26     return self._imshowdata
27
28 @imshowdata.setter
29 def imshowdata(self, val):
30     self._imshowdata = val
31     self._current = 0
32     self._max = len(self._imshowdata) - 1
33     self.imshow_current_changed.emit()
34
35 @property
36 def plotdata(self):

```

```

37     return self._plotdata
38
39     @plotdata.setter
40     def plotdata(self, val):
41         self._plotdata = val
42         self.plot_changed.emit()
43
44     def prev(self):
45         if self._current != self._min:
46             self._current -= 1
47             self.imshow_current_changed.emit()
48
49     def next(self):
50         if self._current != self._max:
51             self._current += 1
52             self.imshow_current_changed.emit()
53
54
55 class DataView(QtGui.QSplitter):
56
57     def __init__(self, model=None):
58         super(DataView, self).__init__()
59         self._model = model
60
61         self.imageview = ImshowCanvas()
62         self.buttons = PrevNextButtons()
63         self.plotview = PlotCanvas()
64
65         self.imagewidget = QtGui.QWidget(parent=self)
66         layout = QtGui.QVBoxLayout()
67         layout.addWidget(self.imageview)
68         layout.addWidget(self.buttons)

```

```

69     self.imagewidget.setLayout(layout)
70
71     self.addWidget(self.imagewidget)
72     self.addWidget(self.plotview)
73
74     self.setStretchFactor(0, 3)
75     self.setStretchFactor(1, 1)
76
77     self.setOrientation(QtCore.Qt.Vertical)
78
79     def imshow_draw(self):
80         self.imageview.update_figure(self._model.imshowdata[self._model.
            _current])
81
82     def plot_draw(self):
83         self.plotview.update_figure(self._model.plotdata)
84
85
86 class Canvas(FigureCanvasQTAagg):
87
88     def __init__(self, parent=None):
89         fig = Figure()
90         self.axes = fig.add_subplot(111)
91         self.axes.hold(False)
92
93         FigureCanvasQTAagg.__init__(self, fig)
94
95         FigureCanvasQTAagg.setSizePolicy(self,
96             QtGui.QSizePolicy.Expanding,
97             QtGui.QSizePolicy.Expanding)
98         FigureCanvasQTAagg.updateGeometry(self)
99

```

```

100     self.setParent (parent)
101
102     def update_figure():
103         pass
104
105
106     class ImshowCanvas (Canvas):
107
108         def update_figure(self, data):
109             self.axes.imshow(data, origin='lower', cmap='gist_heat')
110             self.draw()
111
112
113         class PlotCanvas (Canvas):
114
115             def update_figure(self, data):
116                 self.axes.plot (data)
117                 self.draw()
118
119
120         class PrevNextButtons (QtGui.QWidget):
121
122             def __init__(self, parent=None):
123                 super(PrevNextButtons, self).__init__()
124
125                 self.prev = QtGui.QPushButton("<- Prev")
126                 self.next = QtGui.QPushButton("Next ->")
127
128                 buttons = QtGui.QHBoxLayout()
129                 buttons.addWidget (self.prev)
130                 buttons.addStretch()
131                 buttons.addWidget (self.next)

```

132

133

```
self.setLayout(buttons)
```

The `procpane.py` module provides an interface to the data processing widgets. It holds references to all the processor objects and renders their input widgets in appropriate categories. The processing pane itself is passed to the `mainwindow.py` module for display.

```
1 from PyQt4 import QtCore, QtGui
2
3 from ..analysis.processors import *
4
5
6 class ProcPane(QtGui.QWidget):
7
8     def __init__(self, parent=None):
9         super(ProcPane, self).__init__(parent)
10
11         self.preprocs = [BkgSubtractor(), Rescaler()]
12         self.prebox = ProcBox('Preprocessing', self.preprocs)
13
14         self.fitprocs = [FFTFitter(), PCAFitter(), GPSAFitter()]
15         self.fitbox = ProcBox('Fitting', self.fitprocs)
16
17         self.postprocs = [Normalizer(), GaussFilter()]
18         self.postbox = ProcBox('Postprocessing', self.postprocs)
19
20         self.signal_options = None
21         self.axes_options = None
22
23         layout = QtGui.QVBoxLayout()
24         layout.addWidget(self.prebox)
```



```

25     layout.addWidget(self.fitbox)
26     layout.addWidget(self.postbox)
27     self.setLayout(layout)
28
29     def update_options(self):
30         self.prebox.update_widgets(self.signal_options, self.axes_options)
31         self.fitbox.update_widgets(self.signal_options, self.axes_options)
32         self.postbox.update_widgets(self.signal_options, self.axes_options
33             )
34
35
36 class ProcBox(QtGui.QGroupBox):
37
38     def __init__(self, title, procs, parent=None):
39         super(ProcBox, self).__init__(parent)
40
41         self.setTitle(title)
42
43         self.procs = procs
44         items = [proc.shortdesc for proc in self.procs]
45         self.current_proc = self.procs[0]
46
47         self.stack = QtGui.QStackedWidget()
48         for proc in self.procs:
49             self.stack.addWidget(proc.widget)
50
51         self.options = QtGui.QComboBox()
52         self.options.addItems(items)
53
54         self.go_button = QtGui.QPushButton("Go")
55

```

```

56     self.options.currentIndexChanged.connect (self.stack.
           setCurrentIndex)
57     self.options.currentIndexChanged.connect (self.set_current_proc)
58     self.go_button.clicked.connect (self.compute_called)
59
60     layout = QtGui.QVBoxLayout ()
61     layout.addWidget (self.options)
62     layout.addWidget (self.stack)
63     layout.addWidget (self.go_button)
64     self.setLayout (layout)
65
66     def compute_called(self) :
67         mw = self.parent ().parent ().parent ().parent ()
68         mw.proxy = self.current_proc.compute (mw.proxy)
69         mw.proxy_changed.emit ()
70
71     def set_current_proc(self, idx) :
72         self.go_button.clicked.disconnect (self.compute_called)
73         self.current_proc = self.procs[idx]
74         self.go_button.clicked.connect (self.compute_called)
75
76     def update_widgets(self, signals, axes) :
77         for proc in self.procs:
78             proc.update_widget (signals, axes)

```

A.2.2 Data analysis routines

The `background.py` module provides basic background subtraction capability. The user may select a file to use for background subtraction, which is then

loaded by `libtiff` and subtracted from the selected dataset.

```
1 import numpy as np
2 from libtiff import TIFF as tiff
3
4 def bkg_subtract(data, bkg_loc):
5
6     bkg = tiff.open(bkg_loc, mode='r').read_image().astype('f')
7     return data - bkg
```

The `fft.py` module provides Fast Fourier Transform data extraction capability. The user selects the number of periods represented by the data and the number of data points collected per period; the fit then extracts the absorption, phase, and dark-field data and returns them in that order.

```
1 import numpy as np
2
3 def fft_fit(data, periods, perperiod):
4     data = data[:periods * perperiod]
5     ft = np.fft.rfft(data, axis=0)
6     a0 = np.abs(ft[0] / (periods * perperiod))
7     a1 = np.abs(ft[periods] * 2 / (periods * perperiod))
8     phi = np.angle(ft[periods])
9     return np.array([a0, phi, a1])
```

The `gaussblur.py` module provides a method for attempting to subtract fringe artifacts after the fact. The user selects a size for the Gaussian blur, and the blurred image is then subtracted from the original image, with the hope that not too much genuine data was subtracted along with the artifacts.

```
1 import numpy as np
2 from scipy.ndimage import gaussian_filter
```

```

3
4
5 def gauss_blur(data, blursize):
6
7   for i in range(len(data)):
8     blur = gaussian_filter(data[i], blursize)
9     data[i] -= blur
10
11  return data

```

The `normalize.py` module provides a method for normalizing one processed dataset with respect to a reference dataset. It computes the absorption ratio and the phase difference always; it also attempts to compute the dark-field ratio if possible (not all data processing techniques provide dark-field information).

```

1 import numpy as np
2
3
4 def normalize(data, reference):
5   absorption = data[0] / reference[0]
6   phase = data[1] - reference[1]
7   try:
8     darkfield = data[2] / reference[2]
9     return np.array([absorption, phase, darkfield])
10  except IndexError:
11    return np.array([absorption, phase])

```

The `pca.py` module provides Principal Components Analysis data extraction. It puts the data in the shape required for the PCA calculation, then passes off the calculation to the `matplotlib.mlab` PCA implementation. It returns

the absorption and phase data only (no dark-field).

```
1 """
2 This method computes DPC values according to the algorithm defined in
3
4 J Xu, W Jin, L Chai, Q Xu, "Phase extraction from randomly phase-
5 shifted
6 interferograms by combining principal component analysis and least
7 squares method",
8 Optics Express 19 (21) 2011
9
10 """
11 import numpy as np
12 import matplotlib.mlab as mlab
13
14 def pca_fit(data):
15
16     # Access the data from the Dataset proxy as a numpy array
17     data = data[:]
18
19     means = np.mean(data, axis=0)
20
21     # Put the data into the shape required by the PCA implementation
22     # but keep track of the shape for later
23     shape = data.shape
24     data = (data.reshape((data.shape[0], -1))).T
25
26     pccadata = mlab.PCA(data)
27
28     phi = np.arctan2(pccadata.Y[:, 0], pccadata.Y[:, 1])
29     phi = phi.reshape((shape[1], shape[2]))
```

```
30
31 return np.array([means, phi])
```

The `rescale.py` module provides a very simple function for rescaling an image dataset by a 1D dataset, perhaps of ion chamber counts or recorded CESR current. Usually this will be used to account for a decrease in intensity over the course of the data collection.

```
1 import numpy as np
2
3 def rescale(data, counts):
4
5     return data / counts.reshape((-1, 1, 1))
```

A.2.3 Scripts

The `titan` script allows the software to be invoked from the command line after installation.

```
1 #!/usr/bin/env python
2
3 # Based on the sxfm script from Darren Dale's praxes
4
5 from __future__ import absolute_import
6
7 import sys
8
9 from PyQt4 import QtGui
10
11 from titan.frontend import mainwindow
```

```

12
13 if __name__ == "__main__":
14     qapp = QtGui.QApplication(sys.argv)
15
16     aw = mainwindow.MainWindow()
17     aw.setWindowTitle("TITAN")
18     aw.show()
19
20     sys.exit(qapp.exec_())

```

The `setup.py` script is a standard form for preparing and installing a Python package using the built-in `distutils` distribution utility.

```

1 # This is a slimmed-down version of the setup.py file
2 # written by D. Dale for the praxes project
3 # @ http://github.com/praxes/praxes.
4 # Only the setup() is new.
5
6 from distutils.core import setup
7 from distutils.core import Command
8 from distutils.command.sdist import sdist as _sdist
9 from distutils.command.build import build as _build
10 from distutils.command.bdist_wininst import bdist_wininst as
    _bdist_wininst
11 import multiprocessing
12 import os
13 import subprocess
14 import sys
15
16
17 def convert_ui(args, **kwargs):
18     subprocess.call(args, **kwargs)

```

```

19
20
21 class ui_cvt(Command):
22
23     description = "Convert Qt user interface files to PyQt .py files"
24
25     user_options = []
26
27     boolean_options = []
28
29     def initialize_options(self):
30         pass
31
32     def finalize_options(self):
33         pass
34
35     def run(self):
36         try:
37             to_process = []
38             for root, dirs, files in os.walk('.'):
39                 for f in files:
40                     if f.endswith('.ui'):
41                         source = os.path.join(root, f)
42                         dest = os.path.splitext(source)[0] + '.py'
43                         exe = 'pyuic4'
44                     else:
45                         continue
46
47                     if not os.path.exists(dest):
48                         to_process.append([exe, '-o', dest, source])
49
50     if sys.platform.startswith('win'):

```



```

51     # doing this in parallel on windows will crash your computer
52     [convert_ui(args, shell=True) for args in to_process]
53     else:
54         pool = multiprocessing.Pool()
55         pool.map(convert_ui, to_process)
56     except EnvironmentError:
57         print("""\
58             Warning: PyQt4 development utilities (pyuic4) not found
59             Unable to install graphical user interface
60             """)
61
62
63     class sdist(_sdist):
64
65     def run(self):
66         self.run_command('ui_cvt')
67         _sdist.run(self)
68
69
70     class build(_build):
71
72     def run(self):
73         self.run_command('ui_cvt')
74         _build.run(self)
75
76
77     class bdist_wininst(_bdist_wininst):
78
79     def run(self):
80         self.run_command('ui_cvt')
81         _bdist_wininst.run(self)
82

```

```

83
84 packages = []
85 for dirpath, dirnames, filenames in os.walk('titan'):
86     if '__init__.py' in filenames:
87         packages.append('.'.join(dirpath.split(os.sep)))
88     else:
89         del(dirnames[:])
90
91 scripts = []
92 if sys.platform.startswith('win'):
93     # scripts calling multiprocessing must be importable
94     import shutil
95     shutil.copy('scripts/titan', 'scripts/titan.py')
96     scripts.append('scripts/titan.py')
97 else:
98     scripts.append('scripts/titan')
99
100 setup(name = 'TITAN',
101       author = 'Robin M Baur',
102       author_email = 'rmb62@cornell.edu',
103       cmdclass = {
104         'bdist_wininst': bdist_wininst,
105         'build': build,
106         'sdist': sdist,
107         'ui_cvt': ui_cvt
108     },
109       packages = packages,
110       requires = (
111         'python (>=2.7)',
112         'numpy (>=1.6.1)',
113         'scipy (>=0.10.1)',
114         'h5py (>=2.0.0)',

```

```
115     'praxes (>=0.5.1)'  
116     ),  
117     scripts = scripts,  
118     version = '0.1.0'  
119     )
```

BIBLIOGRAPHY

- [1] Jens Als-Nielsen and Des McMorrow. *Elements of Modern X-Ray Physics*. John Wiley & Sons, Ltd., West Sussex, England, 1st edition, 2001.
- [2] Douglas L Altshuler, William B Dickson, Jason T Vance, Stephen P Roberts, and Michael H Dickinson. Short-amplitude high-frequency wing strokes determine the aerodynamics of honeybee flight. *Proceedings of the National Academy of Sciences of the United States of America*, 102(50):18213–8, December 2005. ISSN 0027-8424. doi: 10.1073/pnas.0506590102. URL <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1312389&tool=pmcentrez&rendertype=abstract>.
- [3] Neil W Ashcroft and N David Mermin. *Solid State Physics*. Brooks-Cole, 1st edition, 1976.
- [4] Simon Baird. Accelerators for Pedestrians Chapter 3: Lattices and Twiss Parameters. URL <http://ps.web.cern.ch/ps/training/pedestrians/chap3.pdf>.
- [5] S. L. Barna, M. W. Tate, S. M. Gruner, and E. F. Eikenberry. Calibration procedures for charge-coupled device x-ray detectors. *Review of Scientific Instruments*, 70(7):2927, 1999. ISSN 00346748. doi: 10.1063/1.1149852. URL <http://link.aip.org/link/RSINAK/v70/i7/p2927/s1&Agg=doi>.
- [6] S. Berujon, H. Wang, I. Pape, K. Sawhney, S. Rutishauser, and C. David. X-ray submicrometer phase contrast imaging with a Fresnel zone plate and a two dimensional grating interferometer. *Optics Letters*, 37(10):1622, May 2012. ISSN 0146-9592. doi: 10.1364/OL.37.001622. URL <http://www.opticsinfobase.org/abstract.cfm?URI=ol-37-10-1622>.

- [7] G Larry Bretthorst. *Bayesian Spectrum Analysis and Parameter Estimation*. Springer-Verlag, Berlin, 1st edition, 1988.
- [8] Clifford Bueno. Scintillating Fiber Optic Faceplate Performance Data. URL <http://www.collimatedholes.com/scindata.html>.
- [9] P D Burns. Slanted-edge MTF for digital camera and scanner analysis. *PICS 2000: Image Processing, Image Quality, Image Capture Systems Conference*, pages 135–138, 2000.
- [10] Center for X-ray Optics. X-Ray Interactions With Matter. URL http://henke.lbl.gov/optical/_constants.
- [11] Michael Chabior, Manfred Schuster, Matthias Goldammer, Christian Schroer, and Franz Pfeiffer. Influence of the grating profiles on the image quality in grating-based x-ray imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 683:71–77, August 2012. ISSN 01689002. doi: 10.1016/j.nima.2012.04.068. URL <http://linkinghub.elsevier.com/retrieve/pii/S0168900212004470>.
- [12] Peter Cloetens, J. P. Guigay, C. De Martino, J. Baruchel, and Michel Schlenker. Fractional Talbot imaging of phase gratings with hard x rays. *Optics Letters*, 22(14):1059, July 1997. ISSN 0146-9592. doi: 10.1364/OL.22.001059. URL <http://www.opticsinfobase.org/abstract.cfm?URI=ol-22-14-1059>.
- [13] John M. Cowley. *Diffraction Physics*. Elsevier, Amsterdam, 3rd edition, 1995. ISBN 0444822186.

- [14] Paul Antoine Douissard, Angelica Cecilia, Thierry Martin, Valentin Chevalier, Maurice Couchaud, Tilo Baumbach, Klaus Dupré, Markus Kühbacher, and Alexander Rack. A novel epitaxially grown LSO-based thin-film scintillator for micro-imaging using hard synchrotron radiation. *Journal of synchrotron radiation*, 17(5):571–83, September 2010. ISSN 1600-5775. doi: 10.1107/S0909049510025938. URL <http://www.ncbi.nlm.nih.gov/pubmed/20724778>.
- [15] M Engelhardt, C Kottler, O Bunk, C David, C Schroer, Joachim Baumann, M Schuster, and F Pfeiffer. The fractional Talbot effect in differential x-ray phase-contrast imaging for extended and polychromatic x-ray sources. *Journal of microscopy*, 232(1):145–57, October 2008. ISSN 1365-2818. doi: 10.1111/j.1365-2818.2008.02072.x. URL <http://www.ncbi.nlm.nih.gov/pubmed/19017212>.
- [16] E. J. Friebele and D. L. Griscom. Color Centers in Glass Optical Fiber Waveguides. *MRS Proceedings*, 61, February 1985. ISSN 1946-4274. doi: 10.1557/PROC-61-319. URL http://journals.cambridge.org/abstract/_S1946427400446610.
- [17] Sol M. Gruner, Mark W. Tate, and Eric F. Eikenberry. Charge-coupled device area x-ray detectors. *Review of Scientific Instruments*, 73(8):2815, 2002. ISSN 00346748. doi: 10.1063/1.1488674. URL <http://link.aip.org/link/RSINAK/v73/i8/p2815/s1\&Agg=doi>.
- [18] J. P. Guigay. On Fresnel Diffraction by One-dimensional Periodic Objects, with Application to Structure Determination of Phase Objects. *Journal of Modern Optics*, 18(9):677–682, September 1971. ISSN 0950-0340. doi: 10.1080/713818491. URL <http://www.informaworld.com/openurl?>

genre=article\&doi=10.1080/713818491\&magic=crossref|
|D404A21C5BB053405B1A640AFFD44AE3.

- [19] A P Hammersley. HST: High Speed Tomography Reference Manual V0.4. URL http://www.esrf.eu/computing/scientific/HST/HST_REF/hst.html.
- [20] Geon-Soo Han and Seung-Woo Kim. Numerical correction of reference phases in phase-shifting interferometry by iterative least-squares fitting. *Applied Optics*, 33(31):7321, November 1994. ISSN 0003-6935. doi: 10.1364/AO.33.007321. URL <http://www.opticsinfobase.org/abstract.cfm?URI=ao-33-31-7321>.
- [21] L. Helfen, T. Baumbach, P. Mikulik, D. Kiel, P. Pernot, P. Cloetens, and J. Baruchel. High-resolution three-dimensional imaging of flat objects by synchrotron-radiation computed laminography. *Applied Physics Letters*, 86(7):071915, 2005. ISSN 00036951. doi: 10.1063/1.1854735. URL <http://link.aip.org/link/APPLAB/v86/i7/p071915/s1\&Agg=doi>.
- [22] L. Helfen, T. Baumbach, P. Cloetens, and J. Baruchel. Phase-contrast and holographic computed laminography. *Applied Physics Letters*, 94(10):104103, 2009. ISSN 00036951. doi: 10.1063/1.3089237. URL <http://link.aip.org/link/APPLAB/v94/i10/p104103/s1\&Agg=doi>.
- [23] Abdelaziz Ikhlef and Maurice Skowronek. Application of a Plastic Scintillating Fiber Array for Low-Energy X-Ray Imaging. *Applied Optics*, 37(34):8081, December 1998. ISSN 0003-6935. doi: 10.1364/AO.37.008081. URL <http://www.opticsinfobase.org/abstract.cfm?URI=ao-37-34-8081>.

- [24] ISO. ISO 12233. URL http://www.iso.org/iso/catalogue_detail.htm?csnumber=33715.
- [25] James R. Janesick. *Scientific Charge-Coupled Devices*. SPIE Press, 2001. ISBN 0819436984.
- [26] Alexander Kazimirov and Darren Dale. CHESS West - A2. URL <http://www.chess.cornell.edu/chess/west/A2.htm>.
- [27] Kris Krug, Liisa Porra, Paola Coan, Arie Wallert, Joris Dik, Andrea Coerd, Alberto Bravin, Muthaffar Elyyan, Péter Reischig, Lukas Helfen, and Tilo Baumbach. Relics in medieval altarpieces? Combining X-ray tomographic, laminographic and phase-contrast imaging to visualize thin organic objects in paintings. *Journal of synchrotron radiation*, 15(Pt 1):55–61, January 2008. ISSN 0909-0495. doi: 10.1107/S0909049507045438. URL <http://www.ncbi.nlm.nih.gov/pubmed/18097079>.
- [28] Wah-Keat Lee. Personal communication.
- [29] Leonard Mandel and Emil Wolf. *Optical Coherence and Quantum Optics*. Cambridge University Press, Cambridge, UK, 1st edition, 1995.
- [30] Donald W. Marquardt. An Algorithm for Least-Squares Estimation of Non-linear Parameters. *SIAM Journal on Applied Mathematics*, 11(2):431, 1963. ISSN 00361399. doi: 10.1137/0111030. URL <http://link.aip.org/link/SMJMAP/v11/i2/p431/s1\&Agg=doi>.
- [31] Richard Mead. Personal communication.
- [32] Atsushi Momose, Wataru Yashiro, Yoshihiro Takeda, Yoshio Suzuki, and Tadashi Hattori. Phase Tomography by X-ray Talbot Interferometry for

- Biological Imaging. *Japanese Journal of Applied Physics*, 45(6A):5254–5262, June 2006. ISSN 0021-4922. doi: 10.1143/JJAP.45.5254. URL <http://jjap.jsap.jp/link?JJAP/45/5254/>.
- [33] Atsushi Momose, Wataru Yashiro, and Yoshihiro Takeda. Sensitivity of X-ray Phase Imaging Based on Talbot Interferometry. *Japanese Journal of Applied Physics*, 47(10):8077–8080, October 2008. ISSN 0021-4922. doi: 10.1143/JJAP.47.8077. URL <http://jjap.jsap.jp/link?JJAP/47/8077/>.
- [34] NIST. X-ray Mass Attenuation Coefficients - Gallium, . URL <http://physics.nist.gov/PhysRefData/XrayMassCoef/ElemTab/z31.html>.
- [35] NIST. X-ray Mass Atten. Coef. - Glass, Borosilicate, . URL <http://physics.nist.gov/PhysRefData/XrayMassCoef/ComTab/pyrex.html>.
- [36] Alan V Oppenheim, Ronald W Schafer, and John R Buck. *Discrete-Time Signal Processing*. Prentice-Hall, 2 edition, 1999.
- [37] David M. Paganin. *Coherent X-Ray Optics*. Oxford University Press, Oxford, UK, 1st edition, 2006.
- [38] P. Pavan, G. Zanella, R. Zannoni, and A. Marigo. Spatial resolution in X-ray imaging with scintillating glass optical fiber plates. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 327(2-3):600–604, April 1993. ISSN 01689002. doi: 10.1016/0168-9002(93)90730-6. URL <http://linkinghub.elsevier.com/retrieve/pii/0168900293907306>.

- [39] PCO-Tech. Sensicam QE. URL <http://www.pco-tech.com/categories/sensitive-cameras/sensicam-qe/>.
- [40] F Pfeiffer, M Bech, O Bunk, P Kraft, E F Eikenberry, Ch Brönnimann, C Grünzweig, and C David. Hard-X-ray dark-field imaging using a grating interferometer. *Nature materials*, 7(2):134–7, February 2008. ISSN 1476-1122. doi: 10.1038/nmat2096. URL <http://www.ncbi.nlm.nih.gov/pubmed/18204454>.
- [41] Franz Pfeiffer, Timm Weitkamp, Oliver Bunk, and Christian David. Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources. *Nature Physics*, 2(4):258–261, March 2006. ISSN 1745-2473. doi: 10.1038/nphys265. URL <http://www.nature.com/doifinder/10.1038/nphys265>.
- [42] P. Revesz. Ion Chamber Flux Calculator. URL <http://www.chess.cornell.edu/calculator/ICFluxCalculator.html>.
- [43] P. Revesz, A. Kazimirov, and I. Bazarov. In situ visualization of thermal distortions of synchrotron radiation optics. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 576(2-3):422–429, June 2007. ISSN 01689002. doi: 10.1016/j.nima.2007.02.110. URL <http://linkinghub.elsevier.com/retrieve/pii/S0168900207004573>.
- [44] Brian Rodricks, Boyd Fowler, Chiao Liu, John Lowes, Lucas Koerner, Mark W. Tate, and Sol M. Gruner. A quantum-limited CMOS-sensor-based high-speed imaging system for time-resolved x-ray scattering. *Proceedings of SPIE*, 7449(7449Y):1–11, 2009. doi: 10.1117/12.828329.

- [45] SciPy. `numpy.linspace` – NumPy v1.6 Manual (DRAFT), . URL <http://docs.scipy.org/doc/numpy/reference/generated/numpy.linspace.html>.
- [46] SciPy. `numpy.random.standard_normal` – NumPy v1.6 Manual (DRAFT), . URL http://docs.scipy.org/doc/numpy/reference/generated/numpy.random.standard_normal.html.
- [47] SciPy. `scipy.optimize.leastsq` – SciPy v0.11 Reference Manual (DRAFT), . URL <http://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.leastsq.html>.
- [48] Andrew C. Scott. Observations on the nature and origin of fusain. *International Journal of Coal Geology*, 12(1-4):443–475, June 1989. ISSN 01665162. doi: 10.1016/0166-5162(89)90061-X. URL <http://linkinghub.elsevier.com/retrieve/pii/016651628990061X>.
- [49] Hua Shao, Don W. Miller, and C.Robert Pearsall. Scintillating fiber optics for X-ray radiation imaging. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 299(1-3):528–533, December 1990. ISSN 01689002. doi: 10.1016/0168-9002(90)90837-V. URL <http://linkinghub.elsevier.com/retrieve/pii/016890029090837V>.
- [50] Mitsuo Takeda, Hideki Ina, and Seiji Kobayashi. Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *Journal of the Optical Society of America*, 72(1):156, January 1982. ISSN 0030-3941. doi: 10.1364/JOSA.72.000156. URL <http://www.opticsinfobase.org/abstract.cfm?URI=josa-72-1-156>.

- [51] H.F. Talbot. LXXVI. Facts relating to optical science. No. IV. *Philosophical Magazine Series 3*, 9(56):401–407, December 1836. ISSN 1941-5966. doi: 10.1080/14786443608649032. URL <http://www.tandfonline.com/doi/abs/10.1080/14786443608649032>.
- [52] Takashi Tanaka and Hideo Kitamura. SPECTRA. URL <http://radiant.harima.riken.go.jp/spectra/>.
- [53] Mark W. Tate, Darol Chamberlain, and Sol M. Gruner. Area x-ray detector based on a lens-coupled charge-coupled device. *Review of Scientific Instruments*, 76(8):081301, 2005. ISSN 00346748. doi: 10.1063/1.2001307. URL <http://link.aip.org/link/RSINAK/v76/i8/p081301/s1\&Agg=doi>.
- [54] Alejandro Téllez-Quiñones, Daniel Malacara-Doblado, and Jorge García-Márquez. Phase-shifting algorithms for a finite number of harmonics: first-order analysis by solving linear systems. *Journal of the Optical Society of America A*, 29(4):431, March 2012. ISSN 1084-7529. doi: 10.1364/JOSAA.29.000431. URL <http://www.opticsinfobase.org/abstract.cfm?URI=josaa-29-4-431>.
- [55] A. Temnykh, T. Kobela, A. Lyndaker, J. Savino, E. Suttner, and Y. Li. Compact Undulator for Cornell High Energy Synchrotron Source. *IEEE Transactions on Applied Superconductivity*, 22(3):4100504–4100504, June 2012. ISSN 1051-8223. doi: 10.1109/TASC.2011.2181932. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6112670>.
- [56] Lloyd N Trefethen and David Bau III. *Numerical Linear Algebra*. SIAM, Philadelphia, 1st edition, 1997. ISBN 0-89871-361-7.

- [57] J. Vargas, J. Antonio Quiroga, and T. Belenguer. Phase-shifting interferometry based on principal component analysis. *Optics Letters*, 36(8):1326, April 2011. ISSN 0146-9592. doi: 10.1364/OL.36.001326. URL <http://www.opticsinfobase.org/abstract.cfm?URI=ol-36-8-1326>.
- [58] J. Vargas, J. Antonio Quiroga, and T. Belenguer. Analysis of the principal component algorithm in phase-shifting interferometry. *Optics Letters*, 36(12):2215, June 2011. ISSN 0146-9592. doi: 10.1364/OL.36.002215. URL <http://www.opticsinfobase.org/abstract.cfm?URI=ol-36-12-2215>.
- [59] Z. Wang and David Russell. Effect of Forewing and Hindwing Interactions on Aerodynamic Forces and Power in Hovering Dragonfly Flight. *Physical Review Letters*, 99(14), October 2007. ISSN 0031-9007. doi: 10.1103/PhysRevLett.99.148101. URL <http://link.aps.org/doi/10.1103/PhysRevLett.99.148101>.
- [60] Timm Weitkamp. Moire interferometry formulas for hard x-ray wavefront sensing. In *Proceedings of SPIE*, volume 5533, pages 140–144. SPIE, 2004. doi: 10.1117/12.559695. URL <http://link.aip.org/link/?PSI/5533/140/1\&Agg=doi>.
- [61] Timm Weitkamp, Christian David, Christian Kottler, Oliver Bunk, and Franz Pfeiffer. Tomography with grating interferometers at low-brilliance sources. *Proceedings of SPIE*, pages 63180S–63180S–10, 2006. ISSN 0277786X. doi: 10.1117/12.683851. URL <http://link.aip.org/link/PSISDG/v6318/i1/p63180S/s1\&Agg=doi>.
- [62] Ralf Widenhorn. Temperature dependence of dark current in a CCD. In *Proceedings of SPIE*, volume 4669, pages 193–201. SPIE, 2002. doi:

10.1117/12.463446. URL <http://link.aip.org/link/?PSI/4669/193/1\&Agg=doi>.

- [63] Jiancheng Xu, Weimin Jin, Liqun Chai, and Qiao Xu. Phase extraction from randomly phase-shifted interferograms by combining principal component analysis and least squares method. *Optics Express*, 19(21):20483, October 2011. ISSN 1094-4087. doi: 10.1364/OE.19.020483. URL <http://www.opticsinfobase.org/abstract.cfm?URI=oe-19-21-20483>.
- [64] Wataru Yashiro, Yoshihiro Takeda, and Atsushi Momose. Efficiency of capturing a phase image using cone-beam x-ray Talbot interferometry. *Journal of the Optical Society of America A*, 25(8):2025, July 2008. ISSN 1084-7529. doi: 10.1364/JOSAA.25.002025. URL <http://www.opticsinfobase.org/abstract.cfm?URI=josaa-25-8-2025>.
- [65] Ryosuke Yokota. Color Centers in Alkali Silicate and Borate Glasses. *Physical Review*, 95(5):1145–1148, September 1954. ISSN 0031-899X. doi: 10.1103/PhysRev.95.1145. URL <http://link.aps.org/doi/10.1103/PhysRev.95.1145>.
- [66] Irene Zanette, Christian David, Simon Rutishauser, Timm Weitkamp, Melissa Denecke, and Clive T. Walker. 2D grating simulation for X-ray phase-contrast and dark-field imaging with a Talbot interferometer. pages 73–79, 2010. doi: 10.1063/1.3399260. URL <http://link.aip.org/link/APCPCS/v1221/i1/p73/s1\&Agg=doi>.
- [67] Irene Zanette, Timm Weitkamp, Tilman Donath, Simon Rutishauser, and Christian David. Two-Dimensional X-Ray Grating Interferometer. *Physical Review Letters*, 105(24), December 2010. ISSN 0031-9007. doi: 10.1103/

PhysRevLett.105.248102. URL [http://link.aps.org/doi/10.1103/
PhysRevLett.105.248102.](http://link.aps.org/doi/10.1103/PhysRevLett.105.248102)