

Parameter Extraction and Uncertainty in Terahertz Time-Domain Spectroscopic Measurements

Nicholas Robert Greenall

*Submitted in accordance with the requirements
for the degree of Doctor of Philosophy*

University of Leeds

School of Electronic and Electrical Engineering

September 2017

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

A number of publications have arisen from the work presented in this thesis:

Greenall, N., A. Valavanis, H. J. Desai, D. O. Acheampong, L. H. Li, J. E. Cunningham, A. G. Davies, E. H. Linfield, and A. D. Burnett. "The Development of a Semtex-H Simulant for Terahertz Spectroscopy." *Journal of Infrared, Millimeter, and Terahertz Waves* 38, no. 3 (2017): 325-338.

N.R. Greenall - Design and implementation of algorithms, processing of simulant spectra, writing of the paper. A.D. Burnett - Chemical analysis, measurements of explosives and component database. Design of spectroscopic system, initial idea, revision of the text. H. J. Desai - initial idea, explosive samples. D. O. Acheampong - explosive samples. L. H. Li, E. H. Linfield - Wafers and Emitter fabrication. J. E. Cunningham, A. G. Davies - Supervision and revision of the paper. Information included in Chapter 4 - Creation of Explosive Simulants

Greenall, N. R., L. H. Li, E. H. Linfield, A. G. Davies, J. E. Cunningham, and A. D. Burnett. "Multilayer extraction of complex refractive index in broadband transmission terahertz time-domain spectroscopy." In *Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*, 2016 41st International Conference on, pp. 1-2. IEEE, 2016.

N.R. Greenall - Initial idea, design and implementation of algorithms. N.R. Greenall measurement and processing of samples. writing of the paper. A.D. Burnett - Design of spectroscopic system, revision of the text. L. H. Li, E. H. Linfield - Wafers and Emitter fabrication. J. E. Cunningham, A. G. Davies - Supervision and revision of the paper. Information included and significantly extended in Chapter 5 - Fitting Sample Models

Greenall, N. R., C. D. Wood, C. Russell, L. H. Li, E. H. Linfield, A. G. Davies, J. E. Cunningham, and A. D. Burnett. "Accurate material parameter extraction from broadband terahertz spectroscopy." In *Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*, 2015 40th International Conference on, pp. 1-2. IEEE, 2015.

N.R. Greenall - Initial idea, Design and implementation of algorithms, measurement and processing of samples, writing of the paper. C.D. Wood, Experimental Assistance, supervision and revision of the paper. C. Russell, - Experimental Assistance A.D. Burnett - Initial idea, design of spectroscopic system, revision of the text. L. H. Li, E. H. Linfield - Wafers and Emitter fabrication. J. E. Cunningham, A. G. Davies - Supervision and revision of the paper. Information included in Chapter 5 - Transfer function methods of extraction

Greenall, N. R., E. H. Linfield, A. G. Davies, L. H. Li, J. E. Cunningham, and A. D. Burnett. "Estimation of spectroscopic uncertainty and correlation in terahertz time domain spectroscopy." In *Infrared, Millimeter, and Terahertz waves (IRMMW-THz)*, 2016 41st International Conference on, pp. 1-2. IEEE, 2016.

N.R. Greenall - Initial idea, Design and implementation of algorithms, measurement and processing of samples, writing of the paper. A.D. Burnett - Design of spectroscopic system, revision of the text. L. H. Li, E. H. Linfield - Wafers and Emitter fabrication. J. E. Cunningham, A. G. Davies - Supervision and revision of the paper. Information included in Chapter 6 - Estimation of Uncertainty in complex refractive index

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

©2017 The University of Leeds and Nicholas Greenall.

The right of Nicholas Greenall to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

Acknowledgements

Firstly I would like to thank my PhD supervisor Dr. Andrew Burnett for providing me with considerable support during my PhD and for supporting my research interests. I would also like to offer my sincere gratitude to Prof. John Cunningham for co-supervising me and supervising my Masters project, as well as providing helpful advice and assistance during my PhD. I would also like to thank Prof. Edmund Linfield and Prof. Giles Davies for their constructive comments and assistance during supervision meetings.

I would also like to extend my gratitude to other members of the group. I would like to thank David Bacon for providing new emitters and detectors for me to use, letting me use his system and equipment and performing considerable system alignment. I would also like to thank Mathew Swithenbank for providing help advice, as well as performing measurements and assisting me generally in my research. I would like to extend my sincere gratitude to both Dr Chris Russell, Dr Reshma Mohandas, Binbin Hong and Nicholas Peters, who have given me constructive and extremely helpful feedback on my work. I would like to thank Dr. Joshua Freeman, Dr. Chris Wood and Dr. Alex Valavanis who have both given me considerable advice and help during my PhD.

I would also like to extend my sincere gratitude to everyone in the Polard Institute, for their companionship and help during my degree, including Dr Nic Hunter, Dr Siddhant Chowdhury, Dr Viktor Doychinov, Dr Andy Sills, and Pierluigi Rubino. I would also like to thank Sue Hobson for her administrative support and all members of staff at the School of Electronic and Electrical Engineering.

I am especially grateful to my parents, Peter and Sarah, and my brother, Will, for their love and support. I would particularly like to thank my girlfriend, Anne-Christine Charigault, who's support during the process of completing this degree has been invaluable. Finally I would like to extend my thanks and gratitude to my wider family, including my grandparents, cousins, aunts and uncles.

Abstract

Terahertz (THz) time domain spectroscopy is emerging as a powerful tool to characterise samples both chemically and physically. In this work different methods of estimating spectroscopic parameters of a sample, its thickness and the uncertainty of these estimates is presented. A number of case studies are also examined including paracetamol polymorphs and a method of creating a spectroscopic simulant of Sementex-H is presented.

Approximation of the sample spectroscopic parameters, real refractive index and absorption coefficient were formed by building up a simple model of the samples interaction with THz radiation. Methods of correcting unwrapping error in the real refractive index were developed, including a method to correct in the presence of discontinuities in the refractive index itself. These approximations were then applied to extract parameters of both lactose and paracetamol samples.

An algorithm to generate spectroscopic simulants was developed and applied to Sementex-H. These simulants consisted of simple mixtures of inert compounds, which were measured and found to have similar spectrum to the target sample.

Methods of fitting resonant models to the sample response were developed to extract both the spectroscopic parameters and sample thickness. These were refined by calibrating for the Gaussian beam profile of the THz radiation, which was shown to increase the accuracy of the extracted thickness. The thickness and spectroscopic parameters of a lactose sample were measured with temperature, and it was found that the spectroscopic parameter change was underestimated when thickness was assumed constant.

A resonant model for multilayered samples was then developed and used to characterise IPA in a flowcell measurement. This was then combined with a method of time segmentation of the sample response, to extract spectroscopic parameters and sample thickness simultaneously. This was then applied to a two layer sample, to extract the spectroscopic parameters of a silicon and a quartz layer from a single measurement.

Finally, methods of propagating the uncertainty from the time domain to the spectroscopic parameters were developed. These were based on a multivariate normal statistical model of the measurements and were compared to numerical bootstrap and Monte-Carlo estimates. These were used to develop confidence intervals for

the extracted refractive index, absorption coefficient and thickness. These methods were applied to both a lactose and quartz sample.

Contents

Abbreviations	xix
List of Symbols	xxi
1 Introduction	1
1.0.1 Objectives and structure of this work	2
2 Principles of THz TDS and signal processing	5
2.1 Terahertz time domain spectroscopy instruments	5
2.1.1 Photo-conductive emission and detection	6
2.1.2 Electro-optic sampling	9
2.1.3 Femtosecond lasers	10
2.1.4 Terahertz TDS systems used in this work	11
2.1.5 Time sampling in a terahertz time domain system	14
2.2 Extraction of spectroscopic parameters in terahertz (THz) time do- main spectroscopy (TDS)	18
2.3 The transfer function and Fourier transform	22
2.3.1 Linear time invariant systems	22
2.3.2 Complex representation	23
2.3.3 Fourier series and Fourier transforms	24
2.3.4 Discrete time sampling	26
2.3.5 The discrete Fourier transform	28
2.4 Propagation of EM radiation in a dielectric	32
2.4.1 Electromagnetic wave propagation	32
2.4.2 Electromagnetic wave propagation within dielectrics	33
2.4.3 Boundaries between dielectrics	34
2.4.4 Polarisation and permittivity	35
2.5 Conclusion	37
3 Estimation of sample refractive index and absorption coefficient	39
3.1 Spectroscopic parameters	40
3.2 Sample transfer function model	41
3.3 Overview of the extraction processing	44

3.4	Averaging and windowing	45
3.5	Deconvolution	49
3.6	Model approximations	50
3.7	Limitations of dynamic range	51
3.8	Terahertz spectroscopy of anhydrous and monohydrate forms of α - lactose	52
3.9	Phase unwrapping corrections	59
3.9.1	Unwrapping the transfer function phase	61
3.9.2	Unwrapping refractive index	62
3.9.3	Unwrapping at discontinuities in real refractive index	66
3.10	Polymorphism of paracetamol	71
3.11	Deconvolution errors	76
3.11.1	Time localised error	77
3.11.2	Systematic deconvolution error when system echoes are present	79
3.11.3	Systematic deconvolution error of sample resonance	81
3.12	Effect of thickness uncertainty	83
3.13	Thickness uncertainty of α lactose monohydrate	85
3.14	Conclusion	87
4	Creation of explosive spectrum simulants	89
4.1	Terahertz spectroscopy of Semtex-H and related explosives	90
4.2	A compound database	92
4.3	Modelling mixture spectra	92
4.4	Quantifying similarity between spectra	93
4.5	Selection and optimisation of mixtures	94
4.6	Results	100
4.7	Conclusion	102
5	Transfer function methods of extraction	105
5.1	Overview of fitting resonant sample transfer function models	105
5.2	Fitting a sample transfer function model	106
5.3	Fitting transfer function models to α lactose monohydrate	111
5.4	Sample resonance models	113
5.5	Fitting resonant models	116
5.6	Effect of thickness error on \tilde{n}	117
5.7	Thickness extraction of quartz	120
5.8	Correction for gaussian beam effects	122
5.8.1	Estimating the beam alignment coefficient	127
5.8.2	Partitioning the impulse response	128
5.9	Measuring the beam alignment coefficient	129
5.10	Extraction of z-cut quartz thickness with calibrated extraction	134

5.10.1	Consistent extraction of z -cut quartz	137
5.11	Temperature dependence of α lactose anhydrous	138
5.12	Multilayer sample extraction	144
5.13	General multilayer finite resonance model	145
5.13.1	Incorporation of the beam alignment coefficient	154
5.14	Extraction of the refractive index of IPA	156
5.15	Multi time fitting	163
5.16	Time localised finite resonance models	165
5.17	Time localised extraction of quartz	167
5.18	Time localised extraction of α -lactose anhydrous	169
5.19	Extraction of complex refractive index of a z -cut quartz and high resistivity silicon layered sample using a time localised model	173
5.20	Conclusion	176
6	Estimating uncertainty in complex refractive index and thickness	177
6.1	Introduction	177
6.2	Modelling a THz time scan	179
6.2.1	Estimating the expectation and covariance	181
6.3	Measurements of time domain uncertainty	182
6.4	Estimating covariance of extracted parameters	190
6.4.1	Sensitivity of the arithmetic mean	192
6.4.2	Sensitivity of windowing	193
6.4.3	Uncertainty of the DFT	193
6.4.4	Sensitivity of the complex logarithm, magnitude and argument	195
6.4.5	Sensitivity of deconvolution	196
6.4.6	Sensitivity of complex refractive index approximations	197
6.4.7	Sensitivity of spectroscopic parameter conversions	198
6.5	Uncertainty of α lactose monohydrate measurements	198
6.6	Estimating confidence intervals	205
6.6.1	Non-simultaneous confidence intervals	206
6.6.2	Bonferroni simultaneous confidence intervals	207
6.6.3	Effective degrees of freedom	208
6.7	Confidence intervals of α lactose monohydrate measurements	209
6.8	Fitting a transfer function model	211
6.8.1	Sensitivity matrices of sample models	212
6.8.2	Derivative of the model coefficients	213
6.8.3	Modifying transfer function models to calculate the derivative	214
6.9	Uncertainty of fitting a resonant model to z -cut quartz	216
6.10	Uncertainty of thickness extracted using total variance	221
6.11	Extracted thickness uncertainty of z -cut quartz	223
6.12	Conclusion	225

7 Conclusion and further work	227
7.1 Further work	231
Appendices	235
A Fabrication of powdered sample pellets	237
B Additional beam alignment calibration measurements	239
C Temperature dependence air volume in a 5% α-lactose anhydrous sample	241
D Multilayer Model Examples	245
E Overview of complex multivariate normal distributions	251
E.1 Bivariate normal variables	251
E.2 Multivariate Normal Variables	254
E.3 Complex multivariate gaussian distributions	256

List of Figures

2.1	Diagram of a photoconductive emitter	7
2.2	Electro-optic sampling schematic	9
2.3	The low frequency resolution broadband system	12
2.4	The high frequency resolution broadband system	13
2.5	The Narrowband THz TDS System: (A) Mechanical delay stage, (B) - Photo conductive emitter in transmission geometry, (C) - Sample, (D) - electro-optic sampling detector	13
2.6	Amplitude sampling of THz radiation	15
2.7	Implicit averaging at the detector	16
2.8	Electrically chopped THz	16
2.9	Lock-in system diagram	16
2.10	Relation between sampled and free space THz	17
2.11	Diagram of LTI system	22
2.12	Periodic bandwidth of a DTFT	28
2.13	Two sine waves at slightly different frequencies	30
2.14	Aliasing of the DFT	31
2.15	A relaxation in permittivity	36
3.1	A THz wave passing through a sample	42
3.2	Attenuation of THz wave in a sample	42
3.3	THz radiation passing through sample and passing through displaced air	43
3.4	Flowchart of extraction of n and α	44
3.5	Average of white noise	45
3.6	A Tukey window applied to a reference scan	46
3.7	Windowed reference and sample measurements	47
3.8	Cross-correlation between reference and sample measurements	48
3.9	absorption coefficient of 100% lactose anhydrous and monohydrate	52
3.10	absorption coefficient of 10 % lactose anhydrous and monohydrate	54
3.11	Different frequency resolutions of the 0.53 THz lactose monohydrate peak	55
3.12	Absorption coefficient of 100% and 10% lactose monohydrate	57

3.13	Refractive index of 100% lactose monohydrate	58
3.14	real refractive index of lactose monohydrate calculated from wrapped phase	60
3.15	Incorrect estimate of the real refractive index of high resistivity silicon	60
3.16	Unwrapping of phase	61
3.17	Extrapolation of phase offset	62
3.18	Finite difference of incorrect estimate of n	63
3.19	Corrected estimate of n of lactose monohydrate	64
3.20	Differential of n of high resistivity of silicon	66
3.21	Detection of discontinuities in n	67
3.22	Segmentation and wrapping of n of lactose monohydrate	68
3.23	Correction of n with discontinuity present	68
3.24	Two segments of real refractive index of lactose monohydrate . . .	69
3.25	Corrected estimate of n of 100% lactose monohydrate	70
3.26	Comparison of discontinuities locations to 10% lactose monohydrate	71
3.27	Absorption coefficient of 100% paracetamol forms I, II, III	71
3.28	Referenced absorption coefficient data of Paracetamol forms I, II, III	72
3.29	Comparison between published and this work's absorption coefficient of paracetamol form I.	73
3.30	Absorption coefficient of 10% Paracetamol I and II, and 20% III . . .	74
3.31	Time traces of reference and silicon measurements including reflections	79
3.32	Estimated parameters of silicon with system reflections present . . .	80
3.33	Fourier transform of the finite difference of n and κ of silicon with system reflections present	81
3.34	Time measurements of silicon plate with a sample reflection present	82
3.35	Extracted parameters of silicon in when including sample reflections	82
3.36	Fourier transform of mean centered n and κ with sample reflections present	82
3.37	Hypothetical non-linear relation between thickness and refractive index	84
3.38	σ_n of 10 % α - lactose monohydrate	85
3.39	σ_κ of 10 % α -lactose monohydrate	86
4.1	THz spectra of RDX (green), PETN (red) and Semtex-H (blue). . . .	90
4.2	Normalised THz absorption spectra of a component library	91
4.3	A diagram to illustrate how spectra is evaluated in relation to frequency difference	94
4.4	Flowchart of the genetic algorithm used in this work	95
4.5	Flowcharts of both the first and second stages of the simulant generation algorithm.	99

4.6	The four predicted spectra for the simulant mixtures are shown as solid lines. The dashed lines show the measured spectra of the mixtures once they were made, diluted with PTFE to 25% by mass and pressed into solid pellets. The solid blue line shows the normalised spectra of Semtex-H previously shown in figure 4.1 for comparison. Spectra have been normalised and offset for clarity. Red - Mixture 1, Green - Mixture 2, Magenta - Mixture 3 and Yellow - Mixture 4. . . .	101
5.1	The processing flow diagram of fitting a resonant sample model. . . .	106
5.2	α of 100 % Lactose monohydrate, using a fitted transfer function model	112
5.3	n of 100 % Lactose monohydrate, using a fitted transfer function model	112
5.4	Model of the THz beam resonating in the sample with thickness l	113
5.5	THz radiation passing through and reflecting in a sample	114
5.6	Absorption coefficient of a 0.512 mm z -cut quartz sample	116
5.7	Real refractive index of a 0.512 mm z -cut Quartz sample	117
5.8	Estimated values of n of 0.512 mm z -cut quartz with different assumed thicknesses	120
5.9	Total variation of the real refractive index of z -cut quartz	121
5.10	Total variation of the real refractive index	122
5.11	The THz focus is shifted at the detector	123
5.12	Focal point displacement when resonance occurs	126
5.13	The partitioned impulse response of quartz	129
5.14	Two different estimates of n of a quartz sample	130
5.15	Estimates of β made using different samples	131
5.16	Different estimates of β , centered at 2 THz (grey line).	132
5.17	The centered mean of the measured estimates of β	133
5.18	The fitted quadratic calibration curve (red), and the estimates of β , using a quartz sample	133
5.19	The real refractive index of z -cut quartz	134
5.20	The total variance of the quartz measurement, using different assumed thicknesses.	135
5.21	The real refractive index of z -cut quartz, under different alignment.	135
5.22	The absorption coefficient of z -cut quartz under different alignment.	136
5.23	The total variance of the extraction process for the quartz sample	136
5.24	β for the different measurements	137
5.25	real refractive index, n of two different z -cut quartz samples.	138
5.26	The measured absorption coefficient of 5 % α -lactose anhydrous.	140
5.27	Thickness of a 5 % α -lactose anhydrous sample as a function of temperature	140
5.28	The average real refractive index of α -lactose anhydrous at each temperature.	141

5.29	the temperature dependent absorption coefficient for a 5 % lactose anhydrous sample	143
5.30	The temperature dependent n of 5% α lactose anhydrous	144
5.31	The THz beam paths through a 2 layered sample and the equivalent binary tree.	146
5.32	The THz radiation passing through a M layered sample	147
5.33	The THz beam paths through a two layered sample	148
5.34	The THz beam paths through a 2 layered sample showing sub trees	149
5.35	The THz beam paths through a 2 layered sample, showing sub trees of a sub tree	150
5.36	n of IPA	158
5.37	Absorption coefficient, α of IPA.	158
5.38	n of z -cut quartz, extracted from a air filled flowcell.	160
5.39	α of z -cut quartz, extracted from an air filled flowcell.	160
5.40	The total variance of the extracted n of z -cut quartz	161
5.41	n of IPA, extracted from a IPA filled flowcell.	162
5.42	α of IPA, extracted from a IPA filled flowcell.	163
5.43	Diagram of the multilayer time localised resonance models	165
5.44	Real refractive index of z -cut quartz extracted uction different methods	167
5.45	Absorption coefficient of z -cut quartz extracted using different methods	168
5.46	The extracted thickness of z -cut quartz using different methods . . .	168
5.47	Relative thickness error in extraction without calibration	169
5.48	Extracted α of 5% α lactose anhydrous using different methods . . .	170
5.49	n of 5% α lactose anhydrous using different methods	171
5.50	Extracted thickness of a 5 % lactose anhydrous sample	172
5.51	The windowed segments of the impulse response of a high resistivity silicon and quartz layered sample.	173
5.52	The n of a high resistivity silicon and z -cut quartz layered sample. .	174
5.53	The α of a high resistivity silicon and z -cut quartz layered sample. .	175
6.1	An air time scan and its associated variance	183
6.2	Correlation matrix of the time scan	184
6.3	Magnitude of spectrum and estimates of corresponding uncertainty	185
6.4	Time shift corrected time stamps of scan	186
6.5	Normalised 10 bin histograms for each time sample	187
6.6	Normalised 500 bin histograms for each time sample estimated using a bootstrap	189
6.7	Uncertainty processing flowchart	191
6.8	Different estimates of the relative uncertainty of n of a 10% Lactose monohydrate sample measurement.	199

6.9	Different estimates of the relative uncertainty of α of a 10% lactose monohydrate measurement.	200
6.10	Different estimates of the variance of α of a z-cut quartz measurement.	201
6.11	Different Estimates of correlation of 10% α lactose monohydrate.	202
6.12	Uncertainty propagation estimates of correlation of 10% α lactose monohydrate.	203
6.13	The estimated 95% probability confidence intervals of n of a 10% α -lactose monohydrate sample.	210
6.14	The estimated 95% probability confidence intervals of α of a 10 % α -lactose monohydrate sample.	211
6.15	Different estimates of variance in z-cut quartz sample.	216
6.16	Correlation matrices of n and α of z-cut quartz.	218
6.17	Estimates of the 95% confidence intervals of n of z-cut quartz	219
6.18	Estimates of the 95% confidence intervals of α of z-cut quartz.	219
6.19	Different estimates of variation of total variation in z-cut quartz.	223
6.20	Different estimates of correlation within the total variation of z-cut quartz	224
6.21	The estimated confidence intervals for total variation of z-cut quartz.	225
B.1	Estimates of β made using different samples on the HFRBB THz TDS system.	239
B.2	Estimates of β made using different samples on the HFRBB THz TDS system.	240
B.3	Estimates of β made using different samples on the LFRBB THz TDS system.	240
C.1	The average real refractive index of the sample at each temperature, after applying a Maxwell Garnett effective medium approximation.	244
D.1	Generated beam path diagrams for the model at different stages.	248
D.2	A resonance tree through a 3 layered sample	249
D.3	A resonance tree through a 4 layered sample	249
E.1	Samples of two variables, x_0, x_1 , with different modeled correlations.	252

List of Tables

- 3.1 Table of absorption peaks in 10% lactose anhydrous and monohydrate 56
- 3.2 Table of absorption peaks in Paracetamol forms I, II and III 75

- 4.1 The four algorithm-generated mixtures, their components, mass percentages and the fitness function value of the spectra compared to the normalised spectrum of Semtex-H. Art - Artemisinin, Micro - Microcellulose, Hex - Hexamethylenetetramine, Lid - Lidocaine . . . 100

Abbreviations

DFT discrete Fourier transform

DTFT discrete time Fourier transform

FWHM full width half maximum

GaP Gallium phosphide

LTI linear time invariant

PTFE polytetrafluoroethylene

RMS root mean square

SNR signal to noise ratio

TDS time domain spectroscopy

THZ terahertz

THz terahertz

ZnTe Zinc telluride

List of Symbols

- $a(k)$ integer multiples of 2π between wrapped and unwrapped phase
- $a(0)$ $a(k)$ at $k = 0$. Equivalent to the phase offset.
- a integer multiple of $\frac{\lambda_b}{l}$
- $\angle H$ unwrapped Transfer function phase
- $\angle H'$ wrapped Transfer function phase
- $\angle H''$ unwrapped Transfer function phase with unknown phase offset
- $b(k)$ integer multiples of 2π between wrapped samples of phase
- \mathbb{F} A function of a linear time invariant system
- c free-space velocity of light
- C_{AB}^{Ws} Weighted normalised Cross-correlation between spectra A and B
- $C_{\tilde{E},e}$ Sensitivity matrix of the Fourier transform operation
- $C_{H,E_{ref}}$ Sensitivity matrix of the deconvolution with respect to the reference measurement
- $C_{H,E_{sam}}$ Sensitivity matrix of the deconvolution with respect to the sample measurement
- $C_{\alpha,\kappa}$ Sensitivity matrix for the conversion between κ and α .
- $C_{l,TV}$ Sensitivity matrix of the total variance operation with respect to H
- \log_e Sensitivity matrix of the complex logarithm
- $C_{\vec{z},\dot{z}}$ Sensitivity matrix of the magnitude and phase calculation of a complex spectrum, \dot{z}
- C_{μ,e_k} Sensitivity matrix of the arithmetic mean operation
- $C_{\tilde{n},H}$ Sensitivity matrix of extraction of \tilde{n}
- $C_{TV,n}$ Sensitivity matrix of the total variance operation with respect to n
- $C_{TV,H}$ Sensitivity matrix of the total variance operation with respect to H
- C_w Sensitivity matrix of the windowing operation, with window w
- C Cross-correlation
- \tilde{n} complex refractive index
- \bar{n} complex refractive index extracted by fitting model
- d_h Host (matrix) compound density
- d_s Sample compound density
- $\mathbb{D}(\mathbf{x})$ Diagonal matrix formed from vector \mathbf{x} .
- e natural logarithm base
- E_0 generated THz radiation

$e(k)$ discrete time domain electric field amplitude	h The impulse response.
$\hat{e}(k)$ measured discrete time domain electric field amplitude	H' Transfer function calculated with error in both reference and sample measurement
E_{ref} frequency domain reference measurement of the THz electric field	H_{δ} Transfer function of pulse in sample response with index δ
e_{ref} reference time domain amplitude of THz electric field	\bar{H} sample transfer function model
E_{sam} frequency domain sample measurement of the THz electric field	H_{sys} the system transfer function
e_{sam} sample time domain amplitude of THz electric field	H_T Transfer function of the resonance tree
$e(t)$ time domain electric field amplitude	H_w Vector of partitioned transfer functions
$\varepsilon_{\text{sam},0}$ frequency representation of the sample error shifted in time such that it is at the origin	H_w Vector of models of partitioned transfer functions
$t_{\varepsilon,0}$ time shift from origin of ε	$\hat{H}(\omega)$ Filtered sample transfer function
$t_{\varepsilon,\text{ref}}$ time shift from origin of ε	$\hat{h}(t)$ Filtered sample impulse response
$t_{\varepsilon,\text{sam}}$ time shift from origin of ε_{sam}	$\bar{H}(\omega)$ sample transfer function model
t_{ε} combined time shift from origin of reference and samples ε	i imaginary unit, $i = \sqrt{-1}$
\bar{x} expectation of multivariate variable x	k integer index
\bar{y} expectation of multivariate variable y	L Signal length, in total number of samples
\bar{z} expectation of multivariate variable z	l sample thickness
F The Fitness score between the mixture spectra and target spectra.	l' Sample thickness normalised to sample real refractive index
F_s Sampling frequency	l_c Thickness of the channel in the flowcell
$\mathbb{E}(\omega)$ Fitting error between transfer function and model	L total count. For instance number of measurements, samples, etc.
$G(l')$ Gouy phase shift correction for change as a function of index normalised sample thickness.	l_L Thickness of sample layer L
$G(z)$ Gouy phase shift at point z	l_q Thickness of the z -cut quartz windows in the flowcell
$H(\omega)$ sample transfer function	l_{rt} Thickness at room temperature
	l The shift tolerance in the weighted normalised cross-correlation
	l_t Thickness at temperature t
	\bar{l} assumed sample thickness used during extraction
	M Total number of reflections in a

time window/model	n' real refractive index estimated from wrapped phase
m integer index	n'' real refractive index estimated with unknown phase offset
M Total number of layers in a sample	$R(x)$ Round to nearest integer operation
M_N Modeled mixture spectra	R_x correlation matrix of multivariate variable x
n_a segment of n consisting of one or more samples	$r_{x,y}$ correlation coefficient between variables x and y
n_b segment of n consisting of one or more samples	x_0 scalar stochastic variable
n'_b incorrect estimate of n_b	x_1 scalar stochastic variable
n_c Real refractive index of the liquid in the channel in the flowcell	t Time
N total count. For instance number of measurements, samples, etc.	$T_{a,s}$ The transfer coefficient at the air-sample interface. The radiation is sourced from the air medium.
n_δ real refractive index estimated from H_δ	$T_{1,2}$ Transfer Coefficient for a interface between to mediums 1 and 2
n_q Real refractive index of the z-cut quartz windows in the flowcell	T_h Vector of time constraints
$\mathbb{E}_{\tilde{n}}$ Extraction error between \tilde{n} and \bar{n}	t_k kth time stamp
n_{LB} Lower wrapping bound on the real refractive index during fitting	t_{max} length of the measurement time window
n_{UB} Upper wrapping bound on the real refractive index during fitting	T_s sampling time period
P_a propagation coefficient of air over sample thickness	T_x Time period of $x(t)$
P propagation coefficient across a material	$v(k)$ error at time stamp index k
P_s propagation coefficient of the sample	V_{rt} Volume at room temperature
$R_{sam,air}$ The reflection coefficient at the sample-air interface. The radiation is sourced from the sample medium.	V_t Volume at temperature t
r difference index in cross correlation	e deterministic time scan in vector form
r_{max} cross-correlation maximum location in the difference index	e_k Measured time vector
n real refractive index	\hat{e} measured time amplitude vector
	s Standard uncorrelated multivariate normal variable
	V Matrix describing structure of covariance
	v error vector in time scan
	x multivariate variable
	y multivariate variable
	z complx multivariate variable

\hat{z} columnated complx variable and conjigate of a complx multivariate variable	uid in the channel in the flow-cell
\tilde{z} columnated complx variable and conjigate of a complx multivariate variable in Fourier symetry	α_{max} absorption coefficient max
\hat{z} columnated real and imaginary parts of a complx multivariate variable	α_q Absorption coefficient of the z -cut quartz windows in the flow-cell
w_0 Beam waist of the focused THz radiation.	β System dependent calibration coefficient for the Guassian profile effects of the THz beam in the sample model.
$w(k)$ window function	Γ_E Co-covariance matrix of the measured spectrum
$w(l')$ Amplitude correction for change in beam waist size as a function of index normalised sample thickness.	Γ_z co-covariance matrix of multivariate variable z
$W(r)$ The weighting function used in the weighted normalised cross-correlation.	$\tilde{\epsilon}$ complex relative permittivity
$w(z)$ Amplitude correction for change in beam waist size at distance z .	Σ_E Covariance matrix of the measured spectrum
$x(k)$ variable of discreate time	$\Sigma_{\tilde{E}}$ Fourier covariance matrix of the measured spectrum
$x(t)$ variable of continuous time	$\Sigma_{\tilde{e}}$ covariance of the mean time scan.
$X(\omega)$ The complex spectrum of $x(t)$	Σ_x covariance matrix of multivariate variable x
$\hat{X}(\omega)$ The complex spectrum of $x(k)$	Σ_y covariance matrix of multivariate variable y
$y(t)$ variable of continuous time	Σ_z covariance matrix of multivariate variable z
$Y(\omega)$ The complex spectrum of $y(t)$	$\Sigma_{\hat{z}}$ augmented covariance matrix of multivariate variable z
z_0 Rayleigh distance of the focused THz radiation.	α_{10} decadic absorption coefficient
z Displacement from the focus of the THz radiation at the detector, in the propagation direction.	δ_{det} RMS of the detector noise
ϵ_{ref} Additive error in the reference measurement	Δ The focal point displacement due to a sample being placed in a Gaussian THz beam
ϵ_{sam} Additive error in the sample measurement	Δ' The focal point displacement at the detector due to a sample being placed in a Gaussian THz beam
α absorption coefficient	δ pulse index in the sample response, $\delta = 0$ for initial transmission without sample reflection, $\delta =$
α_c Absorption coefficient of the liq-	

1, for first reflection	ε_k Combined error at index k
δ_s Sample volume ratio within a pellet mixture	δ_f Computational frequency resolution
$\delta_{s,m}$ Sample mass ratio within a pellet mixture.	ε'' imaginary relative permittivity
δt_h THz pulse time shift incurred by the sample	κ extinction coefficient
$\delta_V(t)$ Pellet volume change at temperature t	κ_{max} extinction coefficient max
δ_a Air volume ratio within a pellet mixture	ε' real relative permittivity
ε_0 frequency representation of the error shifted in time such that it is at the origin	σ_κ uncertainty in κ
$\varepsilon_{0,ref}$ frequency representation of the error shifted in time such that it is at the origin	σ_l uncertainty in l
	σ_n uncertainty in n
	$\mu(x)$ arithmetic mean operation
	ω angular frequency
	λ_b free space wavelength at frequencies accros segment n_b
	$\lambda(k)$ free space wavelength at frequency index k

Chapter 1

Introduction

Terahertz (THz) radiation refers to a frequency band of the electromagnetic spectrum between ultrahigh frequency radiation (up to 300 GHz) and infrared radiation. Generally the frequency range considered is between 0.1 and 10 THz. For a long time, the term "THz gap" [1] has been used widely referring to the fact that this part of the spectrum was inaccessible until relatively recently. This inaccessibility was due to the fact that THz radiation is located at frequencies between the easily accessible conventional radio frequencies and optical frequencies. To access this frequency band, electronic devices needed to either become faster and or new optical devices had to be developed.

There are several advantages of the terahertz frequency band as both an imaging and sensing medium. Notably, many materials which are opaque at other frequencies are transparent at THz frequencies [2], THz radiation is non-ionising [3] and is sensitive to both intramolecular and intermolecular vibrational modes [4].

The fact that many materials are transparent to THz radiation and that it can be used safely at low powers has led to applications in security [5], industrial control [6], non-destructive testing [7], [8] and heritage conservation [9].

Of particular interest is THz spectroscopy, the analysis of the interaction of THz electromagnetic radiation with materials. At THz frequencies a number of effects have been identified, including intermolecular and intramolecular vibrational modes [10]. Gas molecules show rotational modes at THz frequencies [11] which has led to interest in using THz gas spectroscopy for atmosphere research [12]. The sensitivity

to intermolecular modes shown by THz spectroscopy means that it is sensitive to crystal structure, so could be used as a complimentary method to X-ray diffraction to sense crystal structure changes [13]. In particular, it is sensitive to different polymorphs [14] and isomers [15] of the same compound, which has significant implications for the pharmaceutical industry, where such solid state modifications of a compound can have vastly different physical properties and impact drug delivery [16].

Another use of THz spectroscopy is the identification of materials. Drugs of abuse [5] and explosives [17] have both been shown to be identifiable by THz spectroscopy, and THz spectroscopy is foreseen as a potential means of detecting restricted compounds in security screening. Water is particularly absorbent of THz radiation, so applications involving monitoring water content have also been suggested [18].

In this work, a transmission terahertz (THz) time domain spectroscopy (TDS) is used, as it provides a coherent measurement of pulsed THz radiation through a sample [17] and is a relatively common "tabletop" method of performing THz spectroscopy. The research in this work is primarily focused on the extraction of sample refractive index, absorption coefficient and thickness, with their corresponding measurement uncertainty. The techniques developed can be applied to several applications and could be adapted to different THz spectroscopic instruments (in particular continuous wave instruments or reflection instruments) by adapting the model of sample interaction with THz radiation.

1.0.1 Objectives and structure of this work

The objective of this work is to develop algorithms which can be used to extract the real refractive index, absorption coefficient and thickness of a sample, and their corresponding uncertainty, of a sample using THz TDS.

This thesis is divided into six main chapters. The first outlines the principles of THz TDS and signal processing techniques. The second chapter covers simple extraction approximations and phase unwrapping corrections, and applies these techniques to some samples of interest. The third chapter demonstrates a method of creating mixtures which mimic the spectrum of an explosive using a mixture

generating algorithm. The fourth chapter covers more complex extraction techniques, which allow for the extraction of the sample thickness, calibration against Gaussian beam effects, characterising an unknown layer in a sample, a new method of extracting thickness and a method of extracting multiple layer properties. The final experimental chapter covers the uncertainty in measurements and a multivariate statistical method of estimating the uncertainty within parameters. Finally, a conclusion and review of this work is presented.

Chapter 2

Principles of terahertz time-domain spectroscopy and signal processing

In this chapter, the principles of THz TDS and how to extract the refractive index and absorption coefficient from a solid sample are discussed. An overview of different detection and emission devices is first discussed with an emphasis on THz time domain transmission spectroscopy. A brief overview of the operation THz TDS systems used in this work is then given. After this, a section reviewing signal processing in THz TDS and section covering propagation within dielectrics is given.

2.1 Terahertz time domain spectroscopy instruments

Development of terahertz spectroscopy has been primarily driven by the development of coherent detectors and emitters in the THz frequency range. Incoherent radiation from mercury arc lamps [19] was first used to perform spectroscopy, which was followed by the development of photo-conductive switches [20] and electro-optic emitters [21] and detectors [22]. These devices are driven by femtosecond lasers, originally dye lasers [20] were used, but have been supplemented by Ti:sapphire lasers [23] and fiber lasers [24]. These systems generate broadband pulses of THz radiation and measure the generated radiation coherently in the time domain.

Beyond the scope of THz TDS instruments, there are several other THz detection and emission technologies including photo-mixers [25], parametric sources [26], Schottky diodes [27], gas lasers [28] and quantum cascade lasers [29]. These instead generate and detect continuous radiation as opposed to the pulsed radiation used in this work. There are also two notable high powered sources of radiation, synchrotrons and free-electron lasers [30], both of which are electron beam facilities. Additionally, bolometers [31] or Golay cells [32] can be used to perform non-coherent detection for many of these sources. In this work a femtosecond laser is used with a beam splitter to both pump the emitter and sample the THz radiation, and this ensures that the sampling laser does not cause time sample jitter. However, by using two separate lasers with different repetition rates it is possible to achieve significantly faster time sampling [24]. Another variation is to use a laser with a variable THz frequency comb to enhance the frequency resolution of the THz TDS measurement [33].

Transmission instrument geometry is used within this work and is where a sample is placed in between an emitter and detector. This has the notable difficulty of being limited by the instruments dynamic range, while being simpler to align and perform reference measurements. Instead, reflection geometry can be used, where a sample is placed in front of both emitter and detector and the reflected radiation is measured. While instruments in this orientation are more capable of measuring absorbent and optically thick samples [34], they require precise alignment of sample and reference mirror.

This work is focused exclusively on the use of THz TDS instruments and in this section an overview of the operation of the emitters, detectors, laser and instrument is given.

2.1.1 Photo-conductive emission and detection

Photo-conductive switches generate THz radiation similarly to the simple Hertzian dipole [20], by creating a short transient current a far field propagating THz electric wave is emitted.

The bandwidth of this electric wave is proportional to the differential of the

current density [20] and thus a transient current of extremely short time period is required. A femtosecond laser pulse is used to generate this current by exciting electrons across the energy band gap within a photo-active semiconductor. An external electrical bias (typically 100-600 V in this work) is then used to rapidly accelerate these carriers, creating the transient current. The semiconductor is designed so that the carriers generated have an extremely short lifetime. In this work low temperature grown gallium arsenide (LT-GaAs) which is then annealed to introduce defects which shorten the carrier lifetime [35] is used. As the laser pulse width is extremely small compared to the carrier lifetime, the transient current time period is predominantly determined by this carrier lifetime and thus so is the THz bandwidth of the produced radiation.




Figure 2.1: A photo-conductive emitter, the THz radiation is shown as blue cones and the femtosecond laser is shown red line.

In Figure 2.1, the schematic of a photo-conductive emitter is shown. The switch consists of two electrical contacts which are biased with a strong DC electric field. A femtosecond laser pulse is used to excite a freespace THz pulse, which radiates from the surface towards the source of the laser and through the substrate. Collecting radiation in reflection orientation (the former) gives a broader bandwidth as the THz radiation has not propagated through the emitter [36], which will act to absorb radiation. However, more power is generally coupled into the substrate material due to its higher permittivity, leading to greater power being transmitted through the substrate [37].

The switch can also be used in reverse to time sample a THz pulse. A femtosecond

laser beam which arrives at the same time as the THz radiation is used to generate carriers within the switch. However, instead of having an external electrical bias, the THz electric field is used as the driving field for the carriers [38]. This creates a current which is proportional to the amplitude of the THz electric field. As the lifetimes of the carriers are much shorter than the THz pulse, this acts as a time sampling mechanism where both the laser pulse and THz are incident on the detector. By changing the time difference between the two pulses it is possible to time sample an entire THz waveform.

Photoconductive switches have the advantage of being relatively compact coherent high bandwidth room temperature emitters and detectors, which can be driven by "turn-key" lasers. They can also be combined with other similar emission and detection technologies, in particular electro-optic detection. For these reasons, they have become a common technology to drive THz TDS systems [34].

The patterned design of the photoconductive emitter can vary considerably and a number of designs have been attempted, including bowtie, log spiral [39], Hertzian dipole [40] and interlaced fingers [41]. However, the primary element of the design is the gap between the two electrodes. There exists a compromise between a large gap and a narrow gap. In the former, larger biases can be applied [42], increasing THz power. In the former, larger bandwidths of radiation can be produced [43].

While photo-conductive emitters and detectors are a relatively robust technology, they do have some limitations. In particular, the laser pulse energy must be absorbed by the device to emit the radiation, resulting in unwanted heating effects [43] and optical saturation of the device [44], and the electric bias can cause device breakdown [45]. Another issue, which is common to other emission and detection schemes, is that the THz radiation can resonate within the devices. This leads to reflections in the time trace which can limit the reflection-free time window.

In this work, LT-GaAs-on-Quartz devices [46] are used. These devices are LT-GaAs devices fabricated on *z*-cut quartz. The use of quartz serves two purposes, firstly it acts to increase the dark resistivity of the device reducing parasitic heating effects. Secondly, it is a mostly THz transparent material of almost arbitrary thickness which can be used to delay the emitter or detector system reflections.

2.1.2 Electro-optic sampling

Electric-optic sampling is a detection scheme for THz radiation. It is driven by a femtosecond laser which is coherent with the THz radiation (this is usually achieved by using the same or a synced laser to generate the radiation [24]). The THz amplitude is encoded into the laser pulse, which has a much shorter pulse width than the THz radiation. This in effect forms a time sample of the THz electric field, which can be changed in time by changing the arrival time of the laser pulse relative to the THz pulse. This can then be used to sample the entire THz pulse.




Figure 2.2: Electro-optic sampling detector. (A) - Electro-optic crystal, with laser and THz radiation incident on it, (B) - quarter wave plate, (C) - Wollaston prism and (D) - balanced photodiodes.

Figure 2.2 shows the electro-optic sampling scheme. In the electro-optic crystal, Figure 2.2(a), the THz radiation and sampling laser pulse co-propagate and the instantaneous amplitude of the THz radiation modulates the laser beam's polarisation by inducing a birefringence via the Pockels effect [47]. Before entering the electro-optic crystal the laser beam is linearly polarised, the electro-optic crystal exhibits birefringence proportional to the THz electric field [47]. This results in a phase change for different orthogonal components aligned with the different crystal axes and shifts the linear polarisation towards ellipticity. This ellipticity will be proportional the THz electric field and if no electric field is present, the laser beam will remain linearly polarised (provided there is no residual birefringence). This is followed by a quarter wave plate, Figure 2.2(b), which adds a polarisation bias, if no electric field is present the laser beam will be circularly polarised. In the presence of an electric field, it will be elongated than in the absence of an electric field. A Wollaston prism is used to split two orthogonal components of the beam, this is

aligned so that one represents the narrowing of the elliptical polarisation and the other represents the lengthening of the polarisation. Each of these beams are measured by a balanced pair of photo-diodes. These beams will have equal intensity if no THz radiation is present, otherwise the difference will correspond to the time sample amplitude of the THz.

The particular disadvantage of using this detection scheme, is that the THz radiation and optical probe interact over the entirety of the electro-optic crystal. This would in principle increase the sensitivity of the detector, if it were not for the different propagation speeds of the laser and THz pulses. This leads to the temporal position of the probe pulse moving relative to the THz pulse, creating a low-pass filter effect. This can be counteracted by using phase matching schemes, either by angling the crystal and the radiation directions [48] or by finding a probe beam frequency which has a similar propagation velocity to the THz radiation [49]. Another way to compensate for this effect is to use thinner crystals [50], however the THz radiation can become resonant within the detector crystal limiting the usable time window [51]. These reflection can be mitigated by using either matched delay crystals [49] or anti-reflection coatings [52]. Another consideration of thin crystals is the difficulties in manufacturing and handling [53].

Electro-optic sampling has been demonstrated with a number of electro-optic crystal materials, most notably lithium-niobate [48] which requires an angular phase matching scheme or zincblende crystals such as Zinc telluride (ZnTe) [48] or Gallium phosphide (GaP) when increased bandwidth is required [54]. For both materials, a collinear phase matching scheme can be employed with a 820 nm probe radiation [49]. ZnTe is most commonly used because of its higher sensitivity, however, it is bandwidth limited to 5 THz due to absorption modes at these frequencies [55]. GaP is lower sensitivity but has a greater detection range and has been shown to be usable up to 7 THz [54].

2.1.3 Femtosecond lasers

Femtosecond lasers are lasers which generate pulses of optical radiation with femtosecond pulse widths. These lasers are used in THz TDS spectroscopy to both gen-

erate and detect the THz radiation. In this work Ti:Sapphire lasers are used, which make use of a Ti:sapphire crystal as both a lasing material and as a non-linear lens to shape the pulses [56].

The short pulse width is required for both the generation and detection schemes. In the former, the THz radiation is generated by the femtosecond laser pulse and the bandwidth of the generated radiation is proportional to the pulse width of the laser used to drive the generation process [36]. In the sampling instances, the pulse is essentially approximated as a delta function and needs to be relatively small compared to the THz radiation to be an effective instantaneous time sample.

The wavelength of the laser is more nuanced and is generally determined by the photon energy required to generate carriers in photo-conductive devices or the phase matching conditions required in in electro-optic sampling. Fortunately, both of these are satisfied at approximately 800 nm in LT-GaAs semiconductors [57], ZnTe electro-optic crystals and GaP electro-optic crystals, which is within the Ti:sapphire lasers operational range of 650–1000 nm [58].

2.1.4 Terahertz TDS systems used in this work

In this work three different THz TDS systems are primarily used, referred to in this work as: the low frequency broadband (LFRBB THz TDS) system, the high frequency resolution (HFRBB THz TDS) system and the narrowband (NB THz TDS) system. Acronyms are used to maintain brevity during this work. These instruments have been developed by members of the THz group and represent a number of technical innovations made by these members.

The LFRBB THz TDS system has a bandwidth from 0.3 THz to 7 THz. The system is driven by a Vitara-HP (Coherent) laser with a < 20 fs pulse duration and a average power of 1 W, which is used to both generate and detect radiation. In both cases approximately 90% was used for pumping the emitter and 10% was used for the probe beam. In Figure 2.3, the LFRBB THz TDS system is shown. (A) shows the linear delay stage used to change the beam path lengths (and thus arrival times), (B) shows where a sample is placed in the system (not present in reference measurements), (C) the emitter in reflection geometry, (D) the electro-optic sampling detection scheme

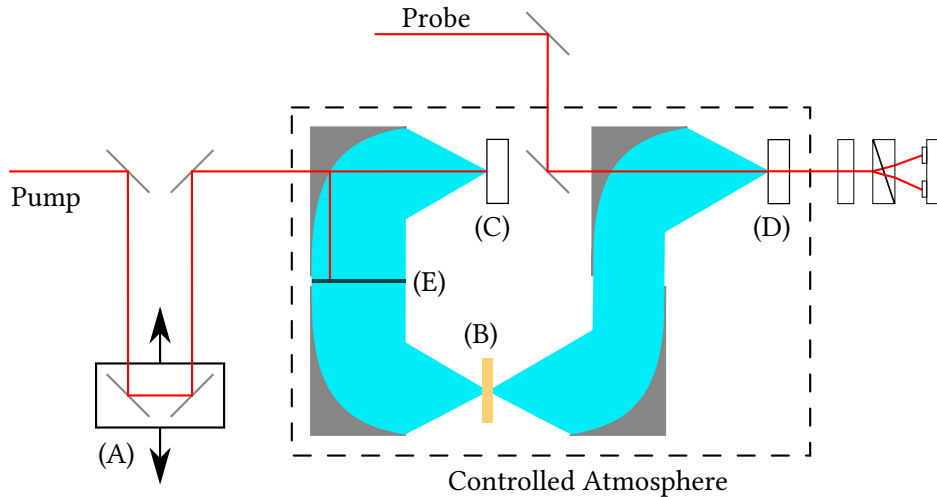


Figure 2.3: The low frequency resolution broadband system: (A) - mechanical delay stage, (B) - sample, (C) - photoconductive emitter in reflection geometry, (D) - electro-optic sampling detector, (E) - silicon beam block for reflected pump laser beam.

and (E) a silicon optical beam block. The beam block is used to block the reflected pump beam (from the emitter and parabolic mirror) from interfering with detection, and is transparent to THz radiation but opaque to the optical beam. The system is encased in a controlled atmosphere of dry air to remove the effects of water vapour from the system [59] (humidity less than 1 % for all measurements).

Unusually for a THz TDS system, the emitter is set up in reflection geometry. The radiation is collected from the frontside of the emitter and as such has a broader bandwidth. The emitter used in the version of the system is a LT-GaAs photoconductive emitter, described in [46]. The detector consists of an electro-optic sampling system, using a 0.15 mm thick GaP electro-optic crystal.

In Figure 2.4 the HFRBB THz TDS system is shown. The system is the same as the LFRBB THz TDS except the emitter and detector have been replaced with LT-GaAs-on-Quartz photoconductive devices [46] and a second delay stage (E) is present in the probe beam. In this work the second delay stage is only used for alignment.

The NB THz TDS system is shown in Figure 2.5. The laser driving both the pump and probe in this system is a Mai Tai (Spectra physics), with a pulse width of 100 fs and power of 2 W. 90% is used to drive the system and 10% of this beam used to probe the detector. The photoconductive emitter used is a LT-GaAs-on-Quartz




Figure 2.4: The high frequency resolution broadband system: (A) - mechanical delay stage, (B) - sample, (C) - Photo-conductive emitter in reflection geometry, (D) - photo-conductive detector, (E) - second delay stage, (F) - silicon beam block for reflection pump laser beam.

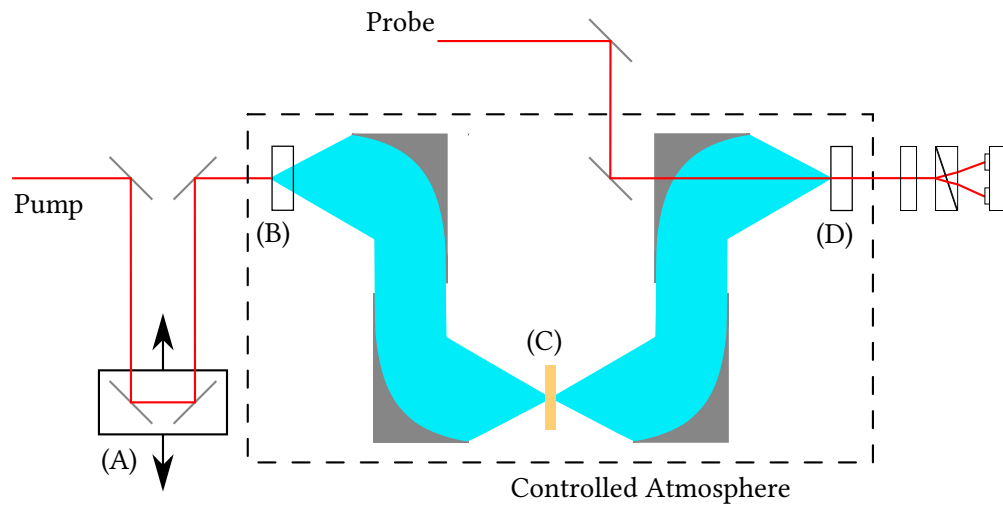


Figure 2.5: The Narrowband THz TDS System: (A) Mechanical delay stage, (B) - Photo conductive emitter in transmission geometry, (C) - Sample, (D) - electro-optic sampling detector

device. The electro-optic sampling system used in this system uses a 2 mm ZnTe electro-optic crystal. This system has a lower bandwidth of 3 THz, which is why it is referred to as the narrowband system in this work.

2.1.5 Time sampling in a terahertz time domain system

The THz TDS system samples a THz waveform, but does this by only sampling different instants of multiple generated waveform and construct a single waveform from these time samples. This relies on the generated waveforms being consistent and that the waveform does not change between samples.

The detection scheme makes use of the laser's short pulse length, effectively treating it as an impulse function, to sample the THz radiation. The probe beam will drive the detector sampling the THz radiation at the instant when the probe arrives. The time instant at which the amplitude of the THz radiation is sampled is dictated by the time difference between the pump and probe pulses being incident on the emitter and detector respectively. This time difference is changed by repositioning the optical delay stage (which will change the path length). In Figure 2.6 the relationship between the optical time difference and sampled waveform is shown. In Figure 2.6 (A), the optical arrival times are shown relative to the THz pulse, with their arrival time difference, δ_t . This relates to the sampled THz waveform in Figure 2.6 (B), where the sample time corresponds to δ_t of the lasers and thus a spatial position of the delay stage. The measured amplitude at the time sample will correspond to the amplitude of the THz waveform at the probe's arrival time.

The stage will be held at a position corresponding to the time sampling instant for a time interval much greater than the repetition rate of the laser. The output at the detector will therefore consist of a signal (optical or electrical depending on detector) which has a repetition rate of the laser. By using a detector with a much lower response time, this is implicitly averaged and will form a constant signal. In Figure 2.7, the averaging effect of the detected signal is shown. Figure 2.7 (A), the output of the detector is shown (assuming stage has constant position at each sampling point) and Figure 2.7 (B) shows the output signal (green) over the sample hold period, T_{hold} . In Figure 2.7 (C) the output of the detector is an average of individual




Figure 2.6: (A) - The amplitude of the THz radiation is sampled at the time instant dictated by the difference between pump and probe arrival times, δ_t . The pump and probe beams are shown as red beams sampling the black THz pulse for clarity. (B) - The time sampled waveform measured. The time instant corresponds to δ_t and the amplitude corresponds to the THz amplitude where the probe arrives.

detected pulses (shown as red lines) and in Figure 2.7 (D) these individual detected signals have a repetition frequency of the laser, F_{laser} , each of which corresponds to an individual generated and detected THz pulse, Figure 2.8 (E).

The bias voltage that is applied to the emitter is electrically chopped with a 7 KHz square wave. This results in the THz and thus detector output being chopped (the detector has a response time comparable to the chopping time period) see Figure 2.8, which is then fed into a lock-in amplifier.

The lock-in amplifier takes a reference from the emitter bias and the detected signal, and demodulates and applies a low-pass filter, which it then samples using an analogue to digital converter, see Figure 2.9. This is sent to a computer which records the THz scan.

To effectively sample the waveform the difference between the pump and probe arrival times can be changed, this is done using mechanical delay stages incorporated into the beam paths. This changes the path lengths of the beams relative to each other, and thus the difference in arrival times of pulses. A simple method of sampling the time waveform is thus to move the stage in a small step, measure, and repeat for the entire measurement time window. An alternate method used in this work, is to move the stage at constant velocity and sample with constant time period.

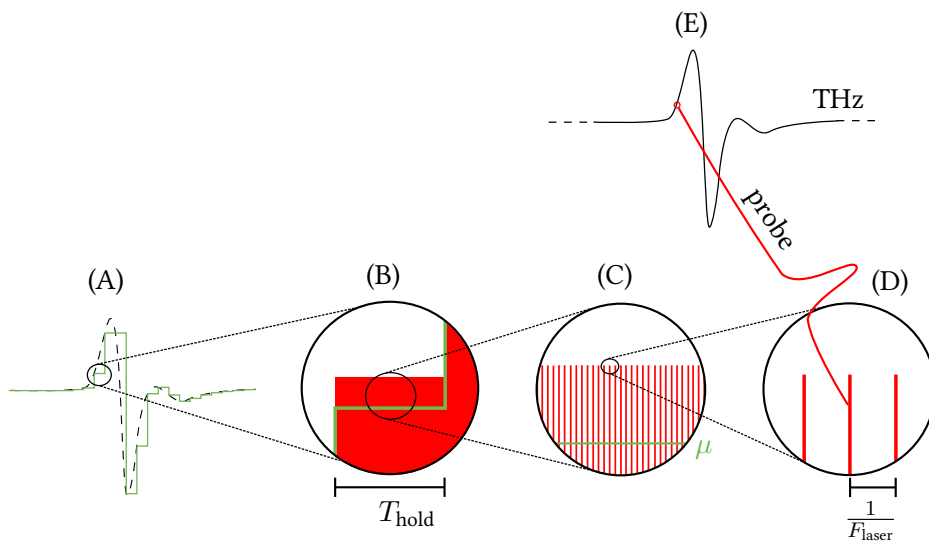


Figure 2.7: (A) - The detected waveform from the detector, (B) - The measured waveform amplitude held for sample hold time T_{hold} , (C) - The measured amplitude (green) is an average of the individual detected signal (red), (D) - The individual detected signals have a repetition frequency of F_{laser} and (E) - The individual detected signals correspond to an individual sampled THz pulses.




Figure 2.8: The driving emitter bias is electrically chopped, resulting in a chopped THz beam and chopped detector output

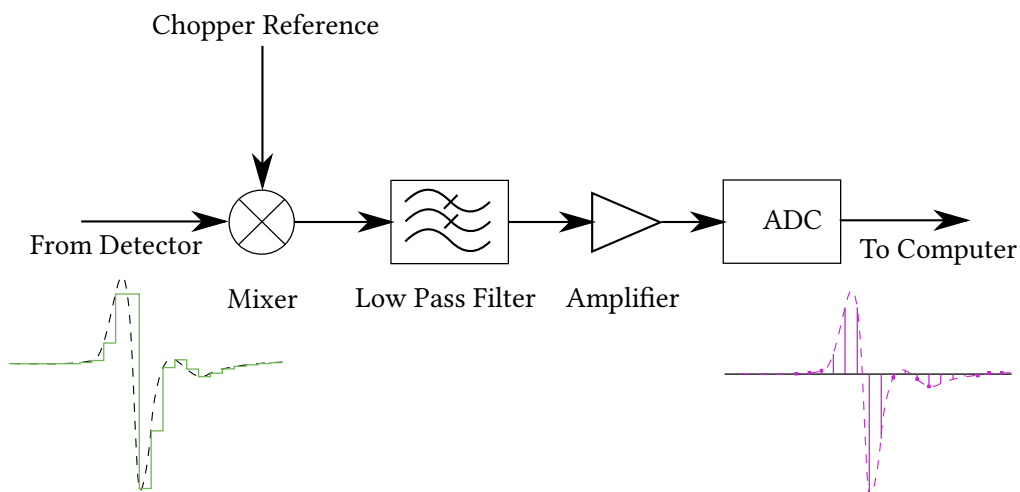


Figure 2.9: The lock-in system diagram: The input is demixed against the reference, this is then low passed filtered using a moving average filter, amplified and is then finally resampled using an ADC.

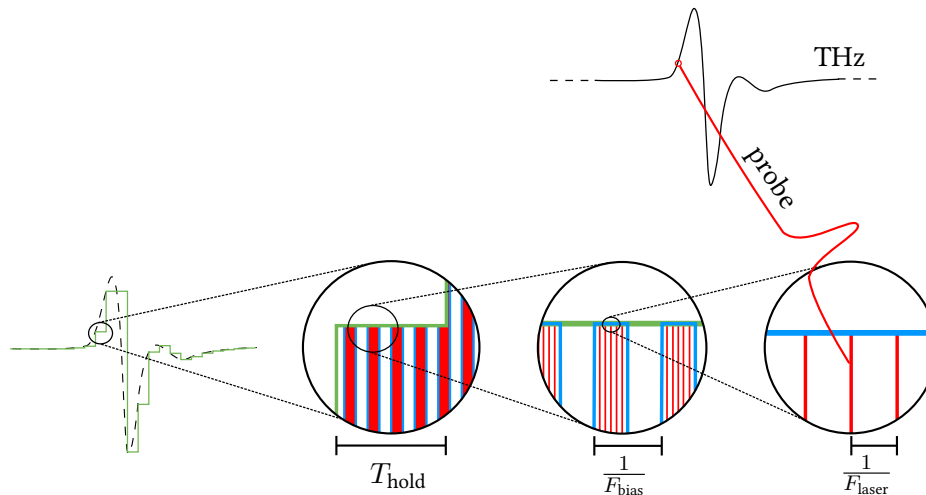


Figure 2.10: Relation between the sampled THz waveform to the free space radiation. At each time sample, the stage position is held constant (or moved slowly) for the sampling hold period, T_{hold} . This consists of a chopped waveform, chopped at F_{bias} . This chopped waveform is an average of individual measured signals with a repetition rate equal to the laser, F_{laser} . Each of the individual measured signals corresponds to one generated and detected THz pulses.

A restriction on the stage hold time or stage velocity, is the lockin's filter time. The lockin uses a moving average filter, the length of which is determined by the time constant used in the lockin settings. If the stage is moving during this time period (regardless of sampling scheme) the filter will act on the THz wave.

The sampling time period (in real time) must therefore be longer than the filter time period, to have no effect. However, the (THz) sampling period must be more than twice the bandwidth of the THz pulse, see § 2.3.4. This is dictated by the spacial separation of the stage at sampling points, so this restriction is on the velocity of the stage (and thus measurement time).

In Figure 2.10, the relation between each THz pulse and the sampled THz waveform is shown. The sampled THz waveform is constructed from millions of individual THz pulses. The sampled waveform (shown using stepped movement for clarity) has a sample time period of T_{hold} . This waveform consists of a chopped (at F_{bias}) waveform. And each of these chopped waveforms is an average of many samples taken at the repetition rate of the laser, F_{laser} .

2.2 Extraction of spectroscopic parameters in THz TDS

The processing of THz TDS signals has developed in parallel to system development. Originally, transmission spectra were used as a means of identifying water vapour in one of the earliest THz TDS systems [59]. This developed into the use of the Beer Lambert law [47] and relating the real refractive index to the change in phase.

Duvillaret *et al* [60] developed a method of treating the sample as a linear system and fitting a planar wave model to the transfer function in the frequency domain. By exploiting this method, the thickness of a sample could be extracted *in situ*. Two different methods of thickness extraction were developed, a method of matching the separated pulses of an optically thick sample [61] and a method of minimising frequency domain etalons in optically thin samples [61]. A method used in this work was developed by Dorney *et al*, who minimised etalons using a direct search of total variance [62]. A variation of this was developed using a Fourier transform to measure these etalons in a quasi space domain instead [63]. Techniques which fit transfer function models are discussed further in Chapter 5.

These methods have been shown to be very effective and are capable of extracting the thickness while mitigating the effect of etalons in the extracted parameters. The two primary issues with these methods has been (i) that there are residual etalons present after extraction and (ii) that a direct search for the thickness is required. For the former, a proposed solution to this has been to use a moving average filter which is dependent of the confidence interval of the measurements [64]. The latter issue will often require a computationally expensive iteration over a range of thicknesses. A fast method of extracting the thickness has been developed using a 3 dimensional optimisation of curve multipliers [65]. This was used to extract the thickness and then a full parameter extraction using a resonant model. While this is reported to be very effective, it relies on a low frequency resolution approximation, extracts a low resolution value and has multiple solutions. This in effect requires a good estimate of thickness and optical parameters before extraction. A method which enables direct extraction of the sample thickness, refractive index and absorption using multiple measurements of the sample with different alignments has been developed [66].

While this enabled direct extraction of the thickness of the sample, it relies on precise knowledge of the sample alignment which is often not trivial.

A similar idea to fitting to the transfer function has been to reconstruct the sample time response and minimise the difference to the measured sample response [67]. While the authors show that this method is accurate, it is noted to be slower and it is not clear what significant benefit this technique has. A similar, but arguably more beneficial idea, has been to combine a transfer function model with a parametric model of refractive index [68], [69]. Parametric methods of extraction are generally more accurate [70] provided that the model accurately represents the data. Unfortunately this typically mandates the use of a non-parametric method to estimate the model. However, a parametric method might be well suited to industrial control applications where the sample of interest is well known. Related to this has been the idea of using a Genetic algorithm to assign properties to a multiple layers of paint [71]. However, this technique would require considerable computational power to operate and is not inherently deterministic.

There have also been several tomographic techniques using the THz pulse to extract sample thickness [72]. These methods rely on using the temporal properties (equivalent to the group phase delay) to extract the thickness and the average refractive index. However, it has been suggested that temporal methods of measuring the thickness are substantially less accurate than frequency domain methods [73]. A suggested method of improving tomography techniques has been to find a sparse representation of the time domain signal, which only encodes information about interfaces [74].

Methods have been developed to extract information in the case of more complex samples. In this case, a singular layer within the sample is unknown but a complete model of the layer structure is used [64]. This is discussed further in Chapter 5. A common method of modelling a layered structure in continuous wave systems is to use a transfer matrix method [75]. This however relies on a standing wave forming through the sample, an effect which does not often occur in THz TDS. One method of dealing with this is to do partial normalisation of the sample response against another related sample. This was done in the case of [76], where a flowcell was used.

However, this is a non-general method which reduces dynamic range. Another method, which has been adapted from finite resonance models, is to estimate a finite resonance model of the multilayer sample using initial estimates [64]. Both of these methods rely on partial information about the sample, in particular only one of the layers can be unknown and all other layers must be specified in the model. Multilayer extraction is discussed further in Chapter 5.

Noise reduction methods have been developed but have primarily focused on white noise reduction. In particular, simple averaging of repeated measurements has been identified as a principle means of reducing the zero mean noise in a measurement [77]. Convolutional Gaussian frequency filters [78] and Wiener deconvolution [79] have been used filter out noise. The primary reason these are not used in this work is that these improve the time domain measurement by filtering the frequency domain. As the analysis is primarily performed in the frequency domain in this work, this is both detrimental and not required. In addition, there are a number of assumptions about the noise when applying a Wiener filter (particularly stationary noise [70]), which were not found to be valid in the measurements in this work, see Chapter 6.

A different technique has been the use of wavelet thresholding [79]. In this technique, a wavelet transform is applied to the time domain measurements and then a subtractive threshold is applied. The theory is that the THz pulse is well compressed into intense coefficients by this process but the noise is not, so applying a subtractive threshold removes more noise than signal. There are several issues with this, the selection of wavelet transform parameters can have a large effect on the performance, there is substantial signal dependent noise (which is well compressed) and the filtering process can introduce distortion into the process. A more recent related method has been empirical mode decomposition [80], which makes use of both a transformed signal and reference noise to perform the filtering. However, since the reference noise is taken in absence of the sample, it likely has similar issues to wavelet thresholding.

A concept related to noise removal is uncertainty analysis. In this case, noise is simply a source of measurement uncertainty which can be characterised. With-

ayachumnankul *et al* [81] used a similar process to the method in this work (see Chapter 6) based on univariate statistics and uncertainty propagation. The primary issue with this method, is that it assumes independence between time and frequency samples, which is not true of the measurements in this work. Krüger *et al* [82] extended this methodology to derive a term for fitting models, which used a similar methodology to that recommended in the Guide to Expression of Uncertainty in Measurement [83], to project the uncertainty onto the error and then onto the extracted parameters. There also exists an older method [61], which is less commonly cited in the literature, which makes use of a calibrated parametric model to estimate the uncertainty in complex refractive index, \tilde{n} . However this is dependent on calibration measurements. This is discussed further in Chapter 6.

While Kramers–Kronig methods are not required to calculate either the real refractive index or absorption of a sample in THz TDS, some authors have devised methods which attempt to exploit approximations of these transforms. Bernier *et al.* [84] devised an iterative method of improving the parameters, by recursively applying the Kramers–Kronig relations to improve extracted parameters. Another more recent use, has been to use Kramers–Kronig relations to extract the thickness of a water channel in a flowcell structure [85], using a similar process to total variance. Peiponen *et al.* has done extensive work on the use of Kramers–Kronig relations in THz spectroscopy, in particular as a method of aligning reflection systems [86], detection of scattering effects [87], consistency checks on extracted parameters [88] and extraction of spectroscopic parameters in the absence of measured thickness [89].

The Fourier transform is known to over parametrise (as it fits to multiple frequencies), so it is common in signal processing to use a time domain model to extract a frequency response which gives an improved frequency response [70]. Tych *et al.* used a continuous time domain model to extract the frequency response of a sample [90]. However, identification of a suitable model can be difficult and will be sample dependent. In particular, as the spectrum of a sample can consist of dozens of absorption modes, extremely high order time domain models may be required.

A different method of using THz radiation is to use empirical analysis to iden-

tify sample properties. In this case, distinguishing features of a THz spectrum are identified for known training samples, this is then used to identify features in future samples. This can be done using spectroscopic parameters [91] or just using the frequency spectrum [92]. In the latter, while simpler, it introduces a dependence on the instrument.

2.3 The transfer function and Fourier transform

Throughout this work, the transfer function of the spectroscopic sample is often referred to. In this section, the concept of an impulse response is introduced. This is followed by the Fourier transform, which is then applied to the impulse response to derive the transfer function. A transfer function is a quantitative measure of a system, which in this case is the spectroscopic sample. The practical aspects of time and frequency sampling are then discussed, which are expanded into the concept of frequency resolution. Frequency resolution is currently a limitation within the THz TDS measurements due to reflections within the time trace.

2.3.1 Linear time invariant systems

A sample within a spectroscopic system can be considered a linear time invariant (LTI) system acting on the THz radiation. A LTI system is a system with a linear relationship between its input and output which does not change over time [93]. In this work, the input is considered to be a reference THz TDS measurement and the output is considered to be a sample THz TDS measurement.




Figure 2.11: A LTI system, which relates input $x(t)$ and output $y(t)$ via function \mathbb{F} . In the diagram above, a THz pulse is the input and a sample THz measurement is the output. The sample is the LTI system which transforms the THz pulse to the sample response.

In Figure 2.11 a LTI system is shown. The input consists of a continuous variable of time, $x(t)$, and the output consists of second variable of time, $y(t)$. The relation

between the two is the function \mathbb{F} which represents the system. In the case of a LTI system this can always be written as convolution operation [93]:

$$y(t) = \mathbb{F}(x(t)) = (x * h)(t) = \int_{-\infty}^{\infty} x(\tau)h(t - \tau)d\tau \quad (2.1)$$

where h is the impulse response of the system and represents the case where the input is an impulse function. A way of quantitatively describing the spectroscopic sample's interaction with the THz radiation is to use its impulse response.

The system is considered linear because convolution is both additive and scalar multiplicative [93]:

$$((ax) * h)(t) = a(x * h)(t) \quad (2.2)$$

$$((x_1 + x_2) * h)(t) = (x_1 * h)(t) + (x_2 * h)(t) \quad (2.3)$$

It is also considered time invariant as any shift in time of the input leads to a corresponding time shift of the output:

$$y(t + a) = (x * h)(t + a) \quad (2.4)$$

2.3.2 Complex representation

A variable is said to be complex if it has both a real and imaginary part:

$$z = a + ib \quad (2.5)$$

where z is a complex complex variable and both a and b are real variables. i is the imaginary unit and is the square root of minus one, $i = \sqrt{-1}$.

The conjugate of a complex variable is defined as:

$$z^* = a - ib \quad (2.6)$$

Thus another definition of a is:

$$a = \Re(z) = \frac{1}{2}(z + z^*) \quad (2.7)$$

and b :

$$b = \Im(z) = \frac{i}{2} (z - z^*) \quad (2.8)$$

The exponential of a complex variable is equal to two complex sinusoids:

$$e^z = |z| e^{i\angle z} = |z| \cos(\angle z) + i |z| \sin(\angle z) \quad (2.9)$$

where $|z|$ is the magnitude and $\angle z$ is the angle or phase of the complex variable.

The magnitude is defined as:

$$|z| = \sqrt{a^2 + b^2} \quad (2.10)$$

and the angle is defined as:

$$\angle z = \tan^{-1} \left(\frac{b}{a} \right) \quad (2.11)$$

A method of representing sinusoids is to use complex numbers:

$$\cos(\theta) = \Re \left(e^{i\theta} \right) = \frac{1}{2} \left(e^{i\theta} - e^{-i\theta} \right) \quad (2.12)$$

2.3.3 Fourier series and Fourier transforms

Fourier transforms are used to convert THz TDS measurements from the time domain to the frequency domain. A brief overview of the Fourier series and the analytical Fourier transform is presented in this subsection. This is then related to the concept of a transfer function of a LTI system.

A function which is periodic can be represented as a summation of sinusoids. Consider the signal $x(t)$, which is periodic, such that:

$$x(t) = x(t + kT_x) \quad (2.13)$$

where k is an arbitrary integer and T_x is the time period of $x(t)$. If this signal has

finite energy, it can be expressed as a Fourier series [70]:

$$x(t) = a_0 + \sum_{k=1}^{\infty} a_k \cos(k\omega_0 t + \theta_k) \quad (2.14)$$

where $\omega_0 = \frac{2\pi}{T_x}$, a_k are the Fourier coefficients and θ_k is a phase offset for each frequency.

The Fourier series can then be expressed in terms of complex exponentials and negative frequencies:

$$x(t) = \sum_{k=-\infty}^{\infty} z_k e^{ik\omega_0 t} \quad (2.15)$$

where z_k is a complex Fourier coefficient and z_{-k} is its conjugate.

z_k can always be found by performing an integration over the period T :

$$z_k = \frac{1}{T_x} \int_{t_0}^{T+t_0} x(t) e^{-ik\omega_0 t} dt \quad (2.16)$$

where t_0 is an arbitrary start point for the integration.

This can be expanded to include the case where $T_x \rightarrow \infty$, i.e. for non periodic functions. Under such case, the series coefficients converges on a continuous frequency *spectrum* [70]:

$$X(\omega) = \mathcal{F}(x(t)) = \int_{-\infty}^{\infty} x(t) e^{-i\omega t} dt \quad (2.17)$$

where \mathcal{F} represents the Fourier transform. Conversely, a function of time can be expressed as the inverse Fourier transform of a spectrum:

$$x(t) = \mathcal{F}^{-1}(X(\omega)) = \int_{-\infty}^{\infty} X(\omega) e^{i\omega t} d\omega \quad (2.18)$$

where \mathcal{F}^{-1} represents the inverse Fourier transform.

The Convolution theorem can be used to represent convolution in the time domain as multiplication in the frequency domain [93]:

$$\mathcal{F}^{-1}(X(\omega)Y(\omega)) = \mathcal{F}^{-1}(\mathcal{F}(x(t))\mathcal{F}(y(t))) = (x * y)(t) \quad (2.19)$$

This can be used with regards to a LTI. The output spectrum $Y(\omega)$ can instead be represented as multiplication of the input spectrum $X(\omega)$ with the transfer function H :

$$Y(\omega) = X(\omega)H(\omega) \quad (2.20)$$

The transfer function is thus the frequency domain equivalent of the impulse response. This can be used to deconvolve an impulse response of a system with known input and output [93]:

$$h(t) = \mathcal{F}^{-1}(H(\omega)) = \mathcal{F}^{-1}\left(\frac{Y(\omega)}{X(\omega)}\right) \quad (2.21)$$

The description of the spectroscopic sample is spectroscopically described in terms of frequency, so applying a Fourier transform is generally required to extract these parameters. Additionally, using Fourier transforms represents a way of deconvolving the impulse response. In this work the frequency domain equivalent, the transfer function, is generally used when describing the sample response.

2.3.4 Discrete time sampling

The THz time trace is not measured continuously, rather it is discretely sampled at regularly intervals. This results in important limitations on the bandwidth of the spectrum, in particular aliasing can occur if the scan is not sampled at a high enough frequency. This sub section briefly outlines discrete time sampling a variable which is a function of time.

It is not realisable to measure a variable over time continuously, instead the variable must be sampled at discrete intervals. This is modelled as being instantaneous samples of a variable $x(t)$:

$$x(k) = x(t_k) \quad (2.22)$$

where $x(k)$ is the time sample of $x(t)$ at time stamp t_k . It should be noted that both k and t_k are infinite sequences for the purposes of analysis. In this work $x(t)$ is

considered bandwidth limited, such that its spectrum:

$$X(\omega) = 0, \quad |\omega| > \omega_c \quad (2.23)$$

where ω_c is the cut-off frequency. Under these circumstances it is preferable to use a periodic sampling:

$$x(k) = x(kT_s) \quad (2.24)$$

where T_s is the sampling period. Using this scheme it is possible to accurately measure bandwidth limited variables. However, the sampling frequency, F_s , must be greater than or equal to twice the highest frequency of $x(t)$ to prevent aliasing [93]:

$$F_s \geq 2f_c \quad (2.25)$$

where $f_c = \frac{\omega_c}{2\pi}$.

For a discrete time signal the integral is replaced with a summation to produce the discrete time Fourier transform (DTFT) [70]:

$$\hat{X}(\omega) = \mathcal{F}(x(k)) = \sum_{k=-\infty}^{\infty} x(k)e^{-i\omega kT_s} \quad (2.26)$$

Where $\hat{X}(\omega)$ is the *continuous* spectrum of the infinite sequence $x(k)$ and is an analytical transform. It is related to the spectrum of $x(t)$, $X(\omega)$ by [70]:

$$\hat{X}(\omega) = F_s \sum_{k=-\infty}^{\infty} X_c(\omega - k2\pi F_s) \quad (2.27)$$

In the case of a bandwidth limited signal this results in the spectrum being periodic with sampling frequency F_s . Shown in Figure 2.12 is an illustration of the periodicity of THz spectrum. The spectrum is both continuous and periodic with the sampling frequency, $F_s = 25$ THz. It should be noted that this is actually a discrete sampled spectrum (see the next section).




Figure 2.12: Periodic bandwidth of a DTFT. In this case, the THz spectrum is periodic with the sampling frequency of 25 THz.

2.3.5 The discrete Fourier transform

A continuous frequency spectrum is not practically realisable for the same reason a continuous time variable is not. Instead, the frequency domain must be sampled. A regularly sampled spectrum is defined as:

$$X(m) = X(m\delta_f) \quad (2.28)$$

where m is the frequency index and δ_f is referred to as the computational resolution (see below). This will produce instantaneous samples in the frequency domain. To synthesize time samples, instead of performing integration, summation is used [70]:

$$\hat{x}(k) = \frac{1}{N} \sum_{j=-\frac{N}{2}}^{\frac{N}{2}-1} X(m) e^{i2\pi\delta_f mk} \quad (2.29)$$

This equation has the same form as equation 2.15 and is thus a finite Fourier series. This also means that the synthesised $\hat{x}(k)$ is periodic with $\frac{1}{\delta_f}$, which equates to N time domain samples.

Within this period of N samples the synthesised samples can be related to the original time samples [70]:

$$\hat{x}(k) = x(k), \quad 0 \leq k < N \quad (2.30)$$

Outside of this range $\hat{x}(k)$ will repeat.

This provides a method of producing finite Frequency samples from finite time samples. For the finite time samples, $x(k)$, of length N , there is the finite frequency spectrum, $X(m)$, which is related to the discrete Fourier transform (DFT) [70]:

$$X(m) = \sum_{k=0}^{N-1} x(k) e^{-i \frac{2\pi}{N} mk} \quad (2.31)$$

It is conventional that the index m extends between 0 and $N - 1$ as the spectrum is periodic, so indices above $\frac{N}{2}$ are equal to the negative indices. There also exists a inverse DFT [93]:

$$x(k) = \sum_{m=0}^{N-1} X(m) e^{i \frac{2\pi}{N} mk} \quad (2.32)$$

In this work, an analytical Fourier transform is not used. Instead a DFT is used on the THz time domain scans. This produces corresponding frequency samples of the THz spectrum. Processing is performed on a set of discrete frequencies, although analytically it is treated as a continuous variable. The measured bandwidth is limited by the sampling frequency and the frequency resolution is limited by the time window used. The latter is often an experimental limitation, due to reflections in the time scan.

It is possible to increase the number of frequency samples, by a process known as zero padding:

$$x_0(k) = \begin{cases} x(k) & \text{if } 0 \leq k < N \\ 0 & \text{if } N \leq k < L \end{cases} \quad (2.33)$$

In this case $x(k)$ has been expanded to a L length sequence by appending zeros. A L length discrete spectrum can then be calculated using the DFT. This process is equivalent to up-sampling and applying a smoothing filter in the frequency domain. Thus no new information is added by this process, instead additional frequency samples are interpolated. It is for this reason that δ_f is referred to as the computational frequency resolution.

The DTFT of a finite $x_w(k)$ of length N is equivalent to infinite zero padding ($L = \infty$). This is distinct from $x(k)$ and is often viewed as sampling $x(k)$ over a

finite length window. In this case $x_w(k)$ is:

$$x_w(k) = x(k)w(k) \quad (2.34)$$

where $w(k)$ is a window function. The simplest window function is a rectangular window function of length N :

$$w(k) = \begin{cases} 1 & , \quad 0 \leq k < N \\ 0 & , \quad \text{otherwise} \end{cases} \quad (2.35)$$




Figure 2.13: Two sine waves at slightly different frequencies

The use of a time window leads to a limited frequency resolution and to frequency leakage (in the time domain this is the sharp truncation of time domain features) [93]. The use of a finite period of time leads to uncertainty in the frequency of a component. In Figure 2.13 two sinusoids are shown with similar periodicity are shown. As time progresses they diverge, and are easily distinguishable. Therefore as the time window length is increased, the uncertainty of the frequency of a sinusoid is decreased. This is fundamentally different from the computational frequency resolution, as it reflects the limit of the time window used to accurately resolve different frequencies. The simplest method of defining frequency resolution is to take the reciprocal of the time window length, $\frac{1}{T_w}$ [93]. This is not a “hard” resolution, it rather reflects the minimum width of the window function in the frequency domain which “smears” the frequency spectrum of the non-windowed time domain signal.

Another consideration is that time signals of finite length can not be enhanced by increasing the time window length.

The DFT has adjustable computational frequency resolution (via zero padding), increasing this will not change the frequency resolution of the time window. However, it is still important to consider the computational frequency resolution because aliasing can occur in the frequency domain [93]. This is most prominent when considering the magnitude or phase of the spectrum.




Figure 2.14: (A) - spectrum magnitude, (B) - spectrum phase and (C) - unwrapped spectrum phase. Blue - zero-padded spectrum and Red - non zero-padded spectrum

The magnitude of the complex spectrum can often consist of peaks which are sampled in the case of the DFT. This sampling however does not have to align with the peaks. This can result in features appearing to be considerably different. In Figure 2.14(A), the non zero-padded spectrum is shown as a red line and the zero-padded spectrum is shown as a blue line. The non zero-padded spectrum appears to be a singular wide peak, however zero-padding has revealed a much narrower peak (frequency resolution limited by the time window used) with considerable frequency leakage (the side lobes).

Another common case of aliasing is with the phase of a spectrum. This is extremely important when considering phase unwrapping. The phase is assumed to be a continuous line, if the sampling period steps over the wrapping discontinuities it will be incorrectly unwrapped. Fortunately aliasing can be prevented by increasing the computational frequency resolution. In Figure 2.14(B) the non zero-padded spectrum is shown as a red line and the zero-padded spectrum is shown as a blue line.

This has been unwrapped in Figure 2.14(C) (see § 3.9) but have produced dramatically different unwrapped phases.

The unwrapping of the phase is particularly important in this work as the spectroscopic parameters extracted in this will be dependent on the unwrapped phase. Thus in this work zero-padding has been used to produce a computational frequency resolution of 10 GHz.

2.4 Propagation of electromagnetic radiation within dielectric materials

In this section the propagation of electromagnetic waves in a dielectric medium is discussed. A simple singular frequency propagation model is discussed in the context of dielectric materials, where the effect of refractive index is discussed. This model is expanded to the idea that there exists a complex multiplier for distance propagated in a dielectric medium. In the subsequent chapters, this concept is used as a basis for propagation within a sample and within air. Boundary conditions are briefly discussed, which provide a simplified basis for the relationships between the reflected, transmitted and incident waves at an interface. Finally, a brief explanation of the features observed in refractive index and absorption is presented. The characteristic shape of a relaxation is also exploited to form an unwrapping correction to the real refractive index, see § 3.9.

2.4.1 Electromagnetic wave propagation

In this work, the propagation of electromagnetic (EM) radiation through a dielectric is considered. The simplest EM wave is a sinusoidal wave [47], which consists of an sinusoidal electric wave with a perpendicular magnetic wave. The magnetic wave is always present, however as only dielectrics are considered in this work, it is assumed to not interact with the materials. Instead the electric wave is of particular importance [47].

The simplest electric wave consists of a sinusoid:

$$e(z, t) = e_0 \Re \left(e^{i\omega t - \frac{2\pi}{\lambda} z + \theta} \right) \quad (2.36)$$

where ω is the angular frequency of the wave, t is the time axis, z is the propagation direction, $e(z, t)$ is the electric wave, e_0 is the amplitude of the wave, θ is the phase of the wave and λ is the wavelength. Using the Fourier transform, an arbitrary shaped wave can be viewed as a superposition of these waves at different frequencies. These waves are often referred to as plane waves, as the phase of the wave is consistent within a plane perpendicular to the propagation direction. In this work, most of the models assume that the waves take this form.

However, later in the work, § 5.8, a slightly more complex propagation mode is considered. This is the Gaussian propagation mode, which has a well defined spatial distribution. In this case, perpendicular to the propagation direction, the intensity drops off with a Gaussian profile (defined by the Gaussian probability distribution of the same name). In the region of the focus, a Gaussian profile beam is well approximated by a plane wave.

2.4.2 Electromagnetic wave propagation within dielectrics

The phase velocity of a wave is:

$$v_e = f \lambda \quad (2.37)$$

Within a dielectric, EM waves propagate at lower velocity. The ratio between the velocity within a material and a vacuum is the real refractive index [47]:

$$n = \frac{c}{v_e} = \frac{\lambda_0}{\lambda} \quad (2.38)$$

where c is the phase velocity in a vacuum (or more concisely the velocity of light) and λ_0 is the wavelength in a vacuum. The frequency does not change within the dielectric, thus the wavelength changes.

Within a dielectric electric waves can be absorbed, this can be modelled by using

damped sine waves as a basis [47]:

$$e(z, t) = e_0 e^{-\frac{\omega \kappa z}{c}} \Re \left(e^{i(\omega t - \frac{2\pi}{\lambda} z + \theta)} \right) \quad (2.39)$$

where the extinction coefficient, κ , is the frequency dependent extinction coefficient. This can be rearranged to include the frequency dependent complex refractive index, \tilde{n} :

$$e(z, t) = e_0 \Re \left(e^{i(\omega t - \frac{\omega \tilde{n}}{c} z + \theta)} \right) \quad (2.40)$$

where $\tilde{n} = n - i\kappa$. This can be used to derive the propagation coefficient in § 3.2.

2.4.3 Boundaries between dielectrics

At an interface between two dielectrics, a boundary condition occurs. The wave inside one material will propagate differently to that in another material. Across the boundary both the electric and magnetic wave must be consistent between the two materials, for a wave incident at the interface, part of the wave is reflected to satisfy this constraint.

The wave impedance is the ratio between the magnetic and electric waves within a dielectric [94]:

$$z = \frac{E}{B} = \frac{1}{\tilde{n}} z_0 \quad (2.41)$$

where z_0 is the wave impedance of free space. The waves at the interface can be related, for the electric field:

$$B_i + B_r = B_t \quad (2.42)$$

and magnetic field:

$$B_i - B_r = B_t \quad (2.43)$$

where B_i and B_r are the electric and magnetic waves incident on the interface, E_t and B_t are the electric and magnetic waves reflected and E_t and B_t are the transmitted electric and magnetic waves. This assumes that the waves are propagating perpendicular to the interface.

Using the wave impedance and these relations, the complex ratios between in-

cident and transmitted, and between the incident and reflected waves can be found [94]. These form the Fresnel reflection and transmission coefficients [47], which are used in § 3.2 to construct models of the sample transfer function.

2.4.4 Polarisation and permittivity

A dielectric can be considered to consist of various charges bound in different ways. For instance, charge will often be distributed around a crystal lattice. When an electric field is applied over a portion of the dielectric, a force is exerted on these charges which will displace the charges from their rest state and an opposite restoring force will occur.

The displacement of these charges relative to the neutral positions induces a electric dipole moment. This can be measured by the induced polarisation [47]:

$$P = xq \quad (2.44)$$

where x is the charge displacement, and q is the charge.

For a dielectric, the polarisation can be related to the electric field by the permittivity [47]:

$$P = (\tilde{\epsilon} - 1) \epsilon_0 E \quad (2.45)$$

where $\tilde{\epsilon}$ is the complex relative permittivity and ϵ_0 is the permittivity of free space. The real part of the permittivity relates to the change in polarisation and the imaginary part relates to the frequency dependent loss of energy within the material.

A propagating electric wave in a dielectric is coupled to the polarisation, thus the permittivity acts to slow the wave down. The permittivity can then be related to the complex refractive index [47]:

$$\tilde{\epsilon} = \tilde{n}^2 \quad (2.46)$$

Thus, the underlying polarisation mechanism (at a particular frequency) will often dictate the value of the permittivity and thus refractive index at that frequency. By modelling the polarisation as a function of frequency, it is possible to identify the

refractive index as a function of frequency.

The bound charges can effectively be modelled by charged masses on the end of a spring, i.e. a simple harmonic oscillator. This can be viewed as being driven by the propagating electric field. The oscillator model can be viewed as a summation of the driving electric force, a restorative force (the metaphorical spring) which is proportional to the displacement, a damping force which is proportional to the mass velocities and Newtons second law which equates force to being proportional to acceleration.

$$m \frac{d^2 x}{dt^2} + a \frac{dx}{dt} + bx = -qE(t) \quad (2.47)$$

This can be converted to polarisation:

$$\frac{d^2 P}{dt^2} + a \frac{dP}{dt} + bP = \frac{q^2}{m} E(t) \quad (2.48)$$

Which can be solved by applying a Fourier transform and rearranging:

$$P(\omega) = \frac{q^2}{m} \frac{1}{\omega_0^2 - \omega^2 + ia\omega} E(\omega) \quad (2.49)$$

where $\omega_0 = \sqrt{b}$. The permittivity can then be said to be [47]:

$$\tilde{\epsilon}(\omega) \propto \frac{1}{\omega_0^2 - \omega^2 + ia\omega} + 1 \quad (2.50)$$




Figure 2.15: A relaxation in permittivity, blue - real permittivity and green - imaginary permittivity

The permittivity (and the underlying polarisation) can be seen as dictating the

complex refractive index as a function of frequency. An underlying oscillator has a natural frequency, which results in a peak in the absorbance of the material (imaginary parts of \tilde{n} and $\tilde{\epsilon}$) which can be fitted with a Lorentzian function and in the real parts a relaxation, a sharp discontinuity within absolute value (this term is not used to refer to dielectric relaxation processes within this work). In Figure 2.15, both the real and imaginary parts of permittivity are shown when this model is used. The imaginary part is shown in green and has a strong Lorentzian peak, and the real part is shown in blue and has a peak followed immediately by a valley.

A dielectric will generally have many polarisation processes and can typically be considered a superposition of different harmonic oscillators. A substantial number of other effects can affect the measured permittivity. In particular, the thermal energy of the material will affect the oscillation mechanisms leading to homogeneous broadening of relaxations [95] and various scattering processes will increase the apparent losses or result in non-homogeneous broadening [96]. In molecular materials at THz frequencies these are generally intermolecular and are non-trivial to assign, requiring detailed modeling and calculation of the absorption modes [97]. While this work does not generally assign physical features to the spectra measured, these relaxation processes can be observed throughout the work and in § 3.9, the shape of a relaxation is exploited to correct the real refractive index.

2.5 Conclusion

In this chapter, a brief overview of the principles required to understand the work in later chapters is presented. The principles of the THz TDS system have been presented, in particular how noise is reduced and time sampling is achieved. The discrete Fourier transform and the transfer function were introduced, both of which are used later to form estimates of material parameters. Finally propagation of the THz radiation within a dielectric is discussed from which a model of THz interaction with a sample is derived.

Chapter 3

Estimation of sample refractive index and absorption coefficient

This chapter consists of a definition of complex refractive index, real refractive index, extinction coefficient, absorption coefficient and dielectric permittivity, and how they interrelate. A brief overview of the processing steps required to approximate them is presented. It is then discussed how these relate to THz propagation through a sample, and a transfer function sample model is then constructed. This model is used to form approximations to calculate these parameters. In the example of α -lactose monohydrate phase unwrapping issues are identified. A simple method of correcting phase unwrapping is then presented, and is demonstrated on the example of α -lactose monohydrate. Three different polymorphic forms of paracetamol are then measured, and the different absorption spectrums between 0.3 and 6 THz are identified. A number of challenges relating to the measurement of samples in THz TDS systems are highlighted, in particular the effect of errors localised in time and of sample thickness uncertainty. Thickness uncertainty is currently the largest source of systematic error in extracted values of refractive index and absorption coefficient from THz TDS measurements [81].

3.1 Spectroscopic parameters

THz radiation which propagates through a sample will be both attenuated and delayed in time, both of which can be related to the optical and dielectric parameters. In this section, the frequency dependence of these parameters is explicitly highlighted. However, in later sections this has merely been implied to maintain brevity.

The parameter used within transfer function models is \tilde{n} , which is a combination of n and κ :

$$\tilde{n}(\omega) = n(\omega) - i\kappa(\omega) \quad (3.1)$$

where ω is the angular frequency and i is the complex unit. The real refractive index, n , is a unitless measure of phase delay (velocity change) with respect to a vacuum, and the extinction coefficient, κ , is a unitless measure of amplitude attenuation per unit length.

The absorption coefficient, α , is used preferentially as a measure of absorption and is related to the extinction coefficient through:

$$\alpha(\omega) = \kappa(\omega) \frac{2\omega}{c} \quad (3.2)$$

where c is the free-space propagation velocity of light. An alternative variation, the decadic absorption coefficient, α_{10} is often used, which can be related to α by [98]:

$$\alpha_{10}(\omega) = \alpha(\omega) \log_{10}(e) \quad (3.3)$$

where e is the natural logarithm base.

Sample models, used later in this work, treat the sample as a dielectric interacting with an electric field. These models are based on \tilde{n} , which can be related to the complex (relative) dielectric permittivity, $\tilde{\epsilon}$, by:

$$\tilde{\epsilon}(\omega) = \tilde{n}(\omega)^2 = (n(\omega) - i\kappa(\omega))^2 \quad (3.4)$$

Like \tilde{n} , this can be split into a real and imaginary parts:

$$\tilde{\varepsilon}(\omega) = \varepsilon(\omega)' - i\varepsilon(\omega)'' = (n(\omega)^2 - \kappa(\omega)^2) - i(2n(\omega)\kappa(\omega)) \quad (3.5)$$

where ε' is the real permittivity and ε'' is the imaginary permittivity. ε' is the frequency dependent ratio between capacitance of a dielectric and capacitance of vacuum and ε'' is a measure of how much energy is lost from the applied electric field in the dielectric material.

3.2 Sample transfer function model

With the aim of calculating the spectroscopic parameters, a model relating the parameters to the measured transfer function was developed, and such models can be found in various guises in the literature [60], [62], [63]. To form this model several assumptions about the sample and the THz radiation are commonly employed: The sample is modelled as a solid non-conductive dielectric slab with uniform refractive index across the medium and the same magnetic permeability as dry air. The THz radiation is modelled as a collimated point beam, which propagates from the air and through the sample, see Figure 3.1(A). Reflections are assumed to occur but are assumed to not be measured in this chapter. To simplify this work, it is assumed that the beam propagates perpendicular to the sample interfaces, however more general models do exist [62]. These assumptions are generally valid when the sample is thin compared to the focal length of the optics used. As thicker samples are used, the effects of the Gaussian beam profile will become more prominent [99]–[101].

At each interface only a portion of the incident THz radiation propagates across the boundary. $T_{1,2}$ is the ratio between the transmitted and incident radiation at the interface. If the refractive index of either material is complex, then this ratio will also be complex indicating an incurred phase change at the interface. This coefficient can be calculated from the materials refractive indexes, respectively \tilde{n}_1 and \tilde{n}_2 :

$$T_{1,2} = \frac{2\tilde{n}_2}{\tilde{n}_1 + \tilde{n}_2} \quad (3.6)$$




Figure 3.1: (A) A THz electric wave, passing through a homogeneous sample. The THz wave is modelled as a point pencil beam perpendicular to the sample interfaces. (B) An interface between two dielectric materials, 1 & 2. In this diagram, material 1 is air and material 2 is the sample. At the interface part of the THz radiation is reflected (red), the other part is transmitted into the second material (smaller blue arrow). There will also be a second reflection at the second interface which is not shown.

In Figure 3.1(B), an interface is shown between air and the sample. $T_{1,2}$ is the ratio between the incident and transmitted radiation. The THz radiation propagating through an attenuating and velocity retarding material will incur a phase shift and attenuation relative to the incident radiation. The complex ratio between the two is the propagation coefficient, P , which can be calculated for a given thickness, l , and \tilde{n} :

$$P = e^{-i\frac{\omega l}{c}\tilde{n}} \quad (3.7)$$

Figure 3.2 shows a wave propagating through a sample. Within the sample, the radiation has shorter wavelength (due to the velocity change) and is attenuated. The wave which leaves the sample has incurred a phase change and has a reduced amplitude compared to the incident radiation. The propagation coefficient is a description of this.

To form a model of the transfer function, the measurements used to perform the deconvolution must be considered. In Figure 3.3, both measurements are shown




Figure 3.2: THz radiation (shown as sine wave for clarity), passing through a sample. The wave passing through the sample has a different wavelength (shown above sample) and will undergo attenuation. This results in a phase change and amplitude relative to the incident radiation which will be sample dependent.




Figure 3.3: (A) Model of the THz radiation passing through a sample of thickness l , (B) Model of the THz radiation in the reference measurement, passing through air which is displaced by the sample.

with respect to the model. In the sample measurement, the THz radiation propagates through the sample. The sample interfaces and propagation through the sample should then be considered when forming the model of E_{sam} . In the reference measurement, instead of propagating through the sample, the THz propagates through air which is displaced from the beam path in the sample measurement. Thus the propagation within air in the absence of a sample must be considered when forming the model of E_{ref} .

For the E_{sam} the following model is constructed:

$$E_{sam} = T_{a,s} P_s T_{s,a} E_0 H_{sys} \quad (3.8)$$

where $T_{a,s}$ and $T_{s,a}$ are the air-sample (left) and sample-air (right) transfer coefficients for the interfaces, P_s is the propagation coefficient for the samples medium over the sample thickness, l , E_0 is the generated THz radiation and H_{sys} is the instrument transfer function. Both H_{sys} and E_0 are consistent between reference and sample measurements, and thus will cancel out within the transfer function model.

A corresponding model for the measured THz radiation can then be constructed for the reference measurement:

$$E_{ref} = P_a E_0 H_{sys} \quad (3.9)$$

where P_a is the propagation coefficient of air over the sample thickness. This is modelled based on the refractive index of dry air, which is assumed to be $\tilde{n}_{air} = 1$.

Thus the modelled transfer function, \bar{H} , is:

$$\bar{H} = \frac{E_{sam}}{E_{ref}} = T_{a,s} \frac{P_s}{P_a} T_{s,a} \quad (3.10)$$

Or in terms of the sample's complex refractive index:

$$\bar{H}(\omega) = \frac{4\tilde{n}}{(\tilde{n} + 1)^2} e^{-i\frac{\omega l}{c}(\tilde{n}-1)} \quad (3.11)$$

3.3 Overview of the extraction processing

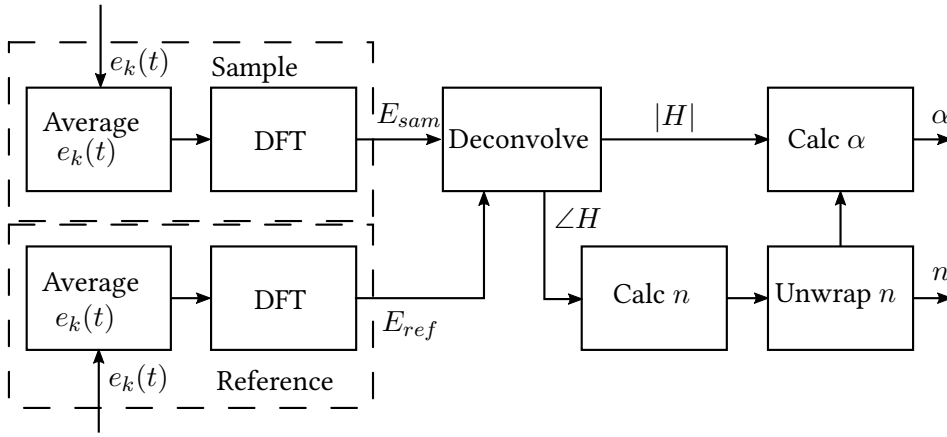


Figure 3.4: Processing flow diagram to extract the n and α of sample. Several measurements of the reference and sample are taken, from which the frequency domain transfer function is deconvolved. From the transfer function, the n and α are calculated.

To extract both n and α of a sample, measurements of a reference (dry air) and of the sample must be taken. These measurements are processed to calculate both parameters, the processing can be viewed as a series of sequential smaller processes. Each of these smaller processes are explained in the following sections.

In Figure 3.4, a diagrammatic overview of the processing chain is shown. For reference and sample measurements, the measurement scans are first averaged (§ 3.4) before a DFT is applied (§ 3.5) to extract the frequency representation of the reference or sample measurement. The transfer function, H , is then calculated by normalising the sample against the reference (§ 3.5). The real refractive index is then calculated (§ 3.6) using the angle (or phase) of the transfer function, $\angle H$, in this work unwrapping corrections are then applied after calculation of n (§ 3.9). The

absorption coefficient is then calculated using the magnitude of the transfer function, $|H|$, and n (§ 3.6).

3.4 Averaging and windowing of time domain data before applying Fourier transforms

Before applying a discrete Fourier transform, the time domain measurements must be processed into a singular windowed time scan. Multiple time scans are taken of a sample or reference, these consist of discrete samples of the amplitude of the THz electric field and time stamps of where they are taken from relative to an origin.

A measured time scan will consist of samples taken regularly in time from a continuous measurement, $e(t)$, of the electric field. The sampled electric field can be written in terms of a time stamp index k , either in terms of the k th time stamp, t_k or in terms of the sampling time period, T_s :

$$e(k) = e(kT_s) = e(t_k) \quad (3.12)$$

where $e(k)$ is the k th sample of $e(t)$.




Figure 3.5: Average of white noise (dashed red) with the multiple scans of white noise (blue circles).

Multiple time scans of the same sample are arithmetically averaged at each time stamp:

$$e(k) = \frac{1}{N} \sum_{m=1}^N e_m(k) \quad (3.13)$$

where m is the index of the measurement, and N is the number of measurements taken. This will asymptotically reduce the power of the zero mean noise with respect to the number of measurements used [93]. In Figure 3.5 an average (dashed red line) of multiple white noise scans is shown (blue circles). The true value should be a constant value with respect to time. The averaged value, tends towards a constant and is considerably less variant than the individual noisy scans. The measurements were sampled at a time period of 0.04 ps.

It is important to perform this step before further processing, in particular the calculation of spectroscopic parameters is non-commutative with arithmetic averaging [102], due to the processing being neither additive or multiplative. This results in the case where averaging multiple extracted \tilde{n} estimates will produce a different result to averaging before extraction and extracting a single \tilde{n} estimate. Importantly, the zero mean noise in the measurement will not be zero mean after extraction, so averaging after extraction is less robust to noise. The effects of averaging, uncertainty and noise are discussed in more detail in Chapter 6.

Each averaged scan is then windowed in time, this is achieved by multiplying with a windowing function. A Tukey window function is used in this work, which is a rectangular window function with cosine roll-off, (see Figure 3.6). The ratio between the roll-off region's and the rectangular regions can be adjusted to minimise window edge discontinuities. This function was used to maximise frequency resolution, while reducing the effects of discontinuities which can be caused by a rectangular window, i.e. frequency leakage.




Figure 3.6: A Tukey window (red dashed line) is applied to a raw reference scan (blue dashed line) to produce a scan without the reflection (green)




Figure 3.7: A reference (blue) measurement and a sample (red) measurement, the identical time windows (dashed lines) are shifted relative to each other.

Windowing a time domain trace is necessary, as the measurement time window is limited by reflections recorded within the time trace. In Figure 3.6, a reference measurement is shown (dashed blue), including a system reflection from the $150 \mu\text{m}$ GaP electro-optic crystal, with the applied time window function (red dashed) and the resulting reflection free time trace (green). These reflections will induce artefacts in the deconvolved transfer function [103]. There are two sources of reflections, one being the sample [60] which the THz field can resonate in, the other being the system [103] in particular the emitter, detector [104] and in the broadband system an optical beam block (high resistivity silicon plate).

If the system reflections have to be removed (so the time position of the reflections is consistent between measurements), the same time window can be applied to all the measurements. The time window is first adjusted and applied to the reference measurement. The sample trace is considered to be approximately an attenuated and time shifted copy of the reference measurement. The same window can then be applied to the sample measurement, by shifting it in time by a corresponding sample time shift. The sample time delay δt_h is estimated by locating the maximum in the cross-correlation between the windowed reference and sample measurement. The cross-correlation between the two measurements gives a measure of the similarity between them with respect to a time shift being applied to the sample measurement. The location of the maximum will therefore correspond to the time shift which needs to be applied to the sample measurement to maximise the similarity to the reference

measurement. By using this technique, the same window can be applied to both reference and sample measurements without incurring additional distortion in either measurement.

Consequently δt_h can be calculated by inverting the shift:

$$\delta t_h = -r_{\max} T_s \quad (3.14)$$

where r_{\max} is the (index) position of the maximum in the cross correlation.




Figure 3.8: The cross-correlation of the reference and a sample measurement. The maximum is marked with a triangle.

The (discrete) cross-correlation between the windowed reference and sample measurement is:

$$C_{ref,sam}(r) = \sum_{k=0}^{k=L-1} e_{ref}(k) e_{sam}(k+r) \quad (3.15)$$

where r is the difference index, e_{ref} is the windowed reference measurement, e_{sam} is the sample measurement and L is the length of the time index k . r_{\max} can then be estimated by locating the maximum of this function.

In Figure 3.7, a time measurement of a sample (red) and a reference (blue) is shown, with the associated windows (dashed lines). The window in each case is the same, but shifted in time by 1.44 ps. In Figure 3.8, the cross-correlation between the two is shown, with the difference given in ps (multiplied by T_s). There is a clear maximum, which was used to estimate the 1.44 ps time shift.

3.5 Deconvolution of a Sample's Transfer Function

To quantify the effect of a sample on the propagating THz electric field, the transfer function of a sample is calculated from reference and sample measurements. In effect, this treats the sample as a linear time invariant system, with known input and output. Since the effect of the sample can be described numerically by convolving the sample's impulse response with an air measurement, the sample's effect can be estimated by performing deconvolution. This is accomplished by normalising the complex frequency response of a sample measurement against that of a reference measurement (air), producing the transfer function of a sample. The impulse response of a system can then be deconvolved by performing the inverse Fourier transform, if required.

A DFT is then applied, to calculate the full complex spectrum of a measurement. Zero padding is used to increase the computational resolution, which refers to the interval between samples in the frequency domain. However, as no new information is introduced to increase the resolution, these new samples are interpolated. This is done for two reasons, firstly the computational resolution (and bandwidth) must be the same in both sample and reference measurements to perform the point by point division, and secondly features of interest may be aliased if the resolution is too low.

The complex frequency spectrum is then truncated down to the measured frequencies. This truncation is not used as a filtering technique, rather the analysis is being restricted to this range. At frequencies outside of this range, there does not exist analysable information, either because of dynamic range limitations, or because of inherent bandwidth limitations of the system. Dynamic range limitations will be discussed further in § 3.7.

The transfer function, H , can then be calculated by normalising the frequency domain sample measurement, E_{sam} , against the frequency domain reference measurement, E_{ref} :

$$H(\omega) = \frac{E_{sam}(\omega)}{E_{ref}(\omega)} \quad (3.16)$$

Using H , the extraction of spectroscopic parameters can then be performed. As with the spectroscopic parameters in later parts of this work, the frequency dependence

is not explicitly stated and the reader should assume this dependence.

3.6 Model approximations

Equation 3.11 is difficult to invert without numerical methods. However, using simple approximations, it is possible to derive approximations to estimate the real refractive index, n , and extinction coefficient κ , individually. These approximations are shown to be accurate when used with samples which are non-resonant and are used as initial points for numerical methods in Chapter 5.

It is assumed that the propagation coefficient dominates the phase of the transfer function (i.e. absorption has negligible effect on the phase at interfaces). Based on this assumption the phase of the transfer function can be equated to the phase of the models propagation coefficient terms:

$$\angle H = \angle \frac{P_s}{P_a} = -\frac{\omega l}{c}(n - 1) \quad (3.17)$$

The rightmost term is not limited to the range $-\pi$ to π , so $\angle H$ must be unwrapped to be accurate. This is explored further in § 3.9, where the effects of wrapped phase are discussed and unwrapping techniques are introduced.

Equation 3.17 can be rearranged into:

$$n = -\frac{c}{\omega l} \angle H + 1 \quad (3.18)$$

Thus, if the thickness of the sample is known, the real refractive index can be estimated. Keeping with this assumption, the transmission coefficient terms can be approximated to be real and dependent solely on n . The transfer function model magnitude can then be approximated as:

$$|H| = \frac{4n}{(n + 1)^2} e^{-\frac{\omega l}{c} \kappa} \quad (3.19)$$

Inverting this, κ , can be calculated:

$$\kappa = \frac{-c}{\omega l} \log_e \left(\frac{(n+1)^2}{4n} |H| \right) \quad (3.20)$$

This can be converted to absorption coefficient using equation 3.2. And thus equation 3.20 is simply a variation of the Beer Lambert law [47].

It should be noted, that the measured absorption coefficient and real refractive index will be affected by additional physical effects which are assumed negligible within the extraction. Effects such as scattering [105] can result in additional losses and phase changes not modelled.

3.7 Limitations of dynamic range

The bandwidth of a measurement is limited by two different aspects, the inherent physical bandwidth of the spectroscopic system and the detection noise of the system. The former refers to the limited ability to emit, focus and detect electromagnetic radiation beyond or below certain frequencies. The latter refers to the inability of the system to detect radiation, due to radiation amplitude being smaller than the detector noise in the system [106]. This detector noise is distinct from measurement noise [107], and refers to the noise in absence of radiation in the system.

When propagating through a sample the THz radiation will be attenuated, the detector noise is however unchanged. This will result in a decrease in bandwidth in the sample measurement, as the emitted radiation is weaker towards the limits of the bandwidth. Inversely there exists a limit on the detectable absorption due to this noise floor.

κ is a measure of absorption, based on a sample measurement normalised to a reference measurement. A maximum value will therefore be related to the reference measurement normalised to the detector noise floor, referred to as the dynamic range of the spectrometer. The dynamic range can be estimated as [106]:

$$DR = \frac{|E_{ref}|}{\delta_{det}} \quad (3.21)$$

where δ_{det} is the root mean square (RMS) of the detector noise. This can be incorporated into equation 3.20, to form an expression for the maximum extinction coefficient, κ_{max} :

$$\kappa_{max} = \frac{c}{\omega l} \log_e \left(\frac{4n}{(n+1)^2} DR \right) \quad (3.22)$$

From which the analysable bandwidth of a sample measurement can be estimated to being the range where $\kappa < \kappa_{max}$.

In this work α is commonly used, so an alternative form of this dynamic range limit is α_{max} [106]:

$$\alpha_{max} = \frac{2}{l} \log_e \left(\frac{4n}{(n+1)^2} DR \right) \quad (3.23)$$

Again, analogously to κ , the analysable bandwidth of the sample measurement can be estimated to be where $\alpha < \alpha_{max}$.

3.8 Terahertz spectroscopy of anhydrous and monohydrate forms of α -lactose

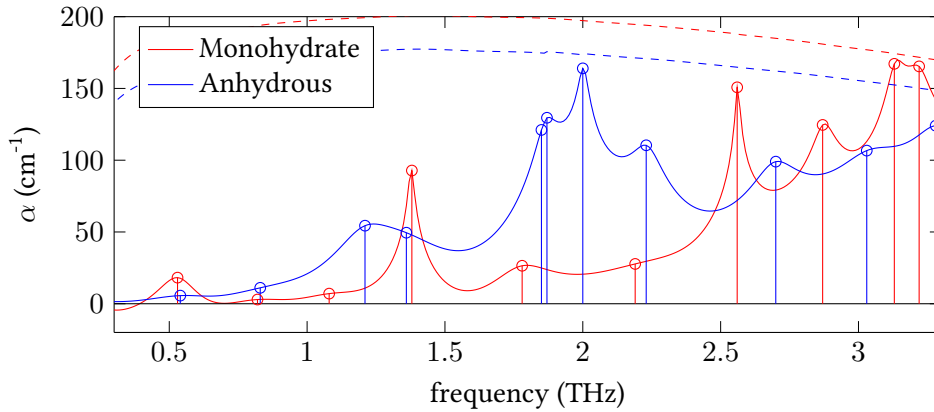


Figure 3.9: absorption coefficient of lactose anhydrous 100% (blue) and lactose monohydrate (red). Peaks have been marked in circles with vertical lines. α_{max} (dashed lines) of lactose anhydrous has been used to determine bandwidth.

α -lactose is a common case study in THz spectroscopy, as it is inexpensive, safe and has distinct features across THz frequency region. α -lactose has monohydrate and anhydrous crystalline forms which have distinctly different absorption spectra and so serve to demonstrate sensitivity to crystalline structure in the THz region.

α -lactose monohydrate is also of interest as a test sample, as it exhibits narrow

absorption features between 0.1 THz and 2 THz, which most THz spectroscopy systems operate over, so it is a practical example to test the available frequency resolution of the instrument. The monohydrate form has 2 sharp features at 0.53 THz and 1.4 THz [108], the former of which is reported to have a full width half maximum (FWHM) of 23 GHz at room temperature [109]. This can be used to demonstrate the instrument's experimental frequency resolution, which is determined by the inverse time window length (in seconds). If the frequency resolution is too low, which can occur when the time window has been limited due to a reflection within the measurement, features will be broadened and merged. It should be noted, that the integrated area of an absorption peak should remain consistent, as the limited frequency resolution has a moving average effect.

α -lactose anhydrous and monohydrate were both measured twice, at 10% and 100% concentrations in pellets with similar thickness and volume. A description of the pellet making process is provided in Appendix A. These were measured at a sampling frequency of 25 THz, using the LFRBB THz TDS system. The emitter was biased with a 7 KHz 150 V square wave and a 220 GHz frequency resolution time window was used.

In Figure 3.9, the absorption coefficient of 100% Lactose monohydrate (red) and anhydrous (blue) is shown. Peaks have been determined by using a peak finding algorithm (peak detection on the second derivative), which have been marked as circles with vertical lines. The estimated α_{max} is shown as dashed lines of the appropriate colour. The bandwidth was limited to 3.3 THz as α of lactose monohydrate meets α_{max} at 3.14 THz. Lactose anhydrous has a large feature at approximately 2 THz, which is being clipped by α_{max} (the curve given is an estimate). This has resulted in it erroneously appearing to be 3 separate peaks, when it is actually just a large singular feature. Peaks which are clipped by α_{max} will be distorted but will have the same approximate place as only the absorption above α_{max} is distorted. This can be verified by looking at a lower concentration (lowering the absorption) or by using a thinner sample (increasing α_{max}).

The spectral features of the two forms are distinctly different, in particular the sub 2 THz, the features of lactose monohydrate is considerably sharper than the

broader features of lactose anhydrous at these frequencies, while the lactose anhydrous has a much larger feature at approximately 2 THz. The lactose monohydrate feature at 0.53 THz is expected to be much narrower than measured here. This is due to the frequency resolution of 220 GHz compared to the feature FWHM of 23 GHz.

At the very low frequencies (below 0.5 THz), the absorption appears to be negative (the sample appears to amplify the radiation). This is likely an artifact introduced by a slight increase in pump laser power, which results in amplification of THz in the sample measurement. This leads to a downward shift in absorption, and where the absorption is negligible, a negative absorption.

Pellets of 10% α lactose monohydrate and anhydrous were made by diluting with polytetrafluoroethylene (PTFE) (which with the exception of a 6 THz peak [17], has very little THz absorption). A full description of the pellet making process is found in Appendix A. This has the effect of decreasing the absorption of the sample, increasing the bandwidth of α below α_{max} . However, additional physical interaction between THz radiation and the mixture will occur, in particularly introducing additional background scattering from the PTFE [105], and asymmetric peak distortion due to the Christiansen effect [96].

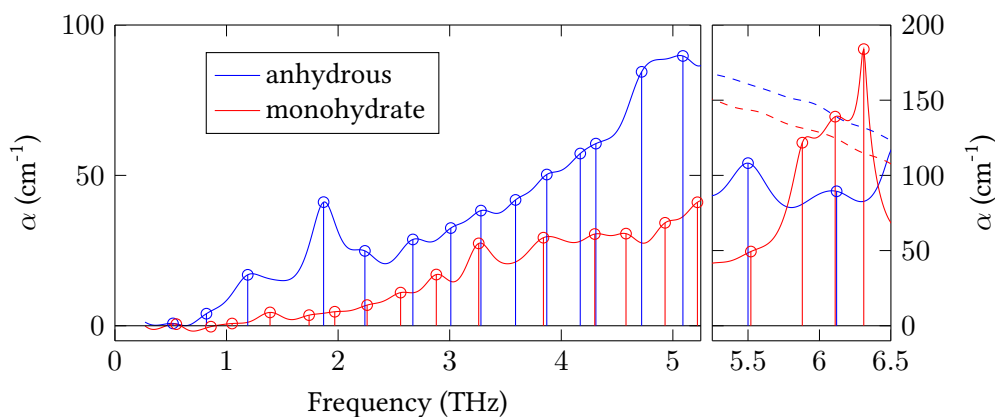


Figure 3.10: absorption coefficient of 10 % lactose anhydrous (blue) and lactose monohydrate (red). The left pane shows 0–5.3 THz, the right pane show 5.3–6.5 THz. Different scales have been used for clarity and the curves do not join between panes. Peaks have been marked in circles with vertical lines. α_{max} (dashed lines) of lactose anhydrous has been used to determine bandwidth.

In Figure 3.10, α of 10% lactose monohydrate (red) and anhydrous (blue) is shown.

Peaks determined are marked as circles with vertical lines, and the bandwidth was determined by α of lactose anhydrous meeting its α_{max} at 6.5 THz. Two different scales have been used for 0–5.3 and 5.3–6 THz as the PTFE peak has dominated the spectrum of the lactose anhydrous measurement. The curves are marked with the same scheme between panes.

Comparing both Figures 3.9 and 3.10, the absorption of diluted lactose is significantly decreased and some slight distortion of low frequency features can be observed. Both noticeably contain a series of peaks across the entire bandwidth. Lactose anhydrous displays strong features at 1.87 and 2.24 THz, and at 4.72, 5.09 and 5.5 THz. While the monohydrate form has a distinct triplet of peaks at 2.88, 3.26 and 3.84. Table 3.1 contains the collected peaks detected in both concentrations and forms of lactose, the peaks have been grouped by frequency. There exist several peaks which are present in the 100% samples but not in the corresponding 10% samples, that are due to a broad peak being “capped” by α_{max} . Apart from these features, there is good agreement between the different concentrations, with variation of less than 50 GHz between the features. There also appears to be several peaks which are at similar frequencies between the monohydrate and anhydrous forms.

The 0.53 THz peak present in lactose monohydrate should have a FWHM of 23 GHz [109], however in Figure 3.9, the feature is comparatively broad. This is due to the low frequency resolution time window used, if a high resolution (longer) time




Figure 3.11: low resolution (red) (220 GHz) and high resolution (blue) (77 GHz) α of the 0.53 lactose monohydrate peak, a Lorentzian curve has been fitted to high resolution data (magenta)

Anhydrous 10%	Anhydrous 100%	Monohydrate 10%	Monohydrate 100%
0.52	0.54	0.55	0.53
0.82	0.83	0.86	0.82
		1.05	1.08
1.19	1.21		
	1.36	1.39	1.38
	1.85		
1.87	1.87	1.74	1.78
	2.00	1.97	
2.24	2.23	2.26	2.19
2.67	2.70	2.56	2.56
3.01	3.03	2.88	2.87
			3.13
3.28	3.28	3.26	3.22
3.59	3.58		
3.87		3.84	
4.17			
4.31		4.30	
4.72		4.58	
5.09		4.93	
		5.22	
5.50		5.52	
		5.88	

Table 3.1: Table of absorption peaks in 10% α lactose anhydrous, 100% α lactose anhydrous, 10% α lactose monohydrate and 100% α lactose monohydrate




Figure 3.12: Absorption coefficient of 100 % lactose monohydrate (blue) and 10 % lactose monohydrate scaled by a factor of 10. Peaks are shown as circles with vertical lines, and the α_{max} of 100 % lactose monohydrate is shown as a dashed line

window is used, the measured FWHM will decrease. In Figure 3.11, the 0.53 THz feature is shown, calculated using a frequency resolution of 220 GHz (red) and a frequency resolution of 77 GHz (blue). To do this a time window including system and sample reflections was used, which lead to a number of artifacts in the spectra (see § 3.11). The feature measured using the 77 GHz resolution is much sharper, and a Lorentzian curve has been fitted to it (magenta curve). The FWHM of this curve is 27 GHz, which is close to the reported value of 23 GHz. An even higher frequency resolution is possible using a longer time window. However, a number of additional reflections are included which will have an adverse effect on the peak shape.

Figure 3.12 shows the different concentrations of lactose monohydrate, with 100% shown in blue, and 10% shown in red. The 10% has been scaled up by 10, which gives an approximation of the α of 100% Lactose monohydrate over a larger bandwidth. This is based on the assumption that scattering is negligible and that PTFE is non-absorbent at THz frequencies. This simple approximation is used in § 4 to model mixtures of different compounds.

In Figure 3.12, the 10% and 100% clearly shows similar features, which demonstrate where the 100% has been capped by α_{max} at 3.26 THz. There appears to be a feature at 1.97 THz in the 10% which is not present in the 100%. This is probably an artifact in the processing (for instance slight frequency leakage from the window function), rather than due to the inclusion of PTFE, which lacks significant features

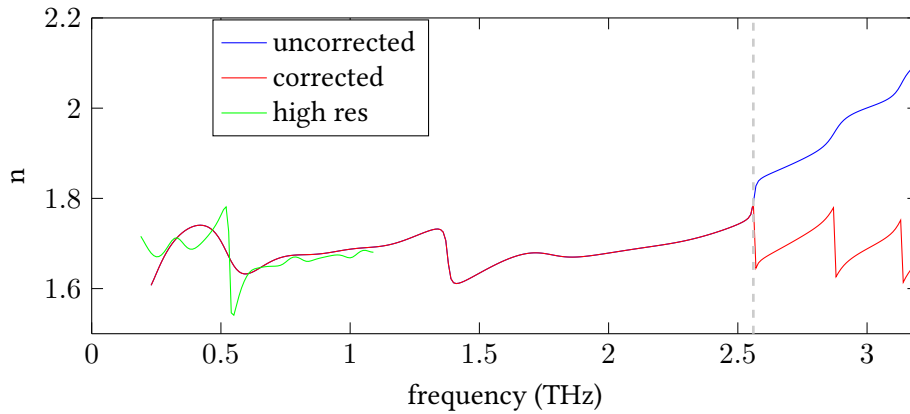


Figure 3.13: n of 100 % Lactose monohydrate (blue), a correction was applied to resolve phase unwrapping issues (red). The high resolution of the 0.53 THz relaxation is shown in green.

at this frequency.

In Figure 3.13, the real refractive index of 100 % Lactose monohydrate is shown. At locations where a feature would be present in the α spectrum, a relaxation is present, resulting in a sharp change in the real refractive index. At 0.53 THz a very broad relaxation is present, which is due to limitation of the frequency resolution. When a high resolution time window (77 GHz) is used, the feature (green) takes the shape of a sharp relaxation. At 1.4 THz there exists a relaxation corresponding to the 1.4 THz peak, and at 2.62 THz there exist a series of sharp discontinuities due to the absorption spectrum being capped by α_{max} .

Sharp features in the α spectrum result in sharp relaxations, as n and κ are inherently part of the same complex function, \tilde{n} . The relationship between the imaginary and real parts of \tilde{n} are the Kramers-Kronig relations [110], which can be used to convert from one to the other. A rule of thumb for the shape of the features in n is to use the differential of κ .

These sharp discontinuities can result in unwrapping errors in n , which if not corrected (blue), lead to n at a higher frequencies than a discontinuity being higher than expected. This can be corrected by applying a correction to the real refractive index (red), to give the correct form to the relaxation. This could be performed by iteratively applying Kramers-Kronig relations to κ and n [84]. However, a far simpler approach is taken in § 3.9.3, which instead uses a model of unwrapping error

to find the flattest n curve.

3.9 Phase unwrapping corrections

To extract the real refractive index and thus the dependent extinction coefficient, the unwrapped phase of the transfer function is required. This is due to equation 3.17, which relates the transfer function phase to the real refractive index:

$$\angle H = \frac{c}{\omega l} (n - 1) \quad (3.24)$$

The right hand side of this expression is not limited to a specific region, however the phase of the transfer function is calculated using an arctan function (equation 2.11), which is only defined in the region between $-\pi$ and π , angles outside this region “wrap” around to the opposite end of this range. The calculated angle is then often a multiple of 2π away from the true angle.

This results in ambiguity such that the true “unwrapped” phase, $\angle H$, is some multiple of 2π away from the wrapped estimate, $\angle H'$:

$$\angle H'(k) = \angle H(k) + a(k)2\pi \quad (3.25)$$

where k is the frequency index and $a(k)$ is an integer multiplier at the frequency sample. $a(k)$ is cumulative, the value at k is dependent on the value at $k - 1$. As real refractive index requires the unwrapped phase, this ambiguity will propagate to the refractive index. This results in two apparent effects with respect to the real refractive index. The first is an error which asymptotically decreases with frequency while the second is discontinuities between values. The former occurs when the initial phase offset, $a(0)$, is unknown, creating an offset across all frequencies ($a(k)$ is cumulative) while the latter occurs when the phase wraps from one value to the other forming discontinuities. These effects are intrinsically the same error, that the phase offset, $a(k)2\pi$, is unknown. Thus to correct wrapping errors, $a(k)$ must be estimated. In this work this is done by assuming that the lowest frequency point is correct, such that $a(0) = 0$, and estimating $a(1)$ – $a(L - 1)$. From the updated




Figure 3.14: Real refractive index of lactose monohydrate calculated from wrapped phase. The value the refractive index wraps around is marked as a dashed grey line.

estimate of $\angle H$ or n a value for $a(0)$ is then formed which can then be applied to recover the unwrapped estimate.

Figure 3.14 show n calculated from the wrapped phase, $\angle H'$, of the 100 % lactose monohydrate sample from § 3.8. A wrapping effect is present which causes discontinuities centered around 1. This is due to the calculated phase of the transfer function wrapping around 0 and the use of equation 3.18 which adds an offset of 1. In Figure 3.15, the incorrect values for n of a 1 mm thick high resistivity silicon plate has been shown (red) along with a corrected estimate (blue). The extracted n has been processed with corrections for discontinuities, however there is a phase offset error which results in an asymptotic error. This is particularly common in samples which have either high sample thickness or a high refractive index.




Figure 3.15: The incorrect estimate (red) of n for high resistivity silicon, and a estimate with phase offset error corrected (blue)




Figure 3.16: The wrapped phase (dashed) is unwrapped (solid).

3.9.1 Unwrapping the transfer function phase

The standard method of unwrapping phase is to assume a linear relationship between frequency points (i.e. that phase is roughly a straight line). At discontinuities between points greater than π , an integer multiple of 2π is added to the next point until the difference between the points is less than π . In Figure 3.16 a diagram showing the wrapped (dashed) and unwrapped (solid) phase is displayed. At discontinuities greater than π an integer multiple of 2π is added, shifting the segments of the phase to form a straight line.

This integer multiple, $b(k)$, can be formulated as the integer argument minimum between samples of $\angle H'$:

$$b(k) = \underset{b(k)}{\operatorname{argmin}} \left(\angle H'(k) - \angle H'(k-1) - b(k)2\pi \right) \quad (3.26)$$

Which can then be related to $a(k)$:

$$a(k) = a(k-1) + b(k) \quad (3.27)$$

which shows that $a(k)$ is cumulative, i.e. the error accumulates across frequencies. Then a new estimate transfer function phase, $\angle H''$ can be formed:

$$\angle H''(k) = \angle H'(k) - a(k)2\pi \quad (3.28)$$

where $a(0) = 0$ is assumed.

This provides a simple method to estimate the unwrapped phase. However, two issues persist, the inherent assumption that the first frequency sample is correct (in

terms of $a(k)$ this is assuming that $a(0) = 0$ and that the unwrapped phase does not have discontinuities greater than π present within it. The first assumption is often incorrect. The phase information at low frequencies can be poor leading to forward propagation of error across frequencies and to the asymptotic error in Figure 3.15. A common method adopted in the literature [101] to circumvent this is to instead truncate at the lower frequencies and extrapolate back towards zero which will give an estimate of the phase close to 0 frequency. This is then used to form an educated prediction as to correct phase offset.

Figure 3.17 shows an estimate of phase with an incorrect phase offset (dashed) and extrapolated offset (solid). The phase before the vertical dashed line is incorrect, leading to an incorrect initial offset. To correct this, this region is instead extrapolated back towards 0. This leads to the correct offset being calculated. This is equivalent




Figure 3.17: The incorrect offset phase (dashed) and the phase with extrapolated offset (solid). Vertical dashed line shows where the phase was extrapolated from.

to locating $a(0)$:

$$a(0) = \underset{a(0)}{\operatorname{argmin}} \left(\angle H''(0) - a(0)2\pi \right) \quad (3.29)$$

Where $\angle H''(0)$ is the extrapolated phase from $\angle H''(k)$ at $k = 0$. The correct phase can then be recovered:

$$\angle H(k) = \angle H''(K) - a(0)2\pi \quad (3.30)$$

3.9.2 Unwrapping refractive index

An alternative solution to the unwrapping problem is to try to resolve it after calculating the real refractive index. This provides a different method of calculating $a(k)$,




Figure 3.18: The finite difference of n' (from Figure 3.14) (blue) and $-\frac{\lambda(k)}{l}$ (red dashed)

which makes use of the assumption that n is constant with respect to frequency. As n is related to the gradient of $\angle H$, this is equivalent to assuming that $\angle H$ is a straight line. However, as this and the previous methods are different, they can converge on different answers, and this can be used as a “sanity” check for other methods and published data. In later sections (§ 3.9.3 & § 5.2), the relations used to perform this wrapping, are exploited to form more accurate extracted values of n and κ .

By combining equations 3.18 and 3.25, the following expression can be derived:

$$n'(k) = n(k) + a(k) \frac{c}{f(k)l} = n(k) + a(k) \frac{\lambda(k)}{l} \quad (3.31)$$

where n' is an estimate of real refractive index based on wrapped phase and will generally take the form of chopped error shown in Figure 3.14. $f(k)$ is the frequency in Hertz at index k and $\lambda(k)$ is the corresponding free space wavelength.

Assuming that the real refractive index is roughly constant, discontinuities can then be removed by subtracting integer multiples of $\frac{\lambda(k)}{l}$. This is equivalent to the method which minimises discontinuities in the raw phase by adding integer multiples of 2π . In Figure 3.18 the finite difference of Figure 3.14 is shown as a blue line, $-\frac{\lambda(k)}{l}$ is also shown as a red-dashed line. The finite difference shows sharp peaks where n has suddenly changed due to the phase wrapping. These will tend to $\frac{\lambda(k)}{l}$ or some integer multiple of, which is why in Figure 3.18 the peaks tend to the red line.

Thus $b(k)$ can be estimated as:

$$b(k) = -R \left(\frac{(n(k) - n(k-1))l}{\lambda(k)} \right) \quad (3.32)$$

where $R(x)$ is the round to nearest integer function. Using equation 3.27, $a(k)$ can then be estimated. As with conventional phase unwrapping, this only resolves the phase ambiguity due to discontinuity between samples, it does not resolve the offset. In other words $a(0)$ is left ambiguous. This can therefore form a new estimate of refractive index, n'' :

$$n''(k) = n'(k) - a(k) \frac{\lambda}{l} \quad (3.33)$$

This assumes that $a(0) = 0$, and thus the new estimate:

$$n''(k) = n(k) + a(0) \frac{\lambda(k)}{l} \quad (3.34)$$

In Figure 3.19, equation 3.34 has been applied to the n' estimate from Figure 3.14. The region at frequencies below 2.5 THz is relatively flat, and displays some expected relaxations. Above frequencies of 2.5 THz, the n appears to deviate to higher values. This is because of underlying discontinuities within the phase which the unwrapping correction fails to correct.




Figure 3.19: The new estimate of n of lactose monohydrate

To recover n , $a(0)$ must be found. This estimate of n is an estimate without knowledge of the initial offset and will take the form of asymptotic error (red curve) shown in Figure 3.15.

By assuming that n is approximately constant, $a(0)$ can be estimated. $\frac{\lambda}{l}$ will increase exponentially at lower frequencies, thus the incorrect real refractive indices will diverge at lower frequencies and have a larger average differential with respect to frequency. By minimising the average differential (find the flattest refractive index curve) the value of refractive index can be found. To achieve this, an error term between the measured and modelled estimate is used:

$$\varepsilon_{n''} = \frac{dn''}{df} - \frac{d\bar{n}''}{df} \quad (3.35)$$

where \bar{n}'' is the model (from equation 3.34) of the calculated n'' . If the differential of n tends to 0 (i.e. n is a constant) then this can be simplified:

$$\varepsilon_{n''} = \frac{dn''}{df} - a(0) \frac{c}{f^2 l} \quad (3.36)$$

A mean square error with respect to frequency can then be formed:

$$\text{MSE}_{n''} = \mu \left(\left(\frac{dn''}{df} - a(0) \frac{c}{f^2 l} \right)^2 \right) \quad (3.37)$$

where $\mu(x)$ is the arithmetic mean operation. An estimate of $a(0)$ can be found by locating where the differential with respect to $a(0)$ is equal to 0, which will be where:

$$a(0) = -R \left(\frac{l}{c} \mu \left(\frac{dn''}{df} f^2 \right) \right) \quad (3.38)$$

This assumes that the differential of the real refractive index tends to 0, that the differential of the bad estimate is accurately known (it is instead estimated from the finite difference) and that the error is constant across frequency. These assumptions are often not true, so weighting the average towards lower frequencies where the error is greater often performs better. Using the weighting $\frac{1}{f^2}$, an estimate of $a(0)$ can be formed, which in a simplified form is:

$$a(0) = -R \left(\frac{lc}{\mu(\lambda^2)} \mu \left(\frac{dn''}{df} \right) \right) \quad (3.39)$$

An unwrapped estimate of n can then be formed using:

$$n(k) = n''(k) - a(0) \frac{\lambda(k)}{l} \quad (3.40)$$




Figure 3.20: The differential of n of silicon (blue) and $\frac{c}{f^2l}$ (red dashed)

In Figure 3.15, the incorrect n of a silicon plate is shown due to the phase offset ($a(0)$) being unknown. To remedy this, equation 3.39 was used to calculate $a(0) = -1$, which has then been used in equation 3.40 to form the correct n shown in Figure 3.15. In Figure 3.20, the differential of incorrect estimate of n of silicon is shown (blue), with $\frac{c}{f^2l}$ overlaid (red dashed). As true n of silicon is almost constant across the bandwidth, the erroneous differential will tend to $\frac{c}{f^2l}$ as $a(0) = -1$.

3.9.3 Unwrapping at discontinuities in real refractive index

These methods which minimise discontinuities within $\angle H$ or n assume that the difference between $\angle H$ frequency samples is always less than π . There exist two conditions where this might not be true. The first is a simple aliasing effect, if the frequency domain is sub sampled (the computational frequency resolution is too low) and the phase is to be unwrapped, discontinuities greater than π are possible and the unwrapping procedure will produce an aliased unwrapped phase, see § 2.3.5 on the DFT. This can be remedied by simply increasing the zero padding (and thus increasing the computational frequency resolution of the discrete Fourier transform).

The second occurs when a fundamental discontinuity is present in the phase which can occur under two conditions. When the signal hits the noise floor (due to

limitation of dynamic range), the phase can be influenced by noise. This dominates in Figure 3.13, where incorrect phase unwrapping has led to n being overestimated above 2.5 THz. This can break the implicit assumption of minimum phase difference, as the measured phase in this region is ambiguous. The phase will therefore consist of random fluctuations and the unwrapping will often incorrectly interpret this. The second cause of discontinuities in the phase, is when there are discontinuities within the real refractive index itself. This can occur when a particularly narrow absorption feature is present, leading to sharp relaxation in real refractive index.

After performing the algorithms in § 3.9.1 or § 3.9.2, these discontinuities are still present in n but are minimised, and are thus detectable and correctable. The wrapping relations with n discussed in § 3.9.2 can be exploited to form the correction but instead of using individual samples, the mean of segments on either side of a discontinuity are used. To detect these discontinuities, the differential squared can be used with a peak detection algorithm. This can then be used to fragment the refractive index into internally consistent (with respect to wrapping) segments.

In Figure 3.21, a refractive index with a discontinuity is shown (left) and is to be fragmented into two segments n_a and n_b . This is done by performing peak detection (differential zero crossing) on the differential squared (right), to the left of the peak will be the first segment (n_a) and to the right the second (n_b).

In Figure 3.22, the extracted n of lactose is shown (blue), with the boundaries of the segments marked by dashed red lines. These occur where the differential peaks (i.e. a sharp change in refractive index) and should occur where either a relaxation or erroneous step is present. However, this can incorrectly segment n which has




Figure 3.21: The real refractive index of a sample contains a discontinuity (left). To break n into n_a and n_b , peak detection is performed on the differential of n , squared (right). The peak is circled, and the corresponding segments are marked by bars.




Figure 3.22: The real refractive index of Lactose monohydrate (blue), with the segmentation points (dashed red).

occurred at 0.25 THz. Division of a larger correct segment into smaller segments is not an issue, as the smaller segments are already consistent with each other. What can cause incorrect solutions, is where there is over segmentation in the region of a true discontinuity. In the neighbourhood of a discontinuity in n , if a segment is too small (due to over segmentation), its average value can be dominated by the discontinuity. Since the correction procedure relies on the segment average being correct, this can lead to incorrect values being produced. Fortunately, by specifying a minimum segment size during the segmentation, this can be avoided.

To extract the correct real refractive index, these segments must be made to be consistent with each other. To do this, there is assumed to be an inherent discontinuity between the segments, but the average value between segments are consistent. This is in essence allowing discontinuities within the real refractive index, as long as the broader frequency response is flat.

To find the correct values for each of the segments, one is assumed to be correct




Figure 3.23: Incorrect n due to a discontinuity (left) and the correct n (right). This correction can be found by minimising $\delta\mu(n)$ between segments n_a and n_b




Figure 3.24: Two segments of the real refractive index of Lactose, n_a (blue), n'_b (red) and the corrected n_b (green). The means of the segments have also been marked as diamonds. The segmentation points have been marked using grey dashed lines.

and the difference between the mean values of segments is minimised. This is shown in Figure 3.23, where two segments on either side of a discontinuity in refractive index have been identified (left). To find the correct real refractive index (right), the difference between segment means $\delta\mu(n)$ has been minimised.

Given two sequential segments, n_a and n_b , which are assumed to be approximately constant (flat) and equal ($n_a \approx n_b$), a relation between n_a and an incorrect estimate of the latter segment, n'_b , can be formed:

$$n_a \approx n'_b + a \frac{\lambda_b}{l} \quad (3.41)$$

where a is an integer constant and λ_b are the wavelengths over segment n_b . n_a is assumed to be correct without any phase unwrapping error. Thus finding the value of a which minimises the difference between segment averages will find the flattest (barring the discontinuity) combinations of the two segments. This can then be done sequentially between pairs of segments, to correct the entire curve.

The squared difference between segment means is:

$$\text{SD} = \left(\mu(n_a) - \mu\left(n'_b + a \frac{\lambda_b}{l}\right) \right)^2 \quad (3.42)$$

which given approximation 3.41, should tend to 0. Thus to estimate a , this must be

minimised. The value of a which minimises this is:

$$a = R \left(l \frac{\mu(n_a) - \mu(n'_b)}{\mu(\lambda_b)} \right) \quad (3.43)$$

This can then be used to form an estimate of n_b which is most consistent with segment n_a :

$$n_b = n'_b + a \frac{\lambda_b}{l} \quad (3.44)$$

In Figure 3.24, the segments between 1.4 and 2.9 THz from Figure 3.22 are shown. Between segments n_a (blue) and n'_b (red) there is a step, and the means of the two (diamonds) are considerably different. Using the assumption that the difference between means should be minimal, a correction has been applied to n'_b to arrive at n_a (green). This has created a discontinuity, but has minimised the difference between its mean and n_a 's. This correction was calculated using equations 3.44 and 3.43.




Figure 3.25: The corrected estimate of n of 100 % Lactose monohydrate. Where discontinuities are present has been marked by vertical dashed red lines.

These corrections are sequentially applied to the segments from low frequency to high frequency to form a new estimate of n . In Figure 3.25, the corrected n is shown, which is significantly flatter but has regular discontinuities. The discontinuities which are due to relaxations have been marked by dashed red lines (the incorrect segmentation at 0.25 THz was present but had no affect on the correction). In Figure 3.26, the extracted n of the 10 % Lactose monohydrate sample (from § 3.8) is shown. The absolute values for both samples will not be consistent but the general

shape of the low concentration should be similar with the high concentration n . In Figure 3.26, the discontinuities positions from Figure 3.25 have been marked as red vertical lines. The majority of these discontinuities lie on features which appear to be smoothed relaxations. Where there does not appear to be consistency is between 3 and 3.5 THz. This is where the absorption of the 100 % sample has reached the limits of the system's dynamic range. At higher frequencies, the discontinuities seem to be consistent. However, this is likely due to the absorption modes at these frequencies being distinct peaks (and thus less affected by reaching the limits of dynamic range), rather than there being a correct feature present.




Figure 3.26: The estimate of n of 10 % Lactose monohydrate. Where discontinuities are present in the estimate of 100 % Lactose monohydrate, has been marked by vertical dashed red lines.

3.10 Polymorphism of paracetamol




Figure 3.27: Absorption coefficient of 100% Paracetamol I (blue), II (red) & III (green). α_{max} has been marked as dashed lines, peaks have been marked by vertical lines with circles.




Figure 3.28: Absorption coefficient of 100 % Paracetamol I (blue), II (red) & III (green), taken from [111]

Paracetamol is a common commercially available analgesic, which has a number of different crystalline forms. It has been previously measured at THz frequencies [4], [112] and its different crystalline forms (I,II,III) have been observed by changing the temperature[111].

Both 100% Paracetamol I and II was measured over 0-3 THz in the LFRBB THz TDS system, and the relevant refractive index and absorption coefficient have been identified. The paracetamol polymorphs (forms I, II and III) were provided by Colin Seaton, Ian Scowen and Tasnim Munshi. The paracetamol form III provided was crystallised using a method which ensured stability at room temperature and included traces of lactose. The measurements were performed at a time sampling frequency of 150 THz with an approximate frequency resolution of 200 GHz. 100 % Paracetamol III was also measured over 0–3 THz in the LFRBB THz TDS at a time sampling frequency of 30 THz, with an approximate frequency resolution of 200 GHz. The emitter was biased with a 7 kHz square wave.

In Figure 3.27, the absorption coefficient of Paracetamol I (blue), II (red) and III (green) is shown. Peaks have been found by applying a peak detection algorithm on the second order differential and marked with circles and vertical lines. α_{max} has been marked for all samples as dashed lines of the appropriate colour. Forms I, II and III reach α_{max} at 2.88, 2.2 and 2.63 THz respectively. The different forms have distinct but similar spectra, and there appear to be several common groupings of peaks between the forms. Form II is relatively flat below 2 THz compared to the




Figure 3.29: Absorption coefficient of 100 % Paracetamol I measured at 415 K (blue) and room temperature (red) taken from [111]. In green the room temperature measurement from this work is shown.

other forms. However, peak detection has found a number of features which can be grouped by frequency with similar peaks in forms I and III at these frequencies. Form I has similar features to form III but at a slightly higher frequency. Features above 2 THz can be seen which are not common to forms II and III, however this might be due to a dynamic range limitation in the measurements of the other forms, rather than an inherent difference in the spectra. Form III shows a relatively strong peak at 0.68 THz, form II has a weaker feature at a higher frequency, but form I does not appear to (although a peak has been detected by the algorithm).

In Figure 3.28, absorption data interpolated from the published figures in [111] is shown. These measurements were taken at different temperatures, at frequencies between 0.4 and 3 THz within an air-tight chamber. The data shows good agreement with the measured data in Figure 3.27. A relative lack of peaks in form II below 2 THz can be seen, with two peaks at 2.3 and 2.8 THz. This suggests that in Figure 3.27 these peaks are being capped by the noise floor. Form III shows weak features corresponding to the 1.32 and 1.79 THz and a strong peak at 0.68 THz, which corresponds to the peaks observed in Figure 3.27. Form I shows much broader features than those observed in Figure 3.27, which could be likely due to the higher sample temperature which will cause homogeneous broadening.

In Figure 3.29, the absorption coefficient of paracetamol form I from [111] is shown with this work's measurement of form I. The blue and red curves show the

referenced absorption data at 415 K and room temperature. The green curve shows the measurement from this work. The difference between the absorption spectra at 415 K is that there are distinct peaks at 1.8 and 2.4 THz, compared to the doublet at 1.84 and 2.12 THz, which is likely due to some associated red shift with temperature and the doublet merging into a broad feature. The two room temperature measurements show similar shape, however the reference measurement shows broader features which might be an artifact from the interpolation process used to extract the data. Notably, these measurements appear to have a systematic difference which increases with frequency. This could be due to a difference in the scattering in the samples [113] or due to thickness uncertainty, see § 3.12.

Lower concentration samples were measured to give a spectral fingerprint between 0 and 6 THz. 10 % concentration form I and II samples were measured at a sampling rate of 150 THz and a 20 % concentration form III sample was measured at a sampling rate of 75 THz. The full powdered sample fabrication process can be found in Appendix A. Both were measured using the LFRBB THz TDS system and with a frequency resolution of approximately 200 GHz. Form III was measured at a higher concentration because it was much less absorbent than the other forms.

In Figure 3.30, the absorption of 10% form I (blue) and II (red), and 20% form III (green). The axis have been limited to under 5.5 THz to prevent the 6 THz PTFE matrix peak dominating the curves. Peak detection has been performed on all the




Figure 3.30: Absorption coefficient of 10 % Paracetamol I (blue) & II (red), and 20 % Paracetamol III (green). α_{max} has been marked as dashed lines, peaks have been marked by vertical lines with circles.

measurements and the peaks have been marked by vertical lines and circles. There exists again similar groupings of peaks between the different forms. However, this becomes weaker at higher frequencies particularly at 2.9 THz where a feature is present in both forms I & II but not in III and at 4.5 THz where forms I and III share a doublet but form II does not. Despite these similarities, all of the forms appear to have different overall shapes which can be seen over the larger bandwidth. Form I displays a consistent series of peaks, while form II has 3 main features (other peaks appear to be fairly weak) and form III takes exhibits a dominant peak at 2.6 THz.

III 100 %	III 20 %	II 100 %	II 10 %	I 100 %	I 10 %
0.68	0.68	0.76	0.6	0.54	0.56
				0.83	0.87
1.01	0.99	1.09	1.1	1.12	1.13
1.32	1.34	1.43	1.4	1.46	1.44
1.79	1.73	1.79	1.73	1.84	1.76
2.37	2.24		2.28	2.12	2.19
	2.63		2.63	2.59	2.55
			2.89		2.88
	3.2		3.44		3.37
	3.72		3.72		3.69
	4.11		4.21		4.19
	4.39				4.44
	4.66				4.66
	4.99		4.99		5
	5.24		5.3		

Table 3.2: Table of absorption peaks in Paracetamol Forms I, II and III.

Table 3.2 shows the peaks detected in the measurements. Where possible the features have been grouped by frequency. From this table, there appear to be several absorption modes exhibited by the different forms at similar features. However, particularly at larger bandwidths, the relative intensities of these features are substantially different giving differently shaped spectra.

These results show the applicability of THz TDS to the detection of different polymorphic forms of pharmaceutical compounds. This is particularly important, as different polymorphs have different physical and chemical properties, which can have both manufacturing, storage and application implications [16].

3.11 Deconvolution errors

From equations 3.16, 3.8 & 3.9, the system transfer function, H_{sys} , should cancel out. However, it was observed that this does not occur as by including a system reflection within both the sample and reference measurements etalons are introduced into \tilde{n} , which would not be present when reflections are excluded. The transfer function is deconvolved from sample and reference measurements, both of which contain some degree of error.

These error terms are defined as the difference between the deterministic value of a measurement and the actual measured value. Random error due to uncertainty within the measurement will propagate to \tilde{n} , which is discussed further in Chapter 6. Other forms are systematic error within the measurement which can be further viewed as systematic measurement errors and systematic modelling errors. In the former, there exists systematic error or bias within the measurement in some form. For instance if there is systematic error in the time base [114], [115] features (in particular reflections) can be different between sample and reference measurements. In the latter, there exists an effect which is not being modelled. A common effect of this type would be resonance within the sample which produces multiple peaks in the time trace. In this case, the resonant sample response can be treated as additive error in addition to the sample response which fits the model. Another source of this type, would be to consider effects of the sample on the focused THz radiation. The sample will affect the focal length of the beam [100], which will in turn affect the focusing on the detector, introducing an effect which is not modelled.

By modifying equation 3.16 to include additive error terms within measurements, an estimate of H including error terms can be formed:

$$H' = \frac{E_{sam} + \varepsilon_{sam}}{E_{ref} + \varepsilon_{ref}} \quad (3.45)$$

where ε_{sam} and ε_{ref} are errors within the sample and reference measurements respectively and H' is transfer function with error in the measurements.

To understand how error affects \tilde{n} , a logarithmic form can be used. Starting with

the transfer function model, equation 3.11, and applying a logarithm to both sides:

$$\log_e(H) = \log_e\left(\frac{4\tilde{n}}{(\tilde{n}+1)^2}\right) - \frac{i\omega l}{c}(\tilde{n}-1) \quad (3.46)$$

This can be rearranged into the form:

$$\tilde{n} = i\frac{c}{\omega l} \left(\log_e(H) + \log_e\left(\frac{(\tilde{n}+1)^2}{4\tilde{n}}\right) \right) + 1 \quad (3.47)$$

Replacing H with H' will lead to two additional logarithmic terms, which are incorporated into the estimate of \tilde{n} with deconvolution errors, \tilde{n}' :

$$\tilde{n}' = i\frac{c}{\omega l} \left(\log_e(H) + \log_e\left(1 + \frac{\varepsilon_{\text{sam}}}{E_{\text{sam}}}\right) - \log_e\left(1 + \frac{\varepsilon_{\text{ref}}}{E_{\text{ref}}}\right) + \log_e\left(\frac{(\tilde{n}+1)^2}{4\tilde{n}}\right) \right) + 1 \quad (3.48)$$

This can be then be written including \tilde{n} :

$$\tilde{n}' = \tilde{n} + i\frac{c}{\omega l} \left(\log_e\left(1 + \frac{\varepsilon_{\text{sam}}}{E_{\text{sam}}}\right) - \log_e\left(1 + \frac{\varepsilon_{\text{ref}}}{E_{\text{ref}}}\right) \right) \quad (3.49)$$

If $\frac{\varepsilon_{\text{sam}}}{E_{\text{sam}}} < 1$ and $\frac{\varepsilon_{\text{ref}}}{E_{\text{ref}}} < 1$, this can be expanded into the form of a Taylor series:

$$\tilde{n}' = \tilde{n} + i\frac{c}{\omega l} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \left(\frac{\varepsilon_{\text{sam}}^k}{E_{\text{sam}}^k} - \frac{\varepsilon_{\text{ref}}^k}{E_{\text{ref}}^k} \right) \quad (3.50)$$

where k is a summation index.

3.11.1 Time localised error

If the error occurs at regions in time, as in the case of error associated with reflections, error terms will have the general form of:

$$\varepsilon = \varepsilon_0 e^{-i\omega t_{\varepsilon,0}} \quad (3.51)$$

where ε_0 is the time centered error spectrum and $t_{\varepsilon,0}$ is the time difference between ε_0 and non time centered spectrum (the shift in time to centre). ε_{sam} and ε_{ref} can be

combined into a singular complex sinusoidal term, in which case a general error at the k th index, ε_k , with time shift, t_ε , can be formed:

$$\varepsilon_k e^{-ik\omega t_\varepsilon} = \frac{\varepsilon_{\text{sam}}^k}{E_{\text{sam}}^k} - \frac{\varepsilon_{\text{ref}}^k}{E_{\text{ref}}^k} \quad (3.52)$$

where ε_k and t_ε are defined as:

$$\varepsilon_k = \left(\frac{\varepsilon_{\text{sam},0}^k}{E_{\text{sam}}^k} e^{-ik\omega \frac{(t_{\varepsilon,\text{sam}} - t_{\varepsilon,\text{ref}})}{2}} - \frac{\varepsilon_{\text{ref},0}^k}{E_{\text{ref}}^k} e^{-ik\omega \frac{(t_{\varepsilon,\text{ref}} - t_{\varepsilon,\text{sam}})}{2}} \right) \quad (3.53)$$

$$t_\varepsilon = \frac{(t_{\varepsilon,\text{sam}} + t_{\varepsilon,\text{ref}})}{2} \quad (3.54)$$

Where $\varepsilon_{\text{sam},0}$ and $\varepsilon_{0,\text{ref}}$ are the corresponding time shifted sample and reference errors, with time shifts $t_{\varepsilon,\text{sam}}$ and $t_{\varepsilon,\text{ref}}$.

By substituting this into Equation 3.49, it can be factored into a summation of complex sinusoids:

$$\tilde{n}' = \tilde{n} + i \frac{c}{\omega l} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \varepsilon_k e^{-ik\omega t_\varepsilon} \quad (3.55)$$

In the case that there is no time localised error in the reference measurement ($\varepsilon_{\text{ref}} = 0$) and there is in the sample measurement (for instance when the THz radiation resonates within the sample and is not accounted for in the processing), then equation 3.49 can be further simplified:

$$\tilde{n}' = \tilde{n} + i \frac{c}{\omega l} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \frac{\varepsilon_{\text{sam},0}^k}{E_{\text{sam}}^k} e^{-ik\omega t_{\varepsilon,\text{sam}}} \quad (3.56)$$

The deconvolved estimates of n and κ including deconvolution error, n' and κ' , can be derived by separating the real and imaginary parts of equation 3.55:

$$n' = n + \frac{c}{\omega l} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \sin(k\omega t_\varepsilon) \varepsilon_k \quad (3.57)$$

$$\kappa' = \kappa - \frac{c}{\omega l} \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \cos(k\omega t_\varepsilon) \varepsilon_k \quad (3.58)$$




Figure 3.31: Time traces of both a 1 mm thick silicon and the associated reference measurement. A system reflection is included in the time window of both measurements

These equations are reached when there is error which is local to a particular point in time, such as error which is associated with a system or sample reflection. In both cases the error occurs across a specific region of time, rather than across the entire time window (i.e. the error has the general form of equation 3.51). The error when propagated to the \tilde{n} , will have a sinusoidal component resulting in etalon artifacts.

3.11.2 Systematic deconvolution error when system echoes are present

Silicon in particular demonstrates the systematic error and produces very visible artifacts which can be seen when calculating with a time window which includes either system or sample reflections. In Figure 3.31, a time scan of a 1 mm thick silicon plate is shown (blue) with its corresponding reference measurement (red). Both of these measurements were measured with the LFRBB THz TDS system. The scans were measured with a bias voltage of 150 V chopped at 7 kHz, and were sampled at 25 THz. In both measurements, a system reflection from the 150 μm thick GaP electro-optic crystal used in the detection scheme is present. The extracted real refractive index and extinction coefficient are shown in Figure 3.32. Both of which show significant etalons when the system reflection is included in the time window. These are not present when both time windows are shortened such that the system reflection is excluded.

This is due to an error which occurs in the measurement of the reflection. This

system reflection is in principle independent of the sample and the transfer function of the system reflection relative to the initial pulse should be consistent. When deconvolving a measurement which includes a system reflection in both windows, it should be removed.

If, however, there is a small systematic change in this system transfer function, artifacts will appear in the extracted parameters. As the error in \tilde{n} is proportional to this systematic error, the error can be very small as long as it is significantly large relative to \tilde{n} .

This error will also occur at the reflections point in time, and as such the etalons should be an expression of equation 3.57. This can be seen by applying a Fourier transform to either n or κ . This technique has been used before in thickness extraction, (see § 5.7), to quantify the change in etalon magnitude [116]. In Figure 3.33, the Fourier transform of the finite difference (to suppress non-sinusoidal features) of n and κ is shown. This is measured on a pseudo time domain scale (this is different to an inverse Fourier transform, but is still measured in time). Features in this domain will correspond to sinusoidal like features in the underlying parameters and what is immediately noticeable is that the both n (blue) and κ (red) have features at the same location and are periodic on the pseudo time scale. An estimate of $t_{\varepsilon,\text{sam}} = t_{\varepsilon,\text{ref}} = 4.4$ ps is formed, by measuring the time difference between the negative peaks between the main pulse and reflection in the silicon measurement in




Figure 3.32: Estimated properties of silicon based on the measurement including a system reflection. Left - real refractive index. Right - extinction coefficient. A measurement which includes a system reflection in both time windows (blue) and without (red)




Figure 3.33: The Fourier transform of the finite difference of n (blue) and κ (dashed red). Both curves show very similar periodic features at approximate intervals of 4.4 ps (vertical blue dashed lines)

Figure 3.31, multiple of this have been marked on Figure 3.33 as blue dashed lines. From this it can be seen that the features in both Fourier transform of n and κ are separated by this interval. This relationship is approximate due to the error being deconvolved with the measurement response, in equation 3.50 this is shown as a division term.

These measurements demonstrate the need to apply windowing to remove system reflections, as small errors localised to the system reflection will propagate through to and dominate the calculated spectroscopic parameters. As these are displaced relative to the initial pulse in the measurement and that the error in \tilde{n} is directly proportional to ε_k , a series of sinusoids is formed in the spectroscopic parameters.

3.11.3 Systematic deconvolution error of sample resonance

When the THz radiation resonates within the sample, reflections (echoes) are present in the sample trace which are not present in the reference trace. These are not accounted for within the sample model (see equation 3.10) and thus can be viewed as a systematic error within the measurement. Like the case with the system reflections, this leads to etalons within n and κ .

In Figure 3.34, the time measurements for a 1 mm thick silicon (blue) plate with its corresponding reference measurement is shown (red). The measurement was




Figure 3.34: Time measurements of a 1 mm thick silicon plate (blue) and the corresponding reference measurement (red). A sample reflection is present in the sample trace, which is not present in the reference measurement




Figure 3.35: The extracted parameters of silicon. Left - real refractive index, right - extinction coefficient. The parameters based on data with sample reflection included have etalons (blue) while the parameters based on data without the reflection do not (dashed red).




Figure 3.36: The Fourier transforms of mean centered n (blue) and κ (dashed red). Both contain sharp periodic features with a period of 22.84 ps (blue dashed lines)

performed using the HFRBB THz TDS, using a 600 V 7 KHz bias voltage, and was sampled at 25 THz. The sample measurement contains a second peak, which corresponds to a reflection within the sample itself. This is not present within the reference measurement.

When processed to extract values of n and κ , this leads to sinusoidal etalons. In Figure 3.35, the n (left) and κ (right) are shown. Both of the measurements based on Figure 3.34 show large etalons with a short oscillation period. If the time window is shortened to exclude the sample reflection, these etalons are removed. This is shown in Figure 3.35 as red dashed lines. In Figure 3.36, the Fourier transform of the finite differences of n (blue) and κ (red) are shown. A sharp set of features are present in both parameter curves and again both are periodic. $t_{\epsilon, \text{sam}}$ was calculated to be 22.84 ps from Figure 3.34 by measuring the time difference between the reflection and the initial peak in the sample scan. In this particular case, the error caused by sample resonance is solely in the sample measurement, see equation 3.56, so $t_{\epsilon, \text{sam}}$ will correspond to the periodicity of the etalons. This has been marked on Figure 3.36 as vertical blue dashed lines, and this matches up to the sharp peak present in the periodic feature from both parameter curves. This is a more exact relation than in the system reflection case, as $\frac{\epsilon_{\text{sam}}}{E_{\text{sam}}}$ will be the transfer function of a reflection within the sample relative to the initial pulse.

This example exhibits a very common effect, as many samples have a refractive index high enough to produce reflections and especially when the sample is thinner, this will produce measurable distortions in the sample parameters. In Chapter 5, this etalon effect is used in conjunction with a sample resonance model to accurately measure sample thickness.

3.12 Effect of thickness uncertainty

In equations 3.18 and 3.20, the sample thickness is required to calculate both n and κ . Any inaccuracy in the thickness will propagate to the n and κ , resulting in uncertainty in these terms. The thickness uncertainty can be propagated through equations 3.18 and 3.20, to form estimates of the uncertainty of n and κ . To do this,




Figure 3.37: (a) - The hypothetical relation between thickness and refractive index is non-linear. (b) In the neighbourhood of the sample thickness (marked as a dashed line on (a)), the relation can be modelled as being approximately linear, leading to a proportional relation between σ_l and σ_n .

these equations are approximated by simpler linear relations.

In Figure 3.37, a hypothetical relation between thickness and refractive index is shown (a). Increasing thickness will asymptotically decrease the value of n towards 0. The relation between uncertainty in thickness, σ_l , and uncertainty in n , σ_n , is therefore non-linear. It is however possible to construct an approximate linear relation at the measured value of thickness by constructing a first order polynomial, see Figure 3.37 (b). The relation between σ_n & σ_l is then approximated by replacing the non-linear relation with a simpler multiplication by the absolute differential at this point:

$$\sigma_n \approx \left| \frac{dn}{dl} \right| \sigma_l \quad (3.59)$$

An approximate relation for σ_n can therefore be derived by differentiating equation 3.18:

$$\sigma_n \approx \left| \frac{c}{\omega l^2} \angle H \right| \sigma_l \approx \frac{\sigma_l}{l} (n - 1) \quad (3.60)$$

A similar relation can be derived for σ_κ :

$$\sigma_\kappa \approx \left| \frac{c}{\omega l^2} \ln \left(\frac{(n+1)^2}{4n} |H| \right) \right| \sigma_l \approx \frac{\sigma_l}{l} \kappa \quad (3.61)$$

In the above approximation, additional terms relating σ_n have been neglected, as these will be much smaller than this primary term.

These can both be arranged into the following forms:

$$\frac{\sigma_n}{n} \approx \frac{\sigma_l}{l} \left(1 - \frac{1}{n} \right) \quad (3.62)$$

$$\frac{\sigma_\kappa}{\kappa} \approx \frac{\sigma_l}{l} \quad (3.63)$$

This leads to the useful observation that the relative uncertainty (uncertainty divided by its measured value) of spectroscopic parameters due to thickness uncertainty will be proportional to the relative uncertainty in thickness. A common technique to increase the α_{max} of a measurement is to use a thinner sample. However, unless the instrument used to measure the thickness of the sample is changed, the absolute uncertainty of the thickness will remain the same. This leads to the relative uncertainty with respect to the thickness increasing and thus the uncertainty in the spectroscopic parameters.

There also exist conditions where the thickness can not be measured without some significant error. In particular when the sample is subject to experimental conditions which effect the thickness, for example temperature change.

3.13 Thickness uncertainty of α lactose monohydrate

A 10% α lactose monohydrate sample was measured, the thickness measured was 0.45 mm to within 0.005 mm using a micrometer. The sample was measured in the




Figure 3.38: σ_n of 10 % α - lactose monohydrate, an estimate based of different calculations of n (green) and an estimate derived from the uncertainty in thickness (red dashed)




Figure 3.39: σ_{κ} of 10 % α -lactose monohydrate, an estimate based of different calculations of κ (green) and an estimate derived from the uncertainty in thickness (red)

LFRRB THz TDS biased with a 7 KHz 150 V square wave and was sampled at 25 THz.

In Figure 3.38, the uncertainty for n of the sample is shown. The green line is an estimate of uncertainty based on different estimates of n calculated using different thicknesses (0.445, 0.45 and 0.455 mm). The dashed red line is an estimate based on equation 3.60, and is calculated based on an uncertainty of 0.005 mm. Both estimates are nearly identical.

In Figure 3.39, the uncertainty for κ of the sample is shown. The green line is an estimate of uncertainty based on different estimates of κ calculated using different thicknesses (0.445, 0.45 and 0.455 mm). The red line is an estimate based on equation 3.61, and is calculated based on an uncertainty of 0.005 mm. At low frequencies the estimates are slightly different, at higher frequencies the estimates converge. This is because the refractive index uncertainty (assumed negligible in equation 3.61) will be present, however the absorption value is very small. At higher frequencies, the estimate in equation 3.61 is more effective as κ is larger.

These estimates have found that the uncertainty in κ is roughly 1% while in n it is closer to 3% due to the relative uncertainty in the thickness measurement and represents a significant systematic error within the measurement. This could of course be mitigated by using a thicker sample, which would decrease α_{max} .

3.14 Conclusion

In this chapter the target spectroscopic parameters were identified with their relevant relation to each other. A sample model was then developed with regards to the THz electric field, which takes \tilde{n} and l as parameters. Approximations for the real and imaginary parts of \tilde{n} , n and κ were derived. These can be used to approximate these parameters from sample and reference measurements.

These were then applied to a α lactose and a discussion of unwrapping corrections for n was presented. The absorption spectra of different polymorphic forms of paracetamol were presented.

The method presented here to estimate both n and κ have two significant issues which were discussed. First there exist deconvolution errors, which arise from error within the measurement. If these errors are localised in time (only occur in particular regions of the time scan), then in \tilde{n} there will be a sinusoidal component errors. Secondly, sources of time localised errors were discussed, system and sample reflections. While both are removable by windowing in the time domain, this will limit the frequency resolution of the measurements.

The methods presented in this chapter are considerably limited by sample reflections and the relative thickness uncertainty. Often it is not possible to simply use a thicker sample to mitigate these effects, and there is an optimal sample thickness for spectroscopic extraction [117]. Another issue is more complex samples, for instance a layered structure, which can produce complex sample resonance. In Chapter 5, more complex methods are discussed which exploit sample resonance and extract sample thickness.

Another issue is the application of THz TDS to specific application. These applications can often involve testing dangerous or expensive samples, so a suitable method of testing and calibrating instrumentation is prudent. In the next chapter, a method of algorithmically creating explosive simulants for security applications is presented.

Chapter 4

Creation of explosive spectrum simulants

This chapter presents an algorithm for generating a mixture of compounds which acts as simulant for Semtex-H in a THz TDS system. The algorithm in this chapter was published in [118]. The methods presented in Chapter 3 were used to extract the absorption coefficient of the simulant mixtures.

THz has gathered considerable interest for security applications, particularly for the detection of concealed explosives [5], [17], [119]. Common explosive compounds are also identifiable by THz spectroscopy [5].

When testing and calibrating an explosive detection system in "real world" conditions, it is preferable not to use an actual explosive and to instead use an inert test sample to maintain restricted access to explosives. In order to practically test explosives at THz frequencies, a non-hazardous analogue is required. This non-explosive simulant should have similar absorption spectra to the target explosive. This does not mean spectra which are objectively the same, but which have spectral features at the same locations in frequency with similar relative intensity to the target explosive.




Figure 4.1: THz spectra of RDX (green), PETN (red) and Semtex-H (blue). The spectra of PETN and Semtex-H have both been offset by 50 and 100 cm^{-1} respectively for clarity

4.1 Terahertz spectroscopy of Semtex-H and related explosives

In Figure 4.1, the absorption spectra of cyclotrimethylenetrinitramine (RDX) (green), pentaerythritol tetranitrate (PETN) (red) and Semtex-H (blue) are shown. Both the RDX and PETN samples were 20 % concentration samples with a PTFE matrix. Semtex-H is a mixture of both RDX and PETN, with a number of additional non THz absorbent components [17]. A simulant of Semtex-H will have to reproduce features at 0.79, 1.44, 1.98, 2.15 and 2.92 THz.

Absorption modes at THz frequencies are often attributable to crystalline structure and this is true for the target Semtex-H [97], [120], creating a simulant chemically will therefore be extremely challenging. A simpler route is to instead find a mixture of inert compounds which mimic select features in the spectra. By finding a combination of these compounds, the full spectra can be mimicked.

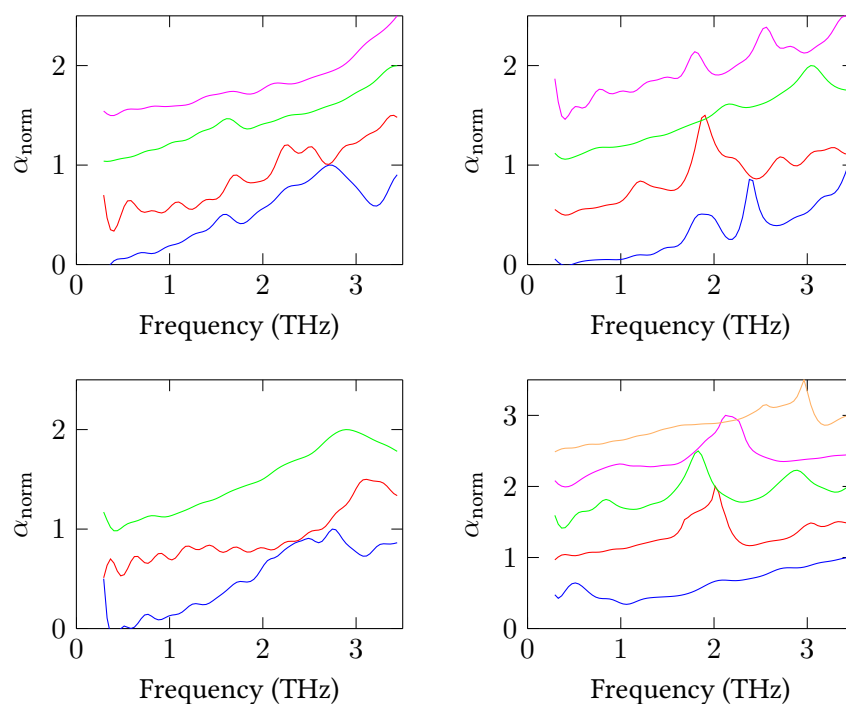


Figure 4.2: Normalised THz absorption spectra of a range of organic molecules. (top-left) amino acids: Leucine (magenta), Isoleucine (green), Glutamine (red) and Asparagine (blue). (top-right) sugars and carbohydrates: Sucrose (magenta), Microcellulose (green), Lactose Anhydrous (red) and Glucose Monohydrate (blue). (bottom-left) a range of organic compounds: Benzocaine (blue), Caffeine (red) and Lidocaine (green). (bottom-right) non-explosive compounds that are chemically similar to known explosives: Cyabouric acid (yellow), Pentaerithritol (magenta), Artemisinin (green), Erithritol (red) and Hexamethylenetetramine (blue). All spectra were normalised to have a maximum absorption of one, and are offset by 0.5 for clarity.

4.2 A compound database

In Figure 4.2, the spectra of a range of sugars and other commonly available organic compounds are shown, which between them have a wide range of spectral features. Having a large number of components increases the number of possible solutions, and therefore the chance of a good solution. This increase in potential solutions also makes the selection of an optimal solution more difficult. The 16 compounds considered in Figure 4.2 were used to form a library of spectra, from which Semtex-H simulants were designed. These compounds were chosen because they had spectral features at similar frequencies to the target spectrum of Semtex-H and because they are both physically and chemically inert when mixed together [118]. A number of the samples shown in Figure 4.2 were diluted using PTFE in order to increase the usable measurement bandwidth of each sample.

4.3 Modelling mixture spectra

To evaluate mixtures of compounds, the combined spectra of a mixture must be considered. The included components of a mixture are assumed to contribute to the overall absorption spectra in proportion to its concentration in the mixture. So that combinations of compounds leads to summation of absorption spectra (scattering is assumed negligible). Thus a mixture spectrum, M_N , can be modelled as a weighted sum of the constituent component spectra:

$$M_N(\omega) = \sum_{j=1}^{j=N} C_j S_j(\omega) \quad (4.1)$$

where N is the number of components, S_j is the pure THz absorption spectrum of each component j taken from the library of spectra and C_j is the concentration (normalised to one) of each component. This method of modelling provides a simple and effective way to locate a viable simulant mixture.

A more complex but generally accurate method of modelling a mixture spectra is to use effective medium approximations, such as Maxwell Garnet [121] or Bruggeman [122]. These methods can produce more accurate mixture spectra but would

introduce additional complexity into the mixture modelling process. These were not used in the simulant generation algorithm, instead multiple viable mixtures were generated using the above method to compensate for a simple mixture model.

This was also used to estimate the pure absorption spectra from a diluted component spectra, the matrix material PTFE was assumed to have negligible absorption. The estimate of a pure spectra based on the diluted spectra was then simply the absorption spectra of the diluted sample divided by the concentration normalised to 1 (i.e. a 10% sample would have its spectra divided by 0.1).

4.4 Quantifying similarity between spectra

Once a mixture is modelled, its spectra are compared to the target spectra using a function to quantise the similarity. Cross-correlation was chosen as the basis for this function. As with equation 3.15, one dataset is shifted relative to the other and a summation is computed giving a quantification of the similarity. This is however performed in the frequency domain, rather than the time domain:

$$C_{AB}(r) = \sum_{k=0}^{k=L} A(k)B(k+r) \quad (4.2)$$

where r is the frequency shift between the spectra A and B , L is the number of samples in each of the spectra, k is the frequency index and C is the cross-correlation. The frequency shift is used to provide a similarity score with a degree of frequency tolerance. A weighted normalised correlation coefficient was used to provide a normalised similarity score between 0 (spectra are completely different) and 1 (spectra are identical) [123]:

$$C_{AB}^{Ws} = \frac{\sum_{r=-l}^{r=l} C_{AB}(r)W(r)}{\sqrt{\sum_{r=-l}^{r=l} C_{AA}(r)W(r)}\sqrt{\sum_{r=-l}^{r=l} C_{BB}(r)W(r)}} \quad (4.3)$$

where $W(r)$ is a weight function, defined as:

$$W(r) = 1 - \frac{|r|}{l} \quad (4.4)$$

Here l is the frequency shift tolerance, i.e. the absolute shift of a spectral feature that will be tolerated by the fitness function before a discriminating penalty is applied. The weight function is used to introduce a preferential score bias to small frequency shifts. A simplified schematic of how this was used can be seen in Figure 4.3. In Figure 4.3.a, the two spectral features, A and B are identical in spectra but shifted by a value r . As r is greater than l , the weighting function is reduced and the final score is poor. In Figure 4.3b, r is less than l and the spectral features are identical; however, r is still quite large. In Figure 4.3c, r is smaller than in Figure 4.3b but the spectral features are not identical. The spectra in Figure 4.3c will always give a better score than Figure 4.3b. C_{AB}^{Ws} gives a score between 1 and 0. Thus, a mixture of compounds that form a suitable Semtex-H simulant must give a resultant spectrum with a high value of C_{AB}^{Ws} .




Figure 4.3: A diagram to illustrate how spectra is evaluated in relation to frequency difference (r) between the target feature (A), dotted) and simulant feature (B), solid), and frequency tolerance (l). (a) r is greater than l , the cross-correlation is truncated using $W(r)$, giving a poor score. (b) The simulant feature has an identical shape, and thus good cross-correlation, but large r , giving a poorer score due to $W(r)$ than (c), where the simulant feature is only partially similar in shape, but has smaller r

4.5 Selection and optimisation of mixtures

Genetic algorithms are a class of minimiser/maximiser algorithms based on the principles of evolution [124]. A genetic algorithm operates on a population of different solutions to a given problem, and operates by replacing poor performing solutions with better solutions generated from the remaining higher ranked population. By iteratively performing this process, an optimal solution can be found.

In Figure 4.4, a flowchart of the genetic algorithm used in this work is shown. The genetic algorithm used in this work operated iteratively on a population of

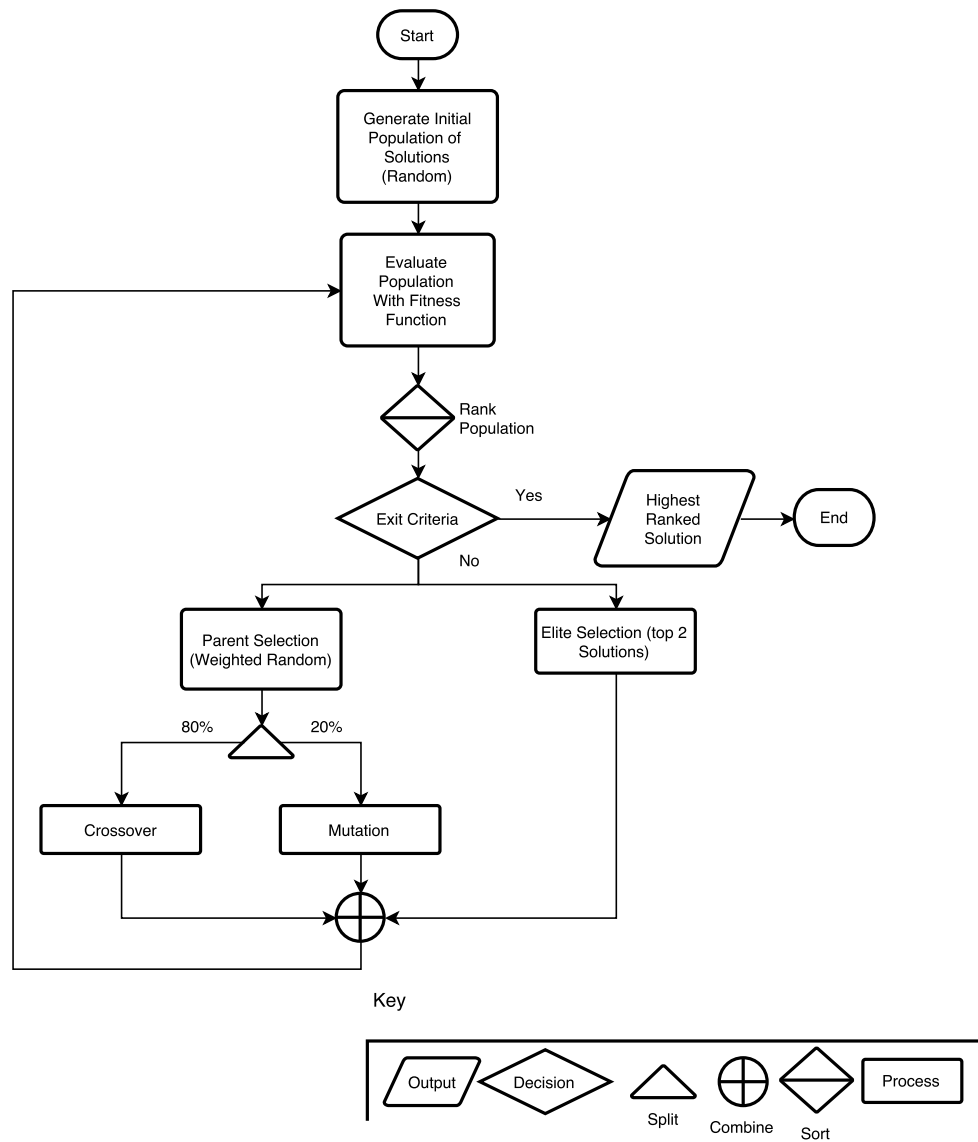


Figure 4.4: Flowchart of the genetic algorithm used in this work

solutions until the top solution was stable over 50 iterations. An initial population is generated randomly with a uniform distribution [125]. With each iteration of the population, individual solutions were evaluated by a fitness function and sorted by score. The top two ranked solutions are selected as “elites”, and pass directly to the next iteration of the population. “Parents” are then selected from the previous iteration of the population to form a new population through either “crossover” or “mutation” [125]. 80 % of the new non-elite population is formed by crossover, a randomised combination of the two parents from the previous iteration. The remaining 20 % are formed by mutation, which is random change to a single parent. The process continues to iterate over the population, until the fitness function score of the top mixture has changed by less than 10^{-6} over 50 iterations. In the first stage of the selection process where chemical components of the mixture are selected (see below), a variation of the genetic algorithm which optimises integer problems was used, during genetic algorithm operations such as "crossover" or "mutation" solutions are forced to take integer values using a penalty function [126]. In the second stage, where the concentration of the fixed components are optimised, a non-linear constraint was used to ensure that all components were positive and their sum was equal to one [127].

To evaluate the performance of each solution, it is evaluated against the target spectrum (the normalised spectrum of Semtex-H) by a fitness function, which provides a fitness score. In this case, the fitness function F is given by:

$$F = 100 (1 - |C_{AB}^{Ws}|) \quad (4.5)$$

F is simply the weighted normalised correlation coefficient, C_{AB}^{Ws} , scaled between 0 and 100, with 0 representing a perfect match for the normalised spectrum of Semtex-H, and 100 representing no correlation. This was done to convert a normalised maximisation problem into a normalised minimisation problem, which is required by the implementation of the genetic algorithm. The genetic algorithm was implemented using MATLAB’s global optimisation toolbox [128]; To create viable simulants, the problem was split into sequential stages: (i) selecting the components of the mixture

and (ii) optimising the relative concentrations of the mixture.

In the component-selection stage, the genetic algorithm was used to form five compound combinations with repetitions, with each component making up 20% of the mixture. This was implemented as optimising a combination of 5 integers with repetitions. This reduces the complexity of the optimisation problem significantly, rather than trying to find an optimum selection of any number of unique compounds with corresponding concentrations which sum to 100 %, the problem becomes a search for the optimal combination of 5 integers between 1 and 16 (the number of compounds). Each component in the spectra library was assigned an integer value and the combinations of integers were converted into a mixture spectra. Repetitions were used to form concentrations which are a multiple of 20 %. For instance, if the algorithm generated a mixture spectra formed from a combination of components 1, 2, 3, 1, and 1; this would be converted to a three component mixture containing components 1, 2, and 3 with concentrations of 60, 20 and 20% by mass, respectively.

The mixture spectra were then modelled using equation 4.1, normalised and then evaluated against the normalised spectrum of Sementex-H using the fitness function, F in equation 4.5. The maximum normalised spectra were used in the fitness function, as a comparison between the relative intensities of the spectrum was desired. A flowchart of the process is shown in Figure 4.5.1 and is shown under component selection.

An initial population of solutions was uniformly and randomly generated, as the process of optimisation is dependent on the initial population, increasing its size (or equivalently repeating the process) will yield a more extensive search for the optimum solution. After evaluating the population solutions, the genetic algorithm then modifies the population of mixture spectra by removing solutions with poorer scoring spectra and generating additional spectra from the remaining population. The evaluation, using fitness scores, and modification of the population was repeated until a consistent best-ranked solution appeared with a change in fitness function less than 10^{-6} over 50 iterations.

This process of selecting a suitable component combination was repeated for a total of 30 times, deliberately making use of the non-deterministic selection process

to produce multiple viable solutions. The highest-ranked mixture spectra from each of the 30 repeats were then compared, duplicates were removed and the remaining mixture spectra ranked. A flow chart of the entire component selection process is shown in Figure 4.5.1.

The top two scoring solutions were then used as a basis for the second stage, in which the genetic algorithm was used to optimise the concentration of each component of the simulant mixture. Randomisation of the two mixture concentrations was used to form the initial population. The genetic algorithm was performed in a similar manner to the first stage, but changes were now made to component concentrations for a fixed component selection. These concentrations were then used to form the spectrum of the mixture, which was then compared to the normalised spectrum of Semtex-H using the fitness function F . This is shown in Figure 4.5.2 as a flowchart under concentration optimisation.

As in the first stage, the process was repeated until a consistent best solution with a change in fitness function of less than 10^{-6} over 50 iterations was achieved. This second stage was then repeated a further four times, and the mean concentrations from each run used as the concentration for the final solution. A flowchart of this second stage is shown in Figure 4.5.2.

The whole process was then repeated ten times, over eight different frequency tolerances and the final solution of each run was collected. At the end of this process four different solutions were found and are summarised in Table 4.1.

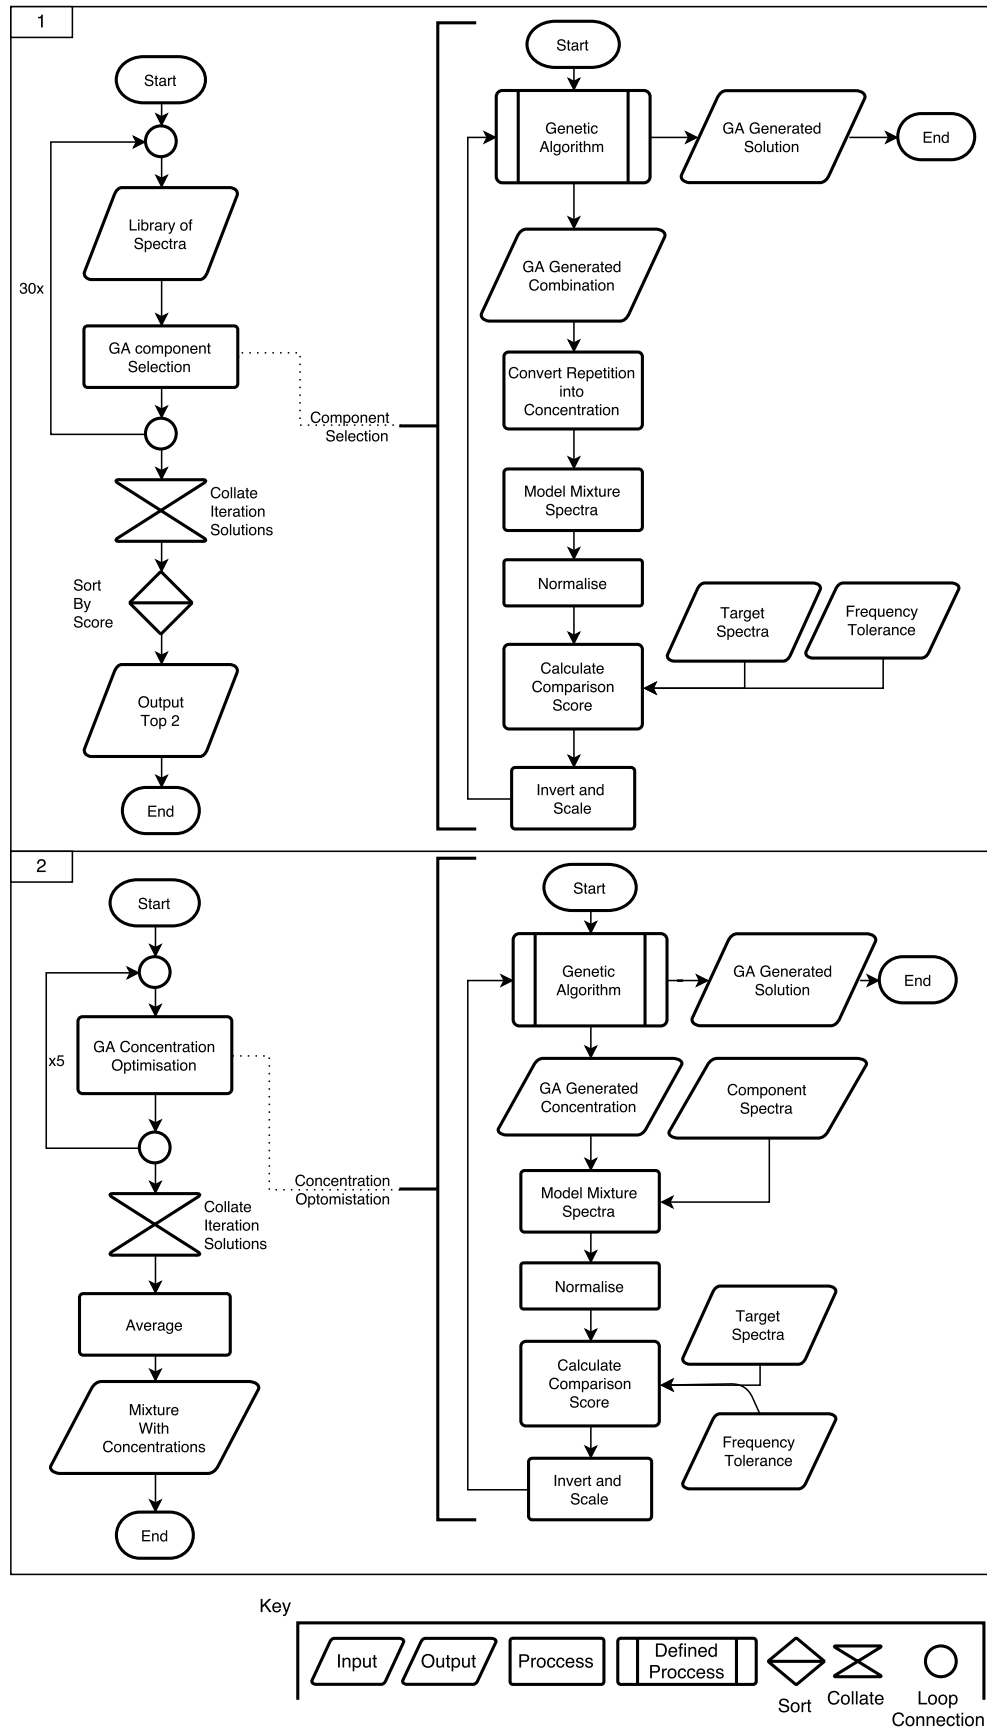


Figure 4.5: Flowcharts of both the first and second stages of the simulant generation algorithm. 1 - The first stage flowchart, the left flowchart shows the process which performs the component selection and the right flowchart shows how the individual component combinations are evaluated. 2 - The second stage flowchart, the left flowchart shows the concentration optimisation and the right flowchart shows how mixtures are evaluated.

4.6 Results

Table 1 shows the components, concentrations and fitness function score (F , where 0 is a perfect match) for each of the four unique mixtures that were obtained using the genetic optimisation algorithm. The process converged on similar mixtures over different frequency tolerances, where components were the same, with preference being given to the mixture at the lowest frequency tolerance. Three mixtures were obtained using the lowest frequency tolerance of 37 GHz (mixtures 1, 2 and 3 in table 1). A further solution (mixture 4) was produced once the frequency tolerance was increased significantly to 620 GHz but which concomitantly resulted in a much lower fitness function score. This lower score is because a greater frequency shift tolerance weakens the constraint of matching the frequency of features for the target spectra, which can give a better fitness score and a greater diversity of solutions. This is at the expense of matching target feature location in frequency. In the case of mixture 4, the feature at 1.7 THz produced a good match with Semtex-H's doublet at 2 THz because of the greater frequency shift tolerance allowed in this case. The fitness score for mixtures obtained using different frequency tolerances are not directly comparable. It was found, however, that the frequency shift tolerance was not required to be increased to achieve multiple low fitness solutions. This is unsurprising given the large full width half maximum (FWHM) of the majority of the spectral features being imitated. All four mixtures contain artemisinin, which closely simulates the low-frequency spectral feature of RDX (0.79 THz) along with a num-

Mixture number	Component name (concentration)	Fitness function, F (0 - 100)	Frequency tolerance (THz)
1	Art (21%), Micro(79%)	1.88	0.037
2	Art (18%), Micro(64%), Hex (18%)	1.8172	0.037
3	Art (17%), Micro(51%), Hex (27%), Asp (5%)	1.8223	0.037
4	Art (23%), Hex (53%), Lid (24%)	0.40923	0.62

Table 4.1: The four algorithm-generated mixtures, their components, mass percentages and the fitness function value of the spectra compared to the normalised spectrum of Semtex-H. Art - Artemisinin, Micro - Microcellulose, Hex - Hexamethylenetetramine, Lid - Lidocaine

ber of higher-frequency spectral features of both RDX and PETN. Microcellulose was also a significant component of three of the mixtures, as it usefully provided a weak contribution to peaks at similar frequencies to those in Semtex-H. These were located at 0.79, 1.44, 1.98, 2.15 and 2.92 THz, as well as a background that increases with frequency, similar to that seen in Semtex-H.




Figure 4.6: The four predicted spectra for the simulant mixtures are shown as solid lines. The dashed lines show the measured spectra of the mixtures once they were made, diluted with PTFE to 25% by mass and pressed into solid pellets. The solid blue line shows the normalised spectra of Semtex-H previously shown in figure 4.1 for comparison. Spectra have been normalised and offset for clarity. Red - Mixture 1, Green - Mixture 2, Magenta - Mixture 3 and Yellow - Mixture 4.

Given the non-deterministic nature of these algorithms, a relaxation of the constraint on component numbers was considered. The number of components in the first part of the process was deliberately kept low in order to make the resultant mixture easier to fabricate. When simulant generation was performed with a larger number of components (up to 20), the resulting mixtures contained very low concentrations and scored no better than those developed using the method described above. If the number of features is compared to that of the target spectra, the reason for this becomes clear. There are only four broad peaks in the target spectra; therefore, produced simulants will only require a similar number of components (assuming a single component produces each feature) to mimic these features. The resulting

predicted spectra for each of the four mixtures are shown as solid lines in Figure 4.6, with each spectra showing four broad spectral features centred roughly at 0.75, 1.5, 2 and 3 THz. Each mixture was then produced and diluted to 25% concentration with PTFE before being pressed into a solid pellet. The THz spectra for each of these samples are shown as dashed lines in Figure 4.6. Figure 4.6 also shows the normalised spectra of Semtex-H (blue solid line) from Figure 4.1 with all spectra being offset of 1 for clarity and comparison.

The experimentally measured spectra of each of the four mixtures are not identical to the predicted spectra; peaks in the spectra are generally weaker and less prominent. These discrepancies are likely caused by a number of factors. First, it has been assumed that the absorption intensity scales linearly with concentration both in the formation of the library and then in equation 4.1 when forming the mixture spectra. Secondly, to construct the library compounds were diluted with PTFE which can cause slight changes in peak positions [120] compared to the spectra of the pure material. Finally, mixing a number of substances will affect the THz scattering [113], particularly if there is a high contrast in refractive index, and in addition may also lead to changes in spectral peak shape [96].

4.7 Conclusion

In this chapter a algorithm was outlined to generate suitable simulants for a target spectrum. The simulants consisted of mixtures of inert compounds which mimicked the normalised spectrum of the target. This was applied to Semtex-H to generate four viable mixtures. Samples of these mixtures were formed and measured. The resulting sample mixtures were compared against the target spectrum and the modelled mixture spectra. It was found that resultant mixtures produced features which mimicked the spectra, however it was limited by the effects of scattering and the combinational effects of a mixture.

This algorithm could be used to generate viable simulant mixtures for a wide range of target samples. This could include samples which are hazardous, restricted (such as drugs of abuse), expensive or unstable at calibration conditions. Additionally,

this technique could be used to decompose a mixture into its constituent components using spectroscopic measurements.

Chapter 5

Transfer function methods of extraction

This chapter discusses different methods of indirectly extracting \tilde{n} of any sample using error minimisation techniques. This is used to extract the thickness of a sample *in-situ* and to extract \tilde{n} of an uncharacterised layer in a multi-layer sample. New algorithms are developed to extract both the thickness and \tilde{n} of a sample. These algorithms have led to two conference presentations [129], [130].

5.1 Overview of fitting resonant sample transfer function models

In order to extract the \tilde{n} of a sample, approximate methods discussed in Chapter 3 can be used. These can be accurate, however, are limited by resonance within the sample and the relative thickness uncertainty. These can both be resolved by fitting a sample transfer function model.

Fitting a transfer function model allows more complex sample models than the one presented in § 3.2. This allows THz resonance within the sample to be considered, allowing the time window to include reflections from the sample. The use of a model which includes this resonance will mitigate the etalons in \tilde{n} , allowing better frequency resolution to be obtained. The reflections from the sample also represent additional information about the sample, this can be used in conjunction with a

resonant model to extract the thickness of the sample *in situ*. This allows the sample's thickness to be measured accurately at the THz focus, under conditions which may prevent accurate thickness measurement by other means.

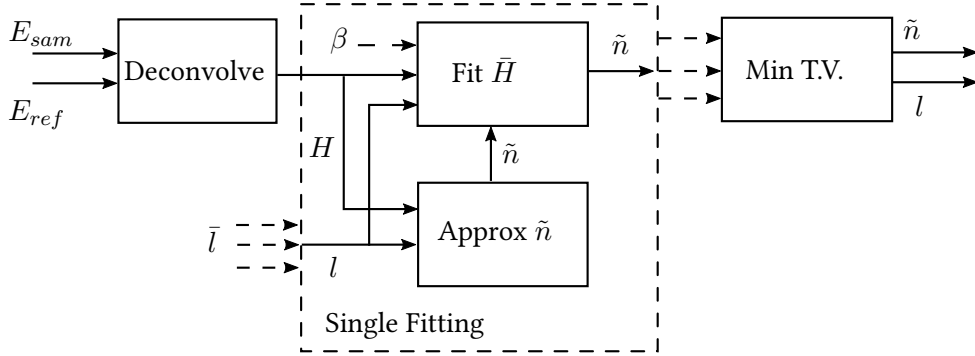


Figure 5.1: The processing flow diagram of fitting a resonant sample model. Deconvolution is followed by fitting a model to the transfer function. This can be repeated using different assumed thicknesses, \bar{l} , to find a minima in total variance, TV., which corresponds to the sample thickness, l .

In Figure 5.1, an overview of the fitting and thickness process is presented. The transfer function of the sample, H , is calculated from reference, E_{ref} , and sample, E_{sam} , measurements. In the dashed central box an extraction for a measured thickness, l , is shown. An initial estimate of \tilde{n} is produced using the methods in § 3.6. Using these as a initial starting point, the sample transfer function model, \bar{H} , is fitted to H by minimising the error between the two, to indirectly extract \tilde{n} . A calibration coefficient, β , can be used to increase the accuracy of the extraction by accounting for a Gaussian beam profile. When using a resonant model, it is possible to extract the thickness by performing a direct search. In this case, the process is repeated for different assumed thicknesses, \bar{l} , and evaluated by minimising total variance, TV., to extract both the l and the \tilde{n} of the sample.

5.2 Fitting a sample transfer function model

In order to fit a sample transfer function model to the calculated sample transfer function, a measure of error must be minimised with respect to \tilde{n} of a sample. This measure of error can be one of several functions, however three will be considered in this work, all of which are a measure of the difference between the modelled and measured transfer function.

In order to minimise the error with respect to \tilde{n} , a minimisation algorithm is used. In the literature both simplex (direct search) [63] and gradient descent [101] have been used (it has also been reported that any error minimisation method works [131]). In this work, the Levenberg–Marquardt algorithm [132], [133] has been used because it was found to generally give the faster convergence. These algorithms find local minima in the error function. This has two important connotations, firstly there may exist a better solution (a minimum which is lower than the one located) and secondly, there can exist residual error after the minimisation process. In practice, however, these can be neglected. The former can be avoided by using techniques which restructure the error to have a singular minimum, and the latter is not significant, as the models are fitted at each frequency point, so there is no compromise between multiple data points.

A minimum (or equivalently maximum) is defined when the differential of the error function with respect to extracted parameter (in this case \tilde{n}) is equal to 0:

$$\frac{d\mathbb{E}}{d\tilde{n}} = 0 \quad (5.1)$$

Where \mathbb{E} is the error. Using the chain rule, this can be shown to be:

$$\frac{d\bar{H}}{d\tilde{n}} \frac{d\mathbb{E}}{d\bar{H}} = 0 \quad (5.2)$$

Where \bar{H} is the transfer function model. This shows that if the error function has a minimum with respect to \bar{H} it also has a minimum with respect to \tilde{n} .

A range of error functions can therefore be used, provided that they have a minimum with respect to \bar{H} . This also inherently means that if different error functions have the same location of minima with respect to \bar{H} , they will produce the same minima with respect to \tilde{n} . In this case, the problem is said to be critically determined, the number of error variables is equal to the number of fitted parameters. The solution to the problem is thus unique as there is no redundancy.

The models used within this work are also non-interdependent with respect to frequency and thus, so are the minima. i.e. finding the minima over all frequencies,

is the same as finding the minima at each frequency [63].

To extract the \tilde{n} of a sample and fit a model to H , the argument minimum is located. This is the value of \tilde{n} which minimises the error between \bar{H} and H :

$$\tilde{n}(\omega) = \underset{\tilde{n}(\omega)}{\operatorname{argmin}} \mathbb{E} (H(\omega), \bar{H}(\tilde{n}, \omega)) \quad (5.3)$$

Where \mathbb{E} is some measure of error between the model and measured transfer functions.

As the extracted \tilde{n} is phase unwrapping sensitive, the error function must be as well. The error functions suggested in the literature are either the absolute difference of complex logarithms [101] or the absolute difference between magnitudes and phases [63]. Alternatively, simply using the square of these error functions produces the same answer (the preferred error minimisation technique, Levenberg-Marquardt, uses the square). The former takes the form:

$$\mathbb{E} = |\log_e(H) - \log_e(\bar{H})| = |\log_e |H| - \log_e |\bar{H}| + i\angle H - i\angle \bar{H}| \quad (5.4)$$

Where H is the experimental transfer function (the calculated transfer function) and \bar{H} is the modelled transfer function. And the latter takes the form:

$$\mathbb{E} = ||H| - |\bar{H}|| + |\angle H - \angle \bar{H}| \quad (5.5)$$

Both of these error functions are similar and will both have a minimum at:

$$|\bar{H}| = |H| \quad (5.6)$$

$$\angle \bar{H} = \angle H \quad (5.7)$$

Which is equatable to:

$$\bar{H} = H \quad (5.8)$$

Phase wrapping can affect the extracted \tilde{n} , this can be observed from equation 5.7, as if $\angle \bar{H}$ is not unwrapped, multiple values of \tilde{n} can produce the same value, leading

to local minima in the overall error.

There exist several solutions to this problem. Arguably the simplest is to try to provide an initial estimate of \tilde{n} which is so close to the actual value, that wrapping issues never occur. This is the approach suggested by Duvillaret *et al.* [60], by using the lower adjacent frequency extracted value as an initial estimate, unwrapping issues can largely be avoided. This method however relies on two things, first an initial point must be extracted correctly, and secondly, that the difference between \tilde{n} at adjacent frequencies is negligible. However, this can lead to error at lower frequencies propagating forward to higher frequencies. This method also enforces a dependence on lower frequencies being calculated first, other than wrapping errors there is no reason to calculate sequentially and if this dependence was removed, the problem becomes poised to being parallel-ised.

A second method would be to form an unwrapped estimate of the phase in equations 5.4 and 5.5. n is approximated as being constant with respect to frequency and the model is evaluated at lower frequencies with the same value of \tilde{n} , starting at 0. The error can then be calculated using the unwrapped phase from both the model and measured transfer functions. However, this is computationally expensive and poorly scaled, at each frequency point all lower frequencies have to be re-evaluated. It also suffers the same issues explored in § 3.9, with conventional unwrapping (i.e. it assumes discontinuities are smaller than 2π). The advantage is this method can be evaluated in parallel as the extraction is not frequency interdependent, and is robust when a poor initial estimate is provided.

An extension of this method, is to extrapolate from a constant n to a phase with constant gradient. The model can then be evaluated at the current frequency, and at an adjacent frequency with the same \tilde{n} . The gradient of the phase can then estimated by taking the difference between the two points. This can then be used to extrapolate back to 0, forming a straight line. The correct phase offset can then be estimated by shifting the line down a multiple of 2π , so that it is within $\pm\pi$ (corresponding to 0 phase at 0). This assumes that the phase of \bar{H} is linear with respect to constant n . However, this is a property of the model used and is not inherently true (in particular, resonant models do not have this property). Thus forming an estimate of the gradient

may fail due to local variation in the model phase.

A efficient method is to form boundaries on the extracted \tilde{n} . As noted in § 3.9, the real refractive index has a relationship with phase wrapping and Sushko *et al.* [134] used this to form boundaries on the extraction process. This however complicates the minima search as a minimum can be formed at the boundary leading to new local minima. To resolve this, wrapping around an initial n value is used. Two boundary conditions are formed:

$$n_{\text{UB}} = n_{\text{init}} + \frac{c}{2fl} \quad (5.9)$$

and:

$$n_{\text{LB}} = n_{\text{init}} - \frac{c}{2fl} \quad (5.10)$$

Where n_{UB} and n_{LB} are the upper and lower bounds respectively. A wrapping error function can then be formed, such that:

$$\mathbb{E}(\tilde{n}) = \varepsilon \left(\tilde{n} - a \frac{c}{fl} \right) \quad (5.11)$$

Where a is an integer, which is calculated from the boundaries:

$$a = \begin{cases} \text{R} \downarrow \left((n - n_{\text{LB}}) \frac{fl}{c} \right) & \text{if } n > n_{\text{UB}} \\ \text{R} \downarrow \left((n - n_{\text{UB}}) \frac{fl}{c} \right) & \text{otherwise} \end{cases} \quad (5.12)$$

Where $\text{R} \downarrow (x)$ rounds x towards 0. After extraction, the n can have wrapping artefacts, however κ will be correct. The estimate of n can then be corrected by using the unwrapping corrections in § 3.9.

This method will therefore converge to values of \tilde{n} in the region of the initial estimate. This is particularly advantageous if discontinuities in the refractive index are present, see § 3.9.3, as other methods of fixing the unwrapping will not generally account for this. This method is also the most computationally simple, with each frequency being non-interdependent, thus simple to parallel-ise. However, it is dependent on the correct value being within these initial boundaries. This will always be true provided that initial estimates are made using the methods in § 3.9,

using the same estimate of thickness.

Using this method of boundary conditions, error functions which are phase insensitive can be used. In this work the difference of real and imaginary parts is used:

$$\mathbb{E} = |\Re(H) - \Re(\bar{H})| + |\Im(H) - \Im(\bar{H})| \quad (5.13)$$

Equivalently the conjugate square could be used:

$$\mathbb{E} = (H - \bar{H})(H - \bar{H})^* \quad (5.14)$$

The advantage of both of these error functions is that they are computationally and analytically simpler than equations 5.4 & 5.5, as neither require logarithms or arctan operations to extract the phase.

To extract \tilde{n} , the error function equation 5.13 is used in conjunction with wrapping around an initial estimate. It should be noted that all the error functions presented in this chapter will often produce the same answer provided correct initial conditions, as the solution is unique.

5.3 Fitting transfer function models to α lactose monohydrate

The measurements of 100% α lactose monohydrate from § 3.8 were re-extracted using a transfer function fitting. The results from § 3.8 were used as both initial points and as a comparison. The model used was the one presented in § 3.2.

In Figure 5.2, the absorption coefficient of the sample calculated using the two different methods is shown. The blue estimate was calculated by fitting the model, and the red dashed line was calculated using equation 3.20, which was also used as an initial estimate for the fitting. Both estimates meet α_{max} at 3.25 THz, which is assumed to be approximately the same for both extractions.

In Figure 5.3, the refractive index of the sample calculated via both methods is shown. The blue line is the estimate based on fitting the model, and the red dashed line is an estimate calculated using equation 3.18. The approximation was used to




Figure 5.2: α of 100 % Lactose monohydrate, using a fitted transfer function model (blue) and a approximation (red, dashed)




Figure 5.3: n of 100 % Lactose monohydrate, using a fitted transfer function model (blue) and a approximation (red, dashed)

create boundary conditions for the error function, see equations 5.9,5.10. Where measurement has been limited by dynamic range, sharp discontinuities are present in the extracted n . If boundary conditions were not used, the n would be extracted incorrectly at these points.

Simply fitting the sample model from § 3.2 has the same limitations as using the approximate methods in § 3.6 as the time window must be shortened to exclude reflections from the sample. It also does not result in some substantial increase in accuracy, this is because the absorption of the sample has negligible effect on the transmission coefficient, see § 3.2, and thus the approximations are very accurate. The benefit of fitting is the ability to model complex sample responses, which can be used to include reflections from the sample. This can be used to extract \tilde{n} at a better frequency resolution and to extract the sample thickness from the sample measurement.

5.4 Sample resonance models

The sample model was developed further in line with the literature to include resonance within the sample [60], [62], [63]. These models assume that reflections occur at the sample interfaces, and that this can result in secondary echoes in the time domain measurement.

To develop this model, the initial model discussed in § 3.2 is used as a starting point. The same assumptions about the sample and the THz beam is used. However internally there is assumed to be resonance between the two sample interfaces. This




Figure 5.4: Model of the THz beam resonating in the sample with thickness l , an angle has been added for clarity. The THz beam reflects and transmits at each interface, creating resonance within the sample.




Figure 5.5: (A) - the THz radiation passing through a sample, (B) - The THz radiation being partially reflected at the sample-air interface

results in a portion of the THz radiation being reflected back and forth creating multiple pulses which can be detected in the time trace.

In Figure 5.4, the modified sample model is shown diagrammatically, with an angle offset for clarity (the THz beam is still assumed to be perpendicular to the sample interfaces). The THz beam passes through the sample, making a single initial pass. Additionally, at the second interface the radiation reflections back into the sample. This cascades with a reflection at the first interface to create a Fabry-Perot resonance effect. Each path through the sample will still partially transmit, leading to multiple pulses of THz radiation being detected.

To form this model analytically, the reflection coefficient is introduced. This relates both the amplitude and phase of the reflected radiation to the incident radiation, at an interface between two materials with refractive indexes \tilde{n}_1 and \tilde{n}_2 , and is:

$$R_{1,2} = 1 - T_{1,2} = \frac{\tilde{n}_2 - \tilde{n}_1}{\tilde{n}_2 + \tilde{n}_1} \quad (5.15)$$

In Figure 5.5 (A), THz radiation passing through a sample is shown. In Figure 5.5 (B), there exists an interface between the sample (1) and air (2), which have different corresponding refractive indexes. The relation between the incident and reflected radiation is the reflection coefficient.

To develop a model which takes into account the resonance between sample air interfaces, the model in § 3.2 is modified to include coefficients which account for the reflection at each interface and propagation back and forth over the sample. A transfer function model which accounts for just one reflection in the sample trace

would take the form:

$$\bar{H} = T_{a,s} \frac{P_s}{P_a} T_{s,a} (1 + R_{s,a}^2 P_s^2) \quad (5.16)$$

Where $T_{s,a}$, $T_{a,s}$, P_s , P_a and $R_{s,a}$ are the transmission coefficient between the sample and air, transmission coefficient between air and the sample, the propagation coefficient of the sample, the propagation coefficient of the displaced air and the reflection coefficient between the sample and air respectively.

To add a second reflection to the model, equation 5.16 can simply be multiplied again by $(1 + R_{sam,air}^2 P_{sam}^2)$, to extend it to three reflections, it simply needs to be multiplied by the same factor. There exists a simple generalisation for M reflections:

$$\bar{H} = T_{air,sam} \frac{P_{sam}}{P_{air}} T_{sam,air} \sum_{\delta=0}^{\delta=M} R_{sam,air}^{2\delta} P_{sam}^{2\delta} \quad (5.17)$$

Where δ is the pulse index, $\delta = 0$ for the initial transmission without reflection, $\delta = 1$ for the first reflection, etc.

M can be approximated reliably by [63]:

$$M = \frac{\text{floor}(\frac{ct_{max}}{nl} - 1)}{2} \quad (5.18)$$

Where t_{max} is the length of the time window and n is in this case an estimate of the average (across frequency) real refractive index.

As $\left| R_{sam,air}^2 P_{sam}^2 \right| < 1$, this model will converge on [60]:

$$\bar{H} = T_{air,sam} \frac{P_{sam}}{P_{air}} T_{sam,air} \frac{1}{1 - R^2 P^2} \quad (5.19)$$

Both are valid approaches to the same modelling problem, however one considers a finite series of reflections, and the other an infinite series. The former will converge on the latter provided a long enough time window or a weak resonance. However, this is not inherently true if the series of reflections is truncated (for instance to remove system reflections).

5.5 Fitting resonant models

It has already been established in § 3.11, that reflections from the sample create sinusoidal artefacts in the extracted \tilde{n} , by fitting a model which models these resonances, these artefacts can be minimised.

z-cut quartz was used as a test case in this work, as it produces a strong resonance and is broadly transparent ($\alpha < 10\text{cm}^{-1}$) at THz frequencies. It was also found to have an extremely flat surface profile, which was profiled using an alpha-step with an accuracy to within $1\ \mu\text{m}$. A consideration for § 5.17 is that the reflections are also well separated which allows time partitioning.

This resonance is difficult to resolve as the reflections are relatively strong compared to other samples, this makes the quartz sample sensitive to inaccuracies within the transfer function model. The strong resonance is due to the large non-dispersive n across the bandwidth of the THz pulse, with an exception at 3.87 THz, where a weak absorption mode was observed.

A 0.512 mm z-cut quartz plate was measured using HFRBB THz TDS system at sampling frequency of 25 THz, with two different windows, the first with a frequency resolution of 30 GHz and the second with a resolution of 200 GHz. The first was used to capture the full reflection train of the quartz sample, the second was used to exclude the reflections from the processing.




Figure 5.6: Absorption coefficient of a 0.512 mm z-cut quartz sample, with $10\ \text{cm}^{-1}$ offsets. Blue - Extracted values using a resonant model, Red - approximate solution including reflections, Green - Approximate solutions excluding reflections. Dashed blue - α_{max} of the measurement.




Figure 5.7: Real refractive index of a 0.512 mm *z-cut* Quartz sample, with 0.05 offsets. Blue - Extracted values using a resonant model, Red - approximate solution including reflections, Green - Approximate solutions excluding reflections

In Figure 5.6, the calculated α with 10 cm^{-1} offsets of the quartz sample is shown, calculated using a sample thickness of 0.512 mm. In Figure 5.7, the corresponding n curves with 0.05 offsets are shown.

The green and red curves were approximated using the methods in § 3.6, using the shorter and longer time windows respectively. The blue curves were extracted by fitting a resonant model to the longer time window measurements. The inclusion of sample reflections has incurred etalons in the extracted parameters. By using a resonant model, these can be reduced. However, this reduction is inconsistent across frequencies and has been found experimentally to be dependent on the thickness used during extraction [62].

5.6 Effect of thickness error on \tilde{n}

The error in the sample thickness used to fit a resonant model will result in etalons being present in the extracted \tilde{n} . This can be derived from the analysis in § 3.11 which derived how time localised error produced etalons.

Using the same method in § 3.11, it is possible to write \tilde{n} as a logarithm of the resonant model:

$$\tilde{n} = \frac{ci}{wl} \left(\log_e(H) - \log_e(T_{a,s}T_{s,a}) - \log_e \left(1 + \sum_{\delta=1}^M R_{s,a}^{2\delta} P_s^{2\delta} \right) \right) + 1 \quad (5.20)$$

Taking the first order Taylor series approximation of a logarithm, the following can be formed:

$$\tilde{n} \approx \frac{ci}{\omega l} (\log_e (H) - \log_e (T_{a,s} T_{s,a}) - R_{s,a}^2 P_s^2) + 1 \quad (5.21)$$

The same approximation can be derived using an infinite resonance model (however they will diverge at higher orders).

If $\bar{H} \approx H$, i.e. when the error has been minimised after fitting, then the error, $\mathbb{E}_{\bar{n}}$, between the extracted refractive index, \bar{n} , and the true value, \tilde{n} , can be estimated as:

$$\mathbb{E}_{\bar{n}} = \bar{n} - \tilde{n} = \frac{ci}{\omega l} (\log_e (T_{a,s} T_{s,a}) - \frac{l}{\bar{l}} \log_e (\bar{T}_{a,s} \bar{T}_{s,a}) + R_{s,a}^{2\delta} P_s^{2\delta} - \frac{l}{\bar{l}} \bar{R}_{s,a}^{2\delta} \bar{P}_s^{2\delta}) \quad (5.22)$$

Where \bar{T} , \bar{R} and \bar{P} are model coefficients formed from \bar{n} and \bar{l} . \bar{l} is the assumed thickness used during the fitting.

Approximating $\frac{1}{\bar{l}} \bar{T}_{s,a} \bar{T}_{a,s} = \frac{1}{l} T_{s,a} T_{a,s}$, an approximation can be formed:

$$\mathbb{E}_{\bar{n}} \approx \frac{ci}{\omega} \left(\frac{1}{\bar{l}} R_{s,a}^2 P_s^2 - \frac{1}{\bar{l}} \bar{R}_{s,a}^2 \bar{P}_s^2 \right) \quad (5.23)$$

For an estimated thickness, in the neighbourhood of the actual sample thickness, a value of \bar{n} is always extractable. However from equation 5.24, there exists a method of measuring $\mathbb{E}_{\bar{n}}$, which can then be used to find the correct thickness.

If the \tilde{n} is roughly flat (i.e. $\frac{d\tilde{n}}{d\omega} \approx 0$), such that its differential with respect to frequency tends to 0, the differential of this error can be extracted:

$$\frac{d\bar{n}}{d\omega} = \frac{d\tilde{n}}{d\omega} + \frac{d\mathbb{E}_{\bar{n}}}{d\omega} \approx \frac{d\mathbb{E}_{\bar{n}}}{d\omega} \quad (5.24)$$

This can then be expressed using the product rule as differentiation of the individual coefficients:

$$\frac{d\bar{n}}{d\omega} \approx \frac{ci}{\omega} \left(\frac{1}{\bar{l}} \left(\frac{dR_{s,a}^2}{d\omega} P_s^2 + R_{s,a}^2 \frac{dP_s^2}{d\omega} \right) - \frac{1}{\bar{l}} \left(\frac{d\bar{R}_{s,a}^2}{d\omega} \bar{P}_s^2 + \bar{R}_{s,a}^2 \frac{d\bar{P}_s^2}{d\omega} \right) \right) \quad (5.25)$$

Assuming that the reflection coefficients are flat ($\frac{dR_s}{d\omega} \approx 0$) and given $\frac{dP_s^2}{d\omega} = \frac{-i2ln}{c} P_s^2$, $\frac{d\bar{n}}{d\omega}$ can be approximated as:

$$\frac{d\bar{n}}{d\omega} \approx \frac{2}{\omega} (\tilde{n}R_{s,a}^2 P_s^2 - \bar{n}\bar{R}_{s,a}^2 \bar{P}_s^2) \quad (5.26)$$

Which leads to the observation that the differential of \bar{n} is equal to the difference of two complex sinusoids, \bar{P}_s^2 and P_s^2 . Using the beat frequency phenomena:

$$\frac{d\bar{n}}{d\omega} \approx \frac{2}{\omega} e^{-i\frac{\omega}{c}(l\bar{n}+\bar{l}\bar{n})} \left(\tilde{n}R_{s,a}^2 e^{i\frac{\omega}{c}(\bar{l}\bar{n}-l\bar{n})} - \bar{n}\bar{R}_{s,a}^2 e^{-i\frac{\omega}{c}(\bar{l}\bar{n}-l\bar{n})} \right) \quad (5.27)$$

Which as $\bar{l} - l$ tends to 0 will approximate to:

$$\frac{d\bar{n}}{d\omega} \approx \frac{4i\tilde{n}}{\omega} e^{-i\frac{\omega}{c}(l\bar{n}+\bar{l}\bar{n})} R_{s,a}^2 \sin\left(\frac{\omega}{c} \Re(\bar{l}\bar{n} - l\bar{n})\right) \quad (5.28)$$

As \bar{l} changes, the differential (and thus error) of the estimated \bar{n} will therefore appear to be a modulated complex sinusoidal. The exception is where the assumed thickness is equal to the actual thickness, in which case equation 5.28 simplifies to 0.

This can be exploited to perform a direct search, by minimising a measure of deviation with respect to assumed thickness, the correct thickness can be located. A common measure in the literature [62], [63] is total variance:

$$\text{TV}(\bar{l}) = \sum_{k=2}^N |\bar{n}_{\bar{l}}(k+1) - \bar{n}_{\bar{l}}(k)| \quad (5.29)$$

Where k is the frequency index and N is the number of frequency domain samples. By finding a minimum of the function with respect to \bar{l} , by directly iterating over different assumed values, the true value of l can be found. This minimum has been observed to be local, both in this work and in the literature [62], which is due to the beat frequency effect interacting with the finite bandwidth. In practice, in this work the real refractive index is used instead of the extinction coefficient, as it is in general flatter and thus produces smoother measures of TV [116].

Another practicality, is that the TV shows an inverse dependence with thickness in samples which have strong features present. This is because the \bar{n} is inversely

proportional to \bar{l} . If there are no features present, this is a constant with respect to frequency so will have little effect on TV. If this is not true, then a background dependence is observed. This can be normalised out by multiplying the TV by \bar{l} .

Alternative methods have been suggested, notably a method which performs a Fourier transform on the real refractive index and minimises the peak present [116], which is similar to the method of analysis in § 3.11.

5.7 Thickness extraction of quartz

In § 5.5, a resonant model was fitted to 0.512 mm z-cut quartz, this has been repeated multiple times with different assumed thicknesses, to produce different extracted values of n . This has then been evaluated with total variance to form a measure of error with respect to thickness and to perform a direct search to find the correct sample thickness.




Figure 5.8: Estimated values of n of 0.512 mm z-cut quartz, extracted using a infinite resonant model with different assumed thicknesses. Each curve is colour coded to its total variance value.




Figure 5.9: Total variation of the real refractive index over assumed thicknesses between 0.35 – 0.65 mm in 0.005 mm steps for the measured z-cut quartz sample. The colour coded markers are the points taken for Figures 5.8, taken in steps of 0.025 mm.

In Figure 5.8, the refractive index values extracted with assumed thicknesses between 0.355–0.63 mm with 0.025 mm steps are shown, each curve has been assigned a colour based on its total variation value corresponding to a marker of the same colour in Figure 5.9. The blue curve is the smoothest curve, and represents the local minimum in the total variance curve, and is measured to be 0.505 mm. This is a slight under estimated compared to measurement with micrometer (taken to be 0.51 mm) and measurement taken with the alpha-step (0.512 mm). Near the smoothest curve, the error appears to be purely sinusoidal, however as the assumed thickness deviates significantly from the true thickness a decaying beat frequency becomes apparent.

In Figure 5.9, the total variance calculated from the real refractive index over assumed thicknesses of 0.35–0.65 mm with 0.005 mm increments is shown (chosen of a computational accuracy of the extracted thickness). The colour coded markers are the points taken for Figures 5.8, taken in steps of 0.025 mm. There is a distinct local minima at 0.505 mm, which corresponds to extracted thickness. There also exists a much weaker collection of local minima. There also exists a global trend to smaller values, which is attributed to the inverse dependence on \bar{l} .

There are residual etalons present in the final extracted parameters and the extracted thickness is a slight underestimate. This can be attributed to the model being slightly incorrect, due to the focused nature of the THz radiation.

5.8 Correction for gaussian beam effects

When modelling the interaction of the THz radiation with a sample, the THz radiation is modelled as being a point pencil beam. This assumption can be valid if the sample is thin relative to the Rayleigh distance of the beam [47]. It has been demonstrated that there exists a systematic bias when using focused radiation [99], [100], when compared to collimated results. Thus integrating corrections into the sample model may improve results.

In this work the beam profile is assumed to be Gaussian, it has been suggested in the literature that mounting lenses on emitters significantly changes the beam profiles [135], so this assumption is not general and will depend on the emitter used. Due to the emitter orientation, lenses are not used in this work.

The THz radiation is focused using a set of off-axis parabolic mirrors, resulting in the THz radiation having a Gaussian beam profile. The sample will act as a primitive lens, shifting the focal point of the beam in the propagation axis. There exist both amplitude and phase changes due to the profile which depend on the distance from the beam focus along the propagation axis. In [100] an inclusion of a calibration parameter, β , into the model from § 3.2 has been suggested to approximate for the effect of Gouy phase shift and the on-axis amplitude variation. To do this, a model of the focal shift in terms of sample thickness and real refractive index is required.

The THz radiation is modelled as being a focused beam, entering a sample. Within the sample, the beam is less tightly focused compared to free space due to Snell's law [47]. This leads to the focal spot being shifted. In Figure 5.10, focused




Figure 5.10: the THz radiation is defocused entering the sample, and re-focused leaving the sample. The radiation which passes through a sample (solid blue line) therefore propagates Δ further than in air (dashed blue line) before converging to a focal point.

THz radiation is shown passing through the sample using simple geometric approximations. Within the sample (solid line), the beam converges slower than in the free space (dotted line). When the beam does converge to a point, if it has passed through the sample, it propagates forward further before converging to a focus. The focus is therefore displaced by Δ .

Approximately, neglecting effects relating to the absorption, angular dispersion and frequency dependence of the sample parameters, the focal point change in the THz profile for propagation over dielectric with given n and l [100]:

$$\Delta = \frac{l}{n} (n - 1) \quad (5.30)$$

This shift relates to the focus at the sample, however the effect at the detector is the cause of the detected amplitude and phase change. In the THz TDS in this work, two parabolic mirrors are used either side of the sample focus, see § 2.1.4. Using the thin lens approximation, the shift at the focus, Δ' , can be related to this value by [100]:

$$\Delta' = \frac{\Delta}{\left(\frac{a}{F} - 1\right)^2} \quad (5.31)$$

Where $\Delta \ll |a + F|$ has been assumed, and a is the distance from the focal spot to the mirror, and F is the focal length of the parabolics (which is the same for both parabolics). The focal point displacement can then be related to a change in phase and amplitude by modelling the THz at the detector as having a Gaussian beam profile.

Assuming that the TDS system is aligned such that the focal point of the refer-




Figure 5.11: The THz focus is shifted at the detector (shown here as a photo-conductive detector), into the detector by Δ'

ence beam is aligned to the detector, this will result in the focal spot being displaced into the detector. In Figure 5.11, the detector is shown (the detector mechanism is arbitrary, the method has been demonstrated on both electro-optic crystals and photo-conductive emitters and detectors [100]), with the focus of the THz radiation aligned in a reference (dashed lines) to the edge of the detector. When a sample is introduced, the focus of the radiation is shifted through the detector (the “detection region” is assumed to be much thinner than the shift). This in effect means the detector detects $-\Delta'$ away from the focus of the radiation when a sample is introduced.

A modified sample model was developed to account for this shift away from the focus of the THz radiation at the detector. Two further coefficients are introduced to model a Gaussian profile beam, the first is an amplitude correction for the beam width, $w(z)$:

$$w(z) = \frac{1}{\sqrt{1 + \left(\frac{z}{z_0}\right)^2}} \quad (5.32)$$

Where z is the displacement from the focus of the beam radiation (at $z = 0$) in the beam propagation direction and z_0 is the Rayleigh distance, which can be related to the beam waist, w_0 , of the beam and the frequency, f :

$$z_0 = \frac{\pi w_0^2}{c} f \quad (5.33)$$

And an additional phase shift, the Gouy phase shift, $G(z)$, is introduced:

$$G(z) = e^{-i\angle \frac{z}{z_0}} \quad (5.34)$$

Where \angle is the angle operation. These terms will be present solely in the sample measurement, as both will equate to 1 in the reference measurement as $z = 0$.

A frequency and system dependent calibration is introduced to group the effects of the parabolics and the Rayleigh distance, referred to in this work as the beam

alignment coefficient, β [100]:

$$\beta = \frac{c}{\omega z_0 \left(\frac{a}{F} - 1\right)^2} \quad (5.35)$$

β is best visualised using a different arrangement:

$$\beta = \frac{1}{2\pi^2} \underbrace{\left(\frac{\lambda}{w_0}\right)^2}_{(i)} \underbrace{\frac{1}{\left(\frac{a}{F} - 1\right)^2}}_{(ii)} \quad (5.36)$$

where (i) is the squared frequency dependent wavelength (λ), beam waist ratio and (ii) is the combined lensing affect of the parabolic mirrors.

This is incorporated into $w(z)$, to form a new coefficient, $w(l')$, in terms of l' [100]:

$$w(l') = \frac{1}{\sqrt{1 + \left(\frac{-\omega\beta l'}{c}\right)^2}} \quad (5.37)$$

Where l' is the path length through the sample normalised to real refractive index, this includes an air displaced term, so that for a non resonant model:

$$l'_0 = \frac{l}{n} - \frac{l}{n_{air}} \quad (5.38)$$

Where $n_{air} = 1$ and l'_0 is the normalised path length for a single pass through the sample, which is equivalent to $-\Delta$.

$G(z)$ can be approximated as $G(l')$ using a first order Taylor series approximation [100]:

$$G(l') \approx e^{+i\frac{\omega\beta l'}{c}} \quad (5.39)$$

The non resonant model from § 3.2 can then be expressed as:

$$\bar{H} = T_{s,a} T_{a,s} \frac{P_s}{P_a} w(l'_0) G(l'_0) \quad (5.40)$$

To consider resonance within the sample, the effect on the focus of the radiation must be considered when the propagation distance through the sample is increased. In Figure 5.12 the propagation of the THz radiation is shown, when a double pass




Figure 5.12: (A) - The THz radiation is defocused entering the sample, and re-focused leaving the sample. The radiation reflects once at each interface to fold in on itself. This results in the focal point shift in the opposite direction to non-resonant case, giving a new displacement to the focal distance, Δ_1 , which can have a drastically different amplitude. (B) - The "unfolded" radiation. In effect the beam has propagated over a sample 3 times as thick, resulting in a longer convergence length to the beam focus. However, because of folding, it results in the propagation outside the sample, being much shorter.

through the sample (internal reflection once at each interface) is shown. This leads to a further shift in the focus of the radiation, in proportion to a double pass through the sample. This leads to Δ_1 being defined as [100]:

$$\Delta_1 = \frac{l}{n} (n - 3) \quad (5.41)$$

This can instead be generalised in terms of l' as, for arbitrary echo index δ :

$$l'_\delta = (1 + 2\delta) \frac{l}{n} - \frac{l}{n_{air}} \quad (5.42)$$

This can then be incorporated into the finite resonance model from § 5.4, to form:

$$\bar{H} = T_{a,s} \frac{P_s}{P_a} T_{s,a} \sum_{\delta=0}^{\delta=M} R_{s,a}^{2\delta} P_s^{2\delta} G(l'_\delta) w(l'_\delta) \quad (5.43)$$

For thin samples with low refractive index, $w(l'_\delta)$ can be assumed to be approximately 1 and the infinite resonance model can be formed:

$$\bar{H} = T_{a,s} \frac{P_s}{P_a} T_{s,a} G(l'_0) \frac{1}{1 - R_{s,a}^2 P_s^2 G(l''_0)} \quad (5.44)$$

Where $l'' = \frac{2l}{n}$.

These modified models however require β to be known for a given THz TDS system to calculate both $w(l')$ and $G(l')$. To do this a method using calibration

samples is required.

5.8.1 Estimating the beam alignment coefficient

To correct for this effect β must be therefore be estimated, in [100] the method used is to measure a calibration sample and then repeat the measurement at a different detector position (the system will require realignment after each calibration measurement). The n of the calibration sample is estimated from each measurement, which is then fed into an equation, which provided a , F , $\mu(n)$ and l are known can be used to calculate β .

A simpler method is to use a calibration sample which produces an echo. By windowing the initial pulse and echo separately, and treating them as separate measurements, different estimations of n can be formed. The initial pulse will only pass through the sample once, see Figure 5.10, while the reflection will pass through the sample multiple times, see Figure 5.11. This results in a change in the detected Gouy phase shift.

A generalisation of equation 3.17 for sequential pulse index δ (0, being the initial pulse) can be formed:

$$n = \frac{-c}{\omega(2\delta + 1)l} \angle H_\delta + \frac{1}{2\delta + 1} \quad (5.45)$$

Where H_δ is the transfer function for the pulse with index δ .

This estimate does not account for the Gouy phase shift, so each estimate will be biased by [100]:

$$n_\delta = n + \beta \frac{n - (2\delta + 1)}{n(2\delta + 1)} \quad (5.46)$$

Where n is the true value of the sample's refractive index, and n_δ is the estimate formed from H_δ .

Taking the difference between the estimate of n for the initial pulse, n_0 , and the echo, n_1 , the following can be simplified:

$$n_0 - n_1 = \beta \frac{n - 1}{n} - \beta \frac{n - 3}{3n} = \beta \frac{2}{3} \quad (5.47)$$

And therefore:

$$\beta = \frac{3}{2}(n_0 - n_1) \quad (5.48)$$

More generally for arbitrary n_a and n_b :

$$\beta = \frac{(2a + 1)(2b + 1)}{2(b - a)}(n_a - n_b) \quad (5.49)$$

This allows β to be calculated relatively trivially from a calibration sample with known thickness, from a single measurement of a sample (using multiple time windows). As this requires sample resonance, an ideal calibration sample must have high real refractive index and negligible absorption. Fortunately there already exist candidates which are ideal for this, high resistivity silicon and z-cut quartz in particular [99].

5.8.2 Partitioning the impulse response

To estimate β , separate estimates of real refractive index are required based on different pulses from the sample. This could be estimated by separately windowing each pulse in the sample trace, and then separately calculating the transfer function of each pulse. This could then be used to individually calculate the real refractive index based on each window.

A superior solution is to partition the impulse response, h , which is the time domain equivalent of the transfer function. These partitions can then be converted to time localised segments of the transfer function, in a similar method to the short time Fourier transform [93]. By separating each pulse into a separate segment, it is possible to calculate the transfer function of each pulse.

The impulse response will generally be susceptible to noise, as the sample and reference measurement out of bandwidth (defined by dynamic range of reference measurement) will usually result in catastrophic cancellation (i.e. $\frac{0}{0}$). This makes it difficult to effectively partition the impulse response. A filter is thus required to provide adequate signal to noise ratio of the impulse response.

It was found that simply using a rectangular window on H was sufficient to form

a filter such that the filtered transfer function \hat{H} in the frequency domain is:

$$\hat{H}(f) = \begin{cases} H(f) & \text{if } -f_c < f < f_c \\ 0 & \text{otherwise} \end{cases} \quad (5.50)$$

The filtered impulse response, \hat{h} was then calculated by performing the inverse discrete Fourier transform (IDFT) on \hat{H} . This is equivalent to filtering the impulse response, h , with an infinite sinc kernel. Alternatively, other filters could be used which might have shorter time responses, however a compromise would be made to the frequency domain response. This could for instance have a negative impact on the analysis after this operation.

The impulse response is then partitioned by applying different windowing functions, with different ending and start points. As with the raw time domain scans, a Tukey window was used to perform each partition. A DFT was then applied to each partitioned section of the impulse response, to form a separated series of transfer function segments, which could give the transfer function for each separate time response of the sample.

5.9 Measuring the beam alignment coefficient




Figure 5.13: The partitioned impulse response of quartz. Gray - raw impulse response, Blue - Windowed initial pulse, Green - windowed second pulse. The peak amplitude normalised (for clarity) windows are shown dashed lines

In Figure 5.13, the impulse response of a 0.512 mm z-cut quartz sample is shown.




Figure 5.14: Two different estimates of n of the quartz sample, from the first impulse window (blue) and the second (green).

The original impulse response (grey), has been segmented into two using windows (dashed lines). The first segment (blue) contains an initial pulse, the second (green) contains the secondary pulse. The sample was measured using the HFRBB THz TDS system, at a sampling frequency of 500 THz with a frequency resolution of 30 GHz. The impulse response was calculated by applying the inverse DFT to the transfer function, calculated by normalising the sample frequency response to that of an air reference. The impulse response has two prominent pulses, corresponding to a transmission through the sample. The first, corresponds to transmission through the sample without reflection; the second corresponds to the first sample reflection. These features have been separated by applying a Tukey window, to isolate the first pulse. This is then subtracted from the original impulse response, and a second window is applied to isolate the second pulse. Each of these windowed transfer functions are converted back to the frequency domain, using separate Fourier transforms, to form two different time localised transfer functions.

In Figure 5.14, n of the quartz sample, based on the initial pulse transfer function (blue) and the reflection pulse (green). These were calculated using equation 5.45, with $\delta = 0$ and $\delta = 1$ respectively. There is a systematic difference between the two estimates from each sample, from which a different estimate of β can be formed.

In Figure 5.15, different estimates of β are shown, based on measurements of different samples. Each of these samples were measured, under the same conditions as the quartz. These measurements were made using thickness measurements per-

formed with an alpha-step surface profiler. The alpha-step has a reported accuracy of less than $1 \mu\text{m}$, and where possible, several different measurements have been made to observe the spatial variation in sample thickness. The blue estimate was calculated using the quartz sample, with a measured thickness of 0.512 mm and no spatial variation was found in these measurements. The green curve was estimated from a 1.015 mm high resistivity silicon sample with an observed spatial standard deviation of 0.0036 mm. The red estimate was made using a 0.58 mm high resistivity silicon sample, unfortunately the sample was damaged before a spatial profile could be measured. The magenta estimate was formed using a TPX sample, with a thickness of 0.969 mm and a spatial deviation of 0.0107 mm.

These different estimates of β were found to have distinct offsets, but have a similar curve across frequency. It was hypothesised that thickness uncertainty was causing different measured offsets, and that centering the curves might result in similar curves. In Figure 5.16, the estimates of β were centered at 2 THz. When this is used, the different estimates of β are shown to be consistent with each other. However, there is noise at both 0.5 THz and above 5 THz. The TPX sample shows the most noise at these extremes, and the silicon samples the least. It is likely then that this is primarily due to the limited signal to noise ratio (SNR) of the reflection pulses which is driving this distortion, samples which produce weaker reflections will give distorted secondary estimates of n .




Figure 5.15: Estimates of β made using different samples. Blue - 0.512 mm z-cut quartz, Green - 1.015 mm High resistivity silicon, Red - 0.58 mm high resistivity silicon, Magenta - 0.969 mm TPX

It is possible to derive the effect of thickness uncertainty on β by extending equation 3.62, from which the following can be derived:

$$\sigma_{\beta} \approx \frac{\sigma_l}{l}(\beta - 1) \quad (5.51)$$

Which considering $\beta \ll 1$ approximates to:

$$\sigma_{\beta} \approx \frac{\sigma_l}{l} \quad (5.52)$$

i.e there is frequency independent uncertainty β with respect to the relative uncertainty in thickness. It is worth noting that this relative error in thickness produces an absolute error in β . Given the measured amplitude of β , uncertainty which is small can be significant if the calibration sample is thin.

Several measurements, of several samples were used to estimate β , a full description of measurements can be found in appendix B. These estimates were found to be alignment dependent, but not detector or emitter dependent (this was tested with the different variations of the THz TDS system, and would require further testing). In Figure 5.17, the centered average of the different estimates of β , with the standard error shown as a dashed line.

This average estimate was observed to be approximately quadratic, and to create a calibration curve, a quadratic curve was fitted to this zero centered average using




Figure 5.16: Different estimates of β , centered at 2 THz (grey line). Blue - 0.512 mm z-cut Quartz, Green - 1.01 mm High resistivity silicon, Red - 0.58 mm high resistivity silicon, Magenta - 1.22 mm TPX




Figure 5.17: The centered mean of the measured estimates of β , the standard error is shown as dashed lines.

linear least squares. It was observed that different measurements of the same z -cut quartz sample produced very consistent estimates of β , which coincides with it having no measured spatial variation in its thickness. Thus the offset from the average of these estimates was added to this fitting, to create an accurate calibration curve. In Figure 5.18, the different estimates of β , based on the same sample, are shown as dashed blue lines. The fitted quadratic curve is shown as a thick red line, only the offset of the curve is based solely on the quartz measurements. A quadratic was used because it fitted the data well, other instruments may produce different values of β and require a different model.




Figure 5.18: The fitted quadratic calibration curve (red), and the estimates of β , using a quartz sample, used to estimate the offset.

5.10 Extraction of z-cut quartz thickness with calibrated extraction




Figure 5.19: The real refractive index of z-cut quartz, the thickness in each case was based on the total variation minimum. Blue - extraction without using a calibration coefficient, Green - with calibration coefficient.

The processing of the 0.512 mm z-cut quartz, extracted in section § 5.5, was repeated with a calibrated model with a fitted β curve taken from § 5.9. In Figure 5.19, the n of the quartz sample is shown. A direct search was performed to locate the sample thickness which minimises the total variation. The blue curve was extracted in the absence of a calibration coefficient (equivalent to $\beta = 0$), and the green curve has been extracted with the fitted calibration curve from § 5.9. Different thicknesses were extracted with these measurements, in the former (without β) a thickness of 0.504 mm was extracted and in the latter (with β) a thickness of 0.512 mm was extracted. The inclusion of β has served to calibrate the thickness extraction, however it was observed that larger etalons between 4 and 6 THz were introduced by using a calibration coefficient.

Extracting the calibration thickness from the same sample shows that the calibration procedure is self consistent. It was found that the calibration could be performed with arbitrary assumed thickness and this would be extracted from the measurement using that calibration.

In Figure 5.20, the total variance of the uncalibrated (blue) and calibrated (green) model extractions is shown. The total variance in the calibrated curve is slightly higher, and the location of the local minima has been moved from 0.504 mm to




Figure 5.20: The total variance of the quartz measurement, using different assumed thicknesses. Blue - without calibration, Green - with calibration

0.512 mm, i.e. it has been moved to be in line with the alpha-step thickness measurement. It would however be expected that the total variance would be lower (as the model should be a better fit), a possible explanation for this is that system was aligned differently such that β should be different for this measurement.

In Figures 5.21 and 5.22, both n and α of the same sample of quartz but measured under different alignment. The measurement was performed in the same system but at a sampling frequency of 500 THz. The extracted values are based on calibrated (green) and non-calibrated (blue) models. The calibrated model is smoother, indicating that the model is likely a slightly better fit. In Figure 5.23, the total variance curves are shown, like before the local minimum has shifted in the calibrated extraction, however the total variance curve has lower absolute values. The minimum has




Figure 5.21: The real refractive index of z-cut quartz, under different alignment. Blue - extracted without calibration, Green - extracted with calibration.




Figure 5.22: The absorption coefficient of z-cut quartz under different alignment. Blue - extracted without calibration, Green - extracted with calibration (20 cm^{-1} offset)

moved from 0.505 mm to 0.515 mm, this only slightly disagrees with the non-contact gauge measurement (0.512 mm), which could be down to measurement or sample alignment uncertainty.

In Figure 5.24, the extracted values of β for the two different measurements are shown, along with the fitted curve used to calibrate the measurements. These estimations of β were calculated as described in § 5.9. The blue curve corresponds to the latter above (Figures 5.21,5.22,5.23), and the green curve corresponds to the former above (Figures 5.19,5.20). The red curve is the fitted parameter used during the extraction. All three curves are consistent between 1 and 4 THz, however beyond 4 THz the green curve diverges. This curve profile was collaborated with another sample taken with the same alignment, see appendix B.




Figure 5.23: The total variance of the extraction process for the quartz sample. Blue - Without calibration, Green - With calibration.

This strongly suggests that β is alignment dependent and could be used to access the quality of the alignment. β will show a strong inverse dependence on the alignment properties of the beam and in Figure 5.24, the higher frequencies could be distorted by miss-alignment.




Figure 5.24: β for the different measurements. Blue - estimated from the measurements used for Figures 5.21,5.22,5.23, Green - estimated from the measurements used for Figures 5.19,5.20 and Red - The fitted calibration used for the extraction of n and α

5.10.1 Consistent extraction of z-cut quartz

Regardless of extracted thickness, the n of a material should be consistent with different sample thicknesses. Therefore, post extraction of l , n and α should be consistent between different samples of the same material. It was found that when sample thickness was extracted with calibrated models the extracted estimates of n and α were more consistent between samples of the same material.

Two different samples of *z-cut* quartz were measured with the LFRBB THz TDS system with a sampling frequency of 25 THz with a bias voltage of 100 V. The frequency resolution used was 120 GHz, and the time window included system reflections. These system reflections do not affect the thickness extraction, but will incur etalons in the extracted parameters. The first sample was 0.512 mm and the second was 0.35 mm thick.

The parameter extraction was performed with and without calibration, using a direct search to extract the thickness. Two different corresponding thicknesses were extracted for each sample. In the former sample, thicknesses of 0.505 mm

and 0.512 mm extracted with and without calibration respectively, the same values extracted in § 5.10. In the latter sample, thicknesses of 0.344 mm and 0.349 mm extracted with and without calibration respectively.




Figure 5.25: n of two different z -cut quartz samples. Blue - uncalibrated measurement of 0.512 mm quartz, Red - uncalibrated measurement of 0.35 mm quartz, Green - calibrated measurement of 0.512 mm and Purple - calibrated measurement of 0.35 mm. The unsmoothed estimates are shown as grey curves.

In Figure 5.25, the different extracted values of n are shown. For clarity the different curves have been smoothed with a moving average filter with width 150 GHz to remove the etalons from the system reflections. The blue and red curves show the uncalibrated measurements of the 0.512 mm and 0.35 mm extracted curves. The purple and green curves show the calibrated extracted values. There is a slight difference between the uncalibrated curves, due to systematic error, primarily in the extracted thickness (the systematic error in the thickness is also accountable for the offset in the uncalibrated results). By accounting for this effect in the model, the error in both curves has been removed. This has resulted in a consistent n between the samples.

5.11 Temperature dependence of α lactose anhyrous

A significant advantage of being able to measure the thickness of a sample *in situ*, is to monitor and correct for a thickness change, which would otherwise be assumed constant. This can reveal effects in the spectra which would otherwise be missed.

A 5% α -Lactose anhydrous sample was measured using the system described

by [91], with the emitter being biased at 100 V and chopped at 10 kHz. The sample fabrication process can be found in Appendix A. The emitter used was driven by an Oxford Instruments cryostat and temperature controller were used in conjunction with liquid helium to cool the sample to 4 K and then measure the sample in 10 K increments, starting at 10 K and ending with 290 K. The heating process was performed by lowering the flow of helium through the cryostat and increasing the heater power (via the temperature controller). A room temperature measurement was performed before cooling the sample. The sample was made using a PTFE matrix, with concentrations measured by mass.

Each measurement was made at a sampling frequency of 150 THz, and windowed with two different windows. The first window included system and sample reflections, with a frequency resolution of 110 GHz. The second window used had a lower frequency resolution of 220 GHz, but did not include either sample or system reflections.

The first window included both a sample reflection and a system reflection. It was not possible to separate the system and sample reflections by windowing (as the sample and system reflections overlapped) and the sample reflection was required to extract the thickness of the sample. A calibrated infinite resonance model was fitted to each measurement at each temperature (see § 5.4), using the β from § 5.9.

A direct search was used in conjunction with total variation to extract the thickness of the sample at each temperature, as in § 5.6. As this window includes system reflections, etalons persist which are not reduced by the direct search and as such a lower frequency resolution window was used for analysis. The thickness extraction is not affected by the inclusion of system reflection in the time window, this is especially useful in setups which incur additional system reflections (such as cryostats, flowcells etc.).

In Figure 5.26, the absorption coefficient for the room temperature measurement is shown. The blue curve shows the low frequency resolution estimation, and the red curve shows the fitted absorption. The low frequency has been approximated using the technique from § 3.6, and the high frequency resolution estimate has been fitted using a resonant model, with the thickness optimised to minimise etalons. These are




Figure 5.26: The measured absorption coefficient of 5 % α -lactose anhydrous. Blue - Low frequency resolution estimate. Red - High frequency resolution estimate, based on fitting using resonant model, including system reflections




Figure 5.27: Thickness of a 5 % α -lactose anhydrous sample as a function of temperature. Blue circles - individual measurements. Line - fitted first order polynomial

still present in the absorption despite this due to the inclusion of system reflections. In this sample, the reduction of frequency resolution has little impact on the resultant spectra. This is because the spectral peaks are broad, however this is not inherently the case for all samples. There are also effects which are present which can change peak shape considerably, for instance homogeneous broadening is usually reduced by lowering the temperature of the sample [97] and changing sample concentration can narrow peaks by reducing non-homogeneous broadening [96]. To provide clear data, the low frequency resolution data has been used for analysis. This data has been estimated using the technique from § 3.6, using the extracted thickness from the high resolution data.

Using the first window in conjunction with a resonant fitting, the sample thick-

ness was measured at each temperature. In Figure 5.27 the measured thickness at each temperature is shown. Circles mark individual measurements of thickness, and the flat line shows a first order polynomial fitted to the thickness. From the measurements, there is a clear trend of the sample contracting with lowering the temperature. The room temperature appears to be an outlier in this dataset, however it is noted that it is measured before the cooling process, so it is possible that the sample structure has changed during cooling and reheating, resulting in a slightly thicker sample.

As discussed in § 3.12, the relative error in sample thickness can have a considerable effect on the sample n and α . During the cooling process the sample contracts, resulting in a significant change in n and α . If the sample thickness is assumed constant during cooling, this will instead result in error. The error in n is particularly significant, as it is of similar magnitude to the change due to temperature change. When the thickness change was taken into account during these measurements, it was found that the change in n was underestimated.

In Figure 5.28, the mean real refractive index (with respect to frequency) is shown as a function of temperature. The blue markers show the means with an extracted thickness, and the red markers show the means where the thickness has been assumed to be constant (measured thickness of 0.58 mm, using micrometer). The lines are first order polynomials fitted to the data excluding the room temperature mea-




Figure 5.28: The average real refractive index at each temperature. Blue - estimate with extracted thickness, Red - Estimate using constant thickness of 0.58 mm. Lines are first order polynomials fitted to the data excluding room temperature measurement.

surement. Without thickness extraction, the fitted offset is 1.46 with a negative gradient of 10^{-4} , while with thickness extraction the offset has increased to 1.5 and the negative gradient has doubled to 2×10^{-4} . The offset difference between the two different sets is due to the measured thickness (via micrometer) being a slight overestimate. Both curves show a similar trend with the real refractive index being higher at lower temperatures. However, the absolute change across temperature is doubled when calculated using the extracted thickness.

A plausible explanation for this change in n with temperature, is that the sample is a mixture of lactose, PTFE and air [136]. If the individual components expand at different rates, this results in the ratio of volumes changing and thus the measured \tilde{n} of the mixture changing with temperature. In Appendix C, the thickness change is attributed to an estimated air volume change. This was found to reduce the n change across temperature.

In Figure 5.29, the temperature dependent α is shown for the 5% α lactose anhydrous sample with 15 cm^{-1} offsets between temperatures. The measurements have been colour coded from coldest (blue) to warmest (red). Most features show a decrease of homogeneous broadening increasing with temperature, and some display weak blue (to higher frequency) shift at lower temperatures. In particular, the peak at 1.3 THz, shows a dramatic sharpening as the temperature decreases as well as a slight blue shift. The feature at 2 THz, show sharpening before merging with its higher frequency neighbor, and forming a broad doublet. Other higher frequency features show sharpening and slight blue shift with temperature. There exists a large feature at 6 THz from the PTFE matrix, which has been truncated out (as not to affect the thickness extraction performed at an earlier stage, due to α_{max}), which starts to dominate at 5 THz.

In Figure 5.30 the temperature dependent estimates of n of the 5% α lactose anhydrous samples are shown with 0.2 offsets. The measurements have been colour coded from coldest (blue) to warmest (red). There are clear trends of the relaxations, which correspond to the absorption peaks, becoming sharper as the temperature is lowered. There is also slight blue shift of the location, particularly the feature at 1.3 THz. The exception to this is the relaxation at 2 THz, which broadens at lower




Figure 5.29: the temperature dependent absorption coefficient for a 5 % lactose anhydrous sample, with 15 cm^{-1} offsets. The measurements have been colour coded from coldest (blue) to warmest (red). The measurements were taken at 4 K, then 10K to 290 K with 10 K increments and at room temperature (taken to be 293 K)

temperatures. This is due to the merging of the two features in the absorption, to form a broad doublet, which results in smearing out the relation across frequencies.




Figure 5.30: The temperature dependent n of 5% α lactose anhydrous offset by 0.2 at each measurement. The measurements have been colour coded by temperature, from coldest (blue) to warmest (red). The measurements were taken at 4 K, then 10 K to 290 K with 10 K increments and at room temperature (taken to be 293 K)

5.12 Multilayer sample extraction

Often there exists samples which consist of multiple layers, the most common example of this is a flowcell, which consists of outer window layers with a liquid in between. In these samples, there exists a singular uncharacterised layer with un-

known \tilde{n} , the other layers have known properties. It is the \tilde{n} of this layer which is extracted. To do this, a measurement of the sample is performed and normalised against an air reference measurement to calculate its transfer function. A model for the layered sample is then fitted to the transfer function, where the \tilde{n} of the unknown layer is used to adjust the fit, and the other layer properties are kept constant (relative to the fitting). This unknown layers thickness can then also be extracted using the thickness extraction process in § 5.7.

5.13 General multilayer finite resonance model

This can be challenging to model, as a reflection in one layer can lead to a series of reflections in another layer, which in turn, can each lead to a series of reflections in another layer, etc. If the layers are thin, the series of reflections in the sample are convergent, and the model simplifies considerably to a product of coefficients. A method which makes use of this, by modelling forward and backward standing waves is the transfer matrix method [75]. However, more generally due to the finite time window there are layers whose resonance is only partially encapsulated (if at all). Another method of simplifying the model, is to combine it with clever experiment design, so that the resonance of some of the layers is constrained. This was done in [76], by using a flow-cell with thick windows with a short time window. The flow-cell windows have a resonance period much greater than the time window, so the resonance within them can be ignored. The model used there was modified further, by using an empty flow-cell (air instead of liquid) as a reference. This removed the need to model propagation in the windows (but not the window-liquid interface), while including an additional resonance term (from the air in the channel) and reducing the overall dynamic range of the measurement.

An alternate, more general method, is to instead form a finite resonance model for the sample which fits the time window. This has been performed before by [64], which mapped out the resonance model using an approximate n for the unknown layer. This was then encoded into a “resonance matrix”, which was used to form a finite resonance model consisting of propagation, reflection and transmission coeffi-




Figure 5.31: Left - The THz beam paths through a 2 layered sample (angle offset for clarity). Right - The equivalent binary tree. Segments have been coloured to differentiate neighboring segments and end points have been labeled using letters.

icients. This method of course relies on an initial estimate being accurate enough to form the correct model (fortunately if thicknesses of each layer are reasonably accurately known, this should be the case). Another potential issue, is that it will model redundant (not detected) beam paths which are outside the dynamic range of the measurement. While this will not significantly change results, it will considerably simplify the model at higher frequencies (as longer paths become more redundant).

In this work, the sample model is built dynamically using a simplified ray tracing scheme. In order to do this, the beam paths are viewed as a "resonant tree", a binary tree formed from segments representing propagation over individual layers. In order to do this, the assumptions from § 3.2 are applied to the multilayer sample. This includes that the THz radiation is a culminated point beam propagating perpendicular to each interface in the sample and that the layers are homogeneous dielectrics.

In Figure 5.31 a binary layer sample is shown on the left, with the different THz beam paths through the sample. On the right is the equivalent resonance tree. The beam segments have been coloured for contrast (so neighboring segments have different colours) and end points have been labeled with different letters. Only those which leave the sample towards the detector result in a detected THz pulse, these have been terminated by arrows. There are many beam paths through the sample which are not detected by the system due to dynamic range, propagation toward the emitter and a finite time window, these have not been marked by arrows.

As the THz pulse propagates through the sample, it will be partially reflected at each interface. This results in a reflection beam path and a transmission beam path.

Each of these beam paths will then propagate to other interfaces, further splitting into an additional two paths, which will then further split into two paths and so forth. These different paths will lead to one of two events, a pulse being detected or a pulse not being detected. The former occurs when a beam path leaves the sample structure towards the detector. The latter occurs when a beam path leaves the structure towards the emitter, results in a propagation time longer than the measurement time window or results in a THz pulse amplitude below the inverse dynamic range, see § 3.7.

In Figure 5.32 the propagation of THz radiation in a layered sample measurement is shown. The sample consists of M layers, each with thickness l_L , where L is the layer index (1 is closest to emitter, M is closest to detector). The total thickness of the sample is $\sum_{L=1}^M l_L$. The sample response will consist of a number of pulses, each associated with a different beam path through the sample. These beam paths are in essence the different paths through the resonance tree, thus the transfer function is:

$$E_{sam} = E_0 H_{sys} H_T \quad (5.53)$$

where H_T is the transfer function of the detected resonance tree and is the sum of these paths which are detected.

As in § 3.2, the radiation propagates over the air which is displaced by the sample




Figure 5.32: The THz radiation passing through a M layered sample (angle for clarity). Each layer can have a unique thickness, l_L .

in the sample measurement, thus the air measurement is:

$$E_{ref} = E_0 H_{sys} P_a \left(\sum_{L=1}^M l_L \right) \quad (5.54)$$

The air propagation term is in this case evaluated using the sum of thicknesses. There are only terms for the displaced air in this measurement, no sample is present during the reference used in this work.

The sample transfer function model can therefore be written as:

$$\bar{H} = \frac{H_T}{P_a} \quad (5.55)$$

i.e. to find the transfer function for a layered sample is the same as finding the transfer function for its resonant tree.

This is not a trivial task as there is no way of calculating if a beam path is redundant (not-detected) without exploring it. In Figure 5.33 a simple resonance tree in a two layered sample is shown. The solid blue line shows the detected resonance tree consisting of a single pass through and reflection in the first layer. However, to determine this detected resonance tree the redundant beam paths, shown as dotted lines, had to be explored.

The symmetry at the nodes of binary trees is exploited to form a transfer function. A node (at an interface) in the resonance tree can be considered a point where a reflection sub tree and transmission sub tree are merged. In Figure 5.34 the same




Figure 5.33: The THz beam paths through a two layered sample (angle for clarity). Solid blue lines - beam paths which result in a detection. Dotted blue lines - redundant beam paths through the sample




Figure 5.34: The THz beam paths through a 2 layered sample (angle for clarity). Solid lines - beam paths which result in a detection. Dotted lines - redundant beam paths through the sample. Blue - initial segments of the tree. Green - Reflection sub tree. Red - Transmission sub tree.

bilayer with resonance tree has been shown. The tree has been split into three parts, the blue path is the initial beam path, the red paths are the transmission sub tree at the first interface and the green paths are the reflection sub tree at the first interface. Each of these sub trees can then be further divided at each interface.

This can then be exploited to define H_T in terms of sub tree transfer functions. Using the colour coding from Figure 5.34, the transfer function for the tree can be written as:

$$H_T = T_{0,1} P_1 (R_{1,2} H_{T,R} + T_{1,2} H_{T,T}) \quad (5.56)$$

where $T_{0,1}$ is the transfer coefficient for the air–first–layer interface, P_1 is the propagation coefficient for the first layer, $R_{1,2}$ is the reflection coefficient between the layers, $T_{1,2}$ is the transmission coefficient between layers, $H_{T,R}$ is the reflection sub tree transfer function and $H_{T,T}$ is the transmission sub tree.

This is a general relation and can be applied to arbitrary sub tree of the main tree. In Figure 5.35 the two layered sample is shown, with the resonant tree split into four different colours. The grey paths are not relevant to this example. Starting with the blue segment, a sub tree is traced out which consists of a green reflection sub tree and a red transmission sub tree.

A general form for a resonant sub tree initially propagating over layer with index L in propagation direction D ($D = 1$ or $D = -1$ for propagating towards or away




Figure 5.35: The THz beam paths through a 2 layered sample (angle for clarity). Solid lines - beam paths which result in a detection. Dotted lines - redundant beam paths through the sample. Gray - non-relevant parts of the resonance tree. Blue - initial segments of the sub tree. Green - Reflection sub tree. Red - Transmission sub tree.

from detector respectively) can be thus given by:

$$H_{L,D} = P_L (R_{L,L+D} H_{L,-D} + T_{L,L+D} H_{L+D,D}) \quad (5.57)$$

where P_L is the propagation coefficient for layer L , $T_{L,L+D}$ is the transmission coefficient between layers L & $L+D$, $R_{L,L+D}$ is the reflection coefficient between these layers, $H_{L,D}$ is the transfer function of the resonant tree, $H_{L,-D}$ is the reflection sub tree and $H_{L+D,D}$ is the transmission sub tree. This is defined *recursively*, only the layer index, L , and propagation direction, D , change for each sub-tree. It should be noted that $L = 0$ and $L = M + 1$ are the air pseudo-layers (no thickness) which are used solely for the transmission and reflection coefficients when leaving the sample.

To complete this termination conditions need to be formed (i.e. when beam paths are stopped being explored). This occurs under four conditions, when a beam path leaves the sample towards the director, when a beam path leaves the sample towards the emitter, when the beam is no longer detectable due to the propagation time being longer than the time window and when the beam is no longer detectable due to the limited dynamic range. There is also a fifth condition in the implementation, a recursion counter to prevent infinite recursion (this is a practical constraint and is not required from a theoretical point of view). With the exception of the first condition, these result in a redundant beam path. A sub tree (a collection of beam paths) which is redundant has a transfer function of $H_T = 0$. Beam paths which leave towards the detector have a transfer function of $H_T = 1$.

The constraints for leaving the sample are easy to resolve, they are entirely resolvable using the layer index L . However, to resolve the constraints based on time and dynamic range, each recursion these constraints must be tightened. A new transfer function model based on these constraints is formed:

$$H_{\mathbb{S}} = \begin{cases} P_L (R_{L,L+D}H_{\mathbb{S}_R} + T_{L,L+D}H_{\mathbb{S}_T}) & , \text{ if } \mathbb{S} = \text{non-redundant} \\ 1 & , \text{ if } \mathbb{S} = \text{detected} \\ 0 & , \text{ otherwise} \end{cases} \quad (5.58)$$

where \mathbb{S} is the sub tree state, \mathbb{S}_R is the reflection sub tree state and \mathbb{S}_T is the transmission sub tree state. The abstraction of a state is simply used for brevity in this work and is replaced by the underlying variables and constraints in the implementation. The construction of a resonance tree transfer function is similar to that in [64] (which considered a tree of fixed size based on time constraints) and the construction of transfer function coefficients is largely adapted from [64].

The state is used to encapsulate all the information about the current recursion:

$$\mathbb{S} = \begin{cases} L & \text{layer} \\ D & \text{direction} \\ T_{\text{lim}} & \text{time constraint} \\ N_{\text{lim}} & \text{noise constraint} \end{cases} \quad (5.59)$$

where T_{lim} is the propagation time constraint and N_{lim} is the noise (dynamic range) constraint. These are used to say if the recursion is in the following states based on the conditions:

$$\mathbb{S} = \begin{cases} \text{non-redundant} & \text{if } 0 < L \leq M \ \& \ 0 < T_{\text{lim}} \ \& \ |P_L| > N_{\text{lim}} \\ \text{detected} & \text{if } M < L \ \& \ 0 < T_{\text{lim}} \ \& \ |P_L| > N_{\text{lim}} \\ \text{redundant} & \text{otherwise} \end{cases} \quad (5.60)$$

The state is then updated for the reflection and transmission sub trees:

$$\mathbb{S}_R = \begin{cases} L & \text{layer} \\ -D & \text{direction} \\ T_{\text{lim}} - t_L & \text{time constraint} \\ \frac{N_{\text{lim}}}{|P_L R_{L,L+D}|} & \text{noise constraint} \end{cases} \quad (5.61)$$

$$\mathbb{S}_T = \begin{cases} L + D & \text{layer} \\ D & \text{direction} \\ T_{\text{lim}} - t_L & \text{time constraint} \\ \frac{N_{\text{lim}}}{|P_L T_{L,L+D}|} & \text{noise constraint} \end{cases} \quad (5.62)$$

Where t_L is the propagation time for a layer and is $t_l = \frac{n_L l_L}{c}$. The new sub tree states not only reflect the new current layer and propagation direction, but also the new constraints that the sub tree must satisfy. Each recursion will decrease the time constraint for the sub tree by the propagation time for the current layer and the noise limit is increased in proportion to attenuation caused by the current layer and interface. This has the effect of tightening the constraints for sub trees for each recursion.

This is then used in conjunction with the current layer position to determine if propagation over the current layer is redundant. An analogy would be finding a path through a maze with one end of a string attached to the path finder and the other to the entrance. The path finder can only find a route through the maze which is shorter than the string but can only rule out a route by exploring it to the limit of the string.

By equating redundant sub tree transfer functions to 0, any parent tree will not have any contribution from the redundant sub tree. In appendix D, the bilayer example is worked through on a recursion by recursion process to achieve a transfer function of the bilayer sample.

This can then be combined with initial conditions and a coefficient for the first

interface, to form a dynamic model for the transfer function of the resonant tree:

$$H_T = T_{0,1} H_{\mathbb{S}_I} \quad (5.63)$$

where \mathbb{S}_I is the initial tree state, which is:

$$\mathbb{S}_I = \begin{cases} 1 & \text{layer} \\ 1 & \text{direction} \\ T_{\max} + t_{\text{air}} & \text{time constraint} \\ \frac{|P_a|}{\text{DR}[T_{0,1}]} & \text{noise constraint} \end{cases} \quad (5.64)$$

where T_{\max} is the time length of the window used, is the dynamic range measurement from § 3.7 and t_{air} is the propagation time of the additional air measured in the reference measurement. This can be calculated using:

$$t_{\text{air}} = \frac{n_{\text{air}} \sum_{L=1}^M l_L}{c} \quad (5.65)$$

It should be noted that the noise constraint was found to have negligible impact on the extracted parameters (compared to the time constraint) but will speed up extraction as paths can quickly become redundant through this constraint. However, it was found to be preferable to slightly overestimate the dynamic range (for constraint purposes) as it also acted as a hard limit for extracted values. When $\alpha > \alpha_{\max}$, the extracted value using this constraint would be equal to α_{\max} . The dynamic range is subject to uncertainty and can easily be underestimated, overestimating it by 10% was found to generally guard against this effect.

All of the coefficients used in the model are precalculated and looked up during the calculation of the resonance tree. To calculate these coefficients, the same assumptions made in § 3.2 are applied to each layer. Namely that the layers are homogeneous dielectrics and the interfaces are perpendicular to the propagation direction. In principle it would be possible to generalise the multilayer model to arbitrary incidence angle using the coefficients from [62], each layer would simply

have an associated angle from the normal which would then be used to calculate the coefficients.

This approach to modelling the transfer function has several advantages. From an experimental viewpoint using an air reference and modelling the entirety of the sample model will give greater dynamic range than using a related sample as a reference [76]. Dynamic range was found to be proportional to the certainty in a measurement, see § 6.7.

Another considerable advantage is that this algorithm can be applied to arbitrary number of layers or constraint, see appendix D. This allows it to be applied to almost any layered sample. Although it should be noted, that its computational complexity will increase with the number of layers or constraint size.

Another feature is that the model is dynamic, it is formed as it is calculated. Another approach would be to use initial approximations to form a static model [64]. However, while this simplifies the model, it is reliant on approximations made before fitting.

Related to this, is during the extraction the parameters of at least one layer can change, which can simplify the model during extraction (decreasing computation time). It was found that the noise (dynamic range) constraint would particularly simplify the model, as the absorption of a layer does not have to change much to make sub tree resonances redundant.

A final advantage is that it can be implemented in very concise code using recursive programming techniques. Related to this, but not used in this work, is that each sub tree can be computed concurrently in parallel leading to faster extraction.

5.13.1 Incorporation of the beam alignment coefficient

It was shown in § 5.8 that the Gaussian beam profile will incur additional effects which must be modelled. These can be incorporated into the multilayer model using the additional coefficients from § 5.8, and modifying the sub tree state.

For each layer in the sample there exists a refractive index normalised thickness, l'_L , and an air normalised sum of layer thickness's, $l'_{air} = \frac{1}{n_{air}} \sum_{L=1}^M l_L$, which if $n_{air} = 1$ will just be total sample thickness. The sample transfer function, equa-

tion 5.55, can then be written as:

$$\bar{H} = \frac{H_T}{P_a G(l'_{air})} \quad (5.66)$$

Equation 5.63 can then be modified to include terms for the beam width and Gouy phase shift:

$$H_{\mathbb{S}} = \begin{cases} P_L G_L (R_{L,L+D} H_{\mathbb{S}_R} + T_{L,L+D} H_{\mathbb{S}_T}) & , \text{ if } \mathbb{S} = \text{non-redundant} \\ 1 & , \text{ if } \mathbb{S} = \text{detected} \\ w(z') & , \text{ otherwise} \end{cases} \quad (5.67)$$

where z' is the normalised propagation distance for the beam path and G_L is the precalculated Gouy shift coefficient for layer L , and is calculated by using the normalised thickness for that layer, $G_L = G(l_L)$.

In this model the detected beam path is not equal to one, but rather the amplitude change which will be detected due to the shift in focus. To calculate this, the n normalised beam path must be calculated for each detected path through the sample. This is performed by cumulatively summing the normalised thicknesses during the recursion. In order to do this \mathbb{S} is modified to include z' . This will have no effect on the constraints of the model, but the reflection and transmission states will be updated to reflect the current propagated distance.

$$\mathbb{S}_R = \begin{cases} L & \text{layer} \\ -D & \text{direction} \\ z' + l'_L & \text{normalised propagation distance} \\ T_{\text{lim}} - t_L & \text{time constraint} \\ \frac{N_{\text{lim}}}{|P_L R_{L,L+D}|} & \text{noise constraint} \end{cases} \quad (5.68)$$

$$\mathbb{S}_T = \begin{cases} L + D & \text{layer} \\ D & \text{direction} \\ z' + l'_L & \text{normalised propagation distance} \\ T_{\text{lim}} - t_L & \text{time constraint} \\ \frac{N_{\text{lim}}}{|P_L T_{L,L+D}|} & \text{noise constraint} \end{cases} \quad (5.69)$$

The initial state will then be:

$$\mathbb{S}_I = \begin{cases} 1 & \text{layer} \\ 1 & \text{direction} \\ 0 & \text{normalised propagation distance} \\ T_{\text{max}} + t_{\text{air}} & \text{time constraint} \\ \frac{1}{\text{DR}} & \text{noise constraint} \end{cases} \quad (5.70)$$

5.14 Extraction of the refractive index of IPA

To extract \tilde{n} of a liquid in a free space THz TDS system, a flowcell is typically used. These samples consist of a three layer system; two outer layers of a THz transparent material, and a active liquid layer. The outer layers have already been characterised, such that both the \tilde{n} and thickness of each of them are known. The liquid layer is then desired to be characterised.

A flowcell consisting of two 2.1 mm *z*-cut quartz windows was measured in NB THz TDS system. The emitter was biased with a 350 V 7 KHz square wave. The samples were measured at a sampling frequency of 33.33 THz, with a frequency resolution of approximately 200 GHz. Isopropanol (IPA) and air were measured in the flowcell, with the same 250 μm PTFE spacers (supports placed either side of the liquid layer, which do not interact with THz radiation). Calibration measurements were not available for this system, as such $\beta = 0$ was assumed. This will lead to a small systematic bias in the extraction of parameters, however, as the channel is thin, the active layer's parameters should be only weakly affected.

The flowcell, loaded with a liquid was treated as a conventional sample, and was

measured using a dry air reference. Estimates of the real refractive index were first made by treating the flow cell as a sample with the active layers thickness, estimated using equation 3.18, and then corrected. The absorption was then (this must be done after correcting n) similarly estimated using equation 3.20, and then corrected.

The correction to n was formed by assuming simple additivity of n between layers, to form the correction:

$$n_c = n - 2(n_q - 1) \frac{l_q}{l_c} \quad (5.71)$$

Where n_q and l_q are the real refractive index and thickness of the quartz layers respectively, n_c and l_c is the real refractive index and thickness of the liquid channel, and n is the estimate calculated using equation 3.18. This correction assumes that the transfer function phase is a sum of the individual layer contributions, and that the outer layers are identical.

Similarly, the absorption can be corrected:

$$\alpha_c = \alpha - 2 \frac{\alpha_q l_q}{l_c} \quad (5.72)$$

Where α_q is the absorption coefficient of the quartz layers, α_c is the absorption coefficient of the liquid channel, and α is the estimate calculated from equation 3.20. Below 0.4 THz, the absorption is negative. This is due to the systematic error in the z -cut quartz's \tilde{n} estimate and was improved by using an empty flowcell to extract the quartz's \tilde{n} .

These corrected estimates were based on the results of a separate measurement of z -cut quartz parameters, and were used as initial conditions for a fitting. The generalised resonant model was then fitted to the transfer function, using a time limit of 30 ps, which was estimated from the impulse response of the sample, a dynamic range limit was not used. A refractive index wrapping error function was used to perform the fit, the wrapping behaviour is not affected by using a multilayer model, provided the n of other layers are kept constant during the extraction.

In Figure 5.36, the n of the IPA in the flowcell is shown. The blue curve is the




Figure 5.36: n of IPA. Blue - extracted using a multilayer algorithm, Red dashed - initial estimate made before extraction.




Figure 5.37: α of IPA. Blue - extracted using a multilayer algorithm, Red dashed - initial estimate made before extraction.

extraction based on fitting a transfer function model and the red dashed curve is the initial approximation. Similarly, in Figure 5.37, the α of the IPA is shown. Again, the blue curve shows the fitted estimate and the red dashed curve shows the initial estimate. At 2.6 THz a peak is present, this is thought to be due to a dynamic range limitation (similar to α_{max}).

The IPA has reasonable coupling with the quartz and is absorbent, thus resonance has not occurred (or is extremely weak) in the channel. The time window was chosen, so that no resonance within the outer layers is measured. This has resulted in a measurement where no reflections are measured, and thus the model fitted will simplify to a non-resonant multilayer model. In this case, just as in § 3.6 the approximations perform extremely well. Unlike in previous examples however, it is still possible to extract the thickness of the channel by using the empty flowcell.

A related sample to the flowcell with a liquid in its middle layer, is a flowcell with dry air in its middle layer. It would be perfectly possible to extract the \tilde{n} of the air. However, dry air can be approximated to have a refractive index of 1, so this is not required. What this can be used to do, is provide both a measurement of the quartz \tilde{n} and the thickness of the channel. Generally it is assumed when performing a multilayer fit that there is a singular unknown layer which needs to be characterised. It is however, perfectly possible to work on the assumption that there are multiple layers which have the same unknown \tilde{n} , and extract the same \tilde{n} for multiple layers. Instead of extracting the \tilde{n} of the air layer, the \tilde{n} of both quartz layers is extracted by fitting a model to the transfer function and keeping the air channel refractive index constant during the fitting.

Air has a much lower n than quartz, and displays little absorption, so resonance will occur within the air layer. This will propagate to the extracted values of \tilde{n} of the quartz, and thus by extracting the \tilde{n} in a air flowcell, it is possible to perform a direct search to extract the thickness of the air channel.

In Figure 5.38 the extracted n of the quartz is shown with different assumed thickness's for the air channel. The curves have been coloured by total variance values (blue - high values and red - low values). The curve extracted at the total variance minimum is shown as a green curve. There are etalons which form and




Figure 5.38: n of z -cut quartz, extracted from an air filled flowcell. The values have been extracted with channel thickness ranging from 0.2 mm to 0.3 mm in $5 \mu\text{m}$ steps. Curves have been coloured by total variance values (red is lower, blue higher), and the extracted value, at the total variance minimum, is shown by a green line




Figure 5.39: α of z -cut quartz, extracted from an air filled flowcell. The values have been extracted with channel thickness ranging from 0.2 mm to 0.3 mm in $5 \mu\text{m}$ steps. Curves have been coloured by total variance values (red is lower, blue higher), and the extracted value, at the total variance minimum, is shown by a green line

reduce as the thickness is changed. Interestingly, the bulk n does not vary with the channel thickness, which would happen if the channel refractive index was extracted instead. This can be attributed to the very weak dependence of the extracted \tilde{n} of quartz on the air channel.

In Figure 5.39 the extracted values of α of the quartz is shown with different assumed thickness's of the air channel. The curves have been coloured by total variance values (blue, high values and red, low values). The curve extracted at the total variance minimum is shown as a green curve. Again, etalons have formed which are smoothed out with a change in channel thickness.




Figure 5.40: The total variance of the extracted n of z-cut quartz, measured against channel thickness.

In Figure 5.40, the measured total variance of the quartz n , against channel thickness. The total variation has a minimum at $275 \mu\text{m}$, which corresponds to the channel thickness, when filled with air. This shows that it is possible to extract the thickness of layers which are not being characterised by the fitting. The fact that these etalons form in the quartz \tilde{n} , indicate that etalons from resonance within the sample are additive with the extracted parameters. If this is indeed the case, this suggests that etalons from multiple sources (resonance within different layers) will add together, and can be minimized independently. This would mean that it would be possible to perform a direct search for the thicknesses of multiple layers sequentially.

In Figures 5.41 and 5.42, the extracted n and α of IPA are shown, using the values from Figures 5.41 and 5.42 for the quartz outer layers during the extraction. These values have been coloured by total variance value, and the curve extracted




Figure 5.41: n of IPA, extracted from a IPA filled flowcell. The quartz \tilde{n} used during extraction have been extracted from an air filled flowcell with channel thickness ranging from 0.2 mm to 0.3 mm in $5 \mu m$ steps. Curves have been coloured by total variance values (red is lower, blue higher), and the extracted value, at the total variance minimum, is shown by a green line

at the total variance minimum is shown in green. It has been assumed by using this method, that the channel thickness is the same in both the air and IPA measurements. The flowcell between measurements was not disassembled between measurements, instead the IPA was inserted via a plumbing system built into the flowcell. The etalons have propagated from the quartz \tilde{n} to the extracted IPA \tilde{n} . This has resulted in similar results to what would be expected, if recursion occurred within the IPA layer. It should be noted, that the local minima in total variance does not change during this process, and the etalons are being driven by the resonance in the air filled flowcell measurement, and not because of resonance within the IPA measurement. This presents a method of extracting the thickness of a layer, using multiple related measurements.




Figure 5.42: α of IPA, extracted from a IPA filled flowcell. The quartz \tilde{n} used during extraction have been extracted from an air filled flowcell with channel thickness ranging from 0.2 mm to 0.3 mm in 5 μm steps. Curves have been coloured by total variance values (red is lower, blue higher), and the extracted value, at the total variance minimum, is shown by a green line

5.15 Multi time fitting

The segmentation method from § 5.9 can be combined with a model fitting to extract information. To do this, a model for each segment is fitted to a segmented transfer function. This allows more information to be extracted from a single measurement. In this case, both the measurement and model is in effect split into multiple corresponding measurements and models. This relies on both the model and measurement being separable in the time domain (i.e. pulses do not overlap). By fitting to multiple measurements, with common parameters, more parameters can be extracted at once. In this case, the thickness can be simultaneously extracted with the complex refractive index.

To facilitate this, the transfer function must be segmented, which is performed using the multiple windowing of impulse responses described in § 5.8.2. This was also used to form time constraints, T_w , for each partition. These partitions were chosen to separate features. At each frequency, an error function was minimised with respect to \tilde{n} and l :

$$\tilde{n}, l = \underset{\tilde{n}, l}{\operatorname{argmin}} \mathbb{E} (H_w, \bar{H}_w(\tilde{n}, l)) \quad (5.73)$$

Where \mathbb{E} is an error function, H_w is the vector of partitioned transfer functions at

that frequency, and \bar{H}_w is the corresponding vector of models.

Unlike in previous fittings, where the thickness was constant, the thickness can be a fitting parameter. This means that the real refractive index cannot be wrapped (as the thickness changes), and thus \mathbb{E} must locate the minimum using unwrapped phase. Fortunately, in the simple case where each transfer function segment contains a single pulse, the phase of each segment tends to be very linear. Thus the unwrapped phase can be extrapolated by calculating the model over 2 other neighboring frequencies and fitting a straight line to 0 in the phase.

The error function is defined as:

$$\mathbb{E} = \sum_{j=0}^{M-1} \left(|H_j| - |\bar{H}_j| \right)^2 + \left| \angle H_j - \angle \bar{H}_j \right|^2 \quad (5.74)$$

By using this technique, it was possible to directly extract the thickness of the sample at each frequency. By doing this, the measurement frequency resolution is reduced (approximately determined by the first window). However, any sample response which can be partitioned in this way will not benefit from an improved frequency response. This is because the individual pulses in time must be shorter than the time window used to separate them. Using a longer window would not confer more information. In the frequency domain, a peak would not get any narrower by using a longer time window because the peak width is fundamental to the sample response and is not affected by the lack of frequency resolution.

This has the advantage of directly extracting the thickness without having to perform a numerically intensive search. Additionally, each frequency can be calculated concurrently in parallel, so this could become dramatically faster using implementation on specialist hardware.

Additionally it was found that this technique was very robust to poor initial estimates in the examples tested, particularly to overestimation (in most cases, a 100% overestimate in thickness as an initial condition would converge). This might be useful if the thickness cannot be measured and only initial estimates from tomography of the peaks are possible. A similar technique which used frequency domain filtering (equivalent to time domain windowing) exists [65], however this was reported to

produce multiple solutions (so requires accurate initial estimates).

5.16 Time localised finite resonance models

To fit to multiple transfer functions segmented in time, a model must be created which fits each segment. To do this, the generalised finite resonant model, § 5.13, is adapted to produce multiple time localised models.




Figure 5.43: Multilayer resonance models for each time segment are built simultaneously. The left diagram shows the transfer function of the first time segment and the right diagram shows the transfer function of the second segment.

In Figure 5.43 a resonance tree is shown. It is then split into two different trees, which consist of different beam paths. These two trees are not sub-trees of the main resonance tree, but are instead formed from different sets of beam paths. These sets of beam paths are time localised to different sequential time windows.

These multiple trees can be created by modifying the sub tree transfer function model to use multiple time constraints to calculate a vector of transfer functions. \bar{H} and H_T are modified to be vectors of multiple different time windowed transfer functions.

A set of time constraints is measured from the different time windows of the impulse response, such that:

$$T_h = \left[0, T_{h,1}, \sum_j^2 T_{h,j}, \dots, \sum_j^N T_{h,j} \right] \quad (5.75)$$

Where T_h is a vector of time constraints, and $T_{h,j}$ is the j th time window length. This is the cumulative sum of the time lengths of each window and will be one element

longer than the number of time windows used.

The sub tree transfer function is then defined as:

$$H_{\mathbb{S}} = \begin{cases} P_L G_L (R_{L,L+D} H_{\mathbb{S}_R} + T_{L,L+D} H_{\mathbb{S}_T}) & , \text{ if } \mathbb{S} = \text{non-redundant} \\ 1 & , \text{ if } \mathbb{S} = \text{detected} \\ w(z') \mathbf{W} & , \text{ otherwise} \end{cases} \quad (5.76)$$

where \mathbf{W} is the window section vector. The sub tree state consistent with § 5.13, however the time constraint, T_{lim} , is now a vector and the subtraction is applied to all elements in this vector. The constraint state is however slightly different:

$$\mathbb{S} = \begin{cases} \text{non-redundant} & \text{if } 0 < L \leq M \text{ \& } 0 < T_{\text{lim},N} \text{ \& } |P_L| > N_{\text{lim}} \\ \text{detected} & \text{if } M < L \text{ \& } 0 < T_{\text{lim},N} \text{ \& } |P_L| > N_{\text{lim}} \\ \text{redundant} & \text{otherwise} \end{cases} \quad (5.77)$$

where $T_{\text{lim},N}$ is the last element of T_{lim} (recursion continues until the last time constraint). \mathbf{W} is the window select vector and is calculated from T_{lim} :

$$\mathbf{W}_j = \begin{cases} 1 & T_{\text{lim},j+1} > 0, T_{\text{lim},j} < 0 \\ 0 & \text{otherwise} \end{cases} \quad (5.78)$$

For example, given the vector $T_{\text{lim}} = [-3, -2, 1, 3]$, $\mathbf{W} = [0, 1, 0]$.

The initial state used is then:

$$\mathbb{S}_I = \begin{cases} 1 & \text{layer} \\ 1 & \text{direction} \\ T_h + t_{\text{air}} & \text{time constraint} \\ \frac{1}{\text{DR}} & \text{noise constraint} \end{cases} \quad (5.79)$$

This model can be used to form time localised models of a sample which consist of several layers. This can in turn be fitted to multiple time segmented transfer functions.

5.17 Time localised extraction of quartz

A measurement was performed by partitioning the impulse response of a 0.512 mm z -cut quartz sample from § 5.10. Two time windows were used to partition the impulse response, each with a frequency resolution of approximately 135 GHz. The first contained the main transmitted pulse and the second contained the first reflection from the sample. The time localised model was then fitted to a partitioned transfer function. The calibration curve presented in § 5.9 was used to calibrate the extraction.

Using this multi-time localised fitting, it is possible to extract \tilde{n} and l . This results in lower resolution data than using a full window, but it will be etalon free, if the impulse response is partitioned such that the transmissions through the sample (initial pulse and reflections) are separated. Using this technique this way does not result in additional information being extracted.




Figure 5.44: Real refractive index of z -cut quartz. Blue - real refractive index extracted simultaneously with thickness. Green - real refractive index extracted with an assumed thickness of 0.5 mm. Red - extracted value formed from § 5.10

Using a time localised fitting in this manner does not confer any additional benefits, simply shorting the time window used to process the data will yield similar results. Using a time localised fitting however incurs the ability to extract the thickness of the sample (at each frequency) simultaneously with \tilde{n} . In Figure 5.44, the real refractive index of the sample has been extracted along with the thickness of the sample. The average of the frequency dependent thickness was 0.512 mm, which is the measured thickness of the sample. The blue curve represents the values of n when the thickness has been extracted simultaneously, and the green curve is




Figure 5.45: Absorption coefficient of *z*-cut quartz. Blue - absorption coefficient extracted simultaneously with thickness. Green - extracted assuming a thickness of 0.5 mm (offset of 5 cm^{-1}). Red - extracted value from § 5.10 (10 cm^{-1} offset)

when the thickness has been kept at an assumed thickness of 0.5 mm (this thickness was used as an initial point during thickness extraction). The red curve shows the estimate from § 5.10 extracted using a resonant model. The resonant model has an extremely similar shape and value, the slight difference is due to the different extracted thickness of 0.515 mm. Using this method, a smooth estimate is extracted due to the limited time windows used.

In Figure 5.45, the corresponding absorption coefficient is shown. The blue curve shows the values of the α when simultaneously extracting thickness and green shows the extracted values, assuming a thickness of 0.5 mm, with a 5 cm^{-1} offset. The red curve shows the estimate from § 5.10 extracted using a resonant model with a 10 cm^{-1} offset.




Figure 5.46: The extracted thickness of *z*-cut quartz. Blue - extracted without β calibration, Green - extracted with β calibration, Red - measured thickness of 0.512 mm




Figure 5.47: Blue - The normalised thickness calibration error. Red - β used to calibrate for this error.

In Figure 5.46, the frequency dependent extracted thickness is shown. The blue curve shows the thickness extracted without using calibration ($\beta = 0$), and the green curve shows the calibrated extracted values. As a point of reference, the measured thickness of 0.512 mm is shown as a red line. The calibrated curve is clearly more consistent and closer to the measured thickness of 0.512 mm. It was observed that the calibration error between the two thicknesses, when normalised to the thickness of measured thickness (0.512 mm) was equal to β . In Figure 5.47, the normalised thickness calibration error (difference between extracted thicknesses) normalised to the measured thickness (0.512 mm) is shown in blue. The calibration curve used for β is shown as a red line. It was observed that these values were almost identical, suggesting that β primarily corrects the effect of the sample thickness on the model.

5.18 Time localised extraction of α -lactose anhydrous

It was found that the room temperature measurement of 5% α -lactose anhydrous from § 5.11 was a suitable candidate for time localised extraction of both l and \tilde{n} . Unfortunately, this was not the case for most of the lower temperature measurements, as the spectral features sharpened, the time domain traces became longer and the individual pulses became impossible to separate successfully. This technique relies on the pulses being separable in the time domain, when this is not the case, convergence on a solution does not seem to be possible.




Figure 5.48: Extracted α of 5% α lactose anhydrous, Blue- extracted using segmented transfer functions, Green - extracted using a single window low res window (120 GHz resolution) and Red - extracted using a single window and an infinite resonant model to resolve the first sample reflection.

Two windows were used to segment the impulse response before converting back to transfer functions. The windows used to segment the transfer functions had a frequency resolution of 200 GHz (both windows were the same length). Neither segment appeared to contain etalons, which would be associated with included system reflections in the initial reference and sample windows (which was the case). This is likely due to the re-windowing, which will distort (reducing their prominence) artifacts shared between the multiple windows. It is assumed that this error is still present, but has simply been obscured in this case.

A time localised model was fitted to the different segmented transfer functions to extract both l and \tilde{n} . To prevent unpredictable behaviour during the fitting, a dynamic range limit was not used during the fitting, only time limits were. Using a dynamic range limit can result in discontinuous results due to the reflection being limited, effectively creating a hard constraint on the extraction (this is not the case when not segmenting the transfer function). The extraction appears to be robust to this limit, this is because the dynamic range is based on the amplitude of the measurement. As the additional information (thickness) is strongly phase dependent, the extraction seems to be able to cope when this limit is reached.

In Figure 5.48, the extracted α of the sample is shown. The blue curve shows the values extracted using multiple transfer function segments using a time localised model. The green curve shows the extracted values using a single low resolution (120




Figure 5.49: n of 5% α lactose anhydrous. Blue- extracted using segmented transfer functions, Dashed Green - extracted using a single window low res window (120 GHz resolution) and Dashed Red - extracted using a single window and an infinite resonant model to resolve the first sample reflection.

GHz) window, calculated using the approximations from § 3.6. The red curve shows the extracted values using a single window (70 GHz resolution) which contained both a single reflection from the system and from the sample. A infinite resonance model using the calibration from § 5.9. Both the green and blue curves do not display etalons due to the low frequency resolution window used. The blue curve displays some distortion at lower frequencies (up to 1 THz), which might be due to weak overlap of the features temporally, which results in a loss of frequency resolution of the features. The red curve displays etalons, particularly at higher frequencies. At low frequencies, where the sample reflection is more dominant and produces etalons, the resonant model has mitigated the etalons. However, at higher frequencies where the system reflection etalons are more dominant, this has not had an effect.

In Figure 5.49, the extracted n is shown. The blue curve shows the multi-segment extracted estimate, the green curve shows the low resolution approximation and the red curve shows the estimate formed from a high resolution time window with a resonant model. The blue estimate is smoother in general, apart from frequencies below 1 THz, where the windowing process has introduced some distortion in the form of large etalons. This suggests a partial overlap between the pulses, and the second window including some of the lower frequencies from the first window. Both the multi segment and high frequency resolution estimate appear to agree, this is because both of these calibrate for the slight error in n due to the Gouy phase

shift, which the approximate solution does not. This could be corrected using the approximation in [100].

In Figure 5.50 the extracted thickness is shown (blue), along with the extracted thickness from § 5.11 (red line). The extracted estimate is not perfectly flat and shows a degree of noise at both ends of the bandwidth. The estimate is largely convergent on the same estimate as using a direct search in conjunction with a resonant model. However, no direct search (performing an assumed thickness sweep) is required and no loss of peak width is observed by using the lower frequency resolution.

This extraction technique can be used successfully on samples which have a time response which does not overlap and can be segmented. This segmentation does lower the measurement frequency resolution. However, the frequency resolution is primarily determined by the sample response (not the technique). If a sample produces a long time response (and thus requires good frequency resolution) which does not overlap, this technique will produce sufficient frequency resolution to resolve the frequency peaks because the time window will contain the entire time response of a pulse.

However, this is reliant on the assumption that the sample response does not overlap and that it is possible to observe that this is the case in the time domain. This can be resolved experimentally by using a thicker sample, which will delay the sample reflection allowing a longer initial time window.




Figure 5.50: Extracted thickness of a 5 % lactose anhydrous sample. Blue - extracted using a segmented transfer function. Red - extracted value from § 5.11 using a direct search with total variance.




Figure 5.51: The windowed segments of the impulse response of a high resistivity silicon and quartz layered sample. Blue - The main transmitted pulse. Green - The first reflection from the quartz layer

5.19 Extraction of complex refractive index of a z-cut quartz and high resistivity silicon layered sample using a time localised model

Fitting a time localised model to a segmented transfer function could be used to extract information other than the thickness of a sample. For instance if a layered sample produced resonance, it might be possible to characterise each layer from a single measurement. This is because a second transfer function segment will carry both amplitude and phase information.

A 1.01 mm high resistivity silicon and 0.512 mm z-cut quartz layer sample were measured using the HFRBB THz TDS system, with a bias voltage of 700 V and at a sampling frequency of 500 THz. The layers were mechanically clamped in a lens holder. The silicon was the layer closest to the emitter and the quartz was closest to the detector.

The transfer function was segmented, so that the main pulse (which passes through both layers) was in the first time window and the first reflection was in the second window. In Figure 5.51, the segmented impulse response is shown. The blue pulse is the main transmitted pulse and the green is the first reflection. This reflection is from the quartz layer which is both thinner and has a lower refractive index. Other reflections have been excluded for this measurement.

A time localised model was then fitted to both segments, the thicknesses and order of both layers was kept constant. The \tilde{n} of both layers was optimised during the fitting to minimise the error of both windows. The initial values used were based on separate measurements of each layer.

In Figure 5.52 the extracted real refractive indices are shown. The left pane shows the n of the silicon and the right pane shows the n of the quartz. In dashed lines, the initial estimates are shown for reference. These were approximated based on equations in § 3.6 and measured under the same conditions. The extracted silicon curve appears to be underestimated while the extracted quartz curve appears to be overestimated.

In Figure 5.53 the extracted absorption coefficient for both layers is shown. The blue curves shows the extracted α of the silicon layer while the green curve shows the extracted α of the quartz layer. The initial estimates have been shown as dashed curves of the same colours for reference. The silicon curve has been overestimated while the quartz curve has been underestimated (hence why it is negative). The extracted estimate of the quartz contains the 3.8 THz peak which is present in the initial estimate.

While these extracted values have been over or underestimated respectively in each case, this technique does have viability and could be improved upon to fully characterise a layered sample. If more segments were used, additional phase




Figure 5.52: The n of a high resistivity silicon and z-cut quartz layered sample. Left pane - n of silicon. Right pane - n of quartz. Solid lines - Extracted from a time segmented transfer function. Dashed lines - approximate estimates based on separate measurements of each layer.




Figure 5.53: The α of a high resistivity silicon and z-cut quartz layered sample. Blue - α of silicon. Green - α of quartz. Solid lines - Extracted from a time segmented transfer function. Dashed lines - approximate estimates based on separate measurements of each layer.

information might be available and it might be possible to extract the thickness of each layer, in addition to \tilde{n} of each layer, from a single measurement. Another potential avenue is to combine it with Kramers–Kronig relations, similar to the method in [85] to extract more plausible values of \tilde{n} .

The issue with this approach to extracting the \tilde{n} of each layer is that uncertainty between the different values are negatively correlated. Overestimate the α of one layer and the other layer will compensate to minimise the error for the overall fit. This results in correlated uncertainty between the extracted estimates of \tilde{n} . This might be compounded with experimental uncertainties.

For instance, if the angular alignment of the sample to the THz beam is slightly wrong, this will result in the propagation length in both samples being longer than the thickness each layer. In a singular layer sample, this will primarily act to increase the effective thickness (which can be extracted). In this case, the quartz layer thickness is effectively underestimated more than the silicon's (due to Snell's law). This results in the quartz's refractive index being overestimated, to compensate the silicon's refractive index is overestimated. This results in the transmission and reflection coefficient being different and thus the extracted absorption acts to compensate for this effect.

5.20 Conclusion

In this chapter, the process of extracting \tilde{n} and l of a sample by fitting models to the sample transfer functions were discussed and demonstrated in detail. It was found that when a model, which was calibrated for the Gaussian beam profile, was fitted to the transfer function of a sample, the thickness could be accurately extracted. This is performed by minimising the etalons in the extracted n as a function of assumed thickness. The method of measuring the beam alignment coefficient, β , from a calibration sample was presented. This was used to calibrate other measurements within this work.

A method of modelling multilayered samples was presented and was used to extract the thickness and \tilde{n} of a liquid layer in a flow cell. This model was extended to work on multiple time segments of the transfer function. This was then used to develop a new method of extracting both the \tilde{n} and l of a single layer sample. This method directly extracts the thickness as a function of frequency, so can be computationally efficient compared to directly searching for the correct thickness. It was found that this technique was frequency resolution limited, but that this limitation is due to the sample response. This technique was then applied to a multilayer sample to extract the \tilde{n} of each layer from a single measurement. The extracted values were observed to be negatively correlated, when a parameter was overestimated in one layer, the parameter in the other layer was underestimated. This could be made more viable in the future by accounting for measurement uncertainties and modelling their effect on the extracted parameters.

The use of resonant models in this work leaves residual etalons present in the extracted parameters. These etalons are partially owing to the different measurement uncertainties present in the measurement. By estimating the uncertainty in the extracted complex refractive index, it might be possible to account for the residual etalons in \tilde{n} .

Chapter 6

Estimating uncertainty in complex refractive index and thickness

In this chapter, a method of estimating the covariance matrices of both \tilde{n} and the variance of l is presented. The covariance matrix is a measure of uncertainty and correlation within a measurement, from which confidence intervals can be formed. A model for the probability distribution of noise within THz TDS measurements based on normal (Gaussian) statistics was developed. Using this model a method of propagating the correlated uncertainty to the spectroscopic parameters based on the guide to the expression of uncertainty in measurement [83], [137] is given. Confidence intervals are formed from these uncertainty estimates, which are applied to both the refractive index and absorption coefficient of α lactose monohydrate and z -cut quartz. The uncertainty of the extracted l of the z -cut quartz is then estimated using a resonant model fitting. The methods developed in this work have been published in [138] and are compared against established empirical methods.

6.1 Introduction

So far in this work, measurements have almost exclusively been considered deterministic and repeatable. However, a measurement is subject to two types of error,

measurement and systematic. Systematic error refers to error within the measurement process itself, which could be underlying biases or uncertainty in the procedure or equipment. A common example of this is THz TDS is the sample thickness, which is relied upon to calculate \tilde{n} , see § 3.12. The measurement error refers to random error which changes with each offset, for example electrical noise within the detector equipment.

The purpose of this chapter is to estimate the uncertainty due to measurement error. While measurement error can be very small, it will have an effect on the extracted parameters which is not insignificant. Additionally, systematic errors such as the alignment error and spatial variation in the sample properties can be experimentally measured, and thus treated as measurement error. For example, by reinserting the sample into the system between measurements, multiple measurements of the same sample under different alignments (both spatial and angular) can be gathered, with different measured error.

Previous methods have been developed which estimate the effect of measurement uncertainty, see § 2.2. In this work, the only assumptions made about the measurement is that it has a multivariate Gaussian distribution. The source of the measurement uncertainty is purposely kept general, as such the methodology in this work should work with any THz spectroscopy system (and could also be applied more broadly to many other instruments well beyond the scope of this work).

The noise in a THz TDS system is modelled as a random variable with a multivariate normal (Gaussian) distribution, see Appendix E. It was found that the measured time domain scans approximately exhibit a zero mean multivariate normal distribution, see § 6.3. A normal distribution model was used as it is completely described by its expectation (average) and covariance, and has multiplicative and additive properties [139]. The expectation represents the average taken over infinite observations (or alternatively the value weighted integral of probability density function), and will in this work be an estimate of the deterministic signal. The variance is the expected squared distance from the expectation, or how far on average a value will be from the expectation. This is used as a measure of uncertainty as it will often represent how close a measured value will be to the deterministic value (i.e. measurements

made with low variance noise will tend to be close to the deterministic value). The covariance is a combined measure of correlation and variance and can be considered a measure of the correlated uncertainty. It was found to be particularly important when transforming between the time and frequency domains and for analysing the underlying causes of uncertainty.

6.2 Modelling a THz time scan

There exists a number of different sources of error within a THz TDS time scan, in particular there exists several sources of measurement error. The uncertainty from this measurement error can be estimated from multiple repeated measures of a THz scan. In order to do this a basic model of the uncertainty in the THz scan is required.

A single THz time domain scan consists of many discrete samples, each taken at time t_k . It is assumed in this work, that these samples are taken with regular period, such that $t_k = T_s k$. Each discrete sample consists of an amplitude measurement of the THz electric field, measured at time $T_s k$, which is noted as:

$$e(k) = e(T_s k) \quad (6.1)$$

Each of these discrete samples has an associated error, which contributes error to the measured amplitude, $\hat{e}(k)$:

$$\hat{e}(k) = e(k) + v(k) \quad (6.2)$$

where $v(k)$ is the error at index k . This error is a collective summation of different errors from different error sources. In the literature, several different sources of error have been described, in particular general detector noise [61] (either in measurement amplifiers or in the photo-diodes), time stamp error and laser noise [140]. Time stamp error has several different associated variations, in particular random error [141], registration error of the stage [142] and periodic jitter in the time stamp [115]. Rather than individually modelling each of these sources, each of these sources are assumed to be normally distributed with zero average, thus the

sum of these variables, $v(k)$, should be normally distributed with zero average. This is confirmed by measurement of the marginal empirical distributions over a THz pulse, see Appendix E.

The THz measurement scan, consisting of individual samples, can be grouped into a multivariate variable:

$$\hat{\mathbf{e}} = \begin{bmatrix} e(0) + v(0) \\ e(1) + v(1) \\ \vdots \\ e(N-2) + v(N-2) \\ e(N-1) + v(N-1) \end{bmatrix} \quad (6.3)$$

If \mathbf{e} is deterministic, such that it is constant with repeated measurement, and \mathbf{v} is zero mean, then $\hat{\mathbf{e}}$ will be normally distributed (see appendix E):

$$\hat{\mathbf{e}} \sim \mathcal{N}_N(\mathbf{e}, \Sigma_{\mathbf{v}}) \quad (6.4)$$

This leads to the useful observation, that the expectation will be \mathbf{e} and the covariance will characterise the uncertainty due to noise.

\mathbf{v} is the multivariate normally distributed variable, and can always be expressed as a matrix times a variable, \mathbf{s} , with an uncorrelated standard normal distribution :

$$\mathbf{v} = \mathbf{V}\mathbf{s} \quad (6.5)$$

This will hold true, regardless of the underlying source (or combinations of sources) of error provided \mathbf{v} is approximately normal. The structure of matrix \mathbf{V} will be determined by the sources (and combination of sources) of error and the processing of data. An example of a process which could affect this matrix, is applying a low-pass filter (for example the effect of the lock-in amplifier).

\mathbf{V} will define the covariance matrix:

$$\Sigma_{\mathbf{v}} = \mathbf{V}\mathbf{V}' \quad (6.6)$$

Which is useful to perform a rudimentary analysis of the error sources.

Fortunately, as \hat{e} has expectation e , e can be estimated from multiple measurements of \hat{e} using the arithmetic mean [107]:

$$e = \mathbb{E}(\hat{e}) \approx \mu(\hat{e}) \quad (6.7)$$

6.2.1 Estimating the expectation and covariance

It is not normally possible to know either of the defining statistics of a stochastic variable. However, it is possible to estimate them with some degree of uncertainty from multiple repeated observations.

For variable \mathbf{x} , the expectation can be estimated using the sample mean [139]:

$$\bar{\mathbf{x}} \approx \mu(\mathbf{x}) \quad (6.8)$$

where $\mu(\mathbf{x})$ is the sample mean of \mathbf{x} , which can be calculated from multiple measurements of \mathbf{x} :

$$\mu(\mathbf{x}) = \frac{1}{M} \sum_{k=0}^{M-1} \mathbf{x}_k \quad (6.9)$$

where \mathbf{x}_k is the k th measurement of a set of M measurements. The arithmetic mean is only an estimator of the expectation and is subject to uncertainty in the original measurements. The standard deviation of the mean is often referred to as standard error [137], and can be calculated from the covariance matrix of \mathbf{x} :

$$\Sigma_{\bar{\mathbf{x}}} = \frac{1}{M} \Sigma_{\mathbf{x}} \quad (6.10)$$

Similarly for the co-covariance matrix:

$$\Gamma_{\bar{\mathbf{x}}} = \frac{1}{M} \Gamma_{\mathbf{x}} \quad (6.11)$$

More formally, $\mu(\mathbf{x})$ is distributed normally:

$$\mu(\mathbf{x}) \sim \mathcal{N}_N(\bar{\mathbf{x}}, \Sigma_{\bar{\mathbf{x}}}, \Gamma_{\bar{\mathbf{x}}}) \quad (6.12)$$

Importantly $\mu(\mathbf{x})$ has expectation $\bar{\mathbf{x}}$ and as the number of measurements is increased, the probability distribution gets more constrained about this value. This can be thought of as the uncertainty decreasing, or alternatively the noise being reduced.

$\Sigma_{\mathbf{x}}$ can be estimated from M measurements using the sample covariance matrix [139]:

$$\Sigma_{\mathbf{x}} \approx \frac{1}{M-1} \sum_{k=0}^{M-1} (\mathbf{x}_k - \mu(\mathbf{x})) (\mathbf{x}_k - \mu(\mathbf{x}))^\dagger \quad (6.13)$$

The sample covariance matrix will again be a random variable and have a Wishart probability distribution, or in the univariate case a χ^2 distribution [139]. This is particularly important, given that the covariance is only estimated and should be considered a random variable in itself.

6.3 Measurements of time domain uncertainty

In § 6.2 a model for a time domain scan was discussed. To validate this model a series of reference measurements were taken on the LFRBB THz TDS. These were measured at a sampling frequency of 25 THz, with a bias voltage of 150 V chopped at 7 KHz. 410 different scans were taken, the limitation of the number of scans taken is the time it takes for one scan. To increase this, the delay stage was moved faster, this resulted in measurements with higher than normal variance. This was not considered an issue in this analysis, as it is the noise which is being characterised.

In Figure 6.1 the average time scan is shown in the top pane. Two red lines mark the range which shows the greatest variation, consisting of 15 time samples. These samples are analysed in greater detail at a later point in this section. In Figure 6.1 the variance of the time scans as a function of time is shown in the bottom pane. What is interesting is that it shows large peaks at the edges of the THz pulse. This is indicative of time base error [141], which manifests itself as an amplitude error proportional to the differential of underlying deterministic signal. This uncertainty was found to be highly correlated within the THz scan.

In Figure 6.2 the correlation matrix of the THz scan is shown. This consists of an image coloured from blue (-1), to green (0) and to red (1). As would be expected, regions which show little THz radiation in the scan are mostly not correlated with




Figure 6.1: .

An air time scan. Top - Blue line - average of 410 different measurements. Vertical dashed red lines - time range for further analysis. Bottom - Variance of the time scans with respect to time




Figure 6.2: Correlation matrix of the time scan. Blue - -1 , Green - 0 and Red - $+1$

other parts of the scan, however in the THz pulse there is a strong checked positive and negative correlation pattern. This happens to correspond with the edges of the THz pulse, rising edges are self correlated, as are falling edges, but rising and falling edges are negatively correlated. In practice this means that when a time sample is measured slightly higher on a rising edge, all the other samples on a rising edge tend to move up as well, and all the samples on a falling edge tend to move down.

The correlation between time samples cannot be assumed negligible, as in [81], and is particularly important when applying a Fourier transform. A singular sample in frequency will be dependent on all time samples within the time window. If time samples are correlated, this will have a dramatic effect on the variance at singular frequency points. The reverse is also true, time samples depend on the entirety of the spectrum, thus correlation in the frequency domain should also be considered (even when estimating variance in the frequency domain).

The simplest example is to consider a time domain scan consisting of a single sinusoid at a singular frequency. If the magnitude and phase of this sinusoid is random,

this will affect all points in the time scan and produce variation between repeated measurement. However, because it is the magnitude and phase of the sinusoid which is random and not the individual time samples, the variation across time will be correlated. If this was reversed, by applying a Fourier transform, this correlation would equate to variance in a very narrow region of the frequency spectrum.

This highlights the main criticism of previous work [81], which attempt to use univariate statistics to estimate the frequency domain uncertainty. It was found that univariate models of uncertainty are insufficient and the correlation in *both* frequency and time domains must be accounted for, when estimating uncertainty in the frequency domain.




Figure 6.3: Green - Magnitude of the spectrum. Blue - estimate of variance based on [81], Dashed Cyan - estimate of variance formed from the bootstrap and Red - estimate of variance based on uncertainty propagation.

In Figure 6.3, the magnitude spectrum (green) of the time domain data is shown, along with different estimates of the variance. The blue curve shows an estimate made with the method in [81], the red curve shows the estimate made using uncertainty propagation, see § 6.4.3, 6.4.4, and the dashed cyan curve shows an estimate made using the bootstrap data, see § 6.4. The bootstrap and multivariate uncertainty propagation agree extremely well, while the other estimate has overestimated the variance considerably.

In the case of time base error, it is possible to approximate a correction for the time base by looking at the deviation from the mean. By dividing the deviation from the mean by the finite difference and multiplying by the sampling time period, and estimate of the time stamp error for each time sample can be found. In Figure 6.4,




Figure 6.4: Time scan between the two red dashed lines in Figure 6.1. The time stamps have been corrected, and the scans have been coloured by the deviation from the mean at the falling edge marked by the vertical line

the different samples are shown with time stamp estimates. At points where the time base error is dominant, the samples appear to form a line with the gradient of the THz pulse. At points where the gradient is small, this has not been effective and the time base correction is very unreliable. This is particularly prominent at the maximum and minimum of the scan. A vertical line has been placed on the sample which has been used to determine the colour of the time scans. The time scans have been coloured from red to blue, representing positive or negative deviation from the mean at this point. Thus the colour represents deviation at this falling edge. Interestingly, while the red stamps appear to be shifted left with a time correction, while the blue stamps appear to be shifted right. Rather than the time base error being random at each point, it appears to be an error across the entire pulse. This could be registration error within the stage [142], which has been previously reported in THz TDS measurements, an alternative hypothesis is that error in the stage position will simply persist over a long period of its movement (i.e. this could be local to regions of the THz scan rather than entire scans). This effect will mean that there is an arrival time uncertainty of the pulse, which will manifest itself as an offset uncertainty in the real refractive index.

The simplest method to assess the probability distribution of the noise of the THz scan is to examine the empirical marginal distributions. In § E, it was shown that a marginal distribution of multivariate normal distribution is also normal. This

will also occur with univariate marginal distribution. A simple method to examine the distribution empirically is to examine the histogram of each time sample and compare to the parametric normal distribution of each time sample.

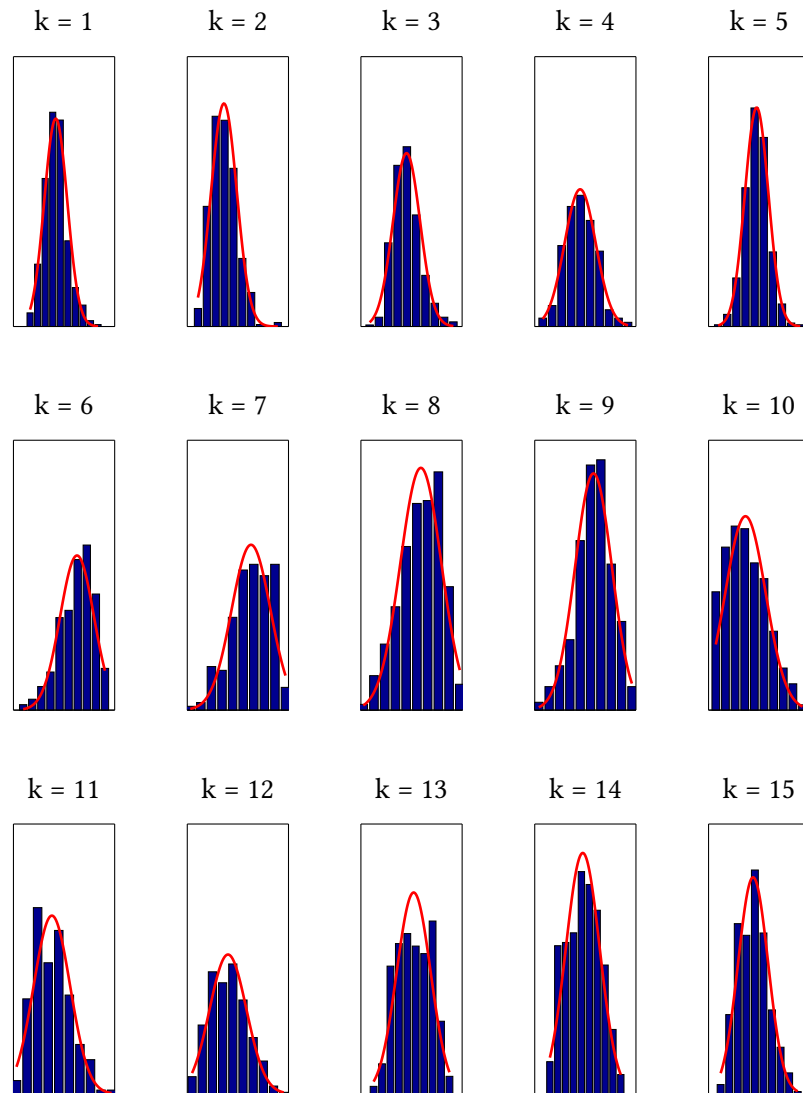


Figure 6.5: Normalised 10 bin histograms for each time sample between the two vertical lines in figure 6.1. The red lines show the parametric normal distribution based on the sample mean and variance

In Figure 6.5, 10 bin histograms are shown for the 15 time samples (k is the index) between the two red lines in Figure 6.1. A histogram can be used as an empirical approximation of the underlying distribution, by comparing them to the parametric model of the noise, it is possible to assess the marginal distributions at each time sample. These histograms have been normalised to the area they cover, so that the area is approximately 1. The red line shows a parametric normal distribution,

specified by the sample mean and variance in each case. In most of the cases, the histograms tend to be similar to each of the parametric normal distributions, although $k = 6 \dots 10$ appear to be slightly skewed and $k = 13, 14$ appears to be more uniform than would be expected if normal. These points occur at peaks within the THz power and are likely down to slight laser power drift over the course of the measurements (several hours) than typical variation in a conventional measurement. Laser power drift will have an effect on the amplitude of the measurement, as the generated THz power is proportional to this noise [81].

A more accurate estimate of the empirical distribution can be formed by using a bootstrap method (see § 6.4) to estimate the distribution of the mean. Using a bootstrap method will be more representative of the empirical distribution of the measurement noise as it is more sensitive to the likelihood of the noise [143] and can be used to form a much larger size of histogram. This was performed with this data set to produce 5000 different estimates of the mean time domain scan. In Figure 6.6 the resulting 500 bin histograms were then normalised and compared to the parametric normal distributions. Due to the large number of bins used in the histogram, the histogram appears to be almost continuous, and is very similar to the parametric curve. Although outliers are present in some cases (see $k = 7, 8, 9$), but considering their relative infrequency, they are assumed negligible.

As the processing is performed on an average of the time scans, it is the distribution of this mean which is most important. Fortunately it appears or tends to be normally distributed, this is generally validated by the processing results when compared to bootstrap estimations. If it were not the case, these would diverge significantly. What these results might indicate is that the distribution of a single scan is less normal. It is preferable from an analysis point of view to continue to treat it as normally distributed, although it is acknowledged that this might not always be true. A consequence which is not considered in this work, is this non-normality might introduce a bias into the sample mean. This is assumed negligible as the variation is typically small relative to the amplitude of the signal.

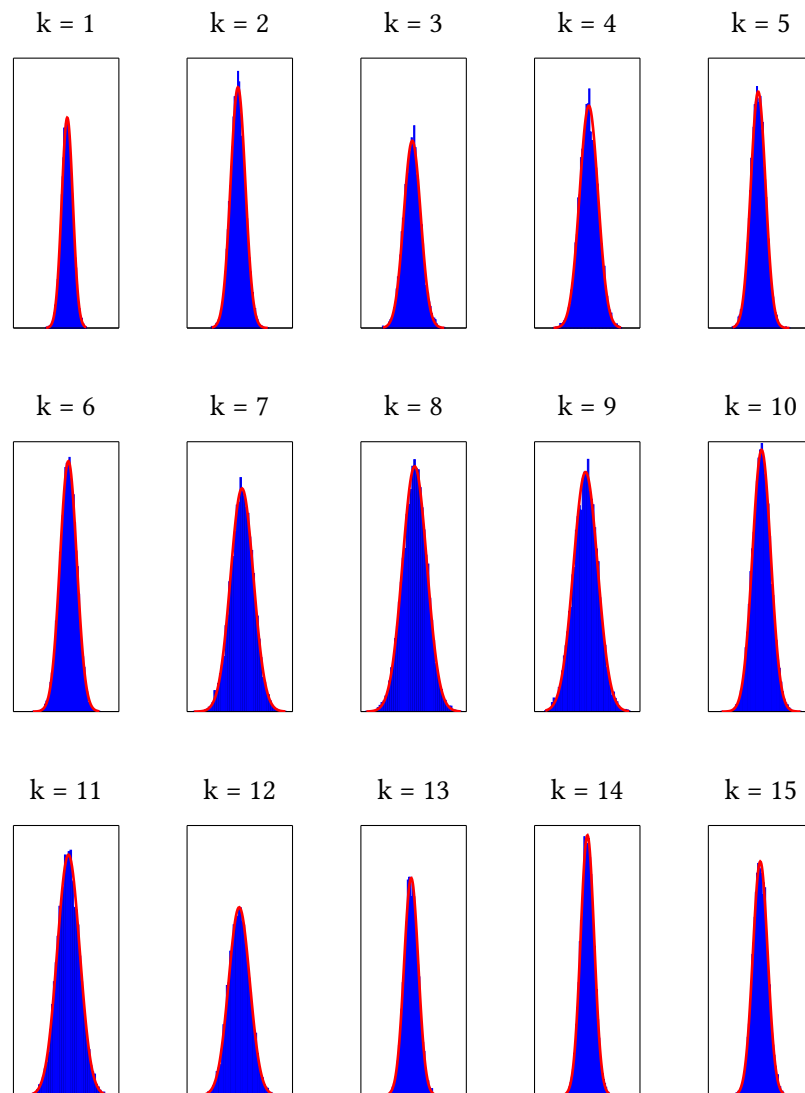


Figure 6.6: Normalised 500 bin histograms for each time sample between the two vertical lines in Figure 6.1 based on a 5000 repetition bootstrap mean. The red lines show the parametric normal distribution based on the sample mean and variance

6.4 Estimating covariance of extracted parameters

Calculating the covariance of the extracted \tilde{n} or l is non-trivial. A naive method of performing this, would be to simply calculate multiple values of \tilde{n} and l based on several measurements. The issue with this approach is that \tilde{n} and l are extracted using the mean of the time domain measurements. This naive estimate of covariance is not representative of these extracted estimates of \tilde{n} and l . There are other issues as well, in particular, it will be more sensitive to outliers in the measurement (averaging in the time domain mitigates this) and that it will be considerably more computationally expensive.

A better method would be to estimate based on different averaged time scans, for instance this could be done by averaging batches of measurements. However, this is not generally required, as a bootstrap method can be used instead. A bootstrap [143] randomly draws (with repetitions) time scans from the original set of measurements to form a new collection of time scans of the same size, which is then averaged. This is then repeated to form multiple averages, from which \tilde{n} and l can be extracted and the covariance can be estimated. The assumption using this method is that the original set of time measurements is representative of the distribution of the measurement. If for instance a low number of measurements is used and there is an outlier, the bootstrap will over represent this outlier in the final statistics. This method is often accurate and requires little information of the underlying distribution. However, a large number (100+) of extracted \tilde{n} values is usually required (which is computationally intensive).

Another method, is to use a Monte-Carlo [83] simulation to generate representative values of \tilde{n} and l . In this case a multivariate normal distribution is assumed and data is simulated with time domain covariance matrices equal to the measured estimates of uncertainty (of the time domain mean). If the underlying distribution is normal, this will be accurate. However, again a large number of extracted \tilde{n} values is usually required.

Instead, uncertainty propagation is used in this work and the above methods are used as a comparison. Uncertainty propagation forms an approximation for

each step of the processing, such that given an estimate of the covariance before the processing step, an estimate of the uncertainty at the output of the step can be made.

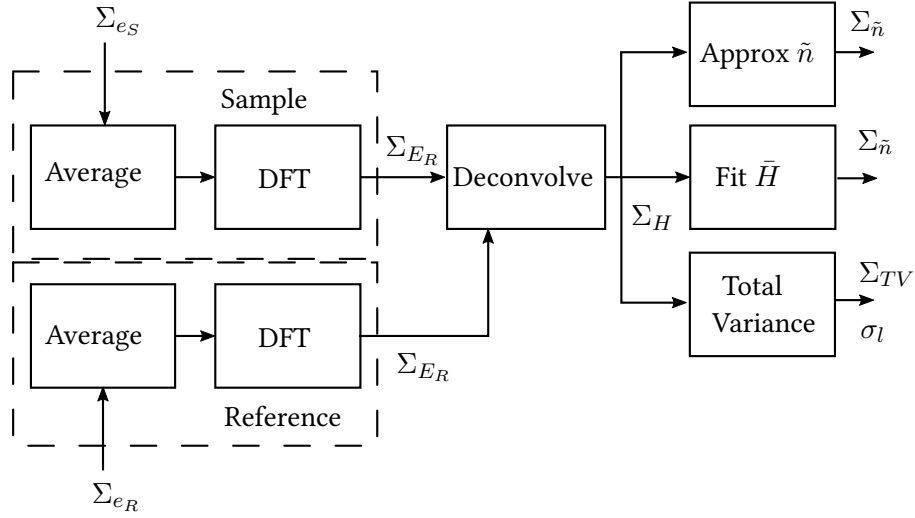


Figure 6.7: Uncertainty processing flowchart. The reference and sample measurements are used to estimate covariance in the time domain, this is then propagated to the extracted parameters using sensitivity matrices for each operation in the flowchart.

This is then used to form a similar processing algorithm to that used to extract \tilde{n} and l . In Figure 6.7, a flowchart of this algorithm is shown to estimate the complex covariance, Σ , and co-covariance (not shown for brevity), Γ , of \tilde{n} and l . For both reference and sample measurements, the covariance matrices of their spectrums, Σ_E , are calculated by performing uncertainty propagation for both averaging and the DFT. An estimate for the covariance matrices of H , Σ_H , is then estimated and then used to estimate the covariances of \tilde{n} , $\Sigma_{\tilde{n}}$, and l , Σ_l . The covariance of \tilde{n} must be estimated in different ways dependent on the method used to extract \tilde{n} . To estimate the covariance of l , the covariance of the total variance is first estimated directly from the transfer function.

A complex representation of variables is used during most of the processing, so a complex multivariate normal distribution should be used for the uncertainty in these variables. Complex variables are essentially two dimensional, thus to measure covariance twice as much information is required. In this work, complementary pairs of complex covariance and co-covariance matrices are used, as a complex variable can be treated as a singular variable during analysis. An overview of complex multivariate normal distributions is found in Appendix E.

Uncertainty propagation is performed for each step by multiplying the covariance of the input covariance by the sensitivity matrix [83] of the operation used. For an arbitrary processing step, variable \mathbf{y} which is function \mathcal{F} of variable \mathbf{x} :

$$\mathbf{y} = \mathcal{F}(\mathbf{x}) \quad (6.14)$$

Then the covariance matrix of \mathbf{y} is:

$$\Sigma_{\mathbf{y}} = \mathbf{C}_{\mathcal{F}} \Sigma_{\mathbf{x}} \mathbf{C}_{\mathcal{F}}^{\dagger} \quad (6.15)$$

where $\Sigma_{\mathbf{y}}$ is the complex covariance matrix of \mathbf{y} , $\Sigma_{\mathbf{x}}$ is the complex covariance matrix of \mathbf{x} , $\mathbf{C}_{\mathcal{F}}$ is the sensitivity matrix of operation \mathcal{F} and \dagger is the conjugate transpose ($\dagger = *'$). The co-covariance has a similar relation:

$$\Gamma_{\mathbf{y}} = \mathbf{A} \Gamma_{\mathbf{x}} \mathbf{A}' \quad (6.16)$$

where $\Gamma_{\mathbf{y}}$ is the complex covariance matrix of \mathbf{y} and $\Gamma_{\mathbf{x}}$ is the complex covariance matrix of \mathbf{x} .

In this work, generally only the sensitivity matrix (\mathbf{C}) is presented as the above relations are used to propagate the covariance and co-covariance matrices. The sensitivity matrix is generally the complex Jacobian of the processing step and can be derived analytically by performing differentiation with respect to the input variable.

From these covariance matrices, confidence intervals are then formed. These represent an interval for a parameter which will cover the deterministic value with some specified likelihood (95 % is used in this work).

6.4.1 Sensitivity of the arithmetic mean

The first operation is to calculate the arithmetic mean of multiple scans, see § 3.4, which will tend towards the underlying deterministic variable. As this is a function of multiple independent normal variables, it will have a normal distribution.

The equation for the arithmetic mean of multiple measurements is given by:

$$\mu(\mathbf{e}) = \frac{1}{M} \sum_{k=0}^{M-1} \mathbf{e}_k \quad (6.17)$$

the sensitivity matrix for of the average, $\mathbf{C}_{\mu, \mathbf{e}_k}$, is the scalar:

$$\mathbf{C}_{\mu, \mathbf{e}_k} = \frac{1}{M} \quad (6.18)$$

Given that each \mathbf{e}_k is identically and independently distributed, the covariance of the mean, $\Sigma_{\bar{\mathbf{e}}}$, can then be written as:

$$\Sigma_{\bar{\mathbf{e}}} = \sum_{k=0}^{M-1} \frac{1}{M^2} \Sigma_{\mathbf{e}} = \frac{1}{M} \Sigma_{\mathbf{e}} \quad (6.19)$$

6.4.2 Sensitivity of windowing

The time domain data is windowed before being converted to the Fourier frequency domain, see § 3.4. This operation can be written in matrix form as:

$$\mathbf{e}_w = \mathbb{D}(\mathbf{w}) \mathbf{e} \quad (6.20)$$

where \mathbf{e}_w is the windowed time domain data, $\mathbb{D}(\mathbf{x})$ indicates the diagonal matrix formed from vector \mathbf{x} and \mathbf{w} is the vector containing the windowing function. The sensitivity matrix of the windowing, \mathbf{C}_w , is then:

$$\mathbf{C}_w = \mathbb{D}(\mathbf{w}) \quad (6.21)$$

This operation can also be used to truncate the data, it should be noted that this truncated data is still normally distributed (as it is a marginal distribution, see Appendix E).

6.4.3 Uncertainty of the DFT

The majority of the remaining operations take place in the complex frequency (Fourier) domain. The DFT is used to convert the real time domain data to com-

plex frequency data. The DFT can be represented as a matrix operation [93]:

$$\tilde{\mathbf{E}} = \mathcal{F} \mathbf{e} \quad (6.22)$$

Where \mathcal{F} is the DFT matrix and $\tilde{\mathbf{E}}$ is the full spectrum frequency domain equivalent of \mathbf{e} , and thus the sensitivity matrix of the Fourier transform, $\mathbf{C}_{\tilde{\mathbf{E}},\mathbf{e}}$, is:

$$\mathbf{C}_{\tilde{\mathbf{E}},\mathbf{e}} = \mathcal{F} \quad (6.23)$$

While this could be used to calculate both the covariance and co-covariance, as for real data $\Sigma = \Gamma$, this is not required. $\tilde{\mathbf{E}}$ is the full frequency spectrum (including frequencies up to sampling frequency, F_s) and will have the same conjugate symmetry as $\tilde{\mathbf{z}}$ in the Appendix E. Thus $\Sigma_{\tilde{\mathbf{E}}}$ is the Fourier covariance matrix discussed in appendix E, and will have the form:

$$\Sigma_{\tilde{\mathbf{E}}} = \mathbf{C}_{\tilde{\mathbf{E}},\mathbf{e}} \Sigma_{\mathbf{e}} \mathbf{C}_{\tilde{\mathbf{E}},\mathbf{e}}^{\dagger} = \begin{bmatrix} \Sigma_{\mathbf{E}} & \overset{\leftrightarrow}{\Gamma}_{\mathbf{E}} \\ \underset{\downarrow}{\Gamma}_{\mathbf{E}}^* & \underset{\downarrow}{\Sigma}_{\mathbf{E}}^* \end{bmatrix} \quad (6.24)$$

where \mathbf{E} is the half spectrum (up to half F_s), and is what is generally used to perform analysis. The Fourier covariance matrix is thus fairly trivial to convert to $\Sigma_{\mathbf{E}}$ and $\Gamma_{\mathbf{E}}$ (and happens to be computationally efficient).

While it is possible to calculate $\Sigma_{\tilde{\mathbf{E}}}$ with the DFT matrix, it is generally more efficient to use a Fast Fourier Transform implementation (FFT):

$$\Sigma_{\tilde{\mathbf{E}}} = \text{FFT} \left(\text{FFT} (\Sigma_{\mathbf{e}})^{\dagger} \right)^{\dagger} \quad (6.25)$$

In this work a sub-set of \mathbf{E} is used for analysis (determined by α_{max} of a measurement), and will have a marginal distribution with the relevant sub matrices, so that:

$$\Sigma'_{\mathbf{E},j,k} = \Sigma_{\mathbf{E},j,k} \quad \text{for} \quad k_{min} \leq j \leq k_{max}, k_{min} \leq k \leq k_{max} \quad (6.26)$$

where $\Sigma'_{\mathbf{E},j,k}$ is the covariance of the frequencies used, k_{min} is the index which

corresponds to the lowest frequency used and k_{max} is the index which corresponds to the highest frequency used.

Just like the DFT, the inverse discrete Fourier transform can be written in matrix form, so that the sensitivity matrix is:

$$\mathbf{C}_{e, \tilde{\mathbf{E}}} = \mathcal{F}^{-1} \quad (6.27)$$

where \mathcal{F}^{-1} is the inverse DFT matrix. This acts on the Fourier covariance, $\Sigma_{\tilde{\mathbf{E}}}$.

6.4.4 Sensitivity of the complex logarithm, magnitude and argument

The magnitude and argument are not complex differentiable functions as they are inherently functions of the underlying real and imaginary parts. To compensate for this, the augmented covariance matrix, see appendix E, can be used in the uncertainty propagation.

The complex logarithm is used in the next section and is closely related to the magnitude and phase of a complex variable. The complex logarithm is complex differentiable and related to the magnitude and argument of complex variable, z :

$$\log_e(z) = \log_e(|z|) + i\angle z \quad (6.28)$$

The sensitivity matrix for the complex logarithm, \mathbf{C}_{\log_e} , is its derivative:

$$\mathbf{C}_{\log_e(z)} = \mathbb{D} \left(\frac{1}{z} \right) = \mathbb{D} \left(\frac{z^*}{|z|^2} \right) \quad (6.29)$$

While it is possible to consider the magnitude and phase as independent, generally this will not be true. It is therefore preferable to consider the jointly distributed multivariate variable of the two:

$$\vec{z} = \begin{bmatrix} |z| \\ \angle z \end{bmatrix} \quad (6.30)$$

This can be related to $\log_e(z)$ by:

$$\vec{z} = \frac{1}{2} \begin{bmatrix} \mathbb{D}(|z|) & \mathbb{D}(|z|) \\ -iI & iI \end{bmatrix} \begin{bmatrix} \log_e(z) \\ \log_e(z)^* \end{bmatrix} \quad (6.31)$$

And thus the sensitivity matrix, $\mathbf{C}_{\vec{z}, \dot{z}}$, is with respect to \dot{z} (see Appendix E):

$$\mathbf{C}_{\vec{z}, \dot{z}} = \begin{bmatrix} \mathbb{D}\left(\frac{z}{2|z|}\right) & \mathbb{D}\left(\frac{z^*}{2|z|}\right) \\ \mathbb{D}\left(\frac{-iz}{2|z|^2}\right) & \mathbb{D}\left(\frac{-iz^*}{2|z|^2}\right) \end{bmatrix} \quad (6.32)$$

6.4.5 Sensitivity of deconvolution

In the frequency domain, the complex transfer can be calculated by division:

$$\mathbf{H} = \frac{\mathbf{E}_{sam}}{\mathbf{E}_{ref}} \quad (6.33)$$

Where \mathbf{H} , \mathbf{E}_{sam} and \mathbf{E}_{ref} are the transfer function, frequency domain sample measurement and frequency domain reference measurement respectively.

The sample and reference measurements are assumed independent, so their contribution to covariance of \mathbf{H} can be formulated as a summation:

$$\Sigma_{\mathbf{H}} = \mathbf{C}_{\mathbf{H}, \mathbf{E}_{sam}} \Sigma_{\mathbf{E}_{sam}} \mathbf{C}_{\mathbf{H}, \mathbf{E}_{sam}}^\dagger + \mathbf{C}_{\mathbf{H}, \mathbf{E}_{ref}} \Sigma_{\mathbf{E}_{ref}} \mathbf{C}_{\mathbf{H}, \mathbf{E}_{ref}}^\dagger \quad (6.34)$$

Where $\mathbf{C}_{\mathbf{H}, \mathbf{E}_{sam}}$, $\mathbf{C}_{\mathbf{H}, \mathbf{E}_{ref}}$ are the sensitivity matrices for the sample and reference measurements respectively. A similar relation can be derived for the complementary covariance matrices, by simply replacing the conjugate transpose with the transpose.

Deriving the sensitivity matrices can be done using the product and quotient rules:

$$\mathbf{C}_{\mathbf{H}, \mathbf{E}_{sam}} = \mathbb{D}\left(\frac{1}{\mathbf{E}_{ref}}\right) \quad (6.35)$$

$$\mathbf{C}_{\mathbf{H}, \mathbf{E}_{ref}} = \mathbb{D}\left(\frac{-\mathbf{E}_{sam}}{\mathbf{E}_{ref}^2}\right) \quad (6.36)$$

When the sensitivity matrices are calculated based on measurement dependent values, the average is used.

6.4.6 Sensitivity of complex refractive index approximations

In § 3.6, the n and κ of a sample is approximated using the magnitude and phase of the transfer function. The derivations for the sensitivity matrix in which case can be found in [81]. It is however preferable to consider the \tilde{n} instead, as it can be related to the complex logarithm.

Using equation 3.47, \tilde{n} can be approximated as:

$$\tilde{n} = i \frac{c}{\omega l} \left(\log_e(\mathbf{H}) + \log_e \left(\frac{(\mathbf{n} + 1)^2}{4\mathbf{n}} \right) \right) + 1 \quad (6.37)$$

where \tilde{n} is the complex refractive index, ω is a vector of angular frequencies, l is the thickness and \mathbf{n} is the real refractive index vector. All operations are performed element by element with the respective vectors.

If the real refractive index is suitably low ($n \lesssim 3$), the transmission terms can be approximated as being negligible, thus this can be further simplified:

$$\tilde{n} = i \frac{c}{\omega l} \log_e(\mathbf{H}) + 1 \quad (6.38)$$

Which is fully complex differentiable, as it wholly a function of the complex logarithm. For high index ($n > 3$) materials the method given in § 6.8 can be applied with a non-resonant model to achieve an accurate estimation.

As this is a function of the complex logarithm, the sensitivity matrix of it, $\mathbf{C}_{\tilde{n}, \mathbf{H}}$, can be calculated using the sensitivity matrix of the logarithm, $\mathbf{C}_{\log_e(\mathbf{H}), \mathbf{H}}$:

$$\mathbf{C}_{\tilde{n}, \mathbf{H}} = \mathbb{D} \left(\frac{ic}{\omega l} \right) \mathbf{C}_{\log_e(\mathbf{H}), \mathbf{H}} = \mathbb{D} \left(\frac{ic}{\omega l} \frac{1}{\mathbf{H}} \right) \quad (6.39)$$

The covariance of n and κ can then be calculated using the following relations:

$$\begin{aligned} \Sigma_n &= \frac{1}{2} \Re(\Sigma_{\tilde{n}} + \Gamma_{\tilde{n}}) \\ \Sigma_{n, \kappa} &= \frac{1}{2} \Im(-\Gamma_{\tilde{n}} + \Sigma_{\tilde{n}}) \\ \Sigma_\kappa &= \frac{1}{2} \Re(\Sigma_{\tilde{n}} - \Gamma_{\tilde{n}}) \end{aligned} \quad (6.40)$$

6.4.7 Sensitivity of spectroscopic parameter conversions

The α is a linear relation of a sample, as such the sensitivity matrix, $C_{\alpha,\kappa}$, between it and the extinction coefficient is:

$$C_{\alpha,\kappa} = \mathbb{D} \left(\frac{2\omega}{c} \right) \quad (6.41)$$

The covariance between n and α can be calculated using the following:

$$\Sigma_{n,\alpha} = \Sigma_{n,\kappa} C_{\alpha,\kappa}' \quad (6.42)$$

The complex permittivity, $\tilde{\varepsilon}$, is simply the square of the \tilde{n} , so the sensitivity matrix between the two is:

$$C_{\tilde{\varepsilon},\tilde{n}} = 2\tilde{n} \quad (6.43)$$

This can then be converted into real and imaginary permittivity using the same form as equation 6.40.

6.5 Uncertainty of α lactose monohydrate measurements

A 10 % α -lactose monohydrate sample was measured in the LFRBB THz TDS, see § 3.8. The sample's n and α were then approximated, see § 3.6. The measurement uncertainty was then estimated in the time domain, and corresponding estimates of the uncertainty of n and α . These estimates were then used to find confidence intervals and a rudimentary analysis of the uncertainty was performed by analysing the correlation and variance.

From the time domain covariance, an estimate of the n and α covariances was formed by sequentially applying the sensitivity matrices of each operation to the covariance and relation matrices. From the final covariance matrices, the variance of both n and α was found. The covariance and mean estimates of both sample and reference measurements were formed from 10 time scans each.

To provide a comparison, three other estimates of the covariance were formed using numerical methods. The first method, referred to as the "naive" method, is to

treat the sample and reference measurements as pairs, and form multiple estimates of n and α . These are then treated as a statistical sample, from which the covariance can be formed. The second method is the bootstrap method [143]. From this multiple estimates of the mean (in the time domain) are formed, and then multiple estimates of n and α can be formed. These are then used to estimate the variance. The final method used for comparison is the "Monte-Carlo" method [83]. This method simulates multiple estimates of the mean time domain scan with covariance and mean equal to the sample estimates of the mean and covariance. From these estimates, multiple estimates of n and α , and their corresponding statistics, were formed.




Figure 6.8: Different estimates of the relative uncertainty of n of a 10% Lactose monohydrate sample measurement. Green line - the naive method, light blue - the bootstrap method, red - the Monte-Carlo method and dashed blue - the propagation of uncertainty method.

In Figure 6.8, the different estimates of the relative standard deviation of n are shown. The green line shows the naive methods estimate, the light blue line shows the bootstrap estimate, the red line shows the Monte-Carlo estimate and the dashed blue shows the estimate propagated from the time domain. These estimates are all fairly similar below 5 THz, and show similar values and shape. Both the Monte-Carlo and Bootstrap methods have been formed using 2000 different estimates. The naive estimate consistently overestimates the variance slightly compared to the other methods, while the other methods track each other constantly. Interestingly, the propagation and Monte-Carlo methods are near identical (below 5 THz), this is to be expected as both are based on the assumption of an underlying normal distribution. The bootstrap does not assume any particular probability distribution,

and instead relies on observed frequency in the sample set. That the bootstrap agree extremely well with the methods based on assumed normality, indicates that this assumption is valid in this case. Above 5 THz, the Monte–Carlo estimate appears to step. This occurs because of unwrapping changes in the different estimates of n . The unwrapping corrections from Chapter 3, are formed based on the measured data (regardless of method) and will therefore be subject to uncertainty as well, which will propagate to n . The uncertainty in the unwrapping correction has not been considered in this work (but is inherently present), as below α_{max} (in α) it tends to be negligible. At 6 THz, the bootstrap also shows this behaviour as well. The reason it does not show unwrapping uncertainty below this frequency is that this uncertainty does not occur within the original sample set. As the Monte–Carlo method models the uncertainty, it can model effects not present within the original data.




Figure 6.9: Different estimates of the relative uncertainty of α of a 10% lactose monohydrate measurement. Green line - naive estimate, light blue - bootstrap estimate, red - Monte–Carlo estimate and dashed blue - propagation of uncertainty estimate

In Figure 6.9, the corresponding different estimates of the relative standard deviation of α are shown. The green line shows the naive estimate, the light blue shows the bootstrap estimate, the red line shows the Monte–Carlo estimate and the dashed blue line shows the uncertainty propagation estimate. The naive estimate is the outlier in this case, while the other three estimates agree extremely well.

In this particular sample measurement, the naive method has appeared to do fairly well compared to the other methods. This is because the measurements in




Figure 6.10: Different estimates of the variance of α of a z-cut quartz measurement. Green line - naive estimate, light blue - bootstrap estimate, red - Monte-Carlo estimate and dashed blue - propagation of uncertainty estimate

this case study did not contain outliers and were extremely low variance. This is not guaranteed to be the case, and sometimes it is not possible to acquire a low noise data set (for instance when measuring a sample in a non-stable state). In Figure 6.10, different estimates of variance for a 0.512 mm z-cut quartz sample's α is shown. Variance was used instead of relative uncertainty, as the relative uncertainty was found to be sensitive to catastrophic collision ($\frac{0}{0}$) and produced extreme values. This sample was measured in the HFRBB THz TDS system, with 20 repetitions of both reference and sample measurements. It was found that one of the reference measurements was an outlier in the frequency domain at higher frequencies and could skew the results. The green (naive) estimate of variance is significantly offset from both the Monte-Carlo (red) and uncertainty propagation methods (dashed blue) estimates. This because it was skewed by the outlier in a way that these estimates were not. Interestingly the bootstrap has also been somewhat skewed by this outlier. The bootstrap is based on the frequency in the sample set, outliers which occur in the original sample set can skew this estimate [143]. This effect will rapidly reduce in larger sample sets, so the issue is rather that the sample size is too small rather than it has outliers in.

In Figure 6.11, the correlation matrices for the 10 % α -lactose monohydrate sample estimated using each method is shown. The top row shows uncertainty propagation estimates, the second shows Monte-Carlo estimates and the third shows

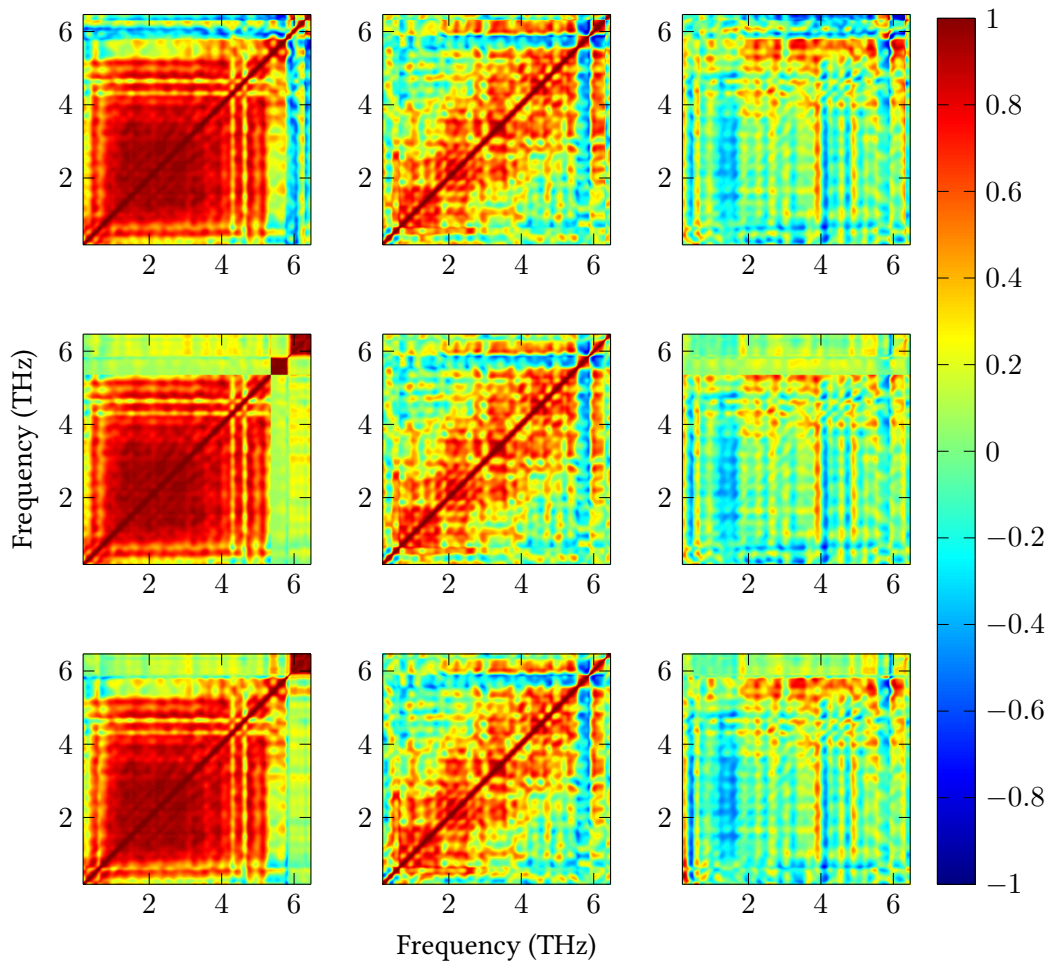


Figure 6.11: Different Estimates of correlation of 10% α lactose monohydrate. Top row - Uncertainty propagation estimates, middle row - Monte-Carlo estimates, bottom row - bootstrap estimates. First column - correlation within n , second column - correlation with α and third column - correlation between n and α

bootstrap estimates. The first column shows the correlation within the n variation, the second shows the correlation with the α variation and the third column shows the correlation between the two. All of the estimates are near identical, with the exception being in the real refractive index correlation. Unwrapping was included in the Monte-Carlo and bootstrap estimates, which has dominated at the higher frequencies making the data appear less correlated than it is in the uncertainty propagation estimate.

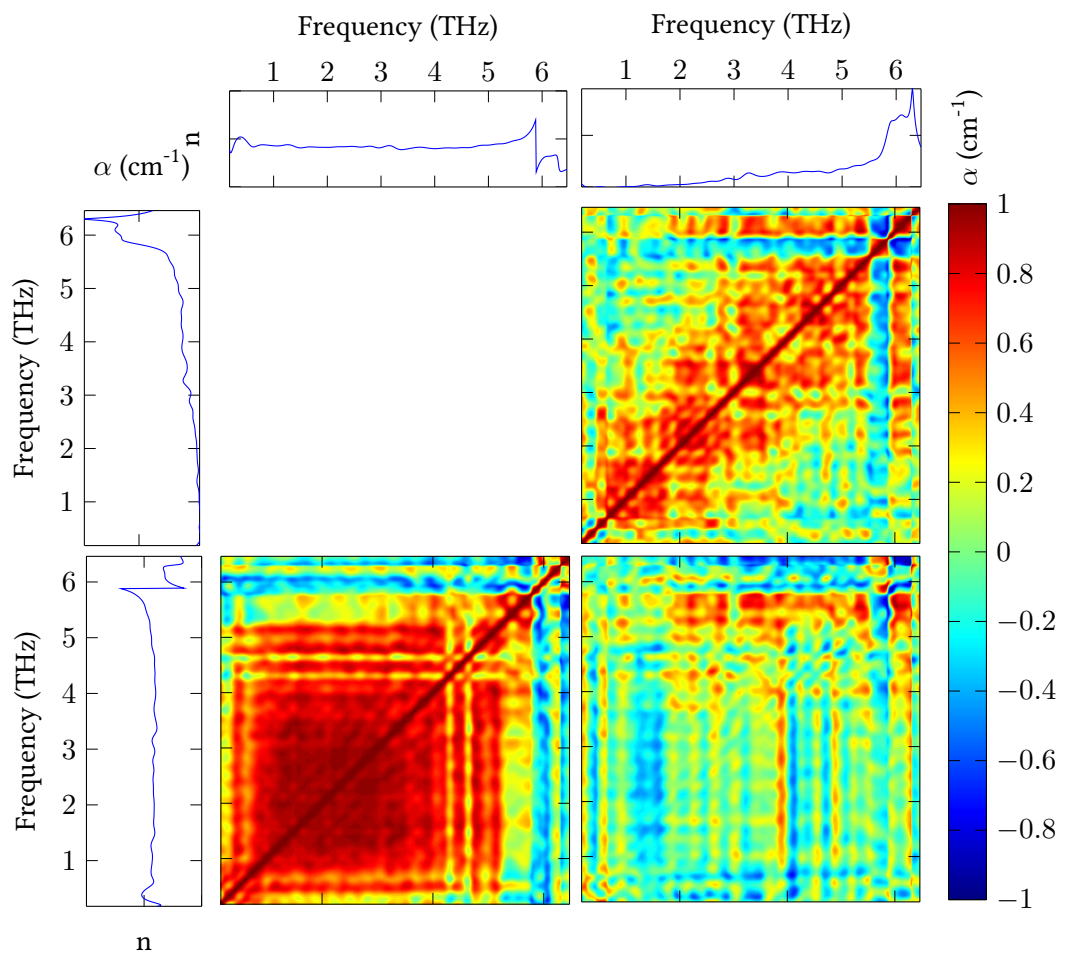


Figure 6.12: Uncertainty propagation estimates of correlation of 10% α lactose monohydrate. Bottom left - n correlation matrix, Bottom right - correlation matrix between n and α , Top right - correlation matrix of α . n and α are shown in the top and left panes for reference.

In Figure 6.12, the uncertainty propagation estimates of correlation are shown. Only the uncertainty propagation estimates are shown in detail, as other estimates produced extremely similar estimates. The correlation of n and α are shown in the bottom left and top right corners respectively. In the bottom right, the correlation

between the two variables are shown. As n and κ (which α is related to) are parts of the same complex variable, they are interdependent. Given this interdependence, it is not unexpected that the variation between the two is correlated, and in this work they are treated as being jointly normal, just as in the case of other complex variables.

The correlation in each variable appears to be broadly frequency independent, i.e. the variation at one frequency seems to be broadly correlated with the variation at every other frequency. This is likely due to the real refractive index offset being uncertain (due to a pulse arrival time uncertainty), which will also result in there being a broad frequency change in α (due to the Kramers–Kronig relations between n and κ [84]).

In both correlation matrices, there exists a strong diagonal band. The main diagonal of a correlation matrix is always one, however in this case it has been expanded to nearby frequencies. This is because of the use of zero padding in the measurement, which in the frequency domain equates to using interpolation. In this case, since several neighbouring frequencies are essentially functions of the same underlying information they must be correlated. Within the n of the sample, there is a large block of positive correlation across a broad range of frequencies. This is a result of the pulse shifting in the measured time window, see § 6.3. There is slightly negative correlation (light blue) at 6 THz. This is due to strong peak at this frequency from the PTFE matrix used in the sample, which will have considerable uncertainty due to the limits of dynamic range at these frequencies. There is also a double band of lower correlation at 4.3 and 4.6 THz, which corresponds to the relaxations at this frequencies. This is a region of negative correlation which is being dominated by the background positive correlation (from the pulse shift in time) and occurs because there is uncertainty in the relaxation positions at these frequencies (it should be noted that the uncertainty is small).

The correlation of α is more localised, which suggests that some variation is due to the peak uncertainty rather than a group shift across frequencies like refractive index. In particular, there are small regions of strong correlation along the diagonal which are associated with singular features, particularly at the 3.26 THz peak, which

indicate localised peak amplitude uncertainty (i.e. the entire peak is changing in unison). However, like n , there is a banding of negative correlation from the edge of the 6 THz peak in PTFE. This suggests that there might be uncertainty in the position of the rising edge (this could be peak position uncertainty or peak width uncertainty due to the limited dynamic range).

The correlation between n and α appears to be much more neutral than the correlation within each of the respective variables. This is partly due to the lack of a dominant diagonal band (these are different variables, so are not self-correlated). Instead there is much more localised peaks in correlation, which correspond to the relaxations and peaks in n and α , which would be expected due to the underlying polarisation mechanisms, see § 2.4.4. Interestingly, the n between 5 and 6 THz appears to correlated broadly across frequencies in α . This might indicate that the variation at these frequencies in n , is driving a more broadband variation in α .

6.6 Estimating confidence intervals

A confidence interval represents a range of values which the deterministic parameter lies within, with some predetermined probability [139]. It should be clear that it is the confidence interval which is random, the parameter which is being estimated is deterministic. Thus performing new measurements will lead to different estimates of the confidence intervals, which will cover the same parameter with the same probability. In the multivariate case, there exists a multivariate confidence region [83]. This confidence region will consist of a N dimensional hyper-ellipsoid which covers the mean vector with some probability. This methodology is not used in this work, as it requires a larger number of observations than dimensions [83] (i.e. to form a confidence region for a 100 dimensional n and α vectors, at least 201 measurements would have to be performed). It would also be difficult to create visualisations of the hyper-ellipsoid which would be analytically useful. Instead, univariate confidence intervals are proposed, these ignore correlation, and solely cover the component with some probability.

6.6.1 Non-simultaneous confidence intervals

A scalar component, x_i , of the multivariate variable, \mathbf{x} , will have univariate distribution. It is then possible to form N confidence intervals based on univariate statistics. These intervals will not consider correlation between different scalar components, but individually will have probability p likelihood of covering the individual component.

To form a confidence interval, consider a measurement of the form:

$$\hat{\mathbf{x}} = \mathbf{x} + \mathbf{y} \quad (6.44)$$

Where \mathbf{y} is a zero mean random variable and \mathbf{x} is a deterministic variable. The distribution for the k th component is distributed normally:

$$\hat{x}_k \sim \mathcal{N}_1(x_k, \sigma_{y_i}^2) \quad (6.45)$$

And for probability p there exists a value a for which the following is true:

$$P(x_k - a \leq \hat{x}_k \leq x_k + a) = p \quad (6.46)$$

Conversely it is possible to specify an interval which will contain x_k based on an observed \hat{x}_k with probability p :

$$P(\hat{x}_k - a \leq x_k \leq \hat{x}_k + a) = p \quad (6.47)$$

This can be seen by considering the case where $x_k - a \leq \hat{x}_k \leq x_k + a$, in which case x_k always lies within a of \hat{x}_k . If \hat{x}_k is not within this range, this is not true. Since \hat{x}_k has probability p of being in this range, the statement in equation 6.47 has the same probability.

For a given probability p , it is possible to determine a value a which satisfies equation 6.47. For a known value of σ_{y_k} , this can be done by using a standard normal

variable, $s \sim \mathcal{N}(0, 1)$, so that $\hat{x}_k = \sigma_{y_k} s + x_k$. There is a constant b for variable s :

$$P(-b \leq s \leq b) = 1 - c = p \quad (6.48)$$

b is then the value equal to the inverse of the normal cumulative distribution function (the integral of the PDF) of s for probability $1 - \frac{c}{2}$:

$$b = \mathcal{N}_s^{-1}\left(1 - \frac{c}{2}\right) \quad (6.49)$$

a can then be calculated by following through to \hat{x}_k :

$$a = b\sigma_{y_k} \quad (6.50)$$

In the case where σ_{y_k} is estimated from the measured data, b must instead be calculated using an inverse of the t cumulative distribution:

$$b = t_v^{-1}\left(1 - \frac{c}{2}\right) \quad (6.51)$$

where v is the degree of freedom in the measurement, and for M measurements is typically $M - 1$. In the case of combinations of variance, an effective degree of freedom should be calculated (see § 6.6.3).

For the multivariate variable $\hat{\boldsymbol{x}}$, there exists the vectors $\hat{\boldsymbol{x}} \pm \boldsymbol{a}$ exist which satisfy the following [139]:

$$P(\hat{x}_k - a_k \leq x_k \leq \hat{x}_k + a_k) = p \quad (6.52)$$

This is only true for individual components, it does not mean that all components of the multivariate intervals cover all components of \boldsymbol{x} simultaneously. It is for that reason, that in this work they are referred to as non-simultaneous confidence intervals.

6.6.2 Bonferroni simultaneous confidence intervals

In the previous sub-section, non-simultaneous intervals were considered. The issue with these intervals is that they only consider the probability of individual compo-

nents of the confidence interval. This will lead to a case where the probability that all the components of the deterministic value lie within the confidence intervals is overestimated (i.e. the confidence interval is too narrow). This can be seen by considering a variable with N components, which are independently and identically distributed, in which case [139]:

$$P(\hat{\boldsymbol{x}} - \boldsymbol{a} \leq \boldsymbol{x} \leq \hat{\boldsymbol{x}} + \boldsymbol{a}) = p^N \quad (6.53)$$

The probabilities multiply, in which case for larger values of N it could be quite likely that individual components of the interval do not cover \boldsymbol{x} .

Instead it is possible to construct a Bonferroni interval which has at least probability p , and specifies that for all the k components of \boldsymbol{x} are within a confidence interval:

$$P(\hat{\boldsymbol{x}} - \boldsymbol{a} \leq \boldsymbol{x} \leq \hat{\boldsymbol{x}} + \boldsymbol{a}) \geq p \quad (6.54)$$

Where a is calculated from a different value of b [139]:

$$b = t_v^{-1} \left(1 - \frac{c}{2A} \right) \quad (6.55)$$

where A is the number of components in the variable \boldsymbol{x} . In the case of n and α , which are essentially the real and imaginary parts of \tilde{n} , this will be $2N$. Both n and α are jointly normal, and can be considered sub-variables of the combination of the two. Thus $2N$ is used so that the intervals have a probability p of covering the deterministic values of both n and α simultaneously.

6.6.3 Effective degrees of freedom

When considering combinations of different normal variables (or functions of combinations of variables) the degrees of freedom is not equal to $M - 1$. Instead the effective degree of freedom needs to be calculated from the underlying variances σ_j^2 ,

degrees of freedom v_j and combination of variances [137]:

$$v_e = \frac{\left(\sum_{j=0}^{L-1} \sigma_j^2\right)^2}{\sum_{j=0}^{L-1} \frac{\sigma_j^4}{v_j}} \quad (6.56)$$

where L is the number of difference sources of variance.

For example, to calculate the effective degrees of freedom for the k th confidence interval of n :

$$v_{e,n_k} = \frac{\sigma_{n_k}^4}{\frac{\sigma_{s,n_k}^4}{M_s-1} + \frac{\sigma_{r,n_k}^4}{M_r-1}} \quad (6.57)$$

where $\sigma_{n_k}^2$ is the k th component of the variance of n , M_s is the number of sample measurements, M_r is the number of reference measurements, σ_{s,n_k}^2 is the sample contribution to $\sigma_{n_k}^2$ and σ_{r,n_k}^2 is the reference contribution. σ_{s,n_k}^2 can be calculated by assuming the variance in the reference measurement is 0, and calculating $\sigma_{n_k}^2$ based solely on the sample variance. The same can be done for σ_{r,n_k}^2 by assuming the variance in the sample measurement is 0. The effective degrees of freedom will need to be estimated for each component of a confidence interval, if the confidence interval is for a parameter which is a function of both reference and sample measurements.

6.7 Confidence intervals of α lactose monohydrate measurements

In Figure 6.13 the n of the 10 % α -lactose monohydrate sample from § 6.5 is shown with confidence intervals. The left pane shows 0.2–5.3 THz and the right pane shows 5.3–6.5 THz. Different scales were used as the uncertainty in relaxations above 5.3 THz tend to dominate the measurement. The green shaded band shows the 95% probability simultaneous confidence interval, and the red shaded band shows the corresponding non-simultaneous interval. Both are based on estimates of the variation made using uncertainty propagation, the simultaneous estimate will be a multiple of the non-simultaneous estimate. The confidence interval is narrow over most of the frequency range, with the exception at 6 THz. It should be noted that while there is some change in the shape of the relaxation features, this measurement is highly




Figure 6.13: The estimated 95% probability confidence intervals of n of a 10% α -lactose monohydrate sample. The left pane shows 0.2–5.3 THz and the right pane shows 5.3–6.5 THz. Different scales have been used to emphasise confidence intervals in these regions. The green shaded region shows the simultaneous interval and the red shaded region shows the non-simultaneous region. The estimate of n is shown as a black line.

correlated (i.e. the uncertainty is in the average refractive index). Below 0.3 THz, the measured phase tends to be extremely noisy and it is expected the confidence region would expand below these frequencies. At higher frequencies, above 5.5 THz, α has reached α_{max} . As such noise will dominate the measurement considerably, which leads to larger confidence bands. The confidence intervals expands around the relaxations at these frequencies, which indicates that there is larger uncertainty around these relaxations. It should be noted, that the unwrapping uncertainty has not been considered, and that the confidence interval will actually be considerably larger at these frequencies if it was considered.

In Figure 6.14 the α of the sample is shown with confidence intervals. The left pane shows 0.2–5.3 THz and the right pane shows 5.3–6.5 THz. Different scales have been used as the uncertainty in relaxations above 5.3 THz tend to dominate the measurement. The green shaded band shows the 95% probability simultaneous confidence interval and the red shaded band shows the corresponding non-simultaneous interval. The black line shows the estimated α and the red line shows α_{max} . The confidence intervals were based on uncertainty propagation of the variation. As would be expected, where α tends to α_{max} , the confidence intervals become extremely large, this is because noise dominates the magnitude measurement at these frequencies. Interestingly, between 3.5 and 6 THz the confidence intervals expand




Figure 6.14: The estimated 95% probability confidence intervals of α of a 10 % α -lactose monohydrate sample. The left pane shows 0.2–5.3 THz and the right pane shows 5.3–6.5THz. Different scales have been used to emphasise confidence intervals in these regions. The green shaded region shows the simultaneous interval and the red shaded region shows the non-simultaneous region. The estimate of α is shown as a black line. The α_{max} of the measurement is shown as a cyan line.

significantly. This occurs because of the decreasing reference measurement amplitude with respect to frequency. If this is considered with regards to the deconvolution sensitivity matrix equations 6.35 and 6.36, it can be seen that the uncertainty in both measurements is inversely proportional to the magnitude of the reference measurement. Thus as it tends to be smaller, the uncertainty in the transfer function (and thus α) tends to be larger.

The dynamic range of the reference, see § 3.7, will therefore have considerable effect on the uncertainty of the extracted \tilde{n} . This could be important when considering system design, for instance designing emitters tuned to a particular range of frequencies might increase the certainty at these frequencies (by decreasing certainty at other frequencies). Another reason this might be relevant is when more unusual references are used, for instance when using self referencing [144] in reflection measurements or when using a similar sample (for instance an empty flowcell [76]).

6.8 Fitting a transfer function model

In Chapter 5 a sample model is fitted to the transfer function. Calculating the sensitivity matrix for such a formulation can be difficult. Instead the assumption is that

the error function has a global minimum at:

$$\bar{H} = H \quad (6.58)$$

Where \bar{H} is the transfer function model. The sensitivity matrix can then be formed as the inverse of the model's Jacobian:

$$C_{\tilde{n},H} = \Delta_{\bar{H},\tilde{n}}^{-1} \quad (6.59)$$

This is a simplification of a more robust analysis [82], which can more generally be applied to uncertainty when using indirect fitting [83].

6.8.1 Sensitivity matrices of sample models

To find the sensitivity matrix of the model, which will be the inverse of its Jacobian with respect to \tilde{n} :

$$C_{\tilde{n},H} = \Delta_{\bar{H},\tilde{n}}^{-1} \quad (6.60)$$

In the case where a model is being fitted to a singular transfer function, at each frequency with a singular parameter (i.e. just a single layer sample's \tilde{n}), the Jacobian will be a diagonal matrix:

$$\Delta_{H,\tilde{n}} = \mathbb{D} \left(\frac{d\bar{H}}{d\tilde{n}} \right) \quad (6.61)$$

It will thus be the derivative of the model at each discrete frequency. Being a diagonal matrix, the inverse is an inversion of the diagonal at each frequency:

$$\Delta_{H,\tilde{n}}^{-1} = \mathbb{D} \left(\frac{d\bar{H}^{-1}}{d\tilde{n}} \right) \quad (6.62)$$

The derivative of the model needs to be calculated. One way of estimating the derivative is to use the finite difference, by making a small step in a parameter, calculating the model and then taking the difference (this is how gradient descent algorithms function without an analytical gradient). The difficulty with this approach, is that the gradient would have to be estimated using a small step size to be accurate for the uncertainty (which experimentally is extremely small), and thus is susceptible

to calculation error. A better approach is to derive the derivative analytically. As all the models used in this work are a product of coefficients, only the derivatives of the coefficients need to be found and the product rule can then be used to find the derivative of the model.

6.8.2 Derivative of the model coefficients

To find the derivative of a sample, the individual model coefficient derivatives are required. These coefficients are functions of the sample's \tilde{n} and can be differentiated with respect to the sample's \tilde{n} .

The propagation coefficient, given by equation 3.7, has the derivative :

$$\frac{dP(\tilde{n})}{d\tilde{n}} = -i\frac{\omega l}{c}P(\tilde{n}) \quad (6.63)$$

The transmission coefficient given by equation 3.6, is a function of two different layer's \tilde{n} , \tilde{n}_1 being the layer sourcing the radiation and \tilde{n}_2 being the layer receiving the radiation. Being a function of two different \tilde{n} 's, there are two different derivatives. The first being with respect to \tilde{n}_2 :

$$\frac{dT_{1,2}}{d\tilde{n}_2} = \frac{2\tilde{n}_1}{(\tilde{n}_1 + \tilde{n}_2)^2} \quad (6.64)$$

The second being with respect to \tilde{n}_1 :

$$\frac{dT_{1,2}}{d\tilde{n}_1} = \frac{-2\tilde{n}_2}{(\tilde{n}_1 + \tilde{n}_2)^2} \quad (6.65)$$

Similarly, there is the reflection coefficient, given by equation 5.15 between two layers, which will have two derivatives. The first being with respect to \tilde{n}_2 :

$$\frac{dR_{1,2}}{d\tilde{n}_2} = \frac{2\tilde{n}_1}{(\tilde{n}_1 + \tilde{n}_2)^2} \quad (6.66)$$

And the second with respect to \tilde{n}_1 :

$$\frac{dR_{1,2}}{d\tilde{n}_1} = \frac{-2\tilde{n}_2}{(\tilde{n}_1 + \tilde{n}_2)^2} \quad (6.67)$$

The Gouy phase correction coefficient is given by equation 5.39, and its derivative with respect to the sample's \tilde{n} is approximately:

$$\frac{dG(n)}{d\tilde{n}} \approx \frac{-i\omega l\beta}{cn^2}G(n) \quad (6.68)$$

Which works if the extinction coefficient is assumed negligible, so that $\tilde{n} \approx n$ in these equations.

Similarly, for the beam waist amplitude correction, see equation 5.37, the derivative is:

$$\frac{dw(n)}{d\tilde{n}} \approx \frac{-\omega^2\beta^2l^2}{c^2n^3\left(1 + \left(\frac{\omega\beta l}{cn}\right)^2\right)}w(n) \quad (6.69)$$

6.8.3 Modifying transfer function models to calculate the derivative

All the models used in this work consist of products of separate coefficients. These coefficients are functions of \tilde{n} and have been shown to be differentiable with respect to \tilde{n} . This means the product rule can be used to calculate the derivative, but also that almost all of the transfer function model itself must be calculated to calculate the derivative (i.e. if one coefficient is factored out of the derivative, the product of the other coefficients and the derivative of the coefficient will be left).

One method of calculating the derivative is to treat it as a joint calculation with the transfer function. Creating a vector, $\bar{\mathbf{H}}$, defined as the transfer function and its derivative:

$$\bar{\mathbf{H}} = \begin{bmatrix} \bar{H} \\ \frac{d\bar{H}}{d\tilde{n}} \end{bmatrix} \quad (6.70)$$

Factoring out a coefficient A , the transfer function can always be written as:

$$\bar{\mathbf{H}} = \begin{bmatrix} A\bar{H}' \\ \frac{dA}{d\tilde{n}}\bar{H}' + A\frac{d\bar{H}'}{d\tilde{n}} \end{bmatrix} \quad (6.71)$$

where $\bar{H}' = \frac{\bar{H}}{A}$. This can be written in matrix form:

$$\bar{H} = \begin{bmatrix} A & 0 \\ \frac{dA}{d\tilde{n}} & A \end{bmatrix} \bar{H}' = \mathbf{A} \bar{H}' \quad (6.72)$$

\mathbf{A} is the coefficient matrix. This can be applied to any model coefficient, and can be applied sequentially. For example, the simplest model can be written as:

$$\bar{H} = \mathbf{T}_{a,s} \mathbf{T}_{s,a} \mathbf{P}_s \mathbf{P}_a^{-1} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad (6.73)$$

Where each of the coefficients is in matrix form with its derivative with respect to the sample's \tilde{n} . It should be noted that in the case of \mathbf{P}_a , the derivative is 0.

This method can be applied to any model within this work to calculate the derivative, simply by replacing the scalar coefficients with matrix forms. There is one exception, which is where an infinite resonance must be considered. In which case, it makes sense to define an temporary Fabry–Perot resonance coefficient:

$$F = 1 - P_s^2 R_{s,a}^2 \quad (6.74)$$

The derivative can then be calculated using the matrix formulation. The resonance matrix, \mathbf{F} , can then be set as:

$$\mathbf{F} = \begin{bmatrix} \frac{1}{F} & 0 \\ -\frac{dF}{d\tilde{n}} \frac{1}{F^2} & \frac{1}{F} \end{bmatrix} \quad (6.75)$$

And the infinite resonance model can be written as:

$$\bar{H} = \mathbf{T}_{a,s} \mathbf{T}_{s,a} \mathbf{P}_s \mathbf{P}_a^{-1} \mathbf{F} \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad (6.76)$$

6.9 Uncertainty of fitting a resonant model to z-cut quartz

A 0.512 mm z-cut quartz sample was measured using the HFRBB THz TDS, as in § 5.10. A infinite resonance transfer function model was fitted to the data with different assumed thicknesses between 505 and 525 μm , with 0.5 μm steps. Total variance was then used to extract the correct thickness, and select the corresponding \tilde{n} . The fitted calibration coefficient, β , from § 5.9 was used to calibrate the extraction.

The uncertainty in n , α and total variance was estimated using uncertainty propagation. To validate this approach, other estimates of the uncertainty based on a Monte–Carlo method and a bootstrap method were used. Both of these methods used 500 repetitions, and took several hours to compute on a desktop machine.




Figure 6.15: Different estimates of variance in z-cut quartz sample. The n estimates are shown on the left and the α estimates are shown on the right. The thick blue lines shows the Monte–Carlo estimates, the green lines show the bootstrap estimates and the dashed red lines show the uncertainty propagation estimate.

In Figure 6.15 the different estimates of variance are shown. The top pane shows the variance of n and the bottom pane shows the variance of α . The thick blue line shows the Monte–Carlo estimate, the green line shows the bootstrap estimate and the dashed red line shows the uncertainty propagation estimate. Estimating the uncertainty over a range of thicknesses, see § 6.11, using Monte–Carlo and bootstrap methods took several hours (due to the high number of repetitions). In comparison, estimation via uncertainty propagation took under 30 seconds (less than a second for a singular thickness).

The different estimates are in extremely good agreement, with the exception of some individual frequencies in n . At certain frequencies, the extraction has converged to an incorrect value in some of the Monte–Carlo and bootstrap estimates. This has resulted in the variance being skewed, resulting in spikes in the variance at these frequencies.

The variance shows etalons across the bandwidth, these occur with the frequency as the residual etalons present in the extracted n and α . It should be noted, that if a non-resonant model is fitted these etalons are not present in the uncertainty. The model used will have a significant effect on the uncertainty as well as extracted \tilde{n} . The variance appears to increase in both variables with frequency, as in § 6.5 the uncertainty is inversely proportional to the dynamic range of the reference and at frequencies where this is smaller, the uncertainty will be larger.

In Figure 6.16 the correlation matrices estimated by uncertainty propagation are shown. In the upper left the correlation matrix of n is shown, in the bottom right the correlation matrix of α is shown and in the top right correlation matrix between the two is shown. The correlation matrices of n and α , show two strong regions of correlation, roughly segmented about 3.8 THz. This happens to correspond with the relaxation in the n . This is indicative of the n offset uncertainty being different either side, creating separate regions of correlation. The lower frequency region is broadly correlated, in the n this is due to pulse arrival time uncertainty in the time domain. In the α a similar effect is observed which is likely down to variation in the pulse amplitude in the time domain (i.e. the entire magnitude spectrum is changing, rather than individual regions).

As in Figure 6.12, there is strong diagonal correlation, however it is much narrower. This is because far less zero padding was used to achieve a similar computational frequency resolution. There also appears to be a lattice like structure superimposed on the correlation matrix, this is especially noticeable in the correlation between the two variables. This lattice like structure occurs because of the positive correlation and negative correlation of associated with the etalons which are present in the uncertainty. These will be associated with uncertainty in the sample reflection. For instance if there is uncertainty in its arrival time this will have a broad effect

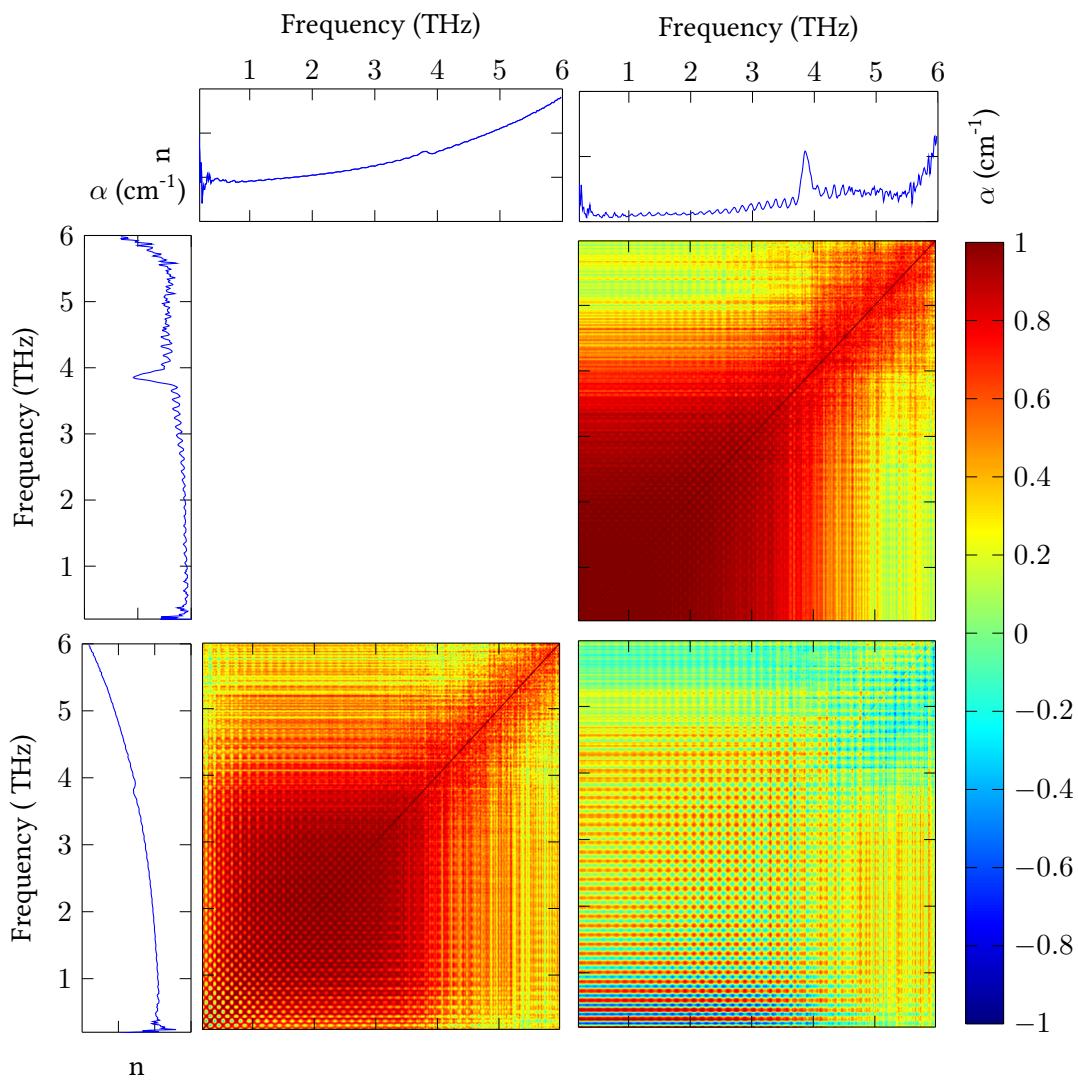


Figure 6.16: Correlation matrices of n and α of z -cut quartz. Bottom left - correlation matrix of n , bottom right - correlation matrix between n and α and upper right - correlation matrix of α . n and α of the sample are shown in the top and left panes for reference.

across the spectrum similar to uncertainty in the main pulse arrival time. However, because the reflection produces etalons in the spectrum, its uncertainty will reflect that and have regions of both positive and negative correlation with similar periodicity.




Figure 6.17: Estimates of the 95% confidence intervals of n . The green shaded region is the simultaneous confidence intervals, and the red shaded region is the non-simultaneous confidence intervals. The black line is the estimate of n .

In Figure 6.17 the extracted n is shown with its corresponding 95% intervals. The green shaded region is the simultaneous confidence intervals and the red shaded region is the non-simultaneous confidence intervals. The confidence intervals were estimated from using the uncertainty propagation method. The black line is the estimate of n . The confidence intervals are fairly narrow across the spectrum, however at higher frequencies the intervals become much broader.




Figure 6.18: Estimates of the 95% confidence intervals of α . The green shaded region is the simultaneous confidence interval and the red shaded region is the non-simultaneous confidence interval. The blue line is the estimate of α .

In Figure 6.18 the estimated α and its corresponding 95% confidence intervals is

shown. The green shaded region is the simultaneous confidence intervals and the red shaded region is non-simultaneous confidence intervals. The black line is the estimated value of α . The confidence interval becomes considerably wider at higher frequencies. The confidence intervals are broad enough to suggest that the residual etalons are at least partially due to a variation within the measurement, particularly due to uncertainty in the sample reflections.

At lower frequencies (between 0.5 and 2.5 THz) the lower confidence interval is smoother than the extracted α . This is due to the uncertainty being greater than the peak and troughs formed by the etalons, indicating that the deterministic value could be considerably smoother. The caveat with this analysis is that the confidence intervals have no indication of correlation and are a random variable (i.e. the confidence interval structure and size are subject to probability).

The uncertainty in these measurements is conditional on the β calibration. However, this calibration curve will be subject to its own uncertainty and has been shown to have a considerable effect on the residual etalons in the extracted parameters, see § 5.10. This uncertainty will likely have a structural effect (i.e. might produce/reduce etalons) on the uncertainty. However, in order to account for this uncertainty, there are several development issues which would need to be resolved (particularly with pooling covariance estimates and with the dimensionality reduction effects of the fitting used).

Another parameter this is conditional on is the extracted thickness, which has been extracted along with its uncertainty in § 6.11 using total variance. This uncertainty will have a considerable effect on the uncertainty of n and α . However, as the uncertainty of the thickness is directly propagated from the uncertainty in n and would require developing a method of estimating a covariance matrix which includes the extracted thickness (i.e. the covariance between n , α and l should be considered). It should be noted that in this sample the uncertainty of the extracted thickness was extremely small and will have negligible effect on the uncertainty of n and α .

6.10 Uncertainty of thickness extracted using total variance

In § 5.6, total variance was used to measure the sample thickness *in situ*, by minimising etalons in the fitted \tilde{n} . The calculation for total variance at assumed thickness l is:

$$TV_l = \mathbf{1}' |D\mathbf{n}_l| \quad (6.77)$$

Where $\mathbf{1}'$ is a N row vector of ones to perform summation, \mathbf{n}_l is the real refractive index extracted using assumed thickness l , and D is a $(N - 1) \times N$ finite difference operator matrix, which is defined as Toeplitz matrix of the form:

$$D = \begin{bmatrix} 1 & -1 & 0 & \cdots & 0 \\ 0 & 1 & -1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 & -1 \end{bmatrix} \quad (6.78)$$

The sensitivity matrix, $C_{TV,n}$, is then in the form of a row vector:

$$C_{TV,n} = sgn(D\mathbf{n}_l)' D \quad (6.79)$$

Where sgn is the sign function:

$$sgn(x) = \begin{cases} x = 1 & : x \geq 0 \\ x = -1 & : x < 0 \end{cases} \quad (6.80)$$

The sensitivity matrix given in equation 6.58 is for a single thickness, and ignores correlation between extracted values of real refractive index using different assumed thickness's. This can be resolved by recognising the inherent dependence on the measurements transfer function and creating a sensitivity matrix, $C_{TV,H}$, which

maps the uncertainty from that to total variance:

$$C_{TV,H} = \begin{bmatrix} C_{TV,n}(l_0)C_{\tilde{n},H}(l_0) \\ C_{TV,n}(l_1)C_{\tilde{n},H}(l_1) \\ \vdots \\ C_{TV,n}(l_{N-1})C_{\tilde{n},H}(l_{N-1}) \end{bmatrix} \quad (6.81)$$

Where C_{TV,n_i} and $C_{\tilde{n},H,l}$ are the sensitivity matrices calculated at thickness l , and l_k is the k th assumed thickness, unto $N - 1$ th thickness.

The covariance matrix of TV is then calculated using the following:

$$\Sigma_{TV} = \frac{1}{2} \mathbb{R} \left(C_{TV,H} \Sigma_H C_{TV,H}^\dagger + C_{TV,H} \Gamma_H C_{TV,H}' \right) \quad (6.82)$$

A subset of TV around the correct local minima can then be used to estimate the uncertainty in thickness. This subset, TV' , will be smaller but still have a marginal normal distribution. A quadratic can be fitted to this reduced subset (alternatively a broad range could be considered by using a higher order polynomial) of the form:

$$TV' = Lb \quad (6.83)$$

where L is design matrix of the form:

$$L = \begin{bmatrix} 1 & l & l^2 \end{bmatrix} \quad (6.84)$$

and b is a vertical vector of polynomial coefficients. The polynomial coefficients can be estimated using linear least squares [139]:

$$\begin{bmatrix} b_0 \\ b_1 \\ b_2 \end{bmatrix} = L^{-1} TV' \quad (6.85)$$




Figure 6.19: Different estimates of variation of total variation. The blue curve shows the Monte–Carlo estimate, the red curve shows the bootstrap estimate and the green curve shows the uncertainty propagation estimate.

where \mathbf{L}^{-1} is the pseudo–inverse of \mathbf{L} . The argument minima of \mathbf{TV}' is:

$$l = \underset{l}{\operatorname{argmin}} \mathbf{TV}' = \frac{-b_1}{2b_2} \quad (6.86)$$

Where l is in this case the extracted sample thickness. The sensitivity matrix, $\mathbf{C}_{l,\mathbf{TV}'}$, between the two is then:

$$\mathbf{C}_{l,\mathbf{TV}'} = \begin{bmatrix} \frac{-1}{2b_2} & \frac{b_1}{2b_2^2} & 0 \end{bmatrix} \mathbf{L}^{-1} \quad (6.87)$$

This will then map the uncertainty in total variance to the uncertainty in the extracted thickness:

$$\sigma_l = \mathbf{C}_{l,\mathbf{TV}'} \Sigma_{\mathbf{TV}'} \mathbf{C}_{l,\mathbf{TV}'}' \quad (6.88)$$

6.11 Extracted thickness uncertainty of *z-cut* quartz

The *z-cut* quartz sample measurement was used to extract the thickness using total variance. Multiple estimates of the uncertainty of total variance were formed, using Monte–Carlo, bootstrap and uncertainty propagation methods.

In Figure 6.19 the estimated variance for the calculated total variance for each of the assumed thicknesses, ranging from 0.5 mm to 0.52 mm with a 0.005 mm step.




Figure 6.20: Different estimates of correlation within the total variation. Left - the Monte-Carlo estimate, middle - the bootstrap estimate and right - the uncertainty propagation estimate.

The blue curve shows the Monte-Carlo estimate, the red curve shows the bootstrap estimate and the green curve shows the uncertainty propagation estimate based on the sensitivity matrix in § 6.10. None of the estimates agree perfectly, as the total variance operation is (purposely) very sensitive to variation.

In Figure 6.21 the different estimates of correlation of the total variation of the measurement. The left image shows the Monte-Carlo estimates, the middle image shows the bootstrap estimate and the right image shows the uncertainty propagation estimate. The estimates are very similar, but small differences are noticeable in the shape of the negative regions. The correlation matrix shows that each side of the minima (see below) is positively correlated with itself but negatively correlated with the other side of the minima. This practically means that if one side of the minima increases the other side will decrease.

In Figure 6.21 the estimated confidence intervals of the total variation. The light shaded region is the simultaneous confidence interval and the darker region is the non-simultaneous confidence interval. The blue line shows the total variance estimate. The estimates have been formed based on uncertainty propagation estimates of variance. The confidence interval is fairly constant across assumed thickness, this is because the variance does not change significantly with assumed thickness.

A quadratic was fitted locally (using linear least squares) to the minima, the region has been marked by vertical lines in Figure 6.21. The fitting has been shown as

two lines, which are based on the simultaneous confidence intervals of the polynomial coefficients. From these polynomial coefficients, 95% confidence intervals for the minima location were estimated to be 512.1 and 512.5 μm . This is the first time, to the authors knowledge, that the thickness uncertainty has been estimated, when the thickness has been extracted using total variance. This is extremely accurate, but as mentioned above, does not account for uncertainty in β , which is likely the biggest uncertainty in this estimate. It has already been shown that the uncertainty in the calibration thickness has a large effect on the measured β , it would be feasible that this would have a large effect on both the uncertainty of β and the uncertainty of the minima location.




Figure 6.21: The estimated confidence intervals for total variation. The lighter region shows the simultaneous confidence interval, the darker region shows the non-simultaneous confidence interval and the blue line shows the estimate of total variation. The purple lines show the confidence intervals for a fitted quadratic between the two vertical dashed lines.

6.12 Conclusion

In this chapter a model for the uncertainty in THz TDS measurements was developed based on multivariate normal statistics. A method of estimating the uncertainty of \tilde{n} was then presented based on the method of uncertainty propagation. This was used to estimate confidence intervals for n and α , using both: (i) simple approximations and (ii) fitting a resonant model to the sample response. It was also shown that the measurement uncertainty had a very small effect on the extracted sample thickness using total variation.

The methods of estimation were shown to be extremely accurate compared to both bootstrap and empirical methods, while being substantially quicker to compute. These methods could be extended to include calibration uncertainty, multilayer sample extraction and time localised fitting.

Chapter 7

Conclusion and further work

In Chapter 5 simple approximate methods of processing the THz TDS measurements were presented and a new method of correcting the refractive index for phase unwrapping issues was developed. These methods of extracting the absorption coefficient and refractive index were found to be sufficient to process some measurements of lactose and paracetamol, as well as develop spectroscopic simulants in Chapter 4.

These approximate methods are often sufficient and should serve as an initial point for the processing. However, if the sample response is complex or there is significant sample thickness uncertainty, the methods presented in Chapter 5 should be used. These method fit a model to the measured transfer function of the sample and by doing so, can incorporate more complex sample responses and extract more information from a single measurement.

A method of using a direct thickness search (total variance) was developed in conjunction with a resonant model which included a calibration coefficient for the THz Gaussian beam profile. This allowed extremely accurate extraction of the sample thickness *in situ* and was used to determine the thickness of *z*-cut quartz and α lactose anhydrous samples.

A simple method of calibrating for this effect was developed. By segmenting the transfer function in the time domain it was possible to isolate the main pulse and first reflections of a calibration sample with known thickness. Using simple estimates of refractive index it was then possible to directly measure the calibration coefficient.

The calibration procedure developed should be adopted in future experimentation and in particular several calibration measurements should be performed after system realignment. These calibration measurements should be performed with transparent, resonant materials with low relative thickness uncertainty (as this was proportional to the uncertainty in β). It was found that *z*-cut quartz was well suited to this task and a variety of quartz samples with differing thicknesses would make excellent calibration samples. Alternatively, if quartz is not available, thick (> 2 mm) high resistivity silicon could make a suitable calibration sample provided that the thickness uncertainty does not scale with sample thickness (for instance, if it is caused by a slanted surface). By performing these calibration measurements, the accuracy of thickness extraction can be greatly enhanced.

A method of modelling multilayered samples was developed and this was used to extract the complex refractive index of a IPA within a flowcell. To perform this extraction, the two *z*-cut quartz layers had to have known complex refractive index and thickness. By measuring an empty flowcell with the same channel thickness it was possible to both extract the complex refractive index of all layers and extract the channel thickness.

To apply this extraction technique in the future to other layered samples, the experiment should be designed in conjunction with the processing with particular consideration of how all the layers should be characterised and how to measure the required information. It may well be the case that it is possible to extract out additional information by considering related samples to the sample of interest, for instance considering multiple measurements of a layered sample with different layers removed.

A time localised model was developed to fit to a time segmented transfer function. This was then used to extract directly the thickness, refractive index and absorption coefficient of a sample. The advantage of this method is that it is robust and a direct search for the thickness is not required. This could therefore be used as an initial step before performing a direct search using total variance to characterise the uncertainty of the thickness, refractive index and absorption coefficient.

This method was extended to extract out the refractive index and absorption co-

efficient of multiple layers. This was applied to a high resistivity silicon and quartz layered sample with known thicknesses. The extracted estimates were negatively correlated and were systematically incorrect. However, it was capable of uniquely assigning the phonon mode correctly to the quartz, which indicates that this method (in its underdeveloped state) could be used as a method of identifying spectral features, rather than accurate values.

When processing sample measurements, it is the recommendation of the author that initial approximations are initially used. These approximate values will often be similar to the final value after processing and should serve as a means of quickly identifying experimental issues (for instance a lower concentration is required) or if further processing is useful. These approximations can then be used as initial estimates during the transfer function fitting.

To perform thickness extraction, resonance within the sample is required. Unfortunately, this is not always the case and samples may be too absorbent or be too transparent for this to occur. To identify if a sample measurement can be used to extract thickness, there are three different potential methods of identifying resonance: i) a reflection is present within the sample time scan which is not present in the reference time scan, ii) after approximating refractive index and absorption coefficient, there are etalons present at lower frequencies which reduce at higher frequencies and iii) performing a direct search using total variance for the thickness. It should be noted, that the first two are simple to perform (relatively) but are not as sensitive as performing the thickness extraction.

It has been observed in this work that samples mixed with PTFE (the matrix material) do often produce reflections which can be exploited to extract out the thickness of the sample. In future work, it would be advisable to adapt existing dilution procedures to take this into account (i.e. dilute not just for bandwidth but for resonance). By measuring pure PTFE (or another matrix material) pellets with different thicknesses, it would be possible to assign an approximate range of pellet mixture masses for low concentrations which produce resonance and whose thickness can be measured. This could then be used in experiments to perform *in situ* thickness extraction.

In Chapter 6, a method of estimating the uncertainty in the refractive index, absorption and thickness were developed. The method developed was found to be extremely accurate when compared to robust numerical methods of estimating the uncertainty. This method is very general and could be adapted to other instrumentation (for instance reflection THz TDS systems) and different processing techniques (for instance fitting Lorentzian peak functions to absorption data).

Before applying these techniques to experimental data, steps should be taken to validate the underlying assumption of normally distributed noise. The simplest method is to examine the bootstrapped histograms of individual time samples and compare them to normal distributions. Secondly, if applying these techniques for the first time or applying a new processing method (for instance a new transfer function model), comparisons between the bootstrap, Monte Carlo and uncertainty propagation estimates should be performed. If they uncertainty propagation agrees with the bootstrap estimate, then it is a suitable method of estimating the uncertainty. The Monte Carlo estimate strictly enforces the condition of normality on the original time domain scans, if the uncertainty propagation estimate agrees with the Monte Carlo estimate but not the bootstrap, the likelihood is that the underlying data is either skewed by an outlier or is not normally distributed.

Experimentally it was found that the uncertainty was inversely proportional to the dynamic range of the reference measurement. Designing experiments so that dynamic range is maximised is thus a method of increasing certainty within the sample measurement. In practice this means that using dry air references are ideal. This proportionality rule is primarily affected by the uncertainty in the reference measurement, so performing more repetitions of the reference measurements will provide a greater reduction in uncertainty than performing repetitions of the sample measurement.

Finally with regards to measurement procedure, it is highly preferable to increase the number of measurement repetitions. This will reduce the uncertainty present by averaging in the time domain but also reduce the variability of the covariance and confidence interval estimates (i.e. taking more measurements increases the certainty of the uncertainty). This should be considered at the expense of noise reduction

within the measurement (i.e. acquiring more data using faster acquisition methods), so that individual measurements are noisier but better estimates of covariance can be gathered. This will not inherently lead to an increase in the uncertainty of the final extracted parameters, as measurement noise reduction is often performed by averaging within the instrumentation rather than in the processing (for instance by using a longer lock in time constant).

7.1 Further work

There are several issues and further avenues of research which can be explored. The foremost limitation still present in the processing are the reflections within the sample trace.

While the primary system reflections have been experimentally delayed by emitter and detector design [46] this might not always be possible due to other experimental constraints (for instance the use of a cryostat). It would therefore be preferable to mitigate the effects of the system reflections in the processing. In order to do this, the underlying cause of the error in these reflections needs to be identified. If it is a systematic measurement error (such as systematic time base error), this can be treated as a systematic uncertainty within the measurement and the work in Chapter 6 could be adapted to calculate the uncertainty in the extracted parameters. If it is an unmodelled effect which is dependent on the sample, it might be more prudent to adapt the sample transfer function model and fit to the transfer function.

Another avenue which has been explored to some extent in the literature is using linear prediction [103], [145], [146]. These methods attempt to model and deconvolve the reflection, which might obscure rather than remove error. Instead it might be more prudent to simply form a time model of the main pulse and fit to a shorter window. Once the model is identified (it will likely be sample dependent), a fitting could be performed over a longer time window. Another avenue would be non-linear processing, for instance the cepstrum [147] is known to separate convolution components in the time domain into a summation. Given that the sample and the reference should have similar convolutional components (from the system

reflection), this could be combined with independent component analysis [148] to remove the common components, which should correspond to the reference and following reflections.

To improve the performance of the simulant generation algorithm, the modelling of mixture spectra could be improved. In particular by incorporating effective medium approximation, such as Maxwell–Garnett [121], and scattering models [96], more accurate models could be used during the optimisation. Another consideration is the predictability and reliability of the produce mixture. In the optimisation, a relative concentration uncertainty could be assumed and used to penalise mixtures which are unlikely to experimentally produce the modelled spectra.

When fitting a transfer function, there are several optimisations which could still be performed. In particular, the current implementation of the optimisation uses a computed finite difference of the error function to descend the gradient. Instead it might be more prudent to incorporate the method of calculating the differential from Chapter 6. This will reduce the calculation complexity significantly.

Fitting a resonant model was shown to mitigate the effects of sample reflections within the trace. This does not however produce an etalon-free estimate of the complex refractive index, which would be expected in most cases. It is unlikely, that it ever will produce a completely smooth estimate without over parametrising the model used to fit, simply because the uncertainty in the complex refractive index displays etalons.

Instead it would be preferable to find the most likely etalon-free estimate for a parameter. This was the approach essentially taken in [63], which used a spatially variant moving average filter. In essence, it smoothed using a moving average filter within confidence intervals on the transfer function. The issue with this approach is it used standard deviation as a confidence interval and treated it as a hard limit (confidence intervals are subject to probability). This approach could be adapted to the confidence intervals formed for the extracted parameters in Chapter 6. However it might instead be better to optimise both the total variance and likelihood of the complex refractive index values, to find the most likely smooth estimate.

The time segmentation methods used to both estimate the beam alignment coef-

efficient and to extract parameters could be improved by performing better separation of the pulses in the impulse response. An issue with the filter used in the deconvolution process is that it has an infinite time response, which leads to slight overlap in the pulses. Therefore, a simple method which might improve the spacial separation is to use a finite impulse response filter when performing the deconvolution. This, however, might incur additional distortion of the frequency domain, so a compromise between the two would have to be made. Another method is to find a sparser approximation [74], which would find an approximation for both the frequency and time responses which is more compact (and thus separable) in the time domain.

The different time segments used during a fitting have different uncertainty. In the case where the fitting is under determined (there is more information than parameters), it would be possible to improve the estimation by weighting each segment in proportion to its certainty. This would result in fitting the transfer function model more towards the main pulse which will have less uncertainty and produce a more accurate estimate of the extracted parameters. This might show particular improvement where the complex refractive index of multiple layers is extracted.

Another consideration when using segmented transfer function models is that it can function on groups of pulses. This could be used when extracting information about multiple layers in a sample, as clusters of pulses can easily be produced (for instance if one layer is optically thick and another optically thin). This could then be used to perform a hybrid of the resonant model fitting and time localised fitting, which could be used to extract the thickness of one layer in the case where spacial separation is not possible.

The uncertainty estimation from Chapter 6 can be considerably expanded. In particular, there is an additional measurement uncertainty from calibration measurements for the beam alignment coefficient, β . This would require developing a methodology of combining multiple measurement uncertainties, as well as propagating to a fitting of β . The same methodology of uncertainty propagation could then be used to estimate the uncertainty in the extracted complex refractive index due to the calibration measurements. This would likely introduce an uncertainty which contains significant etalon like features, so could explain the residual etalon

artefacts present in the complex refractive index when fitting a resonant model.

The uncertainty propagation could be adapted to work with multilayer models and time segmentation methods. In both of these cases it would require extending the calculation of the model differential. In the multilayer case this could then be used to propagate measurement uncertainty of characterised layers to the extracted complex refractive index. In the time segmented fitting, uncertainty propagation of the segmentation would also have to be considered before the uncertainty from different segments could be propagated to the both extracted thickness and complex refractive index.

Finally, the analysis of uncertainty could be expanded. In particular, principle component analysis and factor analysis [139] could be used to analyse the uncertainty within the time domain and within the extracted parameters, to characterise the sources and their effects. Another use of principle component analysis would be to perform dimensionality reduction. This could then be used to form approximate confidence ellipses for the complex refractive index, which would give a quantitative description of the correlated uncertainties effect on the extracted parameters.

Appendices

Appendix A

Fabrication of powdered sample pellets

To create a sample pellet for a given concentration, C , the following steps are taken:

- The sample is measured out by mass to $100C$ mg.
- The sample is ground in a hand mortar and pestle.
- $1\ \mu\text{m}$ PTFE is measured out by mass to $100(1 - C)$ mg.
- The PTFE is mixed with the ground sample to form a consistent mixture.
- A copper ring is placed in the pellet press. This is used to form a support for the pellet.
- 40 mg of the pellet mixture is measured out and placed inside the copper ring.
- The press with the pellet mixture is then fully assembled and put under 8 tonnes of pneumatic pressure for 20 minutes.
- The press is carefully disassembled and the pellet is removed from the press.

40 mg is used to make the pellet as it was found to produce a consistently physically stable, thin pellet. A pellet made using this mass would usually be approximately 0.5 mm thick and leave enough remaining pellet mixture to manufacture a second pellet if required.

Appendix B

Additional beam alignment calibration measurements

This appendix contains different estimates of the beam alignment calibration coefficient, β , used in Chapter 5.9 to perform a quadratic fitting. The measurements were performed in the HFRBB THz TDS and LFRBB THz TDS systems, which used different emitters and detectors but had the same alignment and parabolic. It was verified that measurements from the two systems produced similar estimates of β . In Figures B.1,B.2 several different samples were measured in the HFRBB THz TDS system. In Figure B.3, *z*-cut quartz samples were measured in the LFRBB THz TDS system. It should be noted, that different samples were subject to different thickness uncertainties, which is why the offset of β changes between different samples.




Figure B.1: Estimates of β made using different samples on the HFRBB THz TDS system. Blue - 0.512 mm *z*-cut quartz, Green - 1.015 mm high resistivity silicon, Red - 0.58 mm high resistivity silicon, Purple - 0.969 mm TPX




Figure B.2: Estimates of β made using different samples on the HFRBB THz TDS system. Blue - 0.512 mm *z-cut* quartz, Green - 1.015 mm high resistivity silicon




Figure B.3: Estimates of β made using different samples on the LFRBB THz TDS system. Cyan & Red - 0.512 mm *z-cut* quartz samples, Blue & Green - 0.325 mm (uncertainty not known) *z-cut* quartz samples

Appendix C

Temperature dependence air volume in a 5 % α -lactose anhyrous sample

In § 5.11, the n of a 5 % α -lactose anhyrous sample was found to change significantly with temperature. It was postulated that this might be due to the effective permittivity of the sample changing, as the relative densities of the samples constituents (air, lactose and PTFE) changes with temperature. In this appendix, a correction based on the air volume solely changing with temperature is presented. This was found to significantly reduce the variation with temperature. However, it should be noted that this method was extremely reliant on several assumptions and used a fitting to estimate the air volume content.

This measured \tilde{n} can be related to the effective permittivity of the mixture. From this effective permittivity an estimate of the permittivity of lactose can be formed by using Maxwell-Garnett [121] effective medium approximations. Using Maxwell-Garnett, a relation between the effective permittivity, sample permittivity, host (PTFE matrix) permittivity and air permittivity can be formed [136]:

$$\frac{\varepsilon_{eff} - \varepsilon_h}{\varepsilon_{eff} + 2\varepsilon_h} = \delta_s \frac{\varepsilon_s - \varepsilon_h}{\varepsilon_s + 2\varepsilon_h} + \delta_a \frac{\varepsilon_a - \varepsilon_h}{\varepsilon_a + 2\varepsilon_h} \quad (C.1)$$

where ε_{eff} is the effective (measured) permittivity of the mixture, ε_h is the permittiv-

ity of the host (PTFE matrix), ε_s is the lactose permittivity and ε_a is the permittivity of air (assumed $\varepsilon_a = 1$). Defining the value, $C_{h,a}$, for brevity as:

$$C_{h,a} = \left(\frac{\varepsilon_{eff} - \varepsilon_h}{\varepsilon_{eff} + 2\varepsilon_h} - \delta_a \frac{\varepsilon_a - \varepsilon_h}{\varepsilon_a + 2\varepsilon_h} \right) \quad (C.2)$$

This can be rearranged to:

$$\varepsilon_s = \frac{2\varepsilon_h C_{h,a} + \delta_s \varepsilon_h}{\delta_s - C_{h,a}} \quad (C.3)$$

The sample volume ratio, δ_s , can be formed from the sample mass ratio, $\delta_{s,m}$, sample density, d_s , air volume ratio, δ_a , and host density, d_h :

$$\delta_s = \frac{d_h \delta_{s,m}}{d_h \delta_{s,m} + (1 - \delta_{s,m}) d_s} (1 - \delta_a) \quad (C.4)$$

Similarly, the host ratio:

$$\delta_h = \frac{d_s (1 - \delta_{s,m})}{d_h \delta_{s,m} + (1 - \delta_{s,m}) d_s} (1 - \delta_a) \quad (C.5)$$

This assumes that the sample and host mass ratios are known, but that the air mass is not known (but that the volume ratio is).

If the volume change of each of the components were known as a function of temperature it would be possible to estimate the sample permittivity directly from the measurements. However, since this is not easy to measure *in situ* with the cryogenic measurements.

It is possible to form an estimate based on the assumption that a single component of the sample is responsible for the contraction in sample volume. In this work, it is assumed that the change in air volume within the sample is dominant. This is because the air is a gas and hence more compressible than the solid constituents of the sample.

It is then approximated that the air is solely responsible for the change in sample volume and that the sample contracts uniformly in all directions. From the latter approximation, the thickness change with temperature can be used to estimate the

volume change with temperature:

$$\delta_V(t) = \frac{V_t}{V_{rt}} = \left(\frac{l_t}{l_{rt}}\right)^3 \quad (\text{C.6})$$

Where δ_V is the volume change at temperature t , V_t is the volume at temperature t , V_{rt} is the volume at room temperature, l_t is the thickness at temperature t and l_{rt} is the thickness at room temperature.

From the air assumption, the volumes of the lactose and PTFE host are roughly consistent with temperature. Thus the volume ratio for each as a function of temperature are:

$$\delta_{s,t} = \frac{\delta_s}{\delta_V(t)} \quad (\text{C.7})$$

and:

$$\delta_{h,t} = \frac{\delta_h}{\delta_V(t)} \quad (\text{C.8})$$

Where $\delta_{s,t}$ and $\delta_{h,t}$ are the sample and host volume ratios at temperature t . The air volume ratio can then be formed:

$$\delta_{a,t} = 1 - \delta_{s,t} - \delta_{h,t} \quad (\text{C.9})$$

The density of lactose is 1.525 gcm^{-3} [149], and of PTFE is 2.16 gcm^{-3} [150]. The mass ratio of the lactose to PTFE was 5%, and the average permittivity of PTFE is 2.1025 [151]. An air volume ratio of 0.01 was estimated by minimising the mean square error between 100% lactose anhydrous, see § 3.8, and this sample at room temperature. The extracted thickness was then used to estimate the volume ratio changes in the sample, and thus the permittivity of the sample, from which n and α can be calculated.

It was found that by assuming that the air volume is solely responsible for the sample thickness change significantly reduced the variation in n across temperatures. In Figure C.1, the average of n is shown at each temperature. The blue curve shows the means based on the volume ratios changing, and the red curve shows the extracted values assuming that the volume ratios are constant. The corrected




Figure C.1: The average real refractive index of the sample at each temperature, after applying a Maxwell Garnett effective medium approximation. Blue - Assuming that volume ratios are changing. Red - Assuming that volume ratios are constant.

estimate is significantly flatter across temperature, while the uncorrected estimate shows a relatively large change in n .

Appendix D

Multilayer Model Examples

In this appendix a simplified recursion by recursion example of the multilayer model from § 5.13 is given. To simplify the example considerably (and get a predictable model) a 2 layer sample is used both with thickness 1 m, the angular frequency, ω , has been kept at one and the free space propagation of light, c , has been kept at 1 ms^{-1} . The first layer has a \tilde{n} of 2 and the second a \tilde{n} of 3. By using a $c = 1ms^{-1}$ and the layer thickness of 1 m, this meant that the propagation time for the first layer was 2 s and for the second layer was 3 s. A noise limit of 0 was used (infinite dynamic range) and a time limit of 3 s was used for the model.

The iteration of equation 5.57 is shown with these limits applied. The equation has been coloured to match the programmatically generated image labeled with the same number. The initial equation for the main tree transfer function is:

$$H_T = T_{0,1}H_{1,1}, L = 1, D = 1, T_{\text{lim}} = 5 \quad (\text{D.1})$$

The initial time limit is larger because of the displaced air propagation time. This is expanded into the recursive form (Figure D.1 (1)):

$$H_T = T_{0,1}P_1 (R_{1,2} H_{1,-1} + T_{1,2}H_{2,1}) \quad , L = 1, D = 1, T_{\text{lim}} = 5 \quad (\text{D.2})$$

The current layer index, L , and direction, D , are shown on the right with the current

time limit. The reflection branch is then expanded (Figure D.1(2)):

$$H_T = T_{0,1}P_1 (R_{1,2} P_1 (R_{1,0}H_{1,1} + T_{1,0}H_{0,-1}) + T_{1,2}H_{2,1})$$

$$, L = 1, D = -1, T_{\text{lim}} = 3 \quad (\text{D.3})$$

The non-explored path has been coloured in gray. The propagation direction has flipped but the layer has remained the same. The Transmitted path, which leaves the detector is then explored (Figure D.1(3)):

$$H_T = T_{0,1}P_1 (R_{1,2} P_1 (R_{1,0}H_{1,1} + T_{1,0}H_{0,-1}) + T_{1,2}H_{2,1})$$

$$, L = 0, D = -1, T_{\text{lim}} = 1 \quad (\text{D.4})$$

As $L = 0$, this path is redundant and $H_{0,-1} = 0$. The reflection path is the expanded(Figure D.1(4)):

$$H_T = T_{0,1}P_1 (R_{1,2} P_1 R_{1,0}P_1 (R_{1,2}H_{1,-1} + T_{1,2}H_{2,1}) + T_{1,2}H_{2,1})$$

$$, L = 1, D = 1, T_{\text{lim}} = 1 \quad (\text{D.5})$$

And then the reflection path after that is expanded (Figure D.1(5)):

$$H_T = T_{0,1}P_1 (R_{1,2} P_1 R_{1,0}P_1 (R_{1,2}H_{1,-1} + T_{1,2}H_{2,1}) + T_{1,2}H_{2,1})$$

$$, L = 1, D = -1, T_{\text{lim}} = -1 \quad (\text{D.6})$$

This is redundant because of the time limit, so $H_{1,-1} = 0$. The transmitted path is then expanded (Figure D.1(6)):

$$H_T = T_{0,1}P_1 (R_{1,2} P_1 R_{1,0}P_1 T_{1,2}H_{2,1} + T_{1,2}H_{2,1})$$

$$, L = 2, D = 1, T_{\text{lim}} = -1 \quad (\text{D.7})$$

Which is again redundant because of the time limit. The original transmission is

then expanded, following the beam path through the second layer (Figure D.1(7)):

$$H_T = T_{0,1} P_1 P_2 (R_{2,3} H_{2,-1} + T_{2,3} H_{3,1})$$

$$, L = 2, D = 1, T_{\text{lim}} = 3 \quad (\text{D.8})$$

The reflection at the final interface is then explored (Figure D.1(8)):

$$H_T = T_{0,1} P_1 P_2 (R_{2,3} H_{2,-1} + T_{2,3} H_{3,1})$$

$$, L = 2, D = 1, T_{\text{lim}} = 0 \quad (\text{D.9})$$

This path is redundant because the time limit is 0, thus $H_{2,-1} = 0$. Exploring the transmitted path (Figure D.1(9)):

$$H_T = T_{0,1} P_1 P_2 T_{2,3} H_{3,1}$$

$$, L = 3, D = 1, T_{\text{lim}} = 0 \quad (\text{D.10})$$

As the layer index is 3, this is detected . This means $H_{3,1} = 1$ and H_T can be simplified:

$$H_T = T_{0,1} P_1 P_2 T_{2,3} \quad (\text{D.11})$$

This can be extended to longer time limits and greater dynamic range, to where the model formed is much more complicated, see Figure D.1(10).

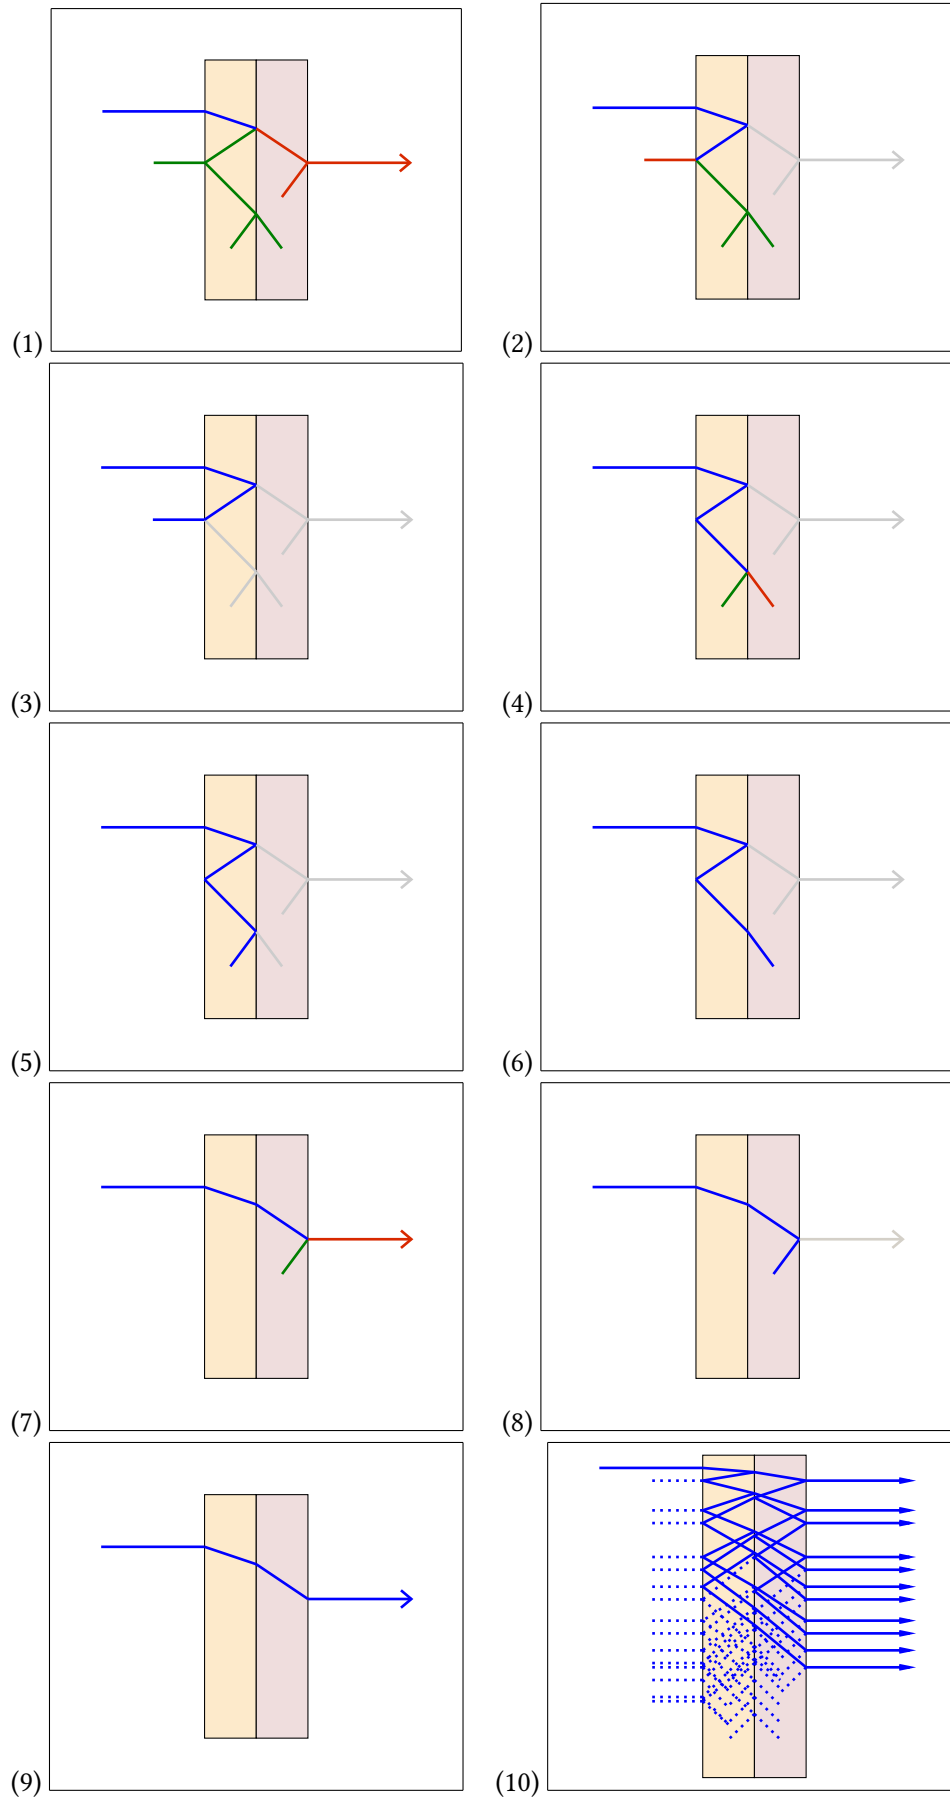


Figure D.1: Generated beam path diagrams for the model at different stages.

This algorithm can be extended to arbitrary number of layers, noise and time constraints. For examples see Figures D.2,D.3, where figures have been generated using a 3 and 4 layer sample.




Figure D.2: A resonance tree through a 3 layered sample




Figure D.3: A resonance tree through a 4 layered sample

Appendix E

Overview of complex multivariate normal distributions

In this appendix, an overview of complex multivariate normal distributions is presented. It is assumed that the reader has an understanding of both complex numbers and univariate normal distributions.

E.1 Bivariate normal variables

The simplest multivariate random variable is the bivariate variable, which is in the case of two variables, x_0, x_1 , which are somehow related and are both normally distributed. In such a case, these two variables can be said to be correlated. By this, there is an association between their measured values, this does not necessarily mean however that there is causation or a strong relation between the variables. In Figure E.1, examples of correlated samples taken from otherwise standard ($\bar{x}_0, \bar{x}_1 = 0$ and $\sigma_{x_0}, \sigma_{x_1} = 1$) normal variables are given. The left figure shows samples taken from two variables which are negatively correlated (so there is a negative association between samples), the middle shows where there is no correlation, and the variables are independent, and the right shows positive correlation between the two variables (so there is a positive association between samples).




Figure E.1: Samples of two variables, x_0, x_1 , with different modeled correlations. Left - Negative correlation, Middle - no correlation (independent) and Right - Positive correlation

To quantify this, the covariance between two variables, σ_{x_0, x_1} , is introduced [152]:

$$\sigma_{x_0, x_1} = \text{cov}(x_0, x_1) = \mathbb{E}((x_0 - \bar{x}_0)(x_1 - \bar{x}_1)) \quad (\text{E.1})$$

This is measure of both the dispersion of each variable and the correlation. This can then be further described as a multiplication of standard deviation and a correlation coefficient, $r_{x,y}$, of the two variables:

$$\sigma_{x_0, x_1} = \sigma_{x_0} r_{x_0, x_1} \sigma_{x_1} \quad (\text{E.2})$$

Conversely, the correlation coefficient can always be calculated by normalising the covariance:

$$r_{x_0, x_1} = \frac{\sigma_{x_0, x_1}}{\sigma_{x_0} \sigma_{x_1}} \quad (\text{E.3})$$

The correlation coefficient will range from -1 to 1 . In this case 1 corresponds to a perfect positive linear relation, -1 corresponds to a negative linear relation, and 0 corresponds to independence (no correlation) [152].

Instead of having separate variables it is preferable to group variables into a single two dimensional vector variable. This bivariate variable, \mathbf{x} , will be a vector of the scalar variables:

$$\mathbf{x} = \begin{bmatrix} x_0 \\ x_1 \end{bmatrix} \quad (\text{E.4})$$

This bivariate variable will have an average, $\bar{\mathbf{x}}$. This expectation is nothing more than a vector of the individual scalar vectors:

$$\bar{\mathbf{x}} = \mathbb{E}(\mathbf{x}) = \begin{bmatrix} \bar{x}_0 \\ \bar{x}_1 \end{bmatrix} \quad (\text{E.5})$$

A covariance matrix [139], $\Sigma_{\mathbf{x}}$, (sometimes referred to as a auto-covariance or variance covariance matrix) can be introduced for a bivariate variable:

$$\Sigma_{\mathbf{x}} = \mathbb{E}((\mathbf{x} - \bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}})') = \begin{bmatrix} \sigma_{x_0}^2 & \sigma_{x_0, x_1} \\ \sigma_{x_1, x_0} & \sigma_{x_1}^2 \end{bmatrix} \quad (\text{E.6})$$

This matrix will contain the variance vector along it's diagonal, and in the upper and lower of diagonal corners, the covariance between the two sub variables. A similar correlation matrix, $R_{\mathbf{x}}$, can be formed:

$$\Sigma_{\mathbf{x}} = \boldsymbol{\sigma}_{\mathbf{x}} R_{\mathbf{x}} \boldsymbol{\sigma}_{\mathbf{x}} \quad (\text{E.7})$$

Where $\boldsymbol{\sigma}_{\mathbf{x}}$ is the standard deviation matrix. The Correlation matrix always has a diagonal containing ones, indicating a sub variable is perfectly correlated with itself. The standard deviation matrix can be calculated by isolating the main diagonal of $\Sigma_{\mathbf{x}}$:

$$\boldsymbol{\sigma}_{\mathbf{x}} = \begin{bmatrix} \sigma_{x_0} & 0 \\ 0 & \sigma_{x_1} \end{bmatrix} \quad (\text{E.8})$$

Being a diagonal matrix, the inverse is particularly easy to calculate:

$$\boldsymbol{\sigma}_{\mathbf{x}}^{-1} = \begin{bmatrix} \frac{1}{\sigma_{x_0}} & 0 \\ 0 & \frac{1}{\sigma_{x_1}} \end{bmatrix} \quad (\text{E.9})$$

And thus the correlation matrix is easily retrievable:

$$R_{\mathbf{x}} = \boldsymbol{\sigma}_{\mathbf{x}}^{-1} \Sigma_{\mathbf{x}} \boldsymbol{\sigma}_{\mathbf{x}}^{-1'} \quad (\text{E.10})$$

The PDF of a multivariate (bivariate is a spacial case) variable can then be written

in terms of vectors and matrices:

$$P(\mathbf{x} = \mathbf{a}) = \frac{1}{\sqrt{2\pi} |\Sigma_{\mathbf{x}}|} e^{-\frac{1}{2}(\mathbf{a}-\bar{\mathbf{x}})' \Sigma_{\mathbf{x}}^{-1} (\mathbf{a}-\bar{\mathbf{x}})} \quad (\text{E.11})$$

where $|\Sigma_{\mathbf{x}}|$ is the determinate of the covariance matrix. In the case of bivariate variable, the bivariate value maps to a scalar probability, and will always form an elliptically (on the x_0, x_1 plane) shaped distribution.

E.2 Multivariate Normal Variables

A multivariate variable is simply a bivariate variable extended to arbitrary number of sub variables [139]. So the N dimensional variable \mathbf{x} is just:

$$\mathbf{x} = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_{N-2} \\ x_{N-1} \end{bmatrix} \quad (\text{E.12})$$

The average vector, $\bar{\mathbf{x}}$ is then:

$$\bar{\mathbf{x}} = \mathbb{E}(\mathbf{x}) = \begin{bmatrix} \bar{x}_0 \\ \bar{x}_1 \\ \vdots \\ \bar{x}_{N-2} \\ \bar{x}_{N-1} \end{bmatrix} \quad (\text{E.13})$$

The covariance matrix is then extended thusly:

$$\Sigma_{\mathbf{x}} = \mathbb{E}((\mathbf{x} - \bar{\mathbf{x}})(\mathbf{x} - \bar{\mathbf{x}})') = \begin{bmatrix} \sigma_{x_0}^2 & \sigma_{x_0,x_1} & \cdots & \sigma_{x_0,x_{N-2}} & \sigma_{x_0,x_{N-1}} \\ \sigma_{x_1,x_0} & \sigma_{x_1}^2 & \cdots & \sigma_{x_1,x_{N-2}} & \sigma_{x_1,x_{N-1}} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \sigma_{x_{N-2},x_0} & \sigma_{x_{N-2},x_1} & \cdots & \sigma_{x_{N-2}}^2 & \sigma_{x_{N-2},x_{N-1}} \\ \sigma_{x_{N-1},x_0} & \sigma_{x_{N-1},x_1} & \cdots & \sigma_{x_{N-1},x_{N-2}} & \sigma_{x_{N-1}}^2 \end{bmatrix} \quad (\text{E.14})$$

The probability density function of a multivariate normal distribution is equation E.11, extended to N dimensions. The notation for a multivariate is the same as in the univariate case, just with vectors and a notation for the dimensionality:

$$\mathbf{x} \sim \mathcal{N}_N(\bar{\mathbf{x}}, \Sigma_{\mathbf{x}}) \quad (\text{E.15})$$

Given a variable, \mathbf{y} , which is a linear function of \mathbf{x} :

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{c} \quad (\text{E.16})$$

where \mathbf{A} is a matrix, and \mathbf{c} is a vector. Both \mathbf{c} and \mathbf{y} have length M , and \mathbf{A} will therefore be a matrix of size $M \times N$. \mathbf{y} will then be normally distributed with L dimensions, it should be noted that $L \neq N$, and can equal 1 (dimensionality reduction is generally possible). The expectation is commutable with the matrix multiplication operation, and thus $\bar{\mathbf{y}}$ is [139]:

$$\bar{\mathbf{y}} = \mathbf{A}\bar{\mathbf{x}} + \mathbf{c} \quad (\text{E.17})$$

And $\Sigma_{\mathbf{y}}$ is then [139]:

$$\Sigma_{\mathbf{y}} = \mathbf{A}\Sigma_{\mathbf{x}}\mathbf{A}' \quad (\text{E.18})$$

Another relevant property is that a marginal distribution of a multivariate normal

variable is always normally distributed [139], thus:

$$\mathbf{y} \subset \mathbf{x}, \quad M < N, \quad \mathbf{y} \sim \mathcal{N}_M(\bar{\mathbf{y}}, \Sigma_{\mathbf{y}}) \quad (\text{E.19})$$

This will commonly be assumed when using a subset of the measured quantities, e.g. when windowing or using select frequencies.

E.3 Complex multivariate gaussian distributions

A large portion of this work models complex variables, in which case the noise in the variables will also be complex. A model for handling the random distribution of complex variables needs to be addressed.

The simplest method of doing this is to simply model the real and imaginary components as separate but correlated variables [153]. This results in univariate variables being transformed to bivariate, and N multivariate variables being transformed to $2N$ multivariate variables. In multivariate notation, given vector \mathbf{z} which a combination of real, \mathbf{x} , and imaginary, \mathbf{y} , parts:

$$\mathbf{z} = \mathbf{x} + i\mathbf{y} \quad (\text{E.20})$$

A new form can then be formed:

$$\hat{\mathbf{z}} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \quad (\text{E.21})$$

Thus $\hat{\mathbf{z}}$ has distribution:

$$\hat{\mathbf{z}} \sim \mathcal{N}_{2N}(\hat{\bar{\mathbf{z}}}, \Sigma_{\hat{\mathbf{z}}}) \quad (\text{E.22})$$

Importantly, the covariance can be written in terms of sub matrices:

$$\Sigma_{\hat{\mathbf{z}}} = \begin{bmatrix} \Sigma_{\mathbf{x}} & \Sigma_{\mathbf{x},\mathbf{y}} \\ \Sigma_{\mathbf{y},\mathbf{x}} & \Sigma_{\mathbf{y}} \end{bmatrix} \quad (\text{E.23})$$

This method is effective, however it would be preferable to perform statistical

analysis in the same representation as the variable being modelled, which makes the analysis less complex. An alternative definition of complex covariance, used more commonly in signal processing applications, is to use the expectation of the conjugate square [154]:

$$\Sigma_z = \mathbb{E} \left((z - \bar{z})(z - \bar{z})^\dagger \right) \quad (\text{E.24})$$

where \dagger represents the conjugate transpose ($\dagger = *'$). This new definition of the covariance matrix is both complex, and often an incomplete description of covariance (there are exceptions [154]). An additional matrix is now required, called the co-covariance matrix (other names proposed are the relation [155] or pseudo covariance matrix [156]). This matrix is:

$$\Gamma_z = \mathbb{E} \left((z - \bar{z})(z - \bar{z})' \right) \quad (\text{E.25})$$

Variable z is then said to have complex normal distribution, with complex average \bar{z} , covariance matrix Σ_z and co-covariance matrix Γ_z :

$$z \sim \mathcal{N}_N(\bar{z}, \Sigma_z, \Gamma_z) \quad (\text{E.26})$$

This complex representation of covariance can be related to real and imaginary covariance matrices by [155]:

$$\Sigma_z = (\Sigma_x + \Sigma_y) + i(\Sigma_{y,x} - \Sigma_{x,y}) \quad (\text{E.27})$$

$$\Gamma_z = (\Sigma_x - \Sigma_y) + i(\Sigma_{y,x} + \Sigma_{x,y}) \quad (\text{E.28})$$

This shows that both matrices are a complex formulation of the same information as the real and imaginary covariance matrices. It is always possible to recover one of the real and imaginary covariance matrices, given both Σ_z and Γ_z .

It is useful to consider the variable \dot{z} which is a vector of z and its conjugate:

$$\dot{z} = \begin{bmatrix} z \\ z^* \end{bmatrix} \quad (\text{E.29})$$

This has a conjugate symmetry (which is shared by the Fourier transform for example) and can be linearly mapped to \hat{z} [155]:

$$\dot{z} = \begin{bmatrix} 1\mathbf{I} & i\mathbf{I} \\ 1\mathbf{I} & -i\mathbf{I} \end{bmatrix} \hat{z} \quad (\text{E.30})$$

$$\bar{z} = \frac{1}{2} \begin{bmatrix} 1\mathbf{I} & 1\mathbf{I} \\ -i\mathbf{I} & i\mathbf{I} \end{bmatrix} \dot{z} \quad (\text{E.31})$$

where \mathbf{I} is an identity matrix of size $N \times N$. If the augmented covariance matrix [154] is then calculated by using equation E.25 on \dot{z} , which is related to Σ_z and Γ_z by:

$$\Sigma_{\dot{z}} = \begin{bmatrix} \Sigma_z & \Gamma_z \\ \Gamma_z^* & \Sigma_z^* \end{bmatrix} \quad (\text{E.32})$$

This also proves that the complex covariance matrices are a complex mapping of the real and imaginary covariance matrices.

\dot{z} can then be said to be normally distributed simply with augmented covariance matrix $\Sigma_{\dot{z}}$:

$$\dot{z} \sim \mathcal{N}_{2N}(\bar{z}, \Sigma_{\dot{z}}) \quad (\text{E.33})$$

However as \dot{z} is a redundant form of z , the form of the normal distribution using covariance and co-covariance matrices is sufficient.

An alternative formulation of \dot{z} is used based on Fourier series conjugate symmetry. A matrix flip transform, $\overset{\leftrightarrow}{\mathbf{I}}$, is specified as the horizontal mirror image of the identity matrix. A new complex variable is defined as:

$$\tilde{z} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \overset{\leftrightarrow}{\mathbf{I}} \end{bmatrix} \dot{z} = \begin{bmatrix} z \\ \updownarrow z^* \end{bmatrix} \quad (\text{E.34})$$

where \uparrow specifies a vertical flip and $\mathbf{0}$ is a matrix of zeros of size $N \times N$. $\tilde{\mathbf{z}}$ will then be in the same form as a discrete Fourier transform of a real multivariate variable. This can be seen as if \mathbf{z} is the first N sequentially ordered discrete frequency components of a Fourier transform with $2N$ frequencies, so that if z_k is the k th discrete complex frequency component:

$$z_k = z \left(\frac{F_s k}{2N} \right) \quad (\text{E.35})$$

The discrete Fourier transform of a real variable has conjugate symmetry so that [93]:

$$z \left(\frac{F_s(2N - k)}{2N} \right) = z \left(\frac{-F_s k}{2N} \right) = z \left(\frac{F_s k}{2N} \right)^* \quad (\text{E.36})$$

i.e. the discrete complex frequency component at frequencies beyond half the sampling frequency (or alternatively at a negative frequency) will be equal to the conjugate of a lower frequency. $\tilde{\mathbf{z}}$ is then the full discrete frequency spectrum:

$$\tilde{z}_k = z \left(\frac{F_s k}{2N} \right) \quad (\text{E.37})$$

The augmented covariance matrix will then be converted to what is referred to in this work as the Fourier covariance matrix:

$$\Sigma_{\tilde{\mathbf{z}}} = \begin{bmatrix} \Sigma_z & \overset{\leftrightarrow}{\Gamma}_z \\ \uparrow \Gamma_z^* & \uparrow \overset{\leftrightarrow}{\Sigma}_z^* \end{bmatrix} \quad (\text{E.38})$$

This represents the “full spectrum” complex covariance of a complex frequency dependent variable. In this work complex representation is used exclusively within the frequency domain, so every complex variable has a Fourier covariance matrix associated with its full spectrum. It is however a redundant form, as a limited selection of frequencies are only ever used (these are still marginally normally distributed) and the full Fourier covariance matrix can always be reconstructed from Σ_z and Γ_z . There are instances when the Fourier covariance matrix is preferred during the analysis as it shares symmetry with the spectrum, see § 2.3.5. It should be noted,

that there is a redundant form of the normal distribution:

$$\tilde{z} \sim \mathcal{N}_{2N}(\tilde{z}, \Sigma_{\tilde{z}}) \quad (\text{E.39})$$

Like \dot{z} , \tilde{z} can be mapped to the real and imaginary sub-variables.

Given $\mathbf{q} = \mathbf{A}\mathbf{z} + \mathbf{c}$ for arbitrary complex matrix \mathbf{A} , the following can be constructed [154]:

$$\Sigma_{\mathbf{q}} = \mathbf{A}\Sigma_z\mathbf{A}^\dagger \quad (\text{E.40})$$

$$\Gamma_{\mathbf{q}} = \mathbf{A}\Gamma_z\mathbf{A}' \quad (\text{E.41})$$

Additionally \mathbf{q} will then be complexly distributed, so that:

$$\mathbf{q} \sim \mathcal{N}_L(\bar{\mathbf{q}}, \Sigma_{\mathbf{q}}, \Gamma_{\mathbf{q}}) \quad (\text{E.42})$$

Another point of interest, is that if z is real:

$$\Sigma_z = \Gamma_z \quad (\text{E.43})$$

This can be used in the case where \mathbf{A} and \mathbf{q} is complex, to form a complex distribution of a function of a real variable.

Bibliography

- [1] A. G. Davies and E. Linfield, "Bridging the terahertz gap," *Physics World*, vol. 17, no. 4, p. 37, 2004.
- [2] J. F. Federici, B. Schulkin, F. Huang, D. Gary, R. Barat, F. Oliveira, and D. Zimdars, "Thz imaging and sensing for security applications—explosives, weapons and drugs," *Semiconductor Science and Technology*, vol. 20, no. 7, p. 266, 2005.
- [3] R. M. Woodward, V. P. Wallace, R. J. Pye, B. E. Cole, D. D. Arnone, E. H. Linfield, and M. Pepper, "Terahertz pulse imaging of ex vivo basal cell carcinoma," *Journal of Investigative Dermatology*, vol. 120, no. 1, pp. 72–78, 2003.
- [4] P. F. Taday, "Applications of terahertz spectroscopy to pharmaceutical sciences," *Philosophical Transactions. Series A, Mathematical, Physical and Engineering Sciences*, vol. 362, no. 1815, pp. 351–63, 2004.
- [5] A. G. Davies, A. D. Burnett, W. Fan, E. H. Linfield, and J. E. Cunningham, "Terahertz spectroscopy of explosives and drugs," *Materials Today*, vol. 11, pp. 18–26, 2008.
- [6] J. A. Zeitler and Y. Shen, "Industrial applications of terahertz imaging," in *Terahertz spectroscopy and imaging*, Springer, 2012.
- [7] D. Zimdars, J. A. Valdmanis, J. S. White, G. Stuk, S. Williamson, W. P. Winfree, and E. I. Madaras, "Technology and applications of terahertz imaging non-destructive examination: Inspection of space shuttle sprayed on foam insulation," vol. 760, no. 1, pp. 570–577, 2005.
- [8] S. Wietzke, C. Jansen, F. Rutz, D. Mittleman, and M. Koch, "Determination of additive content in polymeric compounds with terahertz time-domain spectroscopy," *Polymer Testing*, vol. 26, no. 5, pp. 614–618, 2007.
- [9] J. B. Jackson, J. Bowen, G. Walker, J. Labaune, G. Mourou, M. Menu, and K. Fukunaga, "A survey of terahertz applications in cultural heritage conservation science," *IEEE Transactions on Terahertz Science and Technology*, vol. 1, no. 1, pp. 220–231, 2011.
- [10] J. A. Zeitler, P. F. Taday, D. A. Newnham, M. Pepper, K. C. Gordon, and T. Rades, "Terahertz pulsed spectroscopy and imaging in the pharmaceutical setting—a review," *Journal of Pharmacy and Pharmacology*, vol. 59, no. 2, pp. 209–223, 2007.
- [11] D. M. Mittleman, R. H. Jacobsen, R. Neelamani, R. G. Baraniuk, and M. C. Nuss, "Gas sensing using terahertz time-domain spectroscopy," *Applied Physics B: Lasers and Optics*, vol. 67, no. 3, pp. 379–390, 1998.

- [12] H. M. Pickett, "Thz spectroscopy of the atmosphere," in *Terahertz Spectroscopy and Applications*, International Society for Optics and Photonics, vol. 3617, 1999, pp. 2–7.
- [13] K. Druzicki, J. Mielcarek, A. Kiwilsza, L. Toupet, E. Collet, A. Pajzderska, and J. Wasicki, "Computationally assisted (solid-state density functional theory) structural (x-ray) and vibrational spectroscopy (ft-ir, ft-rs, tds-thz) characterization of the cardiovascular drug lacidipine," *Crystal Growth & Design*, vol. 15, no. 6, pp. 2817–2830, 2015.
- [14] C. J. Strachan, P. F. Taday, D. A. Newnham, K. C. Gordon, J. A. Zeitler, M. Pepper, and T. Rades, "Using terahertz pulsed spectroscopy to quantify pharmaceutical polymorphism and crystallinity," *Journal of Pharmaceutical Sciences*, vol. 94, no. 4, pp. 837–846, 2005.
- [15] E. P. Parrott, J. A. Zeitler, T. Friscic, M. Pepper, W. Jones, G. M. Day, and L. F. Gladden, "Testing the sensitivity of terahertz spectroscopy to changes in molecular and supramolecular structure: A study of structurally similar cocrystals," *Crystal Growth and Design*, vol. 9, no. 3, pp. 1452–1460, 2009.
- [16] D. Singhal and W. Curatolo, "Drug polymorphism and dosage form design: a practical perspective," *Advanced Drug Delivery Reviews*, vol. 56, no. 3, pp. 335–347, 2004.
- [17] W. H. Fan and A. D. Burnett, "Far-infrared spectroscopic characterization of explosives for security applications using broadband terahertz time-domain spectroscopy," *Applied Spectroscopy*, vol. 61, no. 6, pp. 638–643, 2007.
- [18] D. Banerjee, W. von Spiegel, M. Thomson, S. Schabel, and H. Roskos, "Diagnosing water content in paper by terahertz radiation," *Optics Express*, vol. 16, no. 12, pp. 9060–9066, 2008.
- [19] A. Finch, "Chemical applications of far infrared spectroscopy," 1970.
- [20] D. Auston, K. Cheung, and P. Smith, "Picosecond photoconducting hertzian dipoles," *Applied physics letters*, vol. 45, no. 3, pp. 284–286, 1984.
- [21] L. Xu, X. Zhang, and D. H. Auston, "Terahertz beam generation by femtosecond optical pulses in electro-optic materials," *Applied Physics Letters*, vol. 61, no. 15, pp. 1784–1786, 1992.
- [22] Q. Wu and X.-C. Zhang, "Free-space electro-optic sampling of terahertz beams," *Applied Physics Letters*, vol. 67, no. 24, pp. 3523–3525, 1995.
- [23] B. Hu and M. Nuss, "Imaging with terahertz waves," *Optics letters*, vol. 20, no. 16, pp. 1716–1718, 1995.
- [24] C. Janke, M. Först, M. Nagel, H. Kurz, and A. Bartels, "Asynchronous optical sampling for high-speed characterization of integrated resonant terahertz sensors," *Optics Letters*, vol. 30, no. 11, pp. 1405–1407, 2005.
- [25] M. Tani, P. Gu, M. Hyodo, K. Sakai, and T. Hidaka, "Generation of coherent terahertz radiation by photomixing of dual-mode lasers," *Optical and Quantum Electronics*, vol. 32, no. 4, pp. 503–520, 2000.
- [26] K. Kawase, J. Shikata, and H. Ito, "Terahertz wave parametric source," *Journal of Physics D: Applied Physics*, vol. 35, no. 3, 2002.

- [27] N. Karpowicz, H. Zhong, C. Zhang, K. Lin, J. Hwang, J. Xu, and X. Zhang, "Compact continuous-wave subterahertz system for inspection applications," *Applied Physics Letters*, vol. 86, no. 5, p. 054 105, 2005.
- [28] E. R. Mueller, "Terahertz radiation sources for imaging and sensing applications," *Photonics Spectra*, vol. 40, no. 11, p. 60, 2006.
- [29] R. Kohler, A. Tredicucci, F. Beltram, H. E. Beere, E. H. Linfield, A. G. Davies, D. A. Ritchie, R. C. Iotti, and F. Rossi, "Terahertz semiconductor-heterostructure laser," *Nature*, vol. 417, no. 6885, pp. 156–159, 2002.
- [30] G. L. Carr, M. C. Martin, W. R. McKinney, K. Jordan, G. R. Neil, and G. P. Williams, "High-power terahertz radiation from relativistic electrons," *Nature*, vol. 420, no. 6912, p. 153, 2002.
- [31] A. J. Kreisler and A. Gaugue, "Recent progress in high-temperature superconductor bolometric detectors: From the mid-infrared to the far-infrared (thz) range," *Superconductor Science and Technology*, vol. 13, no. 8, p. 1235, 2000.
- [32] M. J. Golay, "A pneumatic infra-red detector," *Review of Scientific Instruments*, vol. 18, no. 5, pp. 357–362, 1947.
- [33] F. Keilmann, C. Gohle, and R. Holzwarth, "Time-domain mid-infrared frequency-comb spectrometer," *Optics Letters*, vol. 29, no. 13, pp. 1542–1544, 2004.
- [34] P. U. Jepsen, D. G. Cooke, and M. Koch, "Terahertz spectroscopy and imaging—modern techniques and applications," *Laser & Photonics Reviews*, vol. 5, no. 1, pp. 124–166, 2011.
- [35] M. Tani, M. Herrmann, and K. Sakai, "Generation and detection of terahertz pulsed radiation with photoconductive antennas and its application to imaging," *Measurement Science and Technology*, vol. 13, no. 11, p. 1739, 2002.
- [36] Y. C. Shen, P. C. Upadhyaya, E. H. Linfield, H. E. Beere, and A. G. Davies, "Ultrabroadband terahertz radiation from low-temperature-grown gas photoconductive emitters," *Applied Physics Letters*, vol. 83, no. 15, p. 3117, 2003.
- [37] J. T. Darrow, X. Zhang, D. H. Auston, and J. D. Morse, "Saturation properties of large-aperture photoconducting antennas," *IEEE Journal of Quantum Electronics*, vol. 28, no. 6, pp. 1607–1616, 1992.
- [38] E. Castro-Camus, J. Lloyd-Hughes, M. Johnston, M. Fraser, H. Tan, and C. Jagadish, "Polarization-sensitive terahertz detection by multicontact photoconductive receivers," *Applied Physics Letters*, vol. 86, no. 25, p. 254 102, 2005.
- [39] D. Dykaar, B. Greene, J. Federici, A. Levi, L. Pfeiffer, and R. Kopf, "Log-periodic antennas for pulsed terahertz radiation," *Applied Physics Letters*, vol. 59, no. 3, pp. 262–264, 1991.
- [40] P. Kirawanich, S. J. Yakura, and N. E. Islam, "Study of high-power wide-band terahertz-pulse generation using integrated high-speed photoconductive semiconductor switches," *IEEE Transactions on Plasma Science*, vol. 37, no. 1, pp. 219–228, 2009.
- [41] G. Matthaus, S. Nolte, R. Hohmuth, M. Voitsch, W. Richter, B. Pradarutti, S. Riehemann, G. Notni, and A. Tunnermann, "Microlens coupled interdigital photoconductive switch," *Applied Physics Letters*, vol. 93, no. 9, p. 091 110, 2008.

- [42] S. Winnerl, F. Peter, S. Nitsche, A. Dreyhaupt, B. Zimmermann, M. Wagner, H. Schneider, M. Helm, and K. Kohler, "Generation and detection of thz radiation with scalable antennas based on gaas substrates with different carrier lifetimes," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 14, no. 2, pp. 449–457, 2008.
- [43] M. R. Stone, M. Naftaly, R. E. Miles, J. R. Fletcher, and D. P. Steenson, "Electrical and radiation characteristics of semilarge photoconductive terahertz emitters," *IEEE Transactions on Microwave Theory and Techniques*, vol. 52, no. 10, pp. 2420–2429, 2004.
- [44] D. S. Kim and D. Citrin, "Dynamics of electric field screening in photoconductive thz sources with spatially patterned excitation," in *Gallium Arsenide and Other Semiconductor Application Symposium, 2005. EGAAS 2005. European*, IEEE, 2005, pp. 417–420.
- [45] L. Hou and W. Shi, "An It-gaas terahertz photoconductive antenna with high emission power, low noise, and good stability," *IEEE Transactions on Electron Devices*, vol. 60, no. 5, pp. 1619–1624, 2013.
- [46] D. R. Bacon, A. D. Burnett, M. Swithenbank, C. Russell, L. Li, C. D. Wood, A. G. Davies, P. Dean, and J. R. Freeman, "Free-space terahertz radiation from a emitter," vol. 24, no. 23, pp. 2779–2781, 2016.
- [47] S. O. Kasap, *Principles of Electronic Materials and Devices*, 3rd ed. McGraw-Hill New York, 2007.
- [48] Q. Wu and X. Zhang, "Ultrafast electro-optic field sensors," *Applied Physics Letters*, vol. 68, no. 12, p. 1604, 1996.
- [49] Q. Wu and X. Zhang, "Design and characterization of traveling-wave electrooptic terahertz sensors," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 2, no. 3, pp. 693–700, 1996.
- [50] P. Y. Han and X. Zhang, "Coherent, broadband midinfrared terahertz beam sensors," *Applied Physics Letters*, vol. 73, no. 21, p. 3049, 1998.
- [51] H. J. Bakker, G. C. Cho, H. Kurz, Q. Wu, and X. Zhang, "Distortion of terahertz pulses in electro-optic sampling," *Journal of the Optical Society of America B*, vol. 15, no. 6, p. 1795, 1998.
- [52] J. Kroll, J. Darmo, and K. Unterrainer, "High-performance terahertz electro-optic detector," *Electronics Letters*, vol. 40, no. 12, pp. 763–764, 2004.
- [53] C. Kubler, R. Huber, S. Tubel, and A. Leitenstorfer, "Ultrabroadband detection of multi-terahertz field transients with gase electro-optic sensors: approaching the near infrared," *Applied Physics Letters*, vol. 85, no. 16, p. 3360, 2004.
- [54] Q. Wu and X. Zhang, "7 terahertz broadband gap electro-optic sensor," *Applied Physics Letters*, vol. 70, no. 14, p. 1784, 1997.
- [55] Q. Wu, M. Litz, and X. Zhang, "Broadband detection capability of znTe electro-optic field detectors," *Applied Physics Letters*, vol. 68, no. 21, p. 2924, 1996.
- [56] U. Keller, "Recent developments in compact ultrafast lasers," *Nature*, vol. 424, no. 6950, pp. 831–838, 2003.
- [57] M. Tani, S. Matsuura, K. Sakai, and S. Nakashima, "Emission characteristics of photoconductive antennas based on low-temperature-grown gaas and semi-insulating gaas," *Applied Optics*, vol. 36, no. 30, pp. 7853–7859, 1997.

- [58] P. F. Moulton, "Spectroscopic and laser characteristics of $\text{Ti:Al}_2\text{O}_3$," *Journal of the Optical Society of America B*, vol. 3, no. 1, pp. 125–133, 1986.
- [59] M. V. Exter, C. Fattinger, and D. Grischkowsky, "Terahertz time-domain spectroscopy of water vapor," *Optics Letters*, vol. 14, no. 20, pp. 1128–1130, 1989.
- [60] L. Duvillaret, F. Garet, and J. Coutaz, "A reliable method for extraction of material parameters in terahertz time-domain spectroscopy," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 2, no. 3, pp. 739–746, 1996.
- [61] L. Duvillaret, F. Garet, and J. Coutaz, "Influence of noise on the characterization of materials by terahertz time-domain spectroscopy," *Journal of the Optical Society of America B*, vol. 17, no. 3, p. 452, 2000.
- [62] T. Dorney, R. Baraniuk, and D. Mittleman, "Material parameter estimation with terahertz time-domain spectroscopy," *Journal of the Optical Society of America A*, vol. 18, no. 7, pp. 1562–1571, 2001.
- [63] I. Pupeza, R. Wilk, and M. Koch, "Highly accurate optical material parameter determination with thz time-domain spectroscopy," *Optics Express*, vol. 15, no. 7, pp. 4335–4350, 2007.
- [64] R. Wilk and I. Pupeza, "Highly accurate thz time-domain spectroscopy of multilayer structures," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 14, no. 2, pp. 392–398, 2008.
- [65] M. Scheller, "Real-time terahertz material characterization by numerical three-dimensional optimization," *Optics Express*, vol. 19, no. 11, pp. 762–769, 2011.
- [66] J. A. Hejase, E. J. Rothwell, and P. Chahal, "A multiple angle method for thz time-domain material characterization," *IEEE Transactions on Terahertz Science and Technology*, vol. 3, no. 5, pp. 656–665, Sep. 2013.
- [67] X. Li, Z. Hong, J. He, and Y. Chen, "Precisely optical material parameter determination by time domain waveform rebuilding with thz time-domain spectroscopy," *Optics Communications*, vol. 283, no. 23, pp. 4701–4706, 2010.
- [68] O. S. Ahmed, M. A. Swillam, M. H. Bakr, and X. Li, "Efficient optimization approach for accurate parameter extraction with terahertz time-domain spectroscopy," *Journal of Lightwave Technology*, vol. 28, no. 11, pp. 1685–1692, 2010.
- [69] G. P. Kniffin and L. M. Zurk, "Model-based material parameter estimation for terahertz reflection spectroscopy," *IEEE Transactions on Terahertz Science and Technology*, vol. 2, no. 2, pp. 231–241, 2012.
- [70] D. Manolakis, V. Ingle, and S. Kogon, *Statistical and Adaptive Signal Processing*. 2000.
- [71] S. Krimi, J. Klier, J. Jonscheit, G. von Freymann, R. Urbansky, and R. Beigang, "Highly accurate thickness measurement of multi-layered automotive paints using terahertz technology," *Applied Physics Letters*, vol. 109, no. 2, p. 021105, 2016.
- [72] T. Yasui, T. Yasuda, K. Sawanaka, and T. Araki, "Terahertz paintmeter for non-contact monitoring of thickness and drying progress in paint film," *Applied Optics*, vol. 44, no. 32, pp. 6849–6856, 2005.

- [73] C. Y. Jen and C. Richter, "Sample thickness measurement with thz-tds: resolution and implications," *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 35, no. 10, pp. 840–859, 2014.
- [74] J. Dong, X. Wu, A. Locquet, and D. S. Citrin, "Terahertz superresolution stratigraphic characterization of multilayered structures using sparse deconvolution," *IEEE Transactions on Terahertz Science and Technology*, 2017.
- [75] P. Harrison and A. Valavanis, *Quantum Wells, Wires and Dots: Theoretical and Computational Physics of Semiconductor Nanostructures*, 4th ed. John Wiley & Sons, 2016.
- [76] A. J. Baragwanath, G. P. Swift, D. Dai, A. J. Gallant, and J. M. Chamberlain, "Silicon based microfluidic cell for terahertz frequencies," *Journal of Applied Physics*, vol. 108, no. 1, p. 013 102, 2010.
- [77] M. Naftaly and R. Dudley, "Methodologies for determining the dynamic ranges and signal-to-noise ratios of terahertz time-domain spectrometers.," *Optics letters*, vol. 34, no. 8, pp. 1213–1215, 2009.
- [78] R. M. Woodward, B. E. Cole, V. P. Wallace, R. J. Pye, D. D. Arnone, E. H. Linfield, and M. Pepper, "Terahertz pulse imaging in reflection geometry of human skin cancer and skin tissue," *Physics in Medicine and Biology*, vol. 47, no. 21, p. 3853, 2002.
- [79] B. Ferguson and D. Abbott, "De-noising techniques for terahertz responses of biological samples," *Microelectronics Journal*, vol. 32, no. 12, pp. 943–953, 2001.
- [80] X. Qiao, X. Zhang, J. Ren, D. Zhang, G. Cao, and L. Li, "Mean estimation empirical mode decomposition method for terahertz time-domain spectroscopy de-noising," *Applied Optics*, vol. 56, no. 25, pp. 7138–7145, 2017.
- [81] W. Withayachumnankul, B. M. Fischer, H. Lin, and D. Abbott, "Uncertainty in terahertz time-domain spectroscopy measurement," *Journal of the Optical Society of America B*, vol. 25, no. 6, p. 1059, 2008.
- [82] M. Krüger, S. Funkner, E. Bründermann, and M. Havenith, "Uncertainty and ambiguity in terahertz parameter extraction and data analysis," *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 32, no. 5, pp. 699–715, 2011.
- [83] Joint Committee for Guides in Metrology (JCGM), "2011 evaluation of measurement data – supplement 2 to the "guide to the expression of uncertainty in measurement" – extension to any number of output quantities," 2011.
- [84] M. Bernier, F. Garet, J. Coutaz, H. Minamide, and A. Sato, "Accurate characterization of resonant samples in the terahertz regime through a technique combining time-domain spectroscopy and kramers-kronig analysis," *IEEE Transactions on Terahertz Science and Technology*, vol. 6, no. 3, 2006.
- [85] H. Son, D. Choi, and G. Park, "Improved thickness estimation of liquid water using kramers–kronig relations for determination of precise optical parameters in terahertz transmission spectroscopy," *Optics Express*, vol. 25, no. 4, pp. 4509–4518, 2017.
- [86] V. Lucarini, Y. Ino, K.-E. Peiponen, and M. Kuwata-Gonokami, "Detection and correction of the misplacement error in terahertz spectroscopy by application of singly subtractive kramers-kronig relations," *Physical Review B*, vol. 72, no. 12, p. 125 107, 2005.

- [87] H. Tuononen, E. Gornov, J. A. Zeitler, J. Aaltonen, and K.-E. Peiponen, "Using modified kramers–kronig relations to test transmission spectra of porous media in thz-tds," *Optics letters*, vol. 35, no. 5, pp. 631–633, 2010.
- [88] K.-E. Peiponen, E. Gornov, Y. Svirko, Y. Ino, M. Kuwata-Gonokami, and V. Lucarini, "Testing the validity of terahertz reflection spectra by dispersion relations," *Physical Review B*, vol. 72, no. 24, p. 245 109, 2005.
- [89] K.-E. Peiponen, E. M. Vartiainen, T. Unuma, J. Axel Zeitler, P. Silfsten, T. Venäläinen, and H. Kishida, "Dispersion relations for evaluating the complex refractive index of medium without the information of its thickness," *Applied Physics Letters*, vol. 102, no. 18, p. 181 110, 2013.
- [90] K. M. Tych, C. D. Wood, and W. Tych, "A simple transfer-function-based approach for estimating material parameters from terahertz time-domain data," *IEEE Photonics Journal*, vol. 6, no. 1, pp. 1–11, 2014.
- [91] A. D. Burnett, W. Fan, P. C. Upadhyaya, J. E. Cunningham, M. D. Hargreaves, T. Munshi, H. G. Edwards, E. H. Linfield, and A. G. Davies, "Broadband terahertz time-domain spectroscopy of drugs-of-abuse and the use of principal component analysis," *The Analyst*, vol. 134, pp. 1658–1668, 2009.
- [92] S. Hadjiloucas, R. K. Galvão, and J. W. Bowen, "Analysis of spectroscopic measurements of leaf water content at terahertz frequencies using linear transforms," *Journal of the Optical Society of America A*, vol. 19, no. 12, pp. 2495–2509, 2002.
- [93] J. S. Orfanidis, *Introduction to Signal Processing*. Prentice-Hall, Inc., 1996.
- [94] D. M. Pozar, *Microwave engineering*. John Wiley & Sons, 2009.
- [95] M. Walther, B. Fischer, M. Schall, H. Helm, and P. U. Jepsen, "Far-infrared vibrational spectra of all-trans, 9-cis and 13-cis retinal measured by thz time-domain spectroscopy," *Chemical Physics Letters*, vol. 332, no. 3, pp. 389–395, 2000.
- [96] M. Franz, B. Fischer, and M. Walther, "The christiansen effect in terahertz time-domain spectra of coarsegrained powders," *Applied Physics Letters*, vol. 92, 2008.
- [97] A. D. Burnett, J. Kendrick, J. E. Cunningham, M. D. Hargreaves, T. Munshi, H. G. M. Edwards, E. H. Linfield, and A. G. Davies, "Calculation and measurement of terahertz active normal modes in crystalline petn," *ChemPhysChem*, vol. 11, pp. 368–378, 2010.
- [98] J. Kendrick and A. D. Burnett, "Pdielec: the calculation of infrared and terahertz absorption for powdered crystals," *Journal of Computational Chemistry*, vol. 37, no. 16, pp. 1491–1504, 2016.
- [99] M. Naftaly, "Metrology issues and solutions in thz time domain spectroscopy: noise, errors, calibration," *IEEE Sensors Journal*, vol. 13, no. 1, pp. 8–17, 2012.
- [100] P. Kuzel, H. Nemeč, F. Kadlec, and C. Kadlec, "Gouy shift correction for highly accurate refractive index retrieval in time-domain terahertz spectroscopy," *Optics Express*, vol. 18, no. 15, pp. 15 338–15 348, 2010.
- [101] V. Apostolopoulos, G. Daniell, and A. Chung, "Complex refractive index determination using planar and converging beam transfer functions," in *Terahertz Spectroscopy and Imaging*, Springer, 2009.

- [102] B. Fischer, "Dynamic range and numerical error propagation in terahertz time-domain spectroscopy," in *Optical Terahertz Science and Technology*, Optical Society of America, 2005, pp. 14–16.
- [103] M. Naftaly and R. E. Miles, "Terahertz time-domain spectroscopy for material characterization," *Proceedings of the IEEE*, vol. 95, no. 8, p. 1658, 2007.
- [104] S. Park, M. Melloch, and A. Weiner, "Analysis of terahertz waveforms measured by photoconductive and electrooptic sampling," *IEEE Journal of Quantum Electronics*, vol. 35, no. 5, pp. 810–819, May 1999.
- [105] L. M. Zurk, B. Orłowski, D. P. Winebrenner, E. I. Thorsos, M. R. Leahy-Hoppa, and L. M. Hayden, "Terahertz scattering from granular material," *Journal of the Optical Society of America B*, vol. 24, no. 9, pp. 2238–2243, 2007.
- [106] P. U. Jepsen and B. M. Fischer, "Dynamic range in terahertz time-domain transmission and reflection spectroscopy," *Optics Letters*, vol. 30, no. 1, p. 29, 2005.
- [107] W. Withayachumnankul and M. Naftaly, "Fundamentals of measurement in terahertz time-domain spectroscopy," *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 35, no. 8, pp. 610–637, 2014.
- [108] J. A. Zeitler, K. Kogermann, J. Rantanen, T. Rades, P. F. Taday, M. Pepper, J. Aaltonen, and C. J. Strachan, "Drug hydrate systems and dehydration processes studied by terahertz pulsed spectroscopy," *International Journal of Pharmaceutics*, vol. 334, pp. 78–84, 2007.
- [109] E. R. Brown, J. E. Bjarnason, M. Fedor, and T. M. Korter, "On the strong and narrow absorption signature in lactose at 0.53 thz," *Applied Physics Letters*, vol. 90, p. 061 908, 2007.
- [110] K. Peiponen and J. J. Saarinen, "Generalized kramers-kronig relations in nonlinear optical- and thz-spectroscopy," *Reports on Progress in Physics*, vol. 72, p. 056 401, 2009.
- [111] J. Sibik, M. J. Sargent, M. Franklin, and J. A. Zeitler, "Crystallization and phase changes in paracetamol from the amorphous solid to the liquid phase.," *Molecular Pharmaceutics*, vol. 11, no. 4, pp. 1326–1334, 2014.
- [112] E. Garcia-Garcia, E. Diez, Y. M. Meziani, J. E. Velazquez-Perez, and J. Calvo-Gallcao, "Terahertz time domain spectroscopy for chemical identification," *IEEE 2013 Spanish Conference on Electron Devices*, pp. 199–202, 2013.
- [113] M. Kaushik, B. Ng, B. Fischer, and D. Abbott, "Reduction of scattering effects in thz-tds signals," *IEEE Photonic Technology*, vol. 24, pp. 155–157, 2012.
- [114] D. Molter, M. Trierweiler, F. Ellrich, J. Jonuschiet, and G. von Frequmann, "Interferometry-aided terahertz time-domain spectroscopy," *Optics Express*, vol. 25, no. 7, 2017.
- [115] A. Rehn, D. Jahn, J. C. Balzer, and M. Koch, "Periodic sampling errors in terahertz time-domain measurements," *Optics Express*, vol. 25, no. 6, pp. 6712–6724, Mar. 2017.
- [116] M. Scheller, C. Jansen, and M. Koch, "Analyzing sub 100 um samples with transmission terahertz time domain spectroscopy," *Optics Communications*, vol. 282, no. 7, pp. 1304–1306, 2009.

- [117] W. Withayachumnankul, B. Fischer, and D. Abbott, "Material thickness optimization for transmission-mode terahertz time-domain spectroscopy," *Optics express*, vol. 16, pp. 7382–7396, 2008.
- [118] N. R. Greenall, A. Valavanis, H. J. Desai, D. O. Acheampong, L. H. Li, J. E. Cunningham, A. G. Davies, E. H. Linfield, and A. Burnett, "The development of a semtex-h simulat for terahertz spectroscopy," *Journal of Infrared, Millimeter and Terahertz Waves*, 2016.
- [119] M. Kemp, P. Taday, B. Cole, J. Cluff, A. Fitzgerald, and W. Tribe, "Security applications of terahertz technology," *AeroSense 2003*, vol. 5070, pp. 44–52, 2003.
- [120] A. D. Burnett, J. Kendrick, C. Russell, J. Christensen, J. E. Cunningham, A. R. Pearson, E. H. Linfield, and A. G. Davies, "Effect of molecular size and particle shape on the terahertz absorption of a homologous series of tetraalkylammonium salts," *Analytical Chemistry*, vol. 85, pp. 7926–7934, 2013.
- [121] J. C. M. Garnett, "Colours in metal glasses, in metallic films, and in metallic solutions. ii," *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, vol. 205, pp. 237–288, 1906.
- [122] M. Scheller, C. Jansen, and M. Koch, "Applications of effective medium theories in the terahertz regime," in *Recent Optical and Photonic Technologies*, InTech, 2010.
- [123] J. Hageman, R. Wehrens, R. de Gelder, W. Leo Meerts, and L. Buydens, "Direct determination of molecular constants from rovibronic spectra with genetic algorithms," *Journal of Chemical Physics*, vol. 113, pp. 7955–7962, 2000.
- [124] M. Mitchell, *An Introduction to Genetic Algorithms*. MIT press, 1998.
- [125] *How the genetic algorithm works*, <http://uk.mathworks.com/help/gads/how-the-genetic-algorithm-works.html>, 2017.
- [126] K. Deep, K. P. Singh, M. L. Kansal, and C. Mohan, "A real coded genetic algorithm for solving integer and mixed integer optimization problems," *Applied Mathematics and Computation*, vol. 212, no. 2, pp. 505–518, 2009.
- [127] A. Conn, N. Gould, and P. Toint, "A globally convergent lagrangian barrier algorithm for optimization with general inequality constraints and simple bounds," *Mathematics of Computation of the American Mathematical Society*, vol. 66, no. 217, pp. 261–288, 1997.
- [128] *Matlab and the global optimisation toolbox release 2015b*, 2015.
- [129] N. R. Greenall, L. H. Li, E. H. Linfield, A. G. Davies, J. E. Cunningham, and A. D. Burnett, "Multilayer extraction of complex refractive index in broadband transmission terahertz time-domain spectroscopy," in *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2016 41st International Conference on*, IEEE, 2016, pp. 1–2.
- [130] N. R. Greenall, C. D. Wood, C. Russell, L. H. Li, E. H. Linfield, A. G. Davies, J. E. Cunningham, and A. D. Burnett, "Accurate material parameter extraction from broadband terahertz spectroscopy," in *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2015 40th International Conference on*, IEEE, 2015, pp. 1–2.

- [131] M. Bernier, F. Garet, and J. Coutaz, "Determining the complex refractive index of materials in the far-infrared from terahertz time-domain data," in *Terahertz Spectroscopy - A cutting Edge Technology*, InTech, 2017.
- [132] K. Levenberg, "A method for the solution of certain non-linear problems in least squares," *Quarterly of Applied Mathematics*, vol. 2, no. 2, pp. 164–168, 1944.
- [133] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *Journal of the society for Industrial and Applied Mathematics*, vol. 11, no. 2, pp. 431–441, 1963.
- [134] O. Sushko, K. Shala, R. Dubrovka, and R. Donnan, "Revised metrology for enhanced accuracy in complex optical constant determination by thz-time-domain spectrometry," *Journal for the Optical Society of America A*, vol. 30, no. 5, pp. 979–86, 2013.
- [135] M. T. Reiten, S. A. Harmon, and R. A. Cheville, "Terahertz beam propagation measured through three-dimensional amplitude profile determination," *Journal for the Optics Society of America B*, vol. 20, no. 10, pp. 2215–2225, 2003.
- [136] E. P. Parrott, J. A. Zeitler, and L. F. Gladden, "Accurate determination of optical coefficients from chemical samples using terahertz time-domain spectroscopy and effective medium theory," *Optical Letters*, vol. 34, pp. 3722–3724, 2009.
- [137] Joint Committee for Guides in Metrology (JCGM), "Evaluation of measurement data: guide to the expression of uncertainty in measurement," 2008.
- [138] N. R. Greenall, E. H. Linfield, A. G. Davies, L. H. Li, J. E. Cunningham, and A. D. Burnett, "Estimation of spectroscopic uncertainty and correlation in terahertz time domain spectroscopy," in *Infrared, Millimeter, and Terahertz waves (IRMMW-THz), 2016 41st International Conference on*, IEEE, 2016, pp. 1–2.
- [139] R. A. Johnson and D. W. Wicharn, *Applied Multivariate Statistical Analysis*, 5th ed. Prentice-Hall New Jersey.
- [140] J. Son, J. V. Rudd, and J. F. Whitaker, "Noise characterization of a self-mode-locked ti:sapphire laser," *Optics Letters*, vol. 17, no. 10, pp. 733–735, 1992.
- [141] D. Jahn, S. Lippert, M. Bisi, L. Oberto, J. C. Balzer, and M. Koch, "On the influence of delay line uncertainty in thz time-domain spectroscopy," *Journal of Infrared, Millimeter, and Terahertz Waves*, vol. 37, no. 6, pp. 605–613, 2016.
- [142] N. Cohen, J. W. Handley, R. D. Boyle, S. L. Braunstein, and E. Berry, "Experimental signature of registration noise in pulsed terahertz systems," *Fluctuation and Noise Letters*, vol. 06, no. 01, pp. L77–L84, 2006.
- [143] B. Efron, "Bootstrap methods: Another look at the jackknife," in *Breakthroughs in Statistics: Methodology and Distribution*, Springer New York, 1992.
- [144] P. U. Jepsen, U. Moller, and H. Merbold, "Investigation of aqueous alcohol and sugar solutions with reflection terahertz time-domain spectroscopy," *Optics Express*, vol. 15, no. 22, pp. 14 717–14 737, 2007.
- [145] O. Hirsch, P. Alexander, and L. F. Gladden, "Techniques for cancellation of interfering multiple reflections in terahertz time-domain measurements," *Microelectronics Journal*, vol. 39, no. 5, pp. 841–848, 2008.

- [146] L. Wang, Y. Liu, Y. Wang, and Z. Zhao, "A method for removing echoes in the terahertz time-domain spectroscopy system," *IEEE International Conference on Microwave and Millimeter Wave Technology*, vol. 1, 2012.
- [147] D. G. Childers, D. P. Skinner, and R. C. Kemerait, "The cepstrum: A guide to processing," *Proceedings of the IEEE*, vol. 65, no. 10, pp. 1428–1443, 1977.
- [148] T. Lee, "Independent component analysis," in *Independent Component Analysis*, Springer, 1998.
- [149] *Lactose pubchem record*, <https://pubchem.ncbi.nlm.nih.gov/compound/440995>.
- [150] *Ptfe manufacturing specifications*, http://www.boedeker.com/ptfe_p.htm.
- [151] P. D. Cunningham, N. N. Valdes, F. A. Vallejo, L. M. Hayden, B. Polishak, J. Zhou Xingand Luo, A. K. Jen, J. C. Williams, and R. J. Twieg, "Broadband terahertz characterization of the refractive index and absorption of some important polymeric and organic electro-optic materials," *Journal of Applied Physics*, vol. 109, no. 4, pp. 043 505–043 505, 2011.
- [152] Y. Rozanov, *Probability Theory: A Concise Course*. Courier Corporation.
- [153] B. D. Hall, "On the propagation of uncertainty in complex-valued quantities," *Metrologia*, vol. 41, no. 3, pp. 173–177, 2004.
- [154] P. J. Schreier and L. L. Scharf, "Second-order analysis of improper complex random vectors and processes," *IEEE Transactions on Signal Processing*, vol. 51, no. 3, pp. 714–725, 2003.
- [155] B. Picinbono, "Second-order complex random vectors and normal distributions," *IEEE Transactions on Signal Processing*, vol. 44, no. 10, pp. 2637–2640, 1996.
- [156] F. Neeser and J. Massey, "Proper complex random processes with applications to information and theory," *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1293–1302, 1993.