

UNITED KINGDOM · CHINA · MALAYSIA

Low Bandwidth Laser Doppler Blood Flowmetry

James Carpenter, MEng.

Thesis submitted to the University of Nottingham for the degree of Doctor of Philosophy

July 2014

Abstract

Laser Doppler Blood Flowmetry (LDBF) has for several decades been applied to measure the flow of red blood cells in living tissue. Laser Doppler Perfusion Imaging (LDPI), a recent advancement which enables full-field blood flow visualisation, is gaining clinical acceptance in fields such as burn diagnostics.

However, video-rate full-field imagers with appropriate sensor and processing capability require large financial and physical resources and this has prompted the development of under-specified systems. These systems may reduce the bandwidth and processing complexity but the question of how they perform compared to their fully specified counterparts remains. The advantages of these cheaper and often highly reconfigurable systems are recognised and so it is beneficial to ask whether any novel processing schemes can reduce the resultant error.

Here a reduced bandwidth LDBF signal processing system has been modelled. Bayesian Inference has been used to show that the Pareto distribution is a likely model for the LDBF power spectrum, despite often being cited as exponential. Methods of evaluating microvascular blood flow have been described and compared. Additionally, one fast algorithm's effectiveness has been explained, and a novel and accurate method using the Hilbert transform has been presented.

By understanding how aliasing modifies the frequency distribution, Bayesian Inference has been used to correct the blood flow output towards gold-standard values. The technique has been shown to correct the output of a low bandwidth CMOS camera imaging a rotating diffuser.

Low bandwidth LDPI systems may be suitable for certain clinical applications where sensitivity to high flow is not required. However, where sensitivity to higher flow than baseline is required, e.g. in burn diagnostics, low bandwidth systems may underestimate the true blood flow leading to misdiagnosis. Nevertheless, low bandwidth systems could be used in this scenario if reliable post-processing is employed, such as that suggested by this thesis.

Acknowledgements

Firstly, I would like all those who helped me complete this work. I particularly want to thank my supervisors, Professors John Crowe and Barrie Hayes-Gill, who have guided me through the last few years and kept me questioning and pursuing this research. I would like to thank David for the many conversations and for helping me set up the optical experiments, and also Shen, for letting me analyse the results from his camera.

I would like to acknowledge and thank all those staff members in the Faculty of Engineering who helped me with ethics, administration and technical expertise.

I also thank those whom I work alongside, who have helped make it a fun few years, and my other friends and family who have been there with me.

Finally I would like to thank my wife, Alyssa, for all her help, support and patience. I could not have done this without her and I dedicate this work to her.

Contents

	Abs	stract		ii		
	Ack	knowledgements				
	List	of Sy	mbols	viii		
	List	of Ab	breviations	ix		
1	Intr	oducti	on	1		
	1.1	Introd	uction	. 1		
	1.2	Laser	Doppler Blood Flowmetry	. 2		
		1.2.1	Background	. 2		
		1.2.2	Light and Tissue	. 3		
		1.2.3	Engineering	. 13		
		1.2.4	A Laser Doppler System	. 21		
		1.2.5	Clinical Applications	. 24		
		1.2.6	Review	. 27		
	1.3	Resear	cch Rationale	. 28		
		1.3.1	State of the Art	. 28		
		1.3.2	Present Research	. 30		
	1.4	Thesis		. 30		
		1.4.1	Aims & Outline	. 30		
		1.4.2	Scope	. 31		
	1.5	Conclu	ision	. 32		
2	The	e First	Moment of the LDBF Photocurrent's Powe	r		
	Spe	ctrum		33		
	2.1	Introd	uction	. 33		
	2.2	The P	hotocurrent	. 34		
		2.2.1	Electric Field	. 34		
		2.2.2	Sample Rates	. 36		

	2.3	Evalua	ating the First Moment	38	
		2.3.1	Frequency Domain Definition	38	
		2.3.2	Frequency Domain Implementation	38	
		2.3.3	Time Domain Equivalence	43	
		2.3.4	Time Domain Implementation	44	
	2.4	Concl	usion \ldots	57	
3	Mo	delling	g the Power Spectrum of the LDBF Photocur-		
	rent	t by B	ayesian Inference	60	
	3.1	Introd	luction	60	
	3.2	Exper	imental Design	61	
		3.2.1	Data Requirements	61	
		3.2.2	Ethical Considerations	61	
		3.2.3	Recording Devices	61	
		3.2.4	In Vivo Data Collection	62	
	3.3	Model	ls	63	
	3.4	Bayes	ian Model Selection	66	
	3.5 Power Spectra Conditioning				
		3.5.1	Overview	70	
		3.5.2	LDBF Theoretical Spectrum	71	
		3.5.3	Properties of Spectral Estimates of White Noise .	71	
		3.5.4	The Laser Doppler Spectrum as Filtered White		
			Noise	73	
		3.5.5	Summary	82	
	3.6	Furthe	er Bayesian Derivation	82	
		3.6.1	Likelihood function	82	
		3.6.2	Model distribution	83	
		3.6.3	Numerical Marginalisation	84	
	3.7	Data I	Processing	86	
		3.7.1	Raw Data Files	86	
		3.7.2	Averaging	95	

		3.7.3	Fitting Process
	3.8	Model	Fit Results
	3.9	Creati	on of Photocurrent
	3.10	Conclu	usion
4	A C	ompar	rison of Processing Algorithms for LDBF 113
	4.1	Introd	uction
	4.2	Theore	etical Accuracy
		4.2.1	Pareto Distribution
		4.2.2	Methodology $\ldots \ldots 115$
		4.2.3	First Moment Results
		4.2.4	LASCA Results
	4.3	Experi	imental Recordings
		4.3.1	Data
		4.3.2	Methodology
		4.3.3	Results
	4.4	Conclu	usion
5	Low	Band	width Correction 154
	5.1	Introd	uction \ldots \ldots \ldots \ldots \ldots \ldots \ldots 154
	5.2	An Int	egrating CMOS Imager
		5.2.1	Overview
		5.2.2	Integrating Pixel
	5.3	Aliasir	ng
		5.3.1	Model
		5.3.2	Model Test
		5.3.3	Anti-Aliasing Comparison
	5.4	Correc	etion using Bayesian Inference
		5.4.1	Bayes' Theorem
		5.4.2	Updated Model
		5.4.3	Methodology

		5.4.4 Results \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots	. 177		
		5.4.5 Implementation	. 185		
	5.5	Conclusion	. 187		
6	$\mathbf{C}\mathbf{M}$	IOS Camera Evaluation	189		
	6.1	Introduction	. 189		
	6.2	CMOS Sensor & FPGA	. 191		
		6.2.1 Overview	. 191		
		6.2.2 Theoretical Response of Diffusing Surface	. 191		
		6.2.3 Experimental Data	. 196		
		6.2.4 Comparison & Correction	. 199		
	6.3	Optronis CL600	. 205		
		6.3.1 Overview	. 205		
		6.3.2 Ethical Considerations	. 205		
		6.3.3 Experimental Method	. 206		
		6.3.4 Results	. 210		
		6.3.5 Discussion	. 213		
	6.4	Conclusion	. 215		
7	Cor	nclusion	216		
	7.1	Introduction	. 216		
	7.2	Summary of Results	. 217		
	7.3	Future Work	. 219		
	7.4	Conclusion	. 220		
	7.5	Summary of Findings	. 221		
Appendix A Derivation of Absolute and Average 222					
Appendix BReduced Bandwidth Images224					
\mathbf{A}	Appendix CCorrected Flow Values241				
R	efere	ences	245		

List of Symbols

Name	Continuous	Discrete
Photocurrent	i(t)	i[n]
Laser Doppler (LD) Signal	s(t)	s[n]
Fourier transform of LD signal	S(f)	S[k]
Power spectrum of LD signal	P(f)	P[k]
Phase spectrum	$\phi(f)$	$\phi[k]$
Amplitude spectrum	A(f)	A[k]
Frequency estimate	C(f)	C[k]
Frequency estimator	$C\bar(f)$	$C\overline{[}k]$
Theoretical spectrum	B(f)	B[k]
Filter impulse response	h(t)	h[n]
Autocorrelation	$r_{ss}(\tau)$	

Name	Symbol
Model parameters	θ
Models	M
Probability	p()
Variance	σ^2
Fourier transform length	N
Number of averages	L
Number of estimators	K
Signal length	T
Sampling rate	f_s
Integration time	au
Boltzmann constant	$k_{\rm B}$

List of Abbreviations

ADC	-	Analogue to Digital Converter
ASIC	-	Application Specific Integrated Circuit
CCD	-	Charge-Coupled Device
CMOS	-	Complementary Metal Oxide Semiconductor
DFT	-	Discrete Fourier Transform
FFT	-	Fast Fourier Transform
FPGA	-	Field Programmable Gate Array
FIR	-	Finite Impulse Response
IIR	-	Infinite Impulse Response
IFFT	-	Inverse Fast Fourier Transform
LDBF	-	Laser Doppler Blood Flowmetry
LDPI	-	Laser Doppler Perfusion Imaging
LDPM	-	Laser Doppler Perfusion Monitoring
LASCA	-	Laser Speckle Contrast Analysis
MAC	-	Multiply Accumulate
NICE	-	National Institute for Health and Care Excellence
PDF	-	Probability Distribution Function
RBC	-	Red Blood Cell
SNR	-	Signal to Noise Ratio

Chapter 1

Introduction

1.1 Introduction

In 1842, Christian Doppler, when observing the colours of stars, proposed an explanation for the apparent frequency shift that occurs when the distance between a source of waves and an object changes with time. The Doppler Effect, as it became known, is readily observable in everyday life with sound waves, such as when an emergency vehicle passes in the street. The sirens seems to change pitch as the vehicle moves towards and then away from the observer. The classic sounds of motor racing owe their trademark slide in pitch to the Doppler Effect, with the high speed movement of the cars shifting the pitch of the engine as the car moves in front of an observer.

The effect can also be observed on moving red blood cells. The frequency of a light source illuminating an area of skin is shifted by the erythrocytes as they are pumped around the body. The discovery of this phenomenon in 1975 [117] has prompted nearly four decades of clinical and engineering research and the subject area forms the basis of this thesis.

This chapter will lay out the technique of laser Doppler blood flowmetry (LDBF) and introduce the thesis. The basic principles of the Doppler Effect and how it applies when light interacts with the near surface regions of the skin will be described. The chapter will also discuss progress in engineering terms, followed by an overview of laser Doppler systems and an evaluation of clinical applications and requirements. This will be followed by an analysis of the research question, that is, 'what is the effect on accuracy of using low bandwidth imaging sensors and can resultant errors can be corrected post-acquisition'. Finally the objectives of the thesis along with its outline will be detailed.

1.2 Laser Doppler Blood Flowmetry

1.2.1 Background

The technique of LDBF is well established [16,90,91,97]. It was first suggested in 1972 by Riva et al. [101] with the first measurements *in vivo* implemented in 1975 by Stern et al. [117]. The research field has enjoyed increased interest over the last four decades, the usefulness of the technique becoming ever more apparent in a variety of clinical situations.

LDBF research can subdivided into three main strands. The first is an understanding of the underlying physical phenomena which are used to determine blood flow. In 1981 Bonner and Nossal [12] published a key theoretical article on the relationship between the speed of the red blood cells (RBCs) and calculation of flow, and it is these same calculations which are used today. Understanding of the physical systems is improving, along with a greater appreciation of what constitutes the Doppler spectrum and the possibility of resolving different velocity components [7, 30, 131]. However, many of the improvements in theoretical understanding have yet to find their way into research devices or commercial systems.

The second strand is the engineering. The overarching improvements in electronics over the last 40 years have paved the way for significant advances and refinements in the way LDBF is implemented. Moving from analogue signal processing to largely digital and from photomultipliers to CMOS imagers, the advancements in the tools required to build flowmeters have allowed development from single point systems to fullfield imagers.

The third strand of research concerns the clinical and research environment. The potential and realised clinical uses are extensive and include burn diagnostics [78], diabetes [17] and post-operative flap monitoring [105].

These three strands will now be described in more detail.

1.2.2 Light and Tissue

This first strand of LDBF research deals with the understanding of light interaction with tissue and RBCs. The Doppler Effect will be considered more generally before being applied to microvascular circulation.

1.2.2.1 The Doppler Effect

As mentioned in Section 1.1, the Doppler Effect has been known and studied for well over 150 years. Srinivasan [115] describes it as 'a change in the observed frequency when there is a relative motion between the source and the detector.' Consider a source object, x, and a receiving object, y, moving with speed, v, directly towards each other (Figure 1.1).



Figure 1.1: A source x emits a wave of frequency f_i . An object y moves towards x with speed, v.

The incident phase angle, $\theta_i(t)$, can be evaluated, where c is the speed of light in the medium and λ_i is the incident wavelength, by:

$$\theta_i(t) = 2\pi \frac{ct}{\lambda_i} + \theta_i(0)$$
(1.1)

Figure 1.2: The wave-fronts are observed by y and reflected. The rate the phase angle changes is increased.



(a) The incident wave with phase angle θ_i (b) The reflected wave $\theta_r = \theta_i + \Delta \theta$

Figure 1.3: As wave-fronts hit an object moving towards the source over displacement D(t), the phase angle increases. The rate of change of θ_r is the new angular frequency.

When the incident wave hits the moving object y (Figure 1.2), the displacement D(t) of y changes the phase angle at time t by:

$$\Delta\theta(t) = 2\pi \frac{D(t)}{\lambda_i} \tag{1.2}$$

In this case, the object is moving at a steady speed and D(t) = vt. The observed phase angle is therefore:

$$\theta_r(t) = \theta_i(t) + \Delta \theta(t) = 2\pi \left(\frac{ct}{\lambda_i} + \frac{vt}{\lambda_i}\right) + \theta_i(0)$$
(1.3)

and the derivative of the phase angle is the angular frequency (so removing $\theta_i(0)$):

$$\omega = \frac{d\theta_r(t)}{dt} = 2\pi \frac{c}{\lambda_i} + 2\pi \frac{v}{\lambda_i}$$
(1.4)

The original frequency is $f_i = \frac{c}{\lambda_i}$ and so the change in frequency, Δf , is:

$$\Delta f = \frac{v}{\lambda_i} \tag{1.5}$$

and the frequency observed by y, f_r , is:

$$f_r = \frac{c+v}{\lambda_i} \tag{1.6}$$

D(t) can be replaced with any expression describing the displacement of y in the direction of x. For example, if y instead moved at a constant speed at an angle ϕ in relation to x, the displacement is $D(t) = vt \cos(\phi)$, the component of the velocity in relation to x is $V(t) = v \cos(\phi)$ and the frequency shift $\Delta f = \frac{v \cos(\phi)}{\lambda}$. For this reason the direction of motion needs to be known in order to determine the Doppler shift. Another example of an expression for displacement is found in the field of vibrometry, where for a surface vibrating with frequency f_v , $D(t) = \sin(2\pi f_v t)$. The Doppler Effect can be used to determine f_v :

$$V(t) = \frac{dD(t)}{dt} = 2\pi f_v \cos(2\pi f_v t)$$
 (1.7)

 $\Delta f = \frac{V(t)}{\lambda_i}$ then forms a frequency modulated (FM) process and f_v is calculated by demodulation. It should be noted that the Doppler equations derived here are non-relativistic.

The case presented above where a source emits a wave and a moving observer detects a shift is modified for LDBF because RBCs do not emit or observe light themselves; i.e. the source and detector are stationary relative to each other, and the emitted waves are Doppler shifted upon interaction with a moving cell. y can therefore be viewed as an emitter of waves (but in reality these are reflected waves) with a single shift from the original $\frac{v}{\lambda_i}$. Then the reflected waves are observed again at x with another shift. The Doppler shift in Equation 1.5 is thus modified by a factor of 2:

$$\Delta f = 2\frac{v}{\lambda_i} \tag{1.8}$$

and any subsequent calculations with alternative velocity expressions are modified similarly.

1.2.2.2 Single Scattering Event

A Doppler shift (Δf) can be detected when a monochromatic light source interacts with a moving RBC. A single scattering event, where an electromagnetic field interacts with an erythrocyte, will first be considered.

As described by Humeau et al. [56], the electric field vector, $\vec{\epsilon_i}$, of an incident wave of amplitude A_i , frequency w, directional vector $\vec{k_i}$, time t and position vector \vec{x} can be written as:

$$\vec{\epsilon_i} = A_i e^{-j(wt - \vec{k_i}\vec{x})} \tag{1.9}$$

As shown in Figure 1.4, the wave will scatter quasi-elastically to vector $\vec{k_s}$, producing a scattering vector \vec{s} . The angle between the incident vector $\vec{k_i}$ and the scattered vector $\vec{k_s}$ is α , and is determined by an appropriate phase function. Commonly used functions describing scattering include the Henyey-Greenstein phase function [4, 5], the Gegenbauer kernel scattering phase function [39] or the Mie phase function [59]. The position vector \vec{x} of the moving particle from the origin is given by $\vec{v}t$ where \vec{v} is the velocity vector. The angle between the scattering vector \vec{s} and the velocity vector \vec{v} is θ .



Figure 1.4: A single scattering event where an electric field vector interacts with a RBC. The Doppler shift is a result of the magnitude of the velocity and scattering vector |v||s| and the angle between the scattering vector and velocity vector, θ .

Assuming a constant electromagnetic magnitude, i.e. $A_i = A_s$, and neglecting the time-invariant phase factor $\vec{k_i} \cdot \vec{x_0} + \vec{k_s} \cdot \vec{x_1}$, the scattered field can be written as:

$$\vec{\epsilon_s} = A_i e^{-j(wt - (\vec{k_i} - \vec{k_s})\vec{v}t)} = A_i e^{-j(wt - \vec{s}\vec{v}t)}$$
(1.10)

Such that:

$$\vec{s}\vec{v}t = |s||v|\cos(\theta)t \tag{1.11}$$

According to Nilsson et al. [88] the scattering vector is:

$$\vec{s} = 2|k|\sin\left(\frac{\alpha}{2}\right) = \frac{4\pi}{\lambda}\sin\left(\frac{\alpha}{2}\right)$$
 (1.12)

where λ is the wavelength of the incident electromagnetic field in tissue. Now, the Doppler shift can be evaluated by:

$$\Delta \omega = \vec{s}\vec{v} = \frac{4\pi}{\lambda}\sin\left(\frac{\alpha}{2}\right)|v|\cos(\theta)$$
 (1.13)

This single shift forms the basis of the Doppler spectrum observed in LDBF in low perfusion sites.

1.2.2.3 The Microvascular Circulation

The epidermis, roughly 120 μ m in the hand [63], forms the surface layer of the skin. Below this is the dermis, which is approximately 1 mm thick. According to Oliveira et al. [93], the majority of these regions are perfused by nutritive capillaries, which have high resistance and relatively low flow [99]. At the boundary between dermis and epidermis, the subpapillary plexus, these nutritive capillaries are known as papillary loops. Below the dermis is the dermal subcutaneous interface [93] and then the hypodermis, which contains thermoregulatory blood flow. The epidermis and dermis are referred to as the dermal layers. Figure 1.5 illustrates the structure of the skin.



Figure 1.5: The structure of the skin. The vessels in the dermis are where the majority of photon-cell interactions occur. The stratum corneum is a layer of protective dead skin cells. It is thicker, for example, on the soles of the feet or palms.

Fredriksson et al. [39] note that 'For LDPI systems and LDF probes with a fibre separation <1.2 mm, skin surface measurements are well within the dermal layers, including only capillaries, arterioles, and venules (the microcirculation).' Skin models in the literature generally deal with layers including the microcirculation alone. For example, at 780 nm wavelength, the measurement depth in the index finger for a fibre separation of 1.2 mm is approximately 0.85 mm [39] (the depth is defined in [39] as where 63% of signal contribution arises from above that point). Thus the majority of photon-cell interactions (where photons then scatter back to the surface and the detector) occur in nutritive vessels and not generally in the larger thermoregulatory vessels at the dermis-hypodermis interface below (although there may be bodily sites where these vessels are closer to the surface and as such LDBF measurements will include a greater influence from them).

1.2.2.4 Light and Tissue Interactions

As photons enter the skin, they 'generally suffer several collisions with somatic cells, connective tissue, blood vessels walls etc. before interacting with a blood cell.' [12]. The authors go on to state that the number of static tissue interactions is large enough that the direction of illumination of the RBCs is effectively random. Figure 1.6 illustrates how light moves in a random fashion throughout the tissue after 10 to 20 scattering events.

Different wavelengths of light are absorbed and scattered in varying proportions: Lindberg and Öberg [74] state that at green wavelengths, absorption in blood is a factor of 10 greater than scattering, but that in the infrared regions scattering dominates. Additionally each skin layer has slightly differing scattering and absorption coefficients. Further information on coefficients at a range of skin depths has been provided by Fredriksson et al. [39].



Figure 1.6: Light enters the tissue (in this case via fibre optics) and is scattered by static and moving cells. Some light is scattered back to the surface and reaches the detector [114].

From the coherent light source, a large number of photons enter the skin. A percentage are absorbed, with a further fraction then scattered by static tissue, whilst others interact with moving RBCs. Some of these photons reach the surface of the tissue and possibly the detector. Thus the power of the frequency shifted light arriving at the detector is small, in the region of 0.06% [47] of the illumination power.

In Bonner and Nossal's theoretical paper [12], the authors describe a scenario where, because of the intensely interwoven and sufficiently dense network of capillaries, venules and arterioles in the microvasculature, the trajectories of the blood vessels form a random network with constant direction over the interaction time. The authors postulated that the velocity distribution was Gaussian. The combination of these hypotheses results in a situation similar to Brownian motion, where the expected value of the velocity vector is zero:

$$\langle V_x \rangle = \langle V_y \rangle = \langle V_z \rangle = 0$$
 (1.14)

(Equation (2) from Binzoni et al. [7]). This simplifies the calculations because Figure 1.4 does not then need to take into account the absolute direction of the RBC on a fixed axis. If the velocity vector's expectation value is zero, the value of interest becomes the expected value of the *speed* distribution, which, of course, does not take into account direction. Bonner and Nossal [12] show, via an expression for the photon autocorrelation function, that the weighted first moment of the photocurrent's spectrum $\frac{M_1}{M_0}$ is proportional to the RMS speed $\langle v^2 \rangle^{\frac{1}{2}}$:

$$\frac{M_1}{M_0} = \frac{\langle v^2 \rangle^{\frac{1}{2}} \beta}{(12\xi)^{\frac{1}{2}a}} f(\bar{m})$$
(1.15)

where, from their paper, a is the radius of an average spherical scatterer, \bar{m} is the average number of collisions, $f(\bar{m})$ is linear for $\bar{m} \ll 1$ and varies as the square root of tissue blood volume for $\bar{m} \gg 1$ (because as blood volume increases, the chance of multiple scatters increases, so M_1 increases). β is an instrumental factor $0 < \beta < 1$, 'which primarily depends upon the optical coherence of the signal at the detector, and... the intermediate scattering function of the Doppler shifted light.' [12] ξ is an 'empirical factor which is related to the shape of the cells.' [12]

As such, this equation is only valid under a set of constraints. There are approximately 4.8 million RBCs per μ L of blood [60]. Therefore, to keep the chances of multiple interactions low ($\bar{m} \ll 1$), the amount of blood as a percentage of tissue must be small (in the microvascular regions, it is roughly 1%). The equation also relies on the assumption of Brownian movement, and movement that is also uniform over the sampled volume in space and uniform over the sampled period in time. These assumptions have been questioned by several authors [7, 56], particularly as full-field imagers are becoming more prevalent [10]. One can imagine the complications of deriving analytical expressions when conjoining sample spaces are imaged at the same time. Binzoni puts the problem thus: 'This can be seen as an infinite number of source-detector couples with the photons simultaneously interacting in a complex manner over all possible paths.' [10] Modelling this presents significant challenges. Additionally, because of the variation in structure over the body and the complications of scattering and absorption, only relative changes in flow can be characterised using this formula [90].

Particularly in the last ten years, research in this field has focussed on two areas: firstly on the interpretation of the spectral moments in an attempt to derive more precise expressions and define more exactly the conditions under which the expressions hold true [6,7,9,67,90]; and secondly to model the microvasculature in more detail, using methods such as Monte-Carlo to analyse the resulting spectra, and thus perhaps to resolve independent velocity components [8, 26, 40, 59, 68, 131].

Of course, there is much overlap in these areas. The first is really concerned with the interpretation of existing methods for calculating blood flow, whereas the second deals mainly in techniques for reliably evaluating tissue samples to determine their structure and photon interactions (e.g. by Monte-Carlo, repeated sampling and parameter searching) and then using the resultant information to very precisely map blood flow through the sample. Needless to say, the second method is extremely time consuming.

However, despite those discussions, the first moment of the photocurrent's spectrum is still used as the gold-standard approach: 'Practically all LDF instruments use the original theory derived by Bonner and Nossal (1981) or some improvement of the same approach.' [7]. In classical laser Doppler theory the unnormalised first moment, M_1 , represents blood flow, while the zeroth moment, M_0 , represents blood volume. But knowledge regarding applicability for imaging is still in its infancy. In 2008, Binzoni and van de Ville acknowledged that the first moment, when 'handled with care', is probably reasonable for full-field imaging, but that it 'does not allow us to obtain information on blood velocity or volume independently', [10] i.e. the zeroth moment is not a valid representation of blood volume in full-field imaging, but still may be determinative in single-point monitoring or scanning imagers.

For the purposes of this thesis, it will be assumed, as is generally accepted, that the first moment of the photocurrent's spectrum represents relative blood flow.

1.2.3 Engineering

In an engineering sense, there are two main streams of development: single point monitoring and area imaging. In this thesis, these will be referred to as laser Doppler perfusion monitoring (LDPM) and laser Doppler perfusion imaging (LDPI) [56]. LDPI can be analysed further as single point scanning, line scanning and full-field imaging. The technical developments have been focused on LDPI, namely how to build up a blood flow map over a region of skin. The drive is towards imagers whose output blood flow maps satisfy the Nyquist criterion, so that any changes in blood flow are accurately displayed (e.g. pulsatile signal), and that also deliver on accuracy and spatial resolution.

1.2.3.1 LDPM

Initially, LDPM monitors were the only laser Doppler blood flowmeters available [2, 114]. The technique involves illuminating a point of tissue with a laser via a fibre optic and probe head, collecting the returned light with another fibre or fibre bundle in the probe, and then connecting those fibres to the active area of a photodiode. Because only one channel is involved, the data collection and processing can be performed in real time and a recording session may last for many days.

The drawback of this method is the susceptibility of the system to movement artefacts, mainly due to the use of optic fibres [112]. Several research groups have attempted to integrate the laser and photodiode into the probe head itself to reduce this noise [27, 53, 112, 120]. The integration of the optics into a probe head has allowed for further development of portable or wireless sensors [113].

Two single point systems currently in production are the Perimed (Sweden) PeriFlux 5000 and the Moor Instruments (Axminster, U.K.) VMS-LDF.

1.2.3.2 LDPI - Single Point Scanning

Single point systems can produce an area image by scanning the region of interest using a moving mirror (Figure 1.7). Only one small area (e.g. 1 mm x 1 mm) is illuminated, sampled and processed at any one moment. This technique is very slow, requiring several minutes to capture one image [129]. It is also very susceptible to movement artefacts: if the subject is moved the image will be corrupted. The advantage is that the processing power required, by today's standards at least, is minimal.



Figure 1.7: LDPI single point scanning. A mirror moves the laser beam over an area, building up an image over several minutes [80].

1.2.3.3 LDPI - Line Scanning

The line scanner is an extension of the point scanner. It simply adds an array of channels in a line and was proposed in 1987 by Fujii et al. [42], who built a system based on 256 channels. A one dimensionally divergent laser beam illuminates a line (e.g. 10 cm x 1 mm) of skin, and the linear photodetector array is focused onto the region. A motorised mirror then moves the laser, scanning the required area, as shown in Figure 1.8. However, it takes between 5 and 10 seconds to capture an image [86]. The processing requirements are clearly larger than the single point scanner, and commercial systems employ multiple analogue to digital converters (ADCs) with FPGA backends [85].



Figure 1.8: LDPI line scanning. A mirror moves the laser line beam over the region of interest [81].

1.2.3.4 LDPI - Full Field

The natural extension of the line scanner is the full-field array. Instead of, for example, a 64x1 line of photodetectors, a 64x64 array of photodetectors can be employed to directly acquire an image without mechanical movement. The advantage is obvious, namely speed, resulting in the reduction of movement artefacts. In 2007, Humeau wrote 'Efforts are made to reduce the acquisition time and approach the ideal situation of real time whole field perfusion imaging' [56]. However, every time a dimension is added, the processing requirements, design considerations and so cost increase exponentially.

The technological requirements can broadly be divided into three categories: illuminating optics, collection optics and processing circuitry. The laser illumination source is of course vital to this system. However, moving from a single point, to a line, to full-field illumination increases the required power by a factor proportional to the area of illumination. The back-scattered light from the skin is of the order of 0.06% of illuminated light [47] and so sufficient radiant power must be provided if the measurements are to be successful and not contaminated by intensity dependent noise.

The collection optics, i.e. the sensor technology, cannot simply be scaled from single point to full-field. Conventional discrete PIN photodiodes have 'insufficient packing density' [106] and thus an array of PINs would not supply the required spatial resolution. CCDs, a popular imaging technology, are not fast enough to supply the required temporal resolution, at least not for the price that LDPIs can command [106]. However, CMOS sensors have emerged as a cheap, fast and high resolution technology, and importantly, are sensitive in the regions of the spectrum employed in LDBF (red and near-infrared) [106]. Therefore the main developments in collection optics have been focused on the utilisation of CMOS technology [52, 65, 73, 105, 110, 111].

The main bottleneck in LDPI is the data processing required to digitize and evaluate the thousands of channels in the array at once. In order to achieve Nyquist appropriate (e.g 10-20 Hz) blood flow maps, trade-offs in the processing are required: for example, reducing the photocurrent sampling rate [73] or employing faster but potentially less accurate processing algorithms [31]. There are several design routes being explored. Firstly custom application specific integrated circuits (ASICs), where an array of CMOS photodetectors are combined with high-speed processing on-chip [48]. There, the processing circuitry can be built directly at the pixel level. The disadvantage is their cost and inflexibility once fabricated, but the advantage is that only the low sampling speed blood flow image data needs to be transferred off-chip. Hybrid solutions of OEM CMOS cameras with FPGA backends are also being explored because they are cheaper and more easily reconfigurable [106]. This requires very fast video transfer between camera and processing platform, for example by using low-voltage differential signalling (LVDS).

This high specification is required if the collection and processing implementation is an extension of the single point, i.e. the unnormalised first moment of the power spectrum of a sufficiently well-sampled photocurrent is taken as representation of flow on a pixel by pixel basis. However, another type of imaging known as Laser Speckle Contrast Analysis (LASCA) offers a way of substantially reducing the temporal resolution required, at the expense of spatial resolution. The theory will be summarised in Section 1.2.3.5 but, briefly, it analyses the standard deviation divided by the mean intensity of a matrix of pixels (e.g. 5x5), where the pixels represent a spatial speckle formation.

Table 1.1 summarises the LDPI state of the art. A range of bandwidths are used with the point and line scanners easily achieving the 40 kHz sampling rate required for a 20 kHz bandwidth. The DOPCAM, the final system in Table 1.1, also achieves this rate because the processing circuitry is on-chip. However CMOS cameras using offchip processing, for example, FPGAs, sample at lower rates so have lower bandwidths. As such CMOS cameras are susceptible to aliasing as they do not usually have any anti-aliasing circuitry built in [106].

Device Name	Manufacture	rYear	Type	Resolution	Speed	Bandwidth
	or Research					
	Group					
LDI2	Moor Instru-	N/S	Point scanner	256x256	5 mins	$20 \mathrm{~kHz}$
	ments					
Periscan PIM3	Perimed	N/S	Point scanner	255 x 255	$5 \mathrm{mins}$	$20 \mathrm{~kHz}$
FLPI	Moor Instru-	N/S	Full-field	576x768	25fps	N/A
	ments		(LASCA)			
CMOS camera	Serov et al.	2005	Full-field	256x256	0.8fps	4 kHz
	[106]					
TOPCAM	Draijer et al.	2009	Full-field	128x128	25fps	$13.5 \mathrm{~kHz}$
	[32]					
LDLS2	Moor Instru-	2009	Line scanner	64x512	6s @ 64x64	20 kHz
	ments					
Pericam PSI	Perimed	2010	Full-field	1386x1036	$0.2 \mathrm{fps}$ (up	N/A
			(LASCA)		to 112)	
FluxExplorer	Schlosser et al.	2010	Full-field	256x256	1fps	$2.5~\mathrm{kHz}$
	[105]					
CMOS camera	Leutenegger	2011	Full-field	480x480	12fps	8 kHz
	et al. [73]					
DOPCAM	He et al. [52]	2011	Full-field	32x32	1fps	20 kHz

Table 1.1: LDPI and LASCA devices currently in development, research or commercial market. Some devices' manufacturers do not specify (N/S) the year of introduction, and the bandwidth is not an applicable (N/A) quantity for LASCA imagers.

1.2.3.5 LASCA

Laser Speckle Contrast Analysis (LASCA) is an alternative method of mapping blood velocities by determination of the speckle contrast over small sub-frames in a single image.

When a coherent light source illuminates a static, optically rough

surface, a static grainy pattern is created as wavefronts constructively and destructively interfere. This spatial grainy pattern is referred to as speckle. However when an object moves through the light, the speckle pattern starts to 'boil' as temporal fluctuations in the individual speckles are introduced. In 1996 and 2001, Briers [15, 16] showed that the temporally fluctuating interference pattern generated by an object moving through light was in fact the same as the temporal fluctuations created by the Doppler Effect. Although these two fields of understanding had developed separately, the underlying physics is the same. Basically LDPI is founded on the understanding of Doppler, whilst LASCA is based on the understanding of fluctuating speckle in the presence of moving particles. However, even though the underlying physics is the same, this does not mean that LDPI and LASCA are identical techniques.

Conventionally, LASCA uses one image produced by integrating over a set exposure time. The exposure time is chosen to be consistent with the decorrelation time of the speckle pattern. This means that it is sensitive to the range of frequencies up to the inverse of the exposure time and insensitive to all other imaged frequencies since they average to a constant intensity. It works by analysing the temporal contrast of a collection of speckles with the optical set-up arranged so that the speckle size and pixel size match. A speckle size greater than pixel size also works, but increases the number of pixels needed to form a contrast image. Normally 5x5 or 7x7 pixels are used (a square with a side length of an odd number of pixels means that the speckle contrast can be assigned to the central pixel [69]) and the speckle contrast K is calculated from the ratio of the standard deviation σ to mean intensity \bar{I} of the pixels used (e.g. 25 or 49):

$$K = \frac{\sigma}{\bar{I}} \tag{1.16}$$

K is thus in the range 0 to 1, where 0 represents a constant intensity 'image' with no contrast and 1 a high contrast image with equal standard deviation and mean. Values will always be in this range for a fully evolved speckle pattern, where the 'phases of the interfering electromagnetic fields are uniformly distributed' [11]. For a set exposure time, low contrast occurs where high velocities blur the speckles in the view of the camera, and high contrast where low velocities effectively produce a speckle pattern that is frozen in time. However, as stated, the sensitivity to a particular velocity is determined by the exposure time.

LASCA has the advantage of being a very fast technique, only requiring one image and using exposure times much greater than those a conventional LDPI camera requires (e.g. 5 ms instead of 50 μ s). One of the main disadvantages, and there must be some trade-offs, is the loss of spatial resolution because every output image pixel requires 5x5 or 7x7 input pixels [110].

Another drawback is that the underlying relationship between blood perfusion and K has not been established [33,56] in the same way that it has for LDPI. For example, the Doppler technique measures the power spectrum directly, and thus in theory provides information on blood cell concentration as well as speed [12]. In 2009, Draijer et al. [33] stated that a concentration measure has not been 'shown to be possible' with LASCA. This means that LDPI gives a measure of flow (velocity times concentration) whereas LASCA only gives information about velocity.

Much has been done to try and resolve this problem with the model and so provide some correspondence between values gained by LASCA to those in conventional LDPI. In 2010, Thompson and Andrews [125] linked a multi-exposure time autocorrelation to the Doppler power spectrum. As the exposure time is varied, an estimation of power in certain frequency bands can be gained and a more conventional analysis performed. However, there are complexities involved in taking such variable exposure images, a fact acknowledged by the authors in a later paper [124], and the cost and complexity of these systems may approach that of conventional Laser Doppler imagers.

Hence, the comparisons between LASCA and LDPI have so far mainly been empirical [92, 107, 123, 124], and independent assessments of LASCA, not unlike LDPI, have focused on the quality of the resultant image [11].

Having outlined the underlying relationship between the first moment of the photocurrent's spectrum and blood flow, this thesis will concentrate on this better understood and more accepted method of calculating blood flow. This is not to say that developments in LASCA are less useful, but that this research builds on the slightly firmer LDPI foundation. In future, it is hoped that the speed advantages of LASCA may be combined with the accuracy of LDPI, and indeed hybrid systems are beginning to emerge [33].

1.2.4 A Laser Doppler System

Before the third strand of research, the clinical side, is addressed, an overview of a typical laser Doppler system will be given. Figure 1.9 is a schematic of a full-field LDPI system.

Laser: The laser must deliver enough power with long enough coherence length to provide sufficient signal to noise ratio (SNR). One of the potential disadvantages of full-field LDPI, as compared to line scanning or single point scanning, is that higher powered lasers are required, e.g. a 250 mW laser and a 400 mW laser are used in [106] and [127] respectively. These powers are approaching class 4 (500 mW), which, because of the necessary safety controls, precautions, training and therefore expense, make integration into a clinical setting more difficult.



Figure 1.9: An overview of a typical LDPI full-field system. The sensor, ADCs, and high speed processing in some systems are often combined into a single chip (ASIC), in others a CMOS sensor is combined with a FPGA.

Imaging optics: Coherent light sources produce speckle patterns when illuminating optically rough surfaces. As well as magnifying and focusing the light with an appropriate depth of field, the imaging optics must resolve these speckle patterns in space to optimise the signal to noise ratio.

Image sensor: Single point and line scanners have tended to use conventional photodiodes, in which the current is a continuous representation of the intensity at the photodiode and is continuously converted to a voltage through a transimpedance amplifier. However the newer full-field imaging devices (e.g. CMOS sensors) tend to use integrating detectors. These accumulate charge over a set time period (the integration time) onto a capacitor, which is then read out at the end of the period. According to Serov and Lasser [106], the integrating detectors improve the SNR compared to a non-integrating detector by $\frac{\pi}{2}\sqrt{N}$ (where N is the number of pixels in a read-out sub-frame). However, by the nature of the integration they also low-pass filter the resultant signal. The effect of this will be more fully considered in Chapter 5.

ADC: Often integrated into the sensor, the analogue to digital converter (ADC) converts the pixel voltages into a digital representation. One ADC is assigned to a sub frame of N pixels, so that the larger the sub-frame the slower the readout speed of all the pixels for a given clock frequency. The resolution is often programmable and typically ranges between 8 and 12 bits. High throughput data transfer is required between sensor, ADC and processor. For example, in a system with 256x256 pixels, sampled at 10 kHz and 10 bit resolution, the link between ADC and processor must be able to handle at least 782 MB/s, and sometimes more if the protocol requires byte-rounded data.

High speed processing: The processing converts the 'photocurrent' (with a sampling rate of up to 40 kHz) to a value representing blood flow (usually output at 1-30 fps). This is normally achieved by calculating the first moment of the photocurrent's power spectrum [23]. Methods of doing this are covered more fully in Chapter 2. The transfer of high speed data from the image sensor and the subsequent digitisation and processing forms the major bottleneck in full-field LDPI.

Flow output: An image representing blood flow is transferred to the display. Any post-processing, for example, averaging [105] or colour-map correction may be added.

Display: A user interface is presented to the researcher or clinician with device calibration and acquisition controls. The display may be a combination of blood flow image with monochromatic photograph to enable easier interpretation of results.

1.2.5 Clinical Applications

An overview of the theoretical background to and engineering advancements in LDBF has been presented so far, along with a typical LDBF device set up. The application of LDBF in clinical situations will now be considered.

LDPI and LDPM are elements in a range of tools clinicians have available for assisting in the diagnosis and treatment of disease and injury, firstly of the skin directly [98], but secondly of other conditions which affect the microcirculation [61]. LDBF is potentially useful in a variety of clinical and research situations. Its main advantages are that it is non-invasive, safe (with appropriate precautions) and relatively instantaneous. However its use is certainly not yet widespread. Table 1.2 is an overview of applications for which LDBF is currently used either in clinical situations or in research. It is important to note that for some applications, for example malaria or diabetes, LDBF applicability is in its very early stages of investigation.

One sector which has received a lot of attention is that of burn assessment. Recently, a report issued by the National Institute for Health and Care Excellence (NICE) in the U.K. [87] advised clinical centres of a new product which reduces the cost and eases diagnosis of the severity of a burn. This was the Moor Instruments LDI2-BI point-scanning imager, which shows the healing potential of large areas of skin tissue based on their blood perfusion. In 2012, Stewart et al. [118] found that LDPI had a >90% positive predictive value for the need for skin excision and grafting post-burn, thus making it a worthwhile

Medical Area	Synopsis	Authors and Dates	LDPI/LDPM	LDPI Range
Burn assessment	Burn regions can be imaged to	Droog et al. (2001) [35], van	LDPI	x3-5 baseline
	determine the severity of the	Herpt et al. (2010) [127],		
	burn: superficial burns show	Erba et al. (2012) [37],		
	increased flow from baseline,	Monstrey et al. (2011) [79],		
	whereas severe damage will	Stewart et al. (2012) [118]		
	show decreased flow.			
Reconstructive Surgery	The post-operative reperfusion	Holzle et al. (2006) [54],	LDPI/LDPM	x1.5 baseline
	of a flap can be monitored,	Booi et al. (2008) [13],		
	and complications or potential	Schlosser et al. (2010) [105]		
	partial or total flap loss can be			
	predicted and resolved using			
	both LDPI and LDPM.			
Cancer Treatment	The microcirculation can be	Scheving et al. (1993) [104],	LDPM	
	studied to determine optimum	Haus (2008) [49]		
	conditions for drug delivery or			
	other treatments.			
Amputations	Blood flow levels in limbs can	Gebuhr et al. (1989) [43]	LDPM	
	determine the the most appro-			
	priate amputation location and			
	healing potential of the stump.			
Malaria Treatment	The reduction in deformability	Dondorp et al. (2000) [29]	LDPM	
	of RBCs (which reduces flow)			
	and the resultant anaemia			
	(which increases flow) can be			
	monitored as treatment is			
	provided.			
Neonatal Care	Baseline changes in flow in the	Suiches et al. (1988) [120],	LDPM	
	first week extrauterine indicate	Suiches et al. (1990) [121],		
	changes in microcirculatory	Purcell and Beeby (2009)		
	structure.	[96]		
Sepsis Diagnosis	Changes in the microcircula-	De Backer et al. (2002) [25],	LDPM	
	tion may be able to predict sep-	Sakr et al. (2004) [103]		
	tic shock, or indicate reduced			
	risk.			
Critical Care	Monitoring the microcircula-	Fries et al. (2005) [41]	LDPM	
	tion could help to forsee resus-			
	citation outcomes.			
Diabetes Assessment	There is scope for the assess-	Kim et al. (2008) [61]	LDPM	
	ment and prevention of dia-			
	betic neuropathy by intersite			
	microcirculatory comparisons.			

Table 1.2: Clinical applications for LDPI and LDPM. The LDPI range is an assessment of the expected range of LDPI values encountered in this medical field.

application area on which to focus.

Another application area is that of skin flap monitoring. Surgical flaps are sections of skin removed from one part of the body (e.g. Transverse Rectus Abdominis Myocutaneous, or TRAM, flap from the abdomen) which are then used to repair a damaged area of skin or assist in reconstructive surgery. Because the reperfusion of a flap after application to the new site is a good indicator of the likelihood of successful transplant, LDPM and LDPI could be used to image or monitor perfusion levels to determine flap health and highlight the possible need for remedial action.

These two application areas serve as a focus for assessing imaging requirements. This thesis is partly concerned with the error created in flow values due to low bandwidth and undersampled photocurrents. However, instead of viewing this error abstractly, it is useful to look at it through the lens of these clinical applications to match specifications to requirements. The range of flow values required in these two application areas is gathered from the literature.

1.2.5.1 Burn Assessment

In 2001, Droog et al. assessed the perfusion values associated with LDBF burn imaging. Having found that there are many factors influencing perfusion values (e.g. angle of imaging, tissue flatness, imaging distance etc.), they found that deep burns generally show lower to equal flow values than the baseline in normal skin, whereas superficial burns show increased perfusion (3-5x) compared to normal skin [35]. As previously discussed, blood flow values are relative and many physiological and instrumental factors affect the output flow values. However, with a properly calibrated and consistent device used in a controlled environment, the output readings should be sensitive up to 5 times baseline flow to avoid incorrect diagnosis of burn severity.

1.2.5.2 Flap Monitoring

Flap monitors do not have the same high flow requirements. There are at least two possible situations in this field of surgery where LDBF may be useful. The first is the location of perforator vessels in the pre-operative flap. The study of Schlosser et al. [105] starts from the evidence-based supposition that blood flow is marginally higher by about 20% in the microvasculature above perforator vessels (that supply and return blood to the tissue and hence need to be located pre-operatively). The technique using LDPI to locate these is still not reliable, but in principle blood flow values need to be sensitive to around 1.5 times baseline or normal flow. The second is the detection of regions of ischaemia post-operatively, where the lack of blood flow will cause the tissue to die and hence require more extensive remedial surgery. If ischaemia can be detected early then the required remedy is lessened in severity, reducing treatment time, cost and further complications [13]. Clearly this detection only requires sensitivity to flow values lower than baseline.

1.2.6 Review

With some of the clinical applications of LDBF having been considered, this concludes the background to the thesis. As has been seen, there is a large volume of literature on the laser Doppler technique. However, it should be noted that it is a constantly evolving technology and new application areas will continue to be discovered whilst existing ones are reinforced. Additionally, businesses are vying to make a commercial success of this technology and therefore research will continue to push forward in each sector. This work aims to add to the scholarly foundation laid down over the past thirty-five or so years.
1.3 Research Rationale

1.3.1 State of the Art

With clinical applicability in mind, the present research mainly focuses on signal processing techniques for providing accurate blood flow values. The rationale prompting the research contained in this thesis will now be considered. First of all, the state of the art in LDBF signal acquisition and processing will be described.

Classically, the LDBF photocurrent is sampled at 40 kHz [65] in order to detect frequency shifts up to 20 kHz. These frequency shifts correspond to movements of up 7.5 mm/s (see Equation 1.13, this is wavelength dependent and is in the case of perfect back-scattering and velocity alignment with k_s). However, in the baseline ranges, RBC speeds are around 2-4 mm/s [119, 126]. Higher speeds would only be reached in extraordinary circumstances, for example for superficial burns as discussed earlier, or if temperature provocation or steroidal cream stimulant is applied, remembering also that blood flow is a product of speed and concentration and blood flow values may be affected by a change in either. Thus 20 kHz represents the upper limit of the Doppler shift range, and is used so as to encapsulate all possible scenarios.

Several research groups [32, 73, 105, 106] have developed imagers which sample the LDBF photocurrent at lower frequencies than this (see Table 1.1), even as low as 5 kHz. The rationale for this is that using existing CMOS image sensors or cameras with external processing negates the need to build custom devices, thus substantially reducing the expense and development time required. However, current CMOS technology, in the price ranges suitable for LDBF, cannot sample at the full bandwidth with large enough resolution and transfer or store the data generated and so the sampling rates are reduced. CMOS technology is therefore a compromise between processing accuracy, and cost and convenience.

Two immediate issues present themselves: firstly, whether any antialiasing circuitry is applied before digitisation and, secondly and more fundamentally, how this undersampling and potential aliasing impact upon blood flow output values.

In 2005, Serov et al. [106] acknowledged the aliasing problem: 'Digital image sensors do not usually include anti-aliasing circuitry in their design; therefore the aliasing effect is virtually unavoidable in the imager.' Aliasing is undesirable in all acquisition scenarios because of the inability to distinguish frequencies greater than or equal to the Nyquist frequency (half the bandwidth) from those less than the Nyquist frequency. The effect of high frequencies manifesting as low frequencies in the power spectrum must therefore be included in any assessment of low bandwidth effects on blood flow values.

In 1990, Obeid et al. [90] stated that in order to prevent flow under-estimations, higher frequency cut-offs are a necessity. This was backed up by Binzoni et al. in 2008 [10], who noted that 'A too low temporal sampling rate will miss high-frequency components and introduce distortion in the [concentration] and [flow] estimations.' Binzoni et al. go on to assess the effect of sampling at 7 kHz on a Monte Carlo generated spectrum by setting an upper cut-off of 3.5 kHz. They report that output blood flow is roughly constant as RBC velocity increases for a constant blood volume, essentially because the higher frequencies generated by larger velocities are undetected. However, Binzoni et al.'s investigation is limited to one sampling frequency and one first moment calculation method, and neglects the effect of aliasing, but is useful for the initial characterisation of low sampling frequency behaviour.

1.3.2 Present Research

It has been identified that there is some improvement to be made in understanding how CMOS imagers which undersample an LDBF photocurrent affect blood flow output values and whether this is acceptable. Therefore an aim of the present research is to characterise the effect of undersampling and see whether it can be compensated for.

By creating a model of the frequency distribution of the photocurrent, the effect of undersampling, aliasing and complete processing methodology can be determined. The model can then be used to predict the power spectrum of the LDBF photocurrent in these conditions and thus compensate for them, moving closer to the original blood flow values.

As the research outcome of determining the effect of undersampling is realised, so can the applicability to the two clinical situations described earlier be ascertained.

1.4 Thesis

1.4.1 Aims & Outline

The aim of the thesis is to describe the background to the research and rationale, moving on to evaluate, by a suitable model, current methods of processing the LDBF photocurrent and presenting a technique for compensating for undersampling.

This will be achieved by the following subsidiary aims, each of which forms a chapter of the thesis:

- 1. Introduce the background to the research and the reasons behind it (this chapter).
- 2. Describe the current methods of evaluating the first moment of the photocurrent's spectrum.

- 3. Determine a model of the power spectrum of the LDBF photocurrent using Bayesian Inference techniques.
- 4. Characterise the effect of algorithm selection and bandwidth on a variety of processing methods by using the model.
- Characterise the effect of aliasing and use the model to compensate for an undersampled simulated photocurrent, also using Bayesian Inference techniques.
- Evaluate the performance of two 'low bandwidth' high speed cameras and correct real recorded data using the techniques described in chapter 5.
- Conclude by summarising the thesis and suggesting suitability for differing applications.

1.4.2 Scope

LDBF is an extremely large field of research that extends from fundamental engineering and physics, through to clinical needs, usability, result interpretation, diagnosis, treatment and cost-benefit analysis. It is therefore necessary to define the scope of this thesis, acknowledging that there are present and pressing issues which may impinge on the work that cannot be dealt with here.

Blood flow imaging results are usually presented to the clinician in a colour coded format, which is in some sense arbitrary. For example, flow values in lower ranges are assigned blue, through the spectrum to mid flow ranges in green and orange and high flow ranges in red. Thus adjustment of the colour map ranges or de-linearisation can arbitrarily change result interpretation without any underlying values changing. However there may be good reasons for adjusting colour maps. For example, the Moor Instruments LDI2-BI re-casts its colour map in terms of wound healing potential. What needs to be acknowledged is that the output flow results are interpreted and diagnoses, treatments and prognoses can be based on that interpretation. Thus, the next stage of an analysis into processing error is the impact on clinical interpretation. This thesis will suggest that mis-diagnoses, the extent of which are dependent on the particular clinical field, may result from inaccurate flow calculations, but a full study is outside its scope.

1.5 Conclusion

The LDBF technique has been described, with the three research strands of physiological and photological understanding, engineering development and clinical applications outlined. The justification of the research presented here has been given, along with the scope and outline of the thesis.

It is currently accepted that the first moment of the photocurrent's power spectrum represents microvascular blood flow, but it is noted here that a reduction in the bandwidth of the signal by undersampling may introduce distortion in the output flow values. It is the intent of the thesis to characterise and compensate for this error which will enable an accurate picture of flow to be presented.

Chapter 2

The First Moment of the LDBF Photocurrent's Power Spectrum

2.1 Introduction

This chapter aims to describe techniques for evaluating the first moment of the photocurrent's power spectrum, widely acknowledged to represent microvascular blood flow. One of the bottlenecks in full field LDBF imaging is processing the photocurrent to obtain this value, as the number of operations to evaluate the fast Fourier transform at each pixel is so large. For this reason a variety of computationally efficient methods have been suggested [9, 31, 50, 73, 82, 105] that reduce processing time and hardware space. The chapter will detail derivation of the methods along with issues regarding implementation. As far as is known, a review of all methods for processing LDBF signals has not yet been published.

Some of the derivations are mathematically precise, others are mathematically precise under certain conditions, whilst others are approximations. An evaluation and comparison of the complexity of each method will be provided, but it is acknowledged that each method has a plurality of possible implementations that can be highly optimised for the chosen hardware.

Firstly, the methods will be classified according to domain, although

evaluation in time and frequency are mathematically identical, then routes using other moments of the power spectrum and the analytic signal will be explored. Before this, some properties of the photocurrent will be recalled to enable appropriate implementation.

2.2 The Photocurrent

2.2.1 Electric Field

The electric field vector of the light received at the photodetector has frequencies comprising the incident frequency plus Doppler shifts from scattering erythrocytes, and phases related to the path length through the epidermis and dermis. Because the photons undergo multiple static scattering, the path lengths and so phases can be assumed to be uniformly random (see Section 1.2.2.4). Hence the photocurrent looks noisy but has spectral properties related to the speed of the RBCs.

If an incident wave $\vec{E}_i = A_i e^{j(\omega_i t + \vec{k_i} \vec{x})}$, where $\vec{k_i}$ is the direction vector and ω_i is the incident frequency, interacts with a moving RBC (a single scattering event) an electric field vector:

$$\vec{E_s} = A_s e^{j\omega_i t} e^{-j((\vec{k_i} - \vec{k_s})\vec{v}t)} e^{-j(\vec{k_i}x_0 + \vec{k_s}x_1)}$$
(2.1)

is created [88]. This corresponds to the Doppler shift as described in Figure 1.4. The three exponential terms in Equation 2.1 correspond to the incident wave, a time varying phase factor (the Doppler shift), and a time invariant phase factor respectively. Here information about the size of the detector is neglected, assuming the photocurrent generated by all electric field vectors scales with area at the same rate. N fields corresponding to N scattering events (assuming the proportion of multiple scatters is small, as is reasonable for small tissue volumes) arrive at the detector along with un-shifted light reflected from the skin surface. The electric field at the detector for the un-shifted light is $\vec{E_r} = A_r e^{j(\omega_r t + \vec{k_r} \vec{x})}$ and the total electric field is:

$$ec{E}_{sum} = ec{E_{r}} + \sum_{s=1}^{N} ec{E_{s}}$$
 (2.2)

The photodetector is a square-law detector and the photocurrent i(t)is a result of the square of the electric field [88], with a proportionality constant related to the detector's quantum efficiency for the incident wavelength. The quantum efficiency does not change significantly across the bandwidth of the LDBF signal ($\omega_r \gg \Delta \omega$). The response time of the photodetector τ is much larger than the period of the electric field so that the photocurrent is an integration of the electric field over time τ and it can be written as:

$$i(t) \propto \int_{-\frac{\tau}{2}}^{\frac{1}{2}} \vec{E}_{sum}(t) \vec{E}_{sum}^{*}(t) \mathrm{d}t$$
 (2.3)

The result is an expression with three main terms:

$$i(t) \propto \int_{-\frac{\tau}{2}}^{\frac{\tau}{2}} \vec{E_r}(t)^2 + 2\vec{E_r}(t) \left(\sum_{s=1}^{N} \vec{E_s}(t)\right) + \left(\sum_{s=1}^{N} \vec{E_s}(t)\right) \times \left(\sum_{s=1}^{N} \vec{E_s}(t)\right) dt$$
(2.4)

The photocurrent will contain a DC element (the first term in Equation 2.4), heterodyne elements, a result of the multiplication of $\vec{E_r}$ by each element $\vec{E_s}$ (the second term), and homodyne elements, a result of the multiplication of each element of $\vec{E_s}$ by each other element of $\vec{E_s}$ (the third term). Since the amplitude of the un-shifted light, A_r , is much greater than the amplitude of the shifted light A_s , the heterodyne elements dominate over homodyne elements [88].

The frequencies in the range of ω_r are large compared to the integration time of the detector [34], the detector therefore acts as a low pass filter and the time-varying portion of i(t) will be a result of intensity fluctuations at frequencies $\Delta \omega$, the direct result of Doppler

shifts by moving scatterers. Different detectors will integrate over different time lengths, but normally these are much smaller than the analogue to digital converter (ADC) sample time (e.g. $\tau=1 \ \mu$ s). Some detectors, however, may be designed to integrate for longer (e.g. $\tau=50 \ \mu$ s). The effect of integration time will be discussed in Chapter 5. After the photodetector, the photocurrent is normally converted, through a transimpedance amplifier, to a voltage and is conditioned for the range of the ADC.

Additionally, $A_r = A_r(f)$ with f in low frequency ranges, up to 30 Hz [64]. This is because the intensity of the reflected light is dependent on skin surface movements [90], optical absorption and the photoplethysmographic effect [66]. This low frequency light scales the higher frequencies (>10 Hz) and the LDBF signal requires division by the low frequency signal before digital conversion. This can be achieved by low pass filtering the signal and dividing the signal by the filter output. The resultant signal will hereafter be referred to as s(t).

2.2.2 Sample Rates

The flow calculation signal processing system has two sample rates that define its operation (Figure 2.1): firstly, the input sample rate of the photocurrent (f_s^{pc}) and secondly, the output sample rate of the blood flow information (f_s^{bf}) . Both are determined by the bandwidth of the respective sampled data. Many samples of the photocurrent are needed to estimate a single blood flow value.

It is generally accepted that the bandwidth of a laser Doppler photocurrent signal s(t) ranges from 20 Hz to 20 kHz (see Section 1.3.1). According to the Nyquist criterion ($f_s > 2f_{bw}$, where f_{bw} is the bandwidth of the signal) the sampling frequency f_s^{pc} must therefore be at least 40 kHz. The bandwidth of blood flow variation is much lower. For well perfused skin this change is cardiac synchronous, which equates



Figure 2.1: Input and output sample rates for the blood flow calculation. Each sample rate depends on the bandwidth of its signal.

to a dominant frequency around 1-2 Hz [19] with further components up to 10 Hz [71]. Unless the system can be designed to synchronise with the heart beat (using, for example, an electronic heart rate monitor), the sample rate f_s^{bf} must be at least 5 Hz and ideally above 20 Hz. This may not be required at bodily sites where the influence of the heart rate on the microvascular flow signal is not as strong. It is acknowledged that, whilst laser Doppler perfusion monitors exceed this easily, with sample rates up to 40 Hz, the state of the art in laser Doppler perfusion imaging is only just reaching these speeds [51]. Heart rate artefacts may be noted at these low output sampling rates.

If the photocurrent is not processed in sequential input samples with a moving window, but rather block by block, the number of samples, N, which can be processed is:

$$N \le \frac{f_s^{pc}}{f_s^{bf}} \tag{2.5}$$

If N is less than the maximum value given by this ratio, some input data is discarded after the previous output data sample is calculated. One reason for discarding data may be that, whilst a system's ADC might sample at 40 kHz, the processing electronics cannot calculate blood flow values continuously at this speed. The theoretical minimum length of N is determined by the required frequency resolution and bias and will be discussed further in Chapter 3. In reality, N is also determined by the processing method (e.g. the fast Fourier transform is more efficient for $N = 2^n$) and memory limits.

With values of N and f_s^{pc} chosen¹, the methods of determining M_1 , the blood flow samples, can be examined.

2.3 Evaluating the First Moment

The photocurrent can be processed either in the frequency domain via a transform or directly in the time domain. The frequency case will be considered first, starting with the definition of the first moment.

2.3.1 Frequency Domain Definition

The *i*th moment of a power spectrum P(f) of a signal is defined as:

$$M_i = \int_0^\infty f^i P(f) \, \mathrm{d}f \tag{2.6}$$

When i=0, the total power in the signal, the RBC concentration, is evaluated. When i=1 the first moment, the blood flow, is evaluated.

2.3.2 Frequency Domain Implementation

The signal s(t) is a time domain process so, for frequency domain calculations, the spectrum P(f) must be determined. The photocurrent will be considered as a stationary stochastic process although strictly speaking it is not stationary because the blood flow varies and so the frequency content of the signal changes with time. However the rate of change of blood flow (dominant 1-2 Hz with components up to 10 Hz [71]) is low enough that small sections of the photocurrent may be considered 'stationary'.

 $^{^{1}}f_{s}$ is used in future without qualifier where f_{s}^{pc} is intended

The Fourier transform of a stationary stochastic process, because it has infinite energy, does not strictly exist. The signal must be multiplied by a suitable window function (rectangular, i.e. a finite number of samples, Hamming or Hanning etc.) to limit the energy. Figure 2.2 illustrates how a window function can be used to bound the photocurrent.



Figure 2.2: An infinite stationary stochastic process $s_{in}(t)$ is multiplied with a window function (in this case Hamming as this window reduces spectral leakage from side lobes) to give a resultant bounded signal.

Mathematically the window can be described by:

$$w(t;\zeta,T) = \begin{cases} 0 & t < \zeta \\ 0.54 - 0.46 \cos\left(\frac{2\pi(t-\zeta)}{T - (t-\zeta)/T}\right) & \zeta \le t \le \zeta + T \\ 0 & t > \zeta + T \end{cases}$$
(2.7)

where ζ represents a specific starting point in all time t and T is the length of the window. Given an infinite signal $s_{in}(t)$, the resultant signal is:

$$s(t) = w(t; \zeta, T) \times s_{in}(t)$$
(2.8)

s(t) now represents a continuous function of time with value 0 outside the window region. Different sections of the infinite signal can be selected by moving ζ . The autocorrelation of the real signal s(t) can now be taken:

$$r_{ss}(\tau) = \int_{-\infty}^{\infty} s(t+\tau)s(t) \, \mathrm{d}t$$
 (2.9)

In practice the lags need only be calculated from $\tau=0$ to $\tau=T$ because the autocorrelation is zero valued outside the range $-T < \tau < T$ and is symmetric around $\tau=0$.

By the Wiener-Khinchine theorem [77], the autocorrelation and power spectrum form a Fourier pair, so that the Fourier transform $\mathcal{F}\{r_{ss}\}$ will yield P(f):

$$P(f) = \int_{-\infty}^{\infty} r_{ss}(\tau) e^{-2j\pi\tau f} \,\mathrm{d}\tau$$
(2.10)

However P(f) may also be evaluated directly by $|S(f)|^2$, where:

$$S(f) = \int_{-\infty}^{\infty} s(t)e^{-2j\pi tf} dt$$
 (2.11)

Equation 2.11 forms the basis of the frequency domain methods. From here on s(t) will refer to the signal starting with t=0 at ζ and with length T. As the signal s(t) is sampled with frequency f_s , the discrete Fourier transform (DFT) is required. s[n] represents the sampled signal, where n=0 when t=0 and $N = Tf_s$. With k representing the discrete frequency bins such that $f = \frac{k}{N}f_s$, the DFT is:

$$S[k] = \sum_{n=0}^{N-1} s[n] e^{-2j\pi nk/N}$$
(2.12)

And the power can be evaluated by:

$$P[k] = |S[k]|^2$$
(2.13)

2.3.2.1 N-point FFT

The DFT is not an efficient method for computer calculation of S[k] because its complexity scales with $O(N^2)^2$. The fast Fourier transform (FFT) is a more computationally efficient algorithm for evaluating the DFT, reducing the complexity of the calculation from $O(N^2)$ to $O(N \log_2 N)$ in the Cooley-Tukey implementation [22]. Its output S[k] does not suffer loss of accuracy compared to the DFT. The terms will be used interchangeably, such that when 'FFT' is used it implies the results of a DFT are calculated. MATLAB 7.12 (The Mathworks) and its libraries are employed by this thesis for FFT evaluation.

The discrete form of Equation 2.6 for i=1 is:

$$M_1 = \frac{f_s}{N} \sum_{k=2}^{N/2+1} kP[k]$$
 (2.14)

with the starting frequency bin adjustable for lower cut-off.

This is the discrete form of a definite integral and so it makes use of the rectangle (or midpoint) rule to calculate area under the spectrum. This makes the assumption that the area between the midpoint and both edges of a frequency bin are roughly equivalent (or if not equivalent that the triangular area lost on one side of the midpoint is equal to the triangular area gained on the other side). This is a valid assumption because in most cases N is large enough that the power spectrum shape has approximately constant gradient across the bin width. However for small N the integration becomes less accurate.

The FFT equation is classically employed in LDPM scenarios [85,89, 112] because it is theoretically the closest to the first moment definition.

 $^{^{2}}O()$ here refers to 'Big O' notation [62].

The choice of N and the accuracy of P[k] control the error in M_1 .

The complexity will be determined by the number of multiplyaccumulate (MAC) operations required to produce one output blood flow value from N=1,024 input samples. In this case, the number of real MACs required to perform an N-point FFT is $10N/2\log_2 N$ [95], and a further N/2 MACs are needed for the weighting and accumulation of the frequency spectrum. Therefore the total number of MACs is:

$$MAC_{FFT} = 10 \times 1,024/2 \log_2 1,024 + 1,024/2 = 51,712$$
 (2.15)

2.3.2.2 Frequency Averaging

Equation 2.14 can be modified so that the P[k] are the average of the square of L DFTs taken from consecutive sets of N samples [84]. In Equation 2.14, the total time period of the window was a result of N samples: $T = \frac{N}{f_s}$. With frequency averaging, the total time period $T = \frac{NL}{f_s}$. As such, Equation 2.14 is simply the case for L=1. For T to remain the same whilst introducing averaging, N must be reduced proportionally. The rationale is that to perform many smaller FFTs is less computationally expensive than longer FFTs. The cost is $O(LN \log N)$. As L increases and N decreases proportionally, the cost is reduced [82]. Small DFTs also have other advantages: the memory requirements are minimal and the coefficients, for example, in an 8-point DFT, may be approximated in hardware by substituting $\frac{\sqrt{2}}{2}$ by $\frac{2}{3}$ [82]. If all the coefficients in the DFT are multiplied by 3 after the approximation is made, the new coefficients for an 8 point sinusoid become 0,2,3,2,0,-2,-3,-2. The effect of approximating the DFT coefficients in this way will be tested in Chapter 4.

Another advantage is that the estimates of P[k] at the particular frequency bins in the N-point DFT may be improved by averaging. The reason for this will be discussed in Chapter 3. However, one of the potential disadvantages is that as N is reduced, the use of the midpoint integration rule, which assumes triangular formations between the spectral vertices, loses accuracy because the gradient of the power spectrum is not constant.

The averaging process is expressed by:

$$P_{av}[k] = \frac{1}{L} \sum_{l=1}^{L} P_l[k]$$
(2.16)

where $P_l[k]$ is the *l*th DFT in a set of L consecutive N-point DFTs. The effect of averaging on M_1 will be assessed theoretically in Chapter 3 and tested in simulations in Chapter 4.

The number of MACs to produce one output sample from 1,024 input samples using 128 averages (i.e. 8 point FFTs) is:

$$MAC_{FA} = 128 \times 10 \times 8/2 \log_2 8 + 8/2 = 15,364$$
 (2.17)

2.3.3 Time Domain Equivalence

Despite the FFT delivering improved performance to the DFT, it is still computationally expensive compared to time domain approaches. As the signal is captured in the time domain, so it seems efficient to process it directly and therefore by-pass any costly transformation into frequency domain. Parseval's theorem states that the energy in a signal is equivalent whether evaluated in the time or frequency domain. If the case where i=0 is considered, Parseval's theorem can be readily applied such that M_0 can be evaluated in the time domain:

$$M_0 = \int_{0}^{f_{\text{max}}} P(f) \, df = \int_{0}^{T} |s(t)|^2 \, dt$$
 (2.18)

When i=1, an equivalent process of frequency weighting the power spectrum must be applied in the time domain. The convolution theorem states that convolution in the time domain is equivalent to multiplication in the frequency domain, or vice-versa [36]. Introducing a function h(t) with Fourier transform:

$$\mathcal{F}\{h(t)\} = f^{\frac{1}{2}}$$
 (2.19)

The convolution of h(t) and s(t) becomes:

$$\mathcal{F}\{h(t) * s(t)\} = \mathcal{F}\{h(t)\}.\mathcal{F}\{s(t)\} = f^{\frac{1}{2}}S(f)$$
(2.20)

Squaring this equation gives total energy and Parseval's theorem can again be applied:

$$M_1 = \int_0^\infty fP(f) \, \mathrm{d}f = \int_0^\infty |h(t) * s(t)|^2 \, \mathrm{d}t$$
 (2.21)

The function h(t) is a filter with $f^{\frac{1}{2}}$ weighted amplitude response, as defined in Equation 2.19, and M_1 can be calculated precisely if Equation 2.19 holds true. An example of the impulse response h(t) for an IIR filter is shown later in Figure 2.5. Note that the phase response of the filter is ignored as the phases of frequencies comprising an LDBF signal are uniformly random and the phase is not used in the first moment calculation.

2.3.4 Time Domain Implementation

Again, the implementation is usually discrete as s(t) is digitised and filtering is performed after the ADC stage. In this case h(t) becomes h[m], and the first moment, for a block of N data points and a finite impulse response (FIR) implementation, is calculated by:

$$M_1 = \frac{1}{N-M} \sum_{n=M-1}^{N-1} \left| \left(\sum_{m=0}^{M-1} h[m]s[n-m] \right) \right|^2$$
(2.22)

where M is the length of the filter. In this case the convolution length is finite, i.e. h[m] is not continuously sliding across the data but is restricted to N - M operations, as this is typical in LDBF where continuous data is not generally available.

2.3.4.1 Filter Designs

There are at least three issues to consider when designing a filter for hardware implementation.

The first is the order of the filter. Filters with greater order are more computationally expensive, but the frequency response is usually closer to the requirement. One advantage of infinite impulse response (IIR) filters over FIR filters is that, because they are recursive, they can achieve the desired frequency response in much smaller length.

The second is the impulse response. Usually, whatever the processing method, a LDBF signal is processed in individual windows [85] rather than continuously. Thus, the beginning of each signal window will comprise an impulse response associated with the filter. Essentially this means that N raw points cannot produce N filtered points because there is a settling time associated with the filter response. For IIR filters, the results will only be valid after the impulse response has reduced to the noise level. For FIR filters, the results are valid after the length of the filter has elapsed.

The third is the stability. FIR filters are always stable because no feedback is utilised, but IIR filters are only stable if the poles of the transfer function are located within the unit circle of the z-plane.

Two filter designs, one FIR and one IIR, will be described.

2.3.4.1.1 He's FIR: In 2008, He [50] designed a 4th-order FIR (response in Figure 2.3) using a least-squares method in MATLAB.

The effect of the error in the frequency weighted filter will be tested in Chapter 4. The design requires 5 multiply-accumulations (MACs) per output sample.

The author also employed a device for reducing computational load in



Figure 2.3: Frequency response of He's $f^{\frac{1}{2}}$ weighted 4th-order FIR filter (calculated from Appendix A in [50]).

the process of squaring and accumulating the filtered samples. Instead of squaring each sample, the absolute is taken, accumulated, and one final square is used. Mathematically, Equation 2.22 is modified to become:

$$M_1 = \frac{\pi}{2(N-M)} \left(\sum_{n=M-1}^{N-1} \left| \left(\sum_{m=0}^{M-1} h[m]s[n-m] \right) \right| \right)^2$$
(2.23)

The reason this can be done is that, if the marginal distribution of the signal is Gaussian with zero mean, then the square and average is the variance of the signal:

$$\frac{\sum_{n=0}^{N-1} (s[n])^2}{N} = \sigma^2$$
 (2.24)

The author [50] shows that, under the conditions above, the average of the absolute process (μ' in his terminology) is proportional to the square root of the variance by a factor k:

$$\mu^{\prime 2} = k\sigma^2 \tag{2.25}$$

where k is dependent on the distribution. For a Gaussian distribution

 $k = \frac{2}{\pi}$. This derivation is given fully in Appendix A. The performance of both square and average and absolute and average will be compared in Chapter 4.

The number of MACs required to evaluate a 5 coefficient FIR is:

$$MAC_{FIR} = 5 \times 1,024 + 1,024 = 6,144$$
 (2.26)

where the additional 1,024 is because of the square and accumulation of the filtered samples.

2.3.4.1.2 Morgan's IIR: In 2011, Morgan et al. [82] designed two 2nd-order IIR filters in cascade. Higher order IIR filters become unstable, hence the cascaded design. As can been seen in Figure 2.4, for 12 MACs per output sample, the frequency response is much closer to the ideal.

The impulse response, shown in Figure 2.5, has reduced to being close to zero by 50 samples. Nguyen [84] starts processing from 16 samples, which seems reasonable as the response is at 0.004 by this point.

The equation for the first moment of the photocurrent's power spectrum using the recursive filter is:

$$M_1 = \frac{1}{N - 16} \sum_{n=16}^{N-1} |y[n]|^2$$
(2.27)

where:

$$y[n] = \sum_{m=0}^{M-1} h_b[m]s[n-m] - \sum_{m=1}^{M} h_a[m]y[n-m]$$
 (2.28)

 h_b and h_a are the filter coefficients for the feed forward and feed back components respectively.

The number of MACs needed to evaluate these two cascading 2ndorder IIR filters is:

$$MAC_{IIR} = 12 \times 1024 + 1024 = 13,312$$
 (2.29)



Figure 2.4: Frequency response of Morgan's 2 stage 2nd-order cascaded IIR filter [82].



Figure 2.5: The impulse response of the IIR filter. The response approaches zero as the number of samples approach 50.

2.3.4.2 The Second Moment

Whilst the first moment is generally used in LDBF applications, in 2010 Binzoni et al. [9] suggested that an expression (Equation 2.30) involving the square root of the second moment (M_2) might provide a

good measure of ' $P_{move}\langle ||\vec{V}||\rangle$ ' (the blood volume times average velocity, their syntax):

$$P_{move}\langle ||\vec{V}||\rangle \propto M_0\left(\sqrt{\frac{M_2}{M_0}}\right)$$
 (2.30)

The advantage of using M_2 is that it is easily evaluated in the time domain using a finite difference method:

$$M_2 \approx \frac{A}{N-1} \sum_{n=0}^{N-2} (|s[n+1] - s[n]|)^2$$
 (2.31)

And M_0 is simply the total power in the signal:

$$M_0 = \sum_{n=0}^{N-1} |s[n]|^2$$
(2.32)

In 2008, Binzoni and van de Ville [10] used a Monte Carlo methodology varying both the Brownian and translational perfusion to show that M_1 is approximately linear to $P_{move}\langle ||\vec{V}||\rangle$, as is generally understood. As the authors then showed in 2010 [9], the expression in Equation 2.30 is also approximately linear. One of the reasons for this may be to do with the shape of the power spectrum. Although the authors themselves did not derive a direct link between M_1 and M_2 , it was highlighted by Koelink et al. [64] that an exact relationship exists under certain conditions, although no proof was given in their article. The proof, that an exact relationship between these moments holds true if the power spectrum is a decaying exponential, is shown here. That the spectrum follows the decaying exponential shape is an observation made by several authors [75, 107], although the accuracy of this widespread assumption is considered in Chapter 3.

A decaying exponential with decay rate 1/m and dependent variable x is used to model the power spectrum. L is the limit of the integration. Firstly the normalisation factor, which is the zeroth moment, must be found:

$$M_0 = \int_0^L e^{-\frac{x}{m}} \, \mathrm{d}x = \left[-me^{-\frac{x}{m}} \right]_0^L = -me^{-\frac{L}{m}} + m \tag{2.33}$$

Then the first moment:

$$M_{1} = \int_{0}^{L} x e^{-\frac{x}{m}} \, \mathrm{d}x = \left[-mx e^{-\frac{x}{m}} \right]_{0}^{L} + m \int_{0}^{L} e^{-\frac{x}{m}} \, \mathrm{d}x \tag{2.34}$$

$$M_1 = -Le^{-\frac{L}{m}} + m \ M_0 = -Lme^{-\frac{L}{m}} - m^2 e^{-\frac{L}{m}} + m^2$$
 (2.35)

And finally the second moment:

$$M_2 = \int_{0}^{L} x^2 e^{-\frac{x}{m}} \, \mathrm{d}x = \left[-mx^2 e^{-\frac{x}{m}} \right]_{0}^{L} + 2m \int_{0}^{L} x e^{-\frac{x}{m}} \, \mathrm{d}x \qquad (2.36)$$

$$M_2 = -L^2 m e^{-\frac{L}{m}} + 2m \ M_1 \tag{2.37}$$

As
$$L \to \infty$$
,

$$\lim_{L \to \infty} M_0 = m$$
$$\lim_{L \to \infty} M_1 = m^2$$
$$\lim_{L \to \infty} M_2 = 2m^3$$

To normalise, both first and second moment are divided by the zeroth. The normalised first moment becomes m, the normalised second moment becomes $2m^2$ and the relationship between them is:

$$\frac{M_1}{M_0} = \frac{1}{2}\sqrt{\frac{M_2}{M_0}}$$
(2.38)

Rearranging:

$$M_1 = \frac{M_0}{2} \sqrt{\frac{M_2}{M_0}}$$
 (2.39)

The relationship between second and first moment will be evaluated in the simulations, testing whether, on this type of signal, there is linearity between the moments.

The number of MACs needed for one output sample is vastly reduced for this algorithm:

$$MAC_{M2} = 2 \times 1,024 = 2,048$$
 (2.40)

A further 1,024 MACs would be needed to evaluate M_0 , but this would normally be calculated alongside all the other algorithms too and therefore it is not included in Equation 2.40.

2.3.4.3 Draijer's Method

In 2009, Draijer et al. [31] proposed a fast time-domain algorithm for evaluating the first moment:

$$M_1 \approx 2\pi \left\langle \left| s(t) \frac{\partial s(t)}{\partial t} \right| \right\rangle$$
 (2.41)

such that the first moment is approximately proportional to the absolute value of the product of the signal and its temporal derivative. They derived this expression by starting with Parseval's theorem, expressed in its purest form:

$$\int_{-\infty}^{\infty} X(f) Y^*(f) df = \int_{-\infty}^{\infty} x(t) y^*(t) dt$$
 (2.42)

The frequency weighted element can be introduced into this equation by noting that:

$$\mathcal{F}\{\frac{\mathrm{d}}{\mathrm{d}t}x(t)\} = jf \tag{2.43}$$

And Equation 2.42, with s(t) = x(t) = y(t) and S(f) = X(f) = Y(f), becomes:

$$\int_{-\infty}^{\infty} jfS(f)S^*(f) df = \int_{-\infty}^{\infty} \frac{d(s(t))}{dt}s^*(t) dt$$
 (2.44)

However the problem with this equation, as described in their paper, is that it evaluates to zero. The authors attempt to overcome this by replacing jf in the left hand side (LHS) of Equation 2.44 with |f| (effectively single siding the frequency spectrum) so that the LHS becomes the first moment definition. The equivalent operation on the RHS, they posit (but without mathematical justification, a fact they acknowledge [31]) is to take the absolute value $\left|\frac{d(s(t))}{dt}s^*(t)\right|$. The time average of this expression is then equivalent to the first moment. In discrete form this is:

$$M_1 \approx \frac{2\pi f_s}{N} \sum_{n=0}^{N-1} |s[n](s[n+1] - s[n])|$$
 (2.45)

The authors [31] test their algorithm empirically, comparing with the standard FFT approach on a phantom and living tissue. Their results suggest it is an effective method for calculating M_1 , and although the algorithm cannot be justified mathematically, the next section will illustrate why it is effective.

Like the 2nd moment calculation, the number of MACs needed is small, but has one additional MAC to account for the additional multiplication:

$$MAC_{DRAIJER} = 3 \times 1,024 = 3,072$$
 (2.46)

2.3.4.4 The Analytic Signal

Having described the work of Draijer et al. [31], this section presents novel work which suggests why Draijer's method is effective, but mathematically unjustifiable, but then uses the 'correct' approach to introduce a method of evaluating the first moment of the photocurrent's power spectrum which has not been applied to LDBF before.

One method of overcoming the zero-value problem presented by Equation 2.44 is to set all negative frequency components to zero. This can be achieved by introducing the analytic signal $s_a(t)$ [23] (Equation 2.47). The analytic signal is complex valued, and is composed of the original signal added to its quadrature in the imaginary plane:

$$s_a(t) = s(t) + j\hat{s}(t)$$
 (2.47)

The average frequency (the normalised first moment) can be evaluated by:

$$\langle f \rangle = \int_{-\infty}^{\infty} f |S(f)|^2 \, \mathrm{d}f = \int_{-\infty}^{\infty} \frac{d(s_a(t))}{dt} \frac{1}{j} s_a^*(t) \, \mathrm{d}t \tag{2.48}$$

which is Equation 1.29 in [21].

In 1998, Huang [55] showed that this equation can be developed as a function of the instantaneous frequency $\dot{\theta}(t)$ and amplitude a(t):

$$\langle f \rangle = \int_{-\infty}^{\infty} \left(\dot{\theta}(t) - j \frac{\dot{a}(t)}{a(t)} \right) a^2(t) d(t)$$
 (2.49)

The instantaneous bandwidth $\frac{\dot{a}(t)}{a(t)}$ is small for signals which are approximately monocomponent [76]. The mean frequency becomes:

$$\langle f \rangle \approx \int_{-\infty}^{\infty} \dot{\theta}(t) a^2(t) \, \mathrm{d}(t)$$
 (2.50)

Equations 2.41 (Draijer [31]), 2.48 (Cohen [21]) and 2.50 (Huang [55]) can be represented before integration as a continuous time signal.

In order to demonstrate the mechanisms of the Draijer algorithm, a signal with similar properties to a LDBF signal was simulated using a monotonically decaying exponential (Equation 2.51) as the power spectrum model [64, 107]:

$$P(f) = \left| K e^{\frac{f}{\langle f \rangle}} \right|^2$$
(2.51)

Because of multiple scattering off static tissue and moving RBCs, the path length and so phases of the electromagnetic waves arriving at the detector are stochastic. A phase spectrum $\theta(f)$ was simulated using the 'rand' function in MATLAB to generate a uniform distribution between $-\pi$ and π .

The bandwidth of Equation 2.51 was set to between 30 Hz and 20 kHz, reasonable for LD spectra, and the sampling rate set to 100 kHz. An inverse fast Fourier transform (IFFT) on 8,192 points yielded a noise-like photocurrent with zero mean. This data is similar to that produced from real LDBF systems.

In Equation 2.48, the signal $s_a(t)$ is complex valued, therefore the quadrature of s(t) was calculated by the Hilbert transform [45,55]. A comparison of the three signals over 9 ms is presented in Figure 2.6. The Draijer method is not precisely equivalent mathematically to the analytic approach, but appears to be similar.



Figure 2.6: A comparison of the three continuous time signals. The similarity of Cohen (blue, vertically offset by 1) and Huang (black) is to be expected, but it is shown here that the Draijer (red) algorithm seems to mimic the analytic approach.

This can be explained by considering that, in the Draijer expression:

$$s(t)\frac{\partial s(t)}{\partial t} = \frac{1}{2}\frac{\partial s(t)^2}{\partial t}$$
(2.52)

The signal s(t) has noise-like characteristics (Figure 2.2) with zero mean. If the signal s(t) is squared, it becomes positive unipolar where each waveform is related to its instantaneous energy. When the squared signal is differentiated, zero crossings are introduced at the peaks and troughs. The 'steepness' of the waveform in the squared signal is related to the frequency. Hence when differentiated, the maximum of the new waveform is proportional to the frequency of the half-periods which form the 'humps' of $\frac{\partial s(t)^2}{\partial t}$, and therefore approximates the instantaneous frequency $\dot{\theta}(t)$. The peaks are also related to the energy in the signal and therefore approximate the instantaneous amplitude $a^2(t)$. The rectification (Equation 2.45) therefore produces a positive signal which approximates $\dot{\theta}(t)a^2(t)$.

Figure 2.7 uses narrowband data (100 Hz) to illustrate. The first signal shows the narrowband s(t). The signal is squared and the Draijer signal shows the additional differentiation and rectification. The value of the peaks are proportional to Equation 2.50, which implies that if integrated over a long enough time period, M_1 may be well approximated by the Draijer approach.

The first moment via the analytic signal can be evaluated discretely by:

$$M_1 \propto \frac{1}{N-1} \left(\sum_{n=0}^{N-2} s[n](\hat{s}[n+1] - \hat{s}[n]) - \sum_{n=0}^{N-2} \hat{s}[n](s[n+1] - s[n]) \right)$$
(2.53)

where $\hat{s}[n]$ represents the discrete analytic signal calculated by Hilbert transform, either in the time domain, by Hilbert filter, or via the FFT in the frequency domain [20]. Both methods of calculating the Hilbert transform will be tested in Chapter 4, alongside the Draijer algorithm.



Figure 2.7: Data with 100 Hz bandwidth was simulated to emphasise how the process of squaring and differentiating approximates the analytic approach. The laser Doppler (LD) signal is in red, with its square in black. The Draijer (green, vertical offset by 1) signal is the differential of the square, and its envelope approximates the Huang signal (blue, vertical offset by 1).

The Hilbert method requires different numbers of MACs depending on whether it is evaluated in the time or frequency domain. In the frequency domain, the FFT is followed by two differentiations, two point-wise multiplications and a final accumulation:

$$MAC_{HILBF} = 10 \times 1,024/2 \log_2 1,024 + 2 \times 1,024 + 2 \times 1,024 + 1,024 = 56,320$$
(2.54)

In the time domain, the FFT is replaced by the Hilbert filter with 3 taps:

$$MAC_{HILBT} = 3 \times 1,024 + 2 \times 1,024 + 2 \times 1,024 + 1,024 = 8,192 \quad (2.55)$$

2.4 Conclusion

Various methods of evaluating the first moment of the power spectrum of the LDBF photocurrent have been described and are summarised in Table 2.1. These methods operate in both time and frequency domain. The domains are theoretically equivalent, however the implementation of transforms, filters and accumulation techniques will affect the resultant accuracy. Some of the methodology makes use of particular properties of the photocurrent. For example, the absolute and average technique makes use of the fact that it is a stochastic signal with Gaussian marginal distribution. A graphical representation of how M_1 may be evaluated from the signal s[n] is shown in Figure 2.8. Additionally, a possible explanation for why the 'Draijer' algorithm is effective has been suggested. The explanation makes use of the analytic signal, and this may also be an effective method of determining M_1 . The accuracy of the algorithms described here in evaluation of M_1 will be compared in simulations in Chapter 4, and plotted against the computational complexity as defined by the MAC number.

Name	Domain	Algorithm
DFT [82,85]	F	$\frac{f_s}{N} \sum_{k=2}^{N/2+1} k S_{av}[k] , \qquad L = 1, 2,, 64$
FIR square and average [50]	Т	$\frac{1}{N-M} \sum_{n=M-1}^{N-1} \left \left(\sum_{m=0}^{M-1} h[m]s[n-m] \right) \right ^2$
FIR absolute and average [50]	Т	$\frac{\pi}{2(N-M)} \left(\sum_{n=M-1}^{N-1} \left \left(\sum_{m=0}^{M-1} h[m]s[n-m] \right) \right \right)^2$
IIR [82]	Т	$\frac{1}{N-16}\sum_{n=16}^{N-1} y[n]^2$,
		$y[n] = \sum_{m=0}^{M-1} h_b[m]s[n-m] - \sum_{m=1}^{M} h_a[m]y[n-m]$
2nd moment DFT	F	$\frac{f_s}{N} \sum_{k=2}^{N/2+1} k^2 S_{av}[k]$
Binzoni's 2nd moment [9]	Т	$\frac{A}{N-1}\sum_{n=0}^{N-2}(s[n+1] - s[n])^2$
Draijer [31]	Т	$\frac{2\pi f_s}{N-1} \sum_{n=0}^{N-2} s[n](s[n+1] - s[n]) $
Analytic [23]	Т	$\frac{1}{\frac{1}{N-1} \left(\sum_{n=0}^{N-2} s[n](\hat{s}[n+1] - \hat{s}[n]) - \sum_{n=0}^{N-2} \hat{s}[n](s[n+1] - s[n])\right)}{\frac{1}{N-1} \left(\sum_{n=0}^{N-2} s[n](\hat{s}[n+1] - s[n])\right)}$

Table 2.1: List of algorithms for testing. Domain signifies whether algorithm processes in time domain (T) or frequency domain (F).



Figure 2.8: An overview of the design routes available to evaluate M_1 . The boxes are colour-coded as follows: blue for DFT based methods, yellow for analytic signal, green for Draijer, red for time domain filters, purple for 2nd moment.

Chapter 3

Modelling the Power Spectrum of the LDBF Photocurrent by Bayesian Inference

3.1 Introduction

This chapter describes how the power spectrum of a LDBF photocurrent can be modelled using a probability distribution whose shape is controlled using a set of parameters. Several authors have noted the similarity of LDBF power spectra to an exponential distribution [107]: in 1999 Lohwasser and Soelkner [75] stated that 'The Doppler frequency spectrum is assumed to be given by the functional relation [of the decaying exponential]'. However, as far as is known, model fits based on measurements on human skin tissue have still to be performed; a state rectified in this chapter.

Also introduced is the concept of Bayesian Inference, the machinery used to evaluate, from the data, which is the best-fit model. The methodology of data collection will be explained, along with the spectral conditioning and other processing required to perform the inference.

3.2 Experimental Design

3.2.1 Data Requirements

To model the power spectrum of the photocurrent, real spectra obtained in vivo over a range of flow values are required. It is known that the photocurrent is a stochastic process [12] and as such the Fourier transform will also be stochastic. Modelling the underlying frequency content of this random process necessitates a long run of data to average out variability. In order to obtain such data, an experiment was designed to capture raw blood flow photocurrents with gradually increasing flow.

3.2.2 Ethical Considerations

Ethical approval for the collection of photocurrents from human participants under varying thermal conditions was granted by the University of Nottingham Faculty of Engineering (Application ID 2012-01). Data was collected and stored in accordance with Faculty requirements and no personally identifiable information was stored with the photocurrent data. Three participants were recruited for 20 minute recordings. This allowed collection of a large number of power spectra for the model fit. There were no specific human requirements or exclusion factors.

3.2.3 Recording Devices

A Moor Instruments Laser Doppler Line Scanner (LDLS) was used as one of two recording instruments. A line of approximately 9 cm length and 2 mm width is illuminated by a laser of 785 nm wavelength. The reflected and Doppler shifted light is detected by 64 spatial channels across the line length. Although the device has the capability to move the line by way of a motorised mirror, for the purposes of this experiment it was held steady in one location. The device samples all channels simultaneously at 40 kHz but it can only record 1024 samples per buffer (T=0.0256 s). However, it is able to repeat this measurement at a rate of 10 Hz. Figure 3.1 illustrates the temporal and spatial sampling characteristics of the device.



Figure 3.1: 64 channels sample a line of approximately 9 cm. Each channel samples at 40 kHz and takes 1024 samples, giving T=0.0256s. The photocurrents are not continuous, but under static flow conditions concatenated photocurrents may be considered stationary.

A Moor Instruments VMS-LDF2 single point monitor was used alongside the LDLS providing a blood flow comparison standard. It also used a laser of 785 nm wavelength. The VMS-LDF2 probes also had the capacity to measure skin surface temperature, important in monitoring the reflex vasodilatory response (refer to Section 3.2.4). The device does not give raw photocurrent output, but serves to validate the flow and temperature changes in the experiment. The positioning of the probe head next to the line is shown in Figure 3.2. The output is a calculated blood flow value in arbitrary units (A.U.).

3.2.4 In Vivo Data Collection

Whilst recording LD photocurrents, the microvascular blood flow was altered by first immersing the recording hand (the right hand) up to the wrist in 15°C water for 2 minutes. This causes vasoconstriction of



Figure 3.2: The laser line from the LDLS imaged a section on the back of the hand. The VMS-LDF2 temperature and flow probe was positioned approximately 1 cm behind.

the microvasculature, restricting the movement of blood. The hand was then dried, positioned under the LDLS laser line and the probe attached as in Figure 3.2. The time between the end of the cold immersion and start of the LDLS recording was 1 minute.

In 2003, Daanen noted that 'When one body part is cooled, vasoconstriction also occurs in other parts of the body. This phenomenon is known as reflex vasoconstriction. Similarly, if heat is applied to another part of the body... the vessels open up... This phenomenon is known as reflex vasodilation.' [24] To increase blood flow in the recorded hand, the contralateral hand was immersed to the wrist in 43°C water and held for 20 minutes. During this 20 minute period the LDLS and VMS-LDF2 recorded the raw photocurrent and blood flow and temperature respectively in the non-immersed hand.

3.3 Models

It is widely accepted that the power spectrum from a LDBF device resembles a decaying exponential e.g. [75,107]. The criteria for selection of model spectra were that they should follow this monotonically
decreasing shape with increasing frequency and the distribution should exist on all positive frequencies. A set of standard probability distribution functions fulfilling these criteria was selected. All the models have an x-axis scale and y-axis scale parameter, and additionally some have a shape parameter. If the distributions have a location parameter, this is set to 0 so that the peak of the distribution is centred. Table 3.1 details the definition of the models, with examples displayed in Figure 3.3.

No.	Name	Number of	Definition $P_M(f)$
		Parameters	
1	Exponential	2	$\theta_1^2 e^{-\frac{f}{\theta_2}}$
2	Exponential with powered exponent	3	$\theta_1^2 e^{-\left(\frac{f}{\theta_2}\right)^{\theta_3}}$
3	Cauchy	2	$\frac{\theta_1^2}{\pi\theta_2 \left[1 + \left(\frac{f}{\theta_2}\right)^2\right]}$
4	Type 3 Logistic	2	$\frac{\theta_1^2 e^{-\frac{f}{\theta_2}}}{\theta_2 \left[1 + e^{-\frac{f}{\theta_2}}\right]^2}$
5	Type 2 Logistic	3	$\frac{\theta_1^2 e^{-\frac{\theta_3 f}{\theta_2}}}{\theta_2 \left[1 + e^{-\frac{f}{\theta_2}}\right]^{\theta_3}}$
6	Type 3 Logistic	3	$\frac{\theta_1^2 e^{-\frac{\theta_3 f}{\theta_2}}}{\theta_2 \left[1 + e^{-\frac{f}{\theta_2}}\right]^{2\theta_3}}$
7	Generalised Pareto	3	$\theta_1^2 \left(1 + \frac{f}{\theta_2}\right)^{-\frac{1}{\theta_3} - 1}$
8	Generalised Pareto (with fixed shape)	2	$\theta_1^2 \left(\overline{1+\frac{f}{\theta_2}}\right)^{-\frac{1}{1.74}-1}$
9	Pareto	3	$\frac{\theta_1^2}{\left(\frac{f}{\theta_2}\right)^{\theta_3+1}}$

Table 3.1: Models selected for comparison. All decrease monotonically with increasing frequency. Model 8 is model 7 but with a fixed shape parameter, determined experimentally by the mean of θ_3 .

The exponential distribution has been used to describe the LDBF power spectrum. It is included first with 2 parameters, where increasing θ_2 increases the power at higher frequencies. θ_1 is, as is the case for the other distributions, the y-axis scale. A three parameter exponential has an additional shape parameter, θ_3 , which increases the negative



Figure 3.3: Example plots of the models from table 3.1. All have a monotonically decreasing shape but the precise shape and distribution of power across the frequency range varies.

gradient of the exponential slope. The Cauchy distribution includes two θ_2 x-axis scale parameters. This has the effect of reducing power at lower frequencies and increasing the power at higher frequencies as θ_2 increases, for fixed y-axis scale. The logistic family of distributions, which has a wide variety of applications in sociology and biology [83], is similar in shape to the exponential (being constructed of exponential expressions) but has a larger tail. The Pareto distributions are based on a 1/f distribution. In initial tests, the Generalised Pareto (model 7) was a likely model and it may be that the third parameter can be replaced with a fixed value which then approximates to the correct shape across all spectra. To test this, a Generalised Pareto with fixed shape was also included (model 8) with the fixed value of the third parameter determined by performing the analysis with the three parameter distribution and selecting the mean value of θ_3 . This value was 1.74.

One usual condition of probability distribution functions is that the integral of the distribution (i.e. the sum of probabilities) equals 1. Because the y-axis scale parameter, θ_1 , is included (so that varying total power can be measured), this condition is no longer true. This means that any other parameters which also scale the distribution independent of θ_1 can be ignored. The distributions are then reduced to the two or three parameters which, when searched over, vary the scale or shape of the distribution. For example, the formula for the Generalised Pareto is:

$$p(x) = \frac{1}{\sigma} \left(1 + \frac{\xi(x-\mu)}{\sigma} \right)^{(-\frac{1}{\xi}-1)}$$
(3.1)

The $1/\sigma$ is absorbed because θ_1 is now included, and since $\mu=0$ the peak is at zero. $\frac{\sigma}{\xi}$ becomes θ_2 as σ now only exists once in the expression. Replacing ξ with θ_3 , the formula becomes:

$$p(x) = \theta_1^2 \left(1 + \frac{x}{\theta_2} \right)^{(-\frac{1}{\theta_3} - 1)}$$
(3.2)

This process of reforming the parameters was repeated for all the distributions.

3.4 Bayesian Model Selection

Bayesian model selection can be used to find the most likely model fit data from a range of models independent of the model parameters. It can perform a fair comparison even if the models have different numbers of parameters, penalising more complex models with higher numbers of parameters according to the principle of Occam's Razor [14].

The ordinary least squares regression is a popular and relatively simple data fitting solution [94]. However, in the case of a LDBF photocurrent, which is a wide-bandwidth stochastic signal, the distribution of the power spectra, calculated by DFT, is heteroscedastic because the variance is correlated with the square of the power (Figure 3.4). In 1968, Jenkins and Watts [58] showed this fact and it will also be demonstrated in Section 3.5. This renders ordinary least squares, which



Figure 3.4: An example of heteroscedastic data - the variance of the power spectrum is not constant across the data. Least squares regression requires homoscedastic data, but the power spectra of LDBF is heteroscedastic.

requires constant variance, unsuitable for the task. However, Bayesian model fitting can calculate the probability at each individual frequency point, by using independent variances. It starts with Bayes' theorem:

$$p(B|A, I) = \frac{p(A|B, I)p(B|I)}{p(A|I)}$$
(3.3)

p(B|A, I) is termed the posterior and is defined as the probability of B given some value or condition A and prior information I. Then p(A|B, I) is the probability of A given B and termed the likelihood function. The two probabilities p(A|I) and p(B|I) are the probabilities of A and B independent of any condition (but based on prior information). The model selection problem considered in LDBF is illustrated in Figure 3.5.

Applying Bayes' theorem:

$$p(M|P(f), I) = \frac{p(P(f)|M, I)p(M|I)}{p(P(f)|I)}$$
(3.4)

the goal is to find which M produces the largest value of the posterior



Figure 3.5: Bayesian Model Selection on Laser Doppler data. Given a signal s(t), and thus a power spectrum P(f), p(M|P(f))is the probability that a model M describes that power spectrum. The most likely model is then the one with the highest probability.

probability for a given P(f). I, any prior information known about the data, will be omitted from all future equations because it does not now formally contribute to the mathematical derivations. p(P(f)) is a normalising function, equal to:

$$\int p(P(f)|M)p(M)\,\mathrm{d}M$$

and is constant for a fixed P(f). The main computational term, the likelihood function, is:

$$p(P(f)|M) \tag{3.5}$$

So, given a certain M, p(P(f)) is the probability of obtaining the spectrum P(f) for any parameters. Let M be a model with two parameters θ_1 and θ_2 . Then the probability of obtaining the spectrum P(f) at some set of parameter values is $p(P(f)|\theta_1, \theta_2, M)$. The requirement, however, is to evaluate the probability independent of the parameter values. This can be achieved by the process of marginalisation, essentially integrating over all possible parameter values:

$$p(P(f)|M) = \iint_{\Delta\theta_1 \Delta\theta_2} p(\theta_1|M) p(\theta_2|M) p(P(f)|\theta_1, \theta_2, M) \,\mathrm{d}\theta_1 \,\mathrm{d}\theta_2 \quad (3.6)$$

where $\Delta \theta$ is the range of possible parameter values searched over (later, $\delta \theta$ will refer to the gap between adjacent discrete parameter values). Note that in equation 3.6 two new terms are introduced, the independent probabilities of the parameter values, or priors. As no information about these probabilities is known in advance of the experiment, they are referred to as uninformative priors and are constant over the parameter range. The probability of a particular value of the parameter occurring is therefore

$$p(\theta|M) = \frac{1}{\Delta\theta}$$
(3.7)

It has been the subject of much debate over the years whether 3.7 actually represents equal parameter probability. In fact, where the scale of the parameters is also an unknown (e.g. is the parameter in the range 0.1 to 1 or 1000 to 10,000?), the flat prior does not assign equal probability per decade [46]. In that case, the Jeffrey's prior is more suitable because it uses a logarithmic scale. However, in the case of these results, where the spectra all fall within a roughly similar power level and shape, and where the posterior is finite and discrete, the flat prior is suitable [57].

The effect of equation 3.7 is to penalise probabilities with higher numbers of parameters: it acts as an Occam's Razor and simpler models are preferred. Many of the models tested have 3 parameters instead of 2. For a model with higher number of parameters to be more likely, the quantity $p(P(f)|\theta_1, ..., \theta_i, M)$ needs to be proportionally larger (by an extra $\Delta \theta$ for each additional parameter) so as to compensate for the penalisation.

The difficulties posed in performing these integrations analytically lend favour to a numerical approach. If P(f) is sampled with frequency resolution $\frac{f_s}{N}$ (where f_s is the sampling frequency and N the number of time domain samples), it consists of samples $P(f_1), P(f_2), P(f_3), \dots, P(f_{N/2})$ up to the Nyquist frequency. Expression 3.5 for certain parameter values becomes

$$p(P(f_1), P(f_2), P(f_3), ..., P(f_{N/2})|\theta_1, \theta_2, M)$$
 (3.8)

$$= p(P(f_1)|\theta_1, \theta_2, M), p(P(f_2)|\theta_1, \theta_2, M), p(P(f_3)|\theta_1, \theta_2, M), ..., p(P(f_{N/2})|\theta_1, \theta_2, M)$$
(3.9)

$$=\prod_{n=1}^{N/2} p(P(f_n)|\theta_1, \theta_2, M)$$
 (3.10)

 $p(P(f_n)|\theta_1, \theta_2, M)$ is the probability of the power at a particular frequency f_n occurring given the model M and parameters θ_1 and θ_2 : i.e. it is the probability of an error from the mean, $P_M(f_n)$ (the value calculated from the model function using θ_1 and θ_2).

Before any further derivations are performed, the properties of the power spectrum and error function require discussion.

3.5 Power Spectra Conditioning

3.5.1 Overview

The aim of this section is to show how a LDBF photocurrent may be processed such that Bayesian Inference can be used to model-fit its power spectrum. This is important because a single power spectrum calculated from a LDBF signal is not, by itself, suitable for the model-fitting exercise.

Maximal probability at the correct point is required. That is to say, where the true power of a certain frequency (an unknown) resides, the probability distribution function (PDF) of the error model of the power should be at its maximum. A spectral estimate, calculated from a time domain signal from LDBF, does not, however, form a PDF with this property because it is a white noise-like stochastic process and, importantly, the variance of the spectral estimate is unknown. This section will show that this is the case, and then how the spectral estimates can be manipulated to become suitable for Bayesian analysis.

3.5.2 LDBF Theoretical Spectrum

The LDBF photocurrent can be considered to be a white noise process passed through a linear filter. Given that: there is a large quantity of red blood cells and a small chance of multiple scattering events off RBCs [12]; the frequencies generated are proportional to the velocities of the RBCs [114] and the velocity distribution has a Brownian component [7]; the heterodyning of Doppler frequencies with reflected coherent light dominates the photocurrent [88] and the phase is random according to the path length through the scattering medium (tissue); then this seems reasonable for the noise-like photocurrent. The frequency response of the 'filter', H(f), then determines the theoretical spectrum, B(f), such that:

$$B(f) = \sigma_z^2 |H(f)|^2$$
 (3.11)

where σ_z^2 is the variance of the white noise process. By first considering properties of unfiltered white noise (B(f)=1), the properties of filtered white noise can be derived and checked for consistency.

3.5.3 Properties of Spectral Estimates of White Noise

White noise is a Gaussian process with infinite bandwidth and zero mean and as such, the Fourier transform, and so power spectrum, is stochastic. This means that if consecutive Fourier transforms (squared to produce power) are performed on a white noise process, single spectral estimates selected from each power spectrum will form a PDF i.e. they will not be constant. White noise was simulated by taking the inverse Fourier transform of a very well sampled frequency spectrum (2²² frequencies, with unity power across the bandwidth) and uniformly random phase. This long signal was broken up into smaller processes and the Fourier transforms taken (Figures 3.6 and 3.7). In this analysis C(f) is an estimate of the power of a frequency f upon taking the square of the Fourier transform of a white noise time series. $\overline{C}(f)$ is an estimator of the power of a frequency f after averaging multiple sets of C(f):

$$\bar{C}(f) = \frac{1}{L} \sum_{l=1}^{L} C_l(f)$$
(3.12)





Figure 3.6: A signal of length $T_{\text{Total}} = 2^{22}$ was broken up into K sections of length T, signified by $s_{Ti}(t)$. K is controlled by the length of T as T_{Total} is constant.



Figure 3.7: Each $s_{Ti}(t)$ is broken up into L sections of length N and the Fourier transforms taken.

Therefore one power spectrum estimator distribution is formed for each T giving K estimators in total. By simulating in this way, it is shown in Figure 3.8 than the PDF of one frequency $(C(f_x))$ of the power spectrum of this process matches the theoretical Chi-squared with 2 degrees of freedom (χ_2^2) . Effectively this is the case for when L = 1. But when the power spectra from multiple Fourier transforms are averaged $(\bar{C}(f)$ when L = 32), the PDF's degrees of freedom increase (by twice the number of averages, in this case 64) and a χ_{64}^2 distribution is created (Figure 3.9).

So, by averaging the spectra, not only does the estimate move closer to the mean, but the distribution changes shape and becomes approximately Gaussian.



Figure 3.8: Simulated distribution of $2\frac{C(f)}{B(f)}$ for one frequency of a white noise process, with the theoretical χ_2^2 . The theory of a white noise process predicts that the variance is proportional to the value of $B(f)^2$ [58].

3.5.4 The Laser Doppler Spectrum as Filtered White Noise

The more white noise is filtered, the less Gaussian it becomes. This is intuitively true since an ideal filter with a delta frequency response will leave a perfect sinusoid. However the filtering in this case is relatively wide band, and the signal still has a Gaussian marginal distribution.



Figure 3.9: Simulated distribution of $64\frac{\bar{C}(f)}{B(f)}$ for one frequency of a white noise process after 32 averages. Upon averaging, the degrees of freedom of the χ^2 distribution is increased and the distribution becomes approximately Gaussian.

3.5.4.1 The Three Distributions

The white noise spectrum of the previous simulation was modified to resemble an exponential distribution, as the spectrum of LDBF is of similar shape [107]. The filter H(f) has the effect of slightly reducing the smoothness of the theoretical spectrum B(f) as it changes from being completely flat to decaying. This is important to note when evaluating the bias of the averaged spectral estimator.

There are three separate distributions to consider. The first is the distribution of C(f), the spectral estimator with no averaging (L = 1). It should resemble a χ_2^2 (Figure 3.10). The second is the averaged spectral estimator, which should resemble a χ_{ν}^2 where ν is twice the number of averages (Figure 3.11). The third is the final distribution of the power spectrum across the bandwidth; this is the distribution for the model fitting exercise (the theoretical version of which is Figure 3.12). The distributions are not to be confused; it is a coincidence that the χ_2^2 actually resembles the LDBF power spectrum, in fact the χ_2^2 is a

decaying exponential by definition. However the first two distributions, the χ , are the distributions of spectral estimators at *one frequency only* and it is only the third, made up of multiple estimators that is actually the full power spectrum of the LDBF process.



Figure 3.10: Distribution $\frac{2C(f)}{B(f)}$ of a frequency f, $\frac{1}{8}$ th of the Nyquist frequency $\frac{f_s}{2}$ with the theoretical χ_2^2 .

There are several choices to be made. First, the number of averages L that make up $\bar{C}(f)$, second, the length of N (i.e. how many samples to take) that make up the total time period T = NL and third, the spectral window to apply to the time series.

3.5.4.2 Number of Averages

The value of L determines the shape of the resultant estimator PDF and decreases the variance. The central limit theorem dictates that upon the addition of independent identically distributed samples, such as the addition of multiple χ_2^2 distributions, the distribution of the estimator is forced towards a Gaussian. As L increases, the degrees of freedom increases by 2L. This is because the degree of freedom increases by one as a sine or cosine (independent Normal distributions) of the Fourier transform is added. One estimate C(f) has 2 degrees of freedom as:



Figure 3.11: Distribution $\frac{64\bar{C}(f)}{B(f)}$ of a frequency f, $\frac{1}{8}$ th of the Nyquist frequency $\frac{f_s}{2}$ with the theoretical χ^2_{64} . The number of averages was 32 and the number of samples in the distribution K = 4096. There is a small bias in the result, the reasons for which are discussed in Section 3.5.4.3.



Figure 3.12: Example LDBF theoretical spectrum, B(f), which resembles a decaying exponential.

$$C(f) = A^2 + B^2$$
 (3.13)

where A and B are the sine and cosine components of the Fourier

transform. A and B form normal distributions with zero mean. Squaring them results in a χ_1^2 distribution and the addition of the two results in χ_2^2 under the additive property of the χ_{ν}^2 distribution:

$$\chi^2_{\nu_1+\nu_2} = \chi^2_{\nu_1} + \chi^2_{\nu_2} \tag{3.14}$$

Hence the addition of multiple transforms increases the degrees of freedom. As ν increases, the skewness property $(\sqrt{\frac{8}{\nu}})$ and the kurtosis property $(\frac{12}{\nu})$ decrease and the distribution tends to normality. This is illustrated in Figure 3.13 by theoretical χ^2_{ν} for different ν against a normal PDF with the same mean and variance $(\mu = \nu \text{ and } \sigma^2 = 2\nu)$.

Recalling that $\nu = 2L$, as the number of averages moves above 32 the distribution very much resembles the normal (Figure 3.13e). It is also worth noting that the variance of the estimator $\bar{C}(f)$ is not dependent on ν because the χ^2_{ν} distributions in Figure 3.13 are distributions of $\frac{\nu \bar{C}(f)}{B(f)}$. That is, despite the χ^2 distribution variances increasing by 2ν , the dependence of the variance on ν is eliminated by the multiplication of $\frac{\bar{C}(f)}{B(f)}$ by ν .

3.5.4.3 Data Length and Window

If L is large then this decreases the available number of samples N to produce the estimator $\overline{C}(f)$ in the same period T (T = NL). The choice of N and the spectral window affects the bias as the bias is dependent on the smoothness of the spectrum with respect to the window size in the frequency domain. A larger N therefore reduces the bias as the spectrum looks 'smoother' at smaller Δf . A window with narrow main lobe and attenuation of the side lobes will decrease the bias as the effect of neighbouring frequencies is reduced. The ideal window from a bias point of view is a delta function at the frequency f, but this requires infinite N. As discussed below, the variance decreases by a factor kwith wider windows, and the choice of window is a function of spectral



Figure 3.13: χ^2_{ν} and \mathcal{N} PDFs with same mean and variance highlight how the estimator PDF tends to normal as L is increased.

smoothness. A Hamming window was chosen for this simulation (note that to obtain correct estimates the spectral values must be normalised by the mean square of the spectral window). The effect of the length of N on the bias is shown in the following Figures (Figures 3.14 to 3.16). These figures also illustrate the third distribution, the power spectrum, against its theoretical values.



Figure 3.14: The expected value of the estimator $\overline{C}(f)$ as a function of frequency with the theoretical spectrum B(f) with length N = 32. The number of averages in the expected value function was 4096.

The number of averages of the estimator decreases as N increases (4096 to 2048 to 1024) but this does not affect the estimator value. The estimator converges well before 1024 averages is reached. It can be seen that the bias decreases with increasing N.

However, as these figures show the expected value of the estimator distribution, they do not reveal the variance. Contrary to intuition, increasing N does not decrease the variance and C(f) is not a consistent estimator of the power spectrum. It can be said that $\bar{C}(f)$ is a consistent estimator because as $L \to +\infty$ the variance tends to zero (Equation 3.15) and therefore a trade-off between bias and variance exists for the same T.



Figure 3.15: The expected value of the estimator $\overline{C}(f)$ as a function of frequency with the theoretical spectrum B(f) with length N = 64. The number of averages in the expected value function was 2048.



Figure 3.16: The expected value of the estimator $\bar{C}(f)$ as a function of frequency with the theoretical spectrum B(f) with length N = 128. The number of averages in the expected value function was 1024.

$$\operatorname{Var}[\bar{C}(f)] = k \frac{B(f)^2}{L}$$
(3.15)

k is a factor affected by the choice of spectral window: windows with higher bandwidths have lower k [58]. For the Hamming window, k = 0.7353.



Figure 3.17: An illustration of the variance of the smoothed spectral estimator for frequency bin 3 where N=64 and the number of averages is 32. The bias is 0.007, with $\sigma=0.021$.

Figure 3.17 reveals the distribution of the estimator (essentially the χ_{64}^2 unnormalised by $\frac{\nu}{B(f)}$) with the variance being proportional to $B(f)^2$. The fact that the variance is not independent of the spectral power requires the normalisation of the exponential in the normal distribution by the variance of the power at each frequency. This will be discussed in the next section as the Bayesian derivation continues. For the purposes of the analysis, the spectral estimator from the Gaussian process will be made equivalent to the power spectrum for Bayesian Inference (equation 3.16).

$$P(f) \equiv \bar{C}(f) \tag{3.16}$$

3.5.5 Summary

It has been shown that single spectral estimates from an LDBF photocurrent form a PDF where the point of maximal probability is not equivalent to the true theoretical power spectrum. However, by averaging multiple power spectra from a stationary photocurrent, the error model becomes more Gaussian and the point of maximal probability approaches the true theoretical point. Now the error model can be approximated by a Gaussian distribution.

3.6 Further Bayesian Derivation

3.6.1 Likelihood function

The effect of averaging, by the central limit theorem, is to approximate a Gaussian distribution. Hence the probability term in equation 3.10 can be evaluated by:

$$p(P(f_n)|\theta_1, \theta_2, M) \approx \frac{1}{\sigma_n \sqrt{2\pi}} e^{-\frac{(P(f_n) - P_M(f_n))^2}{2\sigma_n^2}}$$
 (3.17)

Where:

$$\sigma_n^2 = k \frac{P_M(f_n)^2}{L} \tag{3.18}$$

The likelihood function becomes:

$$p(P(f)|\theta_1, \theta_2, M) = \prod_{n=1}^{N/2} \frac{1}{\sigma_n \sqrt{2\pi}} e^{-\frac{(P(f_n) - P_M(f_n))^2}{2\sigma_n^2}}$$
(3.19)

At this stage, it is usually easier to work with the log of likelihood because the terms can be summed instead of multiplied:

$$\ln(p(P(f)|\theta_1, \theta_2, M)) = \sum_{n=1}^{N/2} \left[-\ln(P_M(f_n)) + \frac{1}{2} (\ln(L) - \ln(2k\pi)) - \frac{L \cdot (P(f_n) - P_M(f_n))^2}{2k P_M(f_n)^2} \right]$$
(3.20)

3.6.2 Model distribution

The model function (assuming a two parameter model) is:

$$P_M(f) = g(\theta_1, \theta_2, M) \tag{3.21}$$

This presents a potential problem. By varying θ_i not only does the comparison term in the exponential in equation 3.19 change, but so does the variance (from equation 3.18). This means that, as the parameters are searched, the variance controlling the likelihood probability also causes a change in the probability outcome. Now the distribution of $P(f_n)$ about a particular model value is normal (equation 3.17), but as the likelihood function involves varying the model values with fixed data, it is interesting to see the distribution of model values $P_M(f_n)$ about a particular $P(f_n)$.



Figure 3.18: The distribution of $P_M(f_x)$ around $P(f_x) = 0.5$. The location of the peak is 0.4924, 1.5% difference. The variance constant L/k = 64.

Figure 3.18 illustrates the result of changing variance, that is, a slightly skewed normal distribution with a small bias. The bias is easily

calculable by evaluating:

$$\frac{\mathrm{d}}{\mathrm{dP}_{\mathrm{M}}} \left(-\ln(P_{M}) - \frac{L(P - P_{M})^{2}}{2kP_{M}^{2}} \right) = 0$$
 (3.22)

where the function arguments have been neglected for ease of reading. This will calculate the maximum of the PDF for P_M and allow verification of the difference between P and P_M . Equation 3.22 evaluates to:

$$P_M^2 + \frac{P.L}{k} P_M - \frac{P^2.L}{k} = 0$$
 (3.23)

and solving for P_M :

$$P_M = \frac{P}{2} (\pm \sqrt{\frac{L}{k}(\frac{L}{k}+4)} - \frac{L}{k})$$
 (3.24)

Setting $\alpha = \sqrt{\frac{L}{k}(\frac{L}{k}+4)} - \frac{L}{k}$ and by taking the difference and percentage:

Bias =
$$100 \frac{P - P_M}{S} = 100 \left(1 - \frac{\alpha}{2}\right)$$
 (3.25)

From equation 3.25 the bias is a constant percentage of P controlled by the constant k and the number of averages, decreasing as L increases. This is illustrated in Figure 3.19 and gives another reason for keeping the number of averages high. The small bias at larger numbers of averages should not adversely affect the results of the model selection problem. For the purposes of this model fitting exercise, L was selected to be 32. This number represents a reasonable number of averages considering the length of photocurrent required, and a bias percentage of 2.2% according to Figure 3.19.

3.6.3 Numerical Marginalisation

Equation 3.6 (from Section 3.4) should be implemented numerically on a model by model basis. $p(P(f)|\theta_1, \theta_2, M)$ may be evaluated numerically as in Equation 3.20, replacing $P(f_n)$ with the spectral data P[n] and



Figure 3.19: The bias percentage as a result of changing variance in the likelihood function. At $\frac{L}{k} = 64$ the percentage is 1.5, matching Figure 3.18. At $\frac{L}{k} = 43.5$ (L=32), the percentage bias is 2.2.

 $P_M(f_n)$ with a model function calculated over the chosen parameter ranges. Thus a 2-D or 3-D map of values calculated from equation 3.20 is built up. Visualising a 2-D map, a probability function is created where the maximum gives the coordinates (i.e. parameters) of the most likely distribution from the model. Marginalisation is the process of integrating over this map. The most likely model is the one with highest summed likelihood, not the one with tallest peak. For three parameters this is:

$$p(P[n]|M) = \frac{1}{A'B'C'} \sum_{a=1}^{A} \sum_{b=1}^{B} \sum_{c=1}^{C} \exp(\ln(p(S[n])|\theta_1[a], \theta_2[b], \theta_3[c], M))$$
(3.26)

where A, B and C are the numbers of parameters evaluated, A'=A-1etc. and $\theta_1[a]$, $\theta_2[b]$ and $\theta_3[c]$ are functions which calculate the parameters from the indices a, b and c. $\frac{1}{A'B'C'}$ acts as the Occam factor. This is because the integration is numerical, so that the summation should be multiplied by the gap $\delta\theta$ between the parameters to give the full area. In this case, the midpoint rule for numerical integration is used as the function is approximately linear at small $\delta\theta$ and the function does not vary sinusoidally. The gap for parameter 1, for example, can be evaluated by:

$$\delta\theta_1 = \theta_1[a+1] - \theta_1[a] = \frac{\theta_1[A] - \theta_1[1]}{A-1}$$
(3.27)

and the parameter range is:

$$\Delta \theta_1 = \theta_1[A] - \theta_1[1] \tag{3.28}$$

From equation 3.7, the summation should be divided by equation 3.28 for each parameter. This leaves:

$$\frac{\theta_1[A] - \theta_1[1]}{A - 1} \cdot \frac{1}{\theta_1[A] - \theta_1[1]} = \frac{1}{A - 1} = \frac{1}{A'}$$
(3.29)

Equation 3.26 will be evaluated for each model on each spectrum. It is assumed that the prior probabilities of the models are equal, as are the prior probabilities of the parameters. The posterior can thus be determined purely from the likelihood function.

3.7 Data Processing

This section shows how the data files received from the two Moor Instruments devices are processed in accordance with Bayesian model selection methodology.

3.7.1 Raw Data Files

3.7.1.1 Moor VMS-LDF2

This device was used as a validator of the experimental method, to show an increase of blood flow and skin surface temperature due to reflex vasodilation. The device already processes the photocurrent into a value representing blood flow, therefore this value can be viewed directly. The output sample rate f_s^{bf} is 40 Hz, so that over 20 minutes more than 48,000 blood flow samples were recorded. The blood flow is smoothed for ease of viewing long term variation by applying a 400 point (10 s) moving average filter, and the results for subjects 1, 2 and 3 can be seen in Figures 3.20, 3.21 and 3.22 respectively. The temperature is displayed as recorded by the point probe.



Figure 3.20: VMS-LDF2 flow and temperature recording on subject 1. The recorded hand was immersed for the first 2 minutes to prompt vasoconstriction, followed by 1 minute of drying and positioning. At 3 minutes, the contralateral hand was immersed in 43°C water. The results show an increase of flow, coinciding with an increase in skin surface temperature.

The flow results are not smooth functions, but have medium term temporal variations. In some cases this was simply due to movement artefacts - if the fibre optic light guide was moved this affected the noise level substantially. There are also respiratory and myogenic variations [66] and resistive mechanisms on top of the heart beat. Stefanovska et al. [116] found five oscillations within the range 0.0095-1.6 Hz by LDBF. These may have contributed to the variations seen in the VMS-LDF2



Figure 3.21: The same experiment as in Figure 3.20 repeated for subject 2. The flow increase is not as large as subject 1 at this site, but the temperature still shows an increase of 4°C.



Figure 3.22: The same experiment as in Figure 3.20 repeated for subject 3. Again the temperature increases by roughly 4°C but the flow remains stationary.

results. An FFT on subject 1 blood flow data (Figure 3.23) illustrates the frequency bands present in the signal.

There is some variation in gradient of flow increase. All show a skin surface temperature increase of 3-5°C. In subject 1, this appears to be accompanied by a flow increase of similar gradient. Subject 2 shows a slight increase, but subject 3 seems to have a relatively flat flow response over the 20 minutes. This could be due to variation in haemodynamic responses to heat provocation between subjects, but it may also be due to probe positioning. The results from the Moor LDLS add some understanding.



Figure 3.23: FFT of four minutes of data from subject 1 shows the low frequency bands present in the blood flow signal. The colours correspond to physiological frequencies as follows: purple is cardiac, yellow is respiratory, green is myogenic, light blue is sympathetic and grey is endothelial. Image was captured using the Moor Instruments VMS-PC software.

3.7.1.2 Moor LDLS

This device, in contrast with the VMS-LDF2, records raw photocurrents. The blocks, with 1,024 samples at f_s =40 kHz, were windowed by a Hamming window, the square of the FFT taken, frequency-weighted and accumulated (as in Section 2.3.2.1) to give a blood flow value.

Figure 3.24 shows the whole of the blood flow as recorded by the LDLS over 18.5 minutes (the LDLS recording started after 2 minutes of cold immersion and 1 minute extra for drying and positioning) across 64 channels. The location of the larger veins are visible by the spatial regions of higher flow. It can be seen here, and to a greater extent in

Figure 3.26, that some spatial locations show increasing flow with others giving flatter responses.



Figure 3.24: The whole of the LDLS recording for subject 1 (the position of the laser line is static and shown in Figure 3.1). The vein locations appear as regions of higher flow, which also increase with time.

Figure 3.25 depicts the raw flow values. Averaging has again been applied, which smooths out the heart rate and noise, to view longer term trends. Figures 3.26, 3.27 and 3.28 show the flow processed by a 100 point (10 s) moving average filter.

So some channels (i.e. skin locations) show an increase in flow, others are more stationary, and there are degrees in between. This situation, interesting in its own right (and also observed by Rendell et al. [100]), is acceptable for the purposes here. Photocurrents at a variety of flow values are now available for model selection.



Figure 3.25: Channel 52 from the Moor LDLS on subject 1 before averaging. The flow gradually increases due to reflex vasodilation.



Figure 3.26: Two channels from LDLS recording on subject 1. The first shows an increase of flow whereas the second shows a flatter flow response. This suggests tissue with varying underlying blood vessel structure exhibits variations in response to temperature provocation.



Figure 3.27: A similar effect to subject 1's increase in flow over time can be observed in subject 2, though not as marked. Some regions of the image (e.g. channel 28) show increasing flow.



Figure 3.28: Subject 3, whose VMS-LDF2 flow recording was relatively stationary, shows that the recording location has an effect on flow response.

3.7.2 Averaging

The data received from the LDLS is in blocks of 1,024 samples. This gives a power spectrum of 512 points up to the Nyquist frequency. In order to use 32 blocks for averaging, 16,384 samples are required. However, the block rate is 10 Hz, so 3.2 seconds of data would need to be collected. The problem with this is that blood flow varies at roughly the heart rate (1-2 Hz) and thus adjacent photocurrents are not necessarily similarly stationary. However, where the flow *and* concentration are similar in the same channel, the spectral shape should also be similar. To have enough photocurrent blocks to average, a method where blocks of like flow and concentration are concatenated was employed. Firstly, the flow and concentration for each block were calculated by the 1,024 point FFT method.

The flow and concentration values were normalised so that the range for each was between 0 and 1. The flow and concentration ranges were then each divided into 75 bands, as shown in Figure 3.29. This number was chosen so that it was large enough to ensure comparable statistics for the grouped photocurrents, but small enough to make available a sufficient number of photocurrents per band. 75 flow bands and 75 concentration bands give a total possible set of $75 \times 75 = 5,625$ bins. Each photocurrent was then assigned to a bin according to its flow and concentration value. Channels at the edges of the image were discarded because the laser line was wider than the hand so no tissue was imaged.

Each bin was then examined to see how many blocks it contained. If it contained at least 32 blocks, the spectra were averaged to create 1 spectrum. If more sets of 32 were available in the band then these sets were also averaged. If fewer than 32 spectra were in the bin (many bins contained no spectra) then the bin was ignored. Figure 3.30 shows that upon averaging stochastic spectra, they become smoother. Thus for each subject, a file of averaged spectra was created. Subject 1 produced 12,305 averaged spectra, subject 2, 14,207, and subject 3, 7,295.



Figure 3.29: The flow and concentration were normalised and 75 equally spaced bands created for each (this figure shows example signals from subject 1). Therefore across the pulsatile signal seen here in the flow graph, approximately 5 bands straddle each pulse.



(b) Averaged spectrum (32 averages)

Figure 3.30: Examples of the effect of averaging on spectra from subject 1. 32 averages produces a much smoother spectrum.

The flow and concentration were recalculated for the averaged spectra. Figures 3.31, 3.32 and 3.33 show the distribution and histograms of flow and concentration for each subject. Any spectra with flat frequency responses (e.g. due to non-tissue imaging) were discarded. Three spectra sampled from a higher region of flow and concentration, a medium region and a lower region respectively for each of the three subjects are shown in Figures 3.34, 3.35 and 3.36.



Figure 3.31: Subject 1: recalculated flow and concentration after spectral averaging. Flat spectra with low flow values were removed.



Figure 3.32: Subject 2: recalculated flow and concentration after spectral averaging. Flat spectra with low flow values were removed.


Figure 3.33: Subject 3: recalculated flow and concentration after spectral averaging. Flat spectra with low flow values were removed.



Figure 3.34: Subject 1: examples of averaged spectra (after 32 averages) at various flow levels. Greater flow is represented by a broader frequency spectrum.



Figure 3.35: Subject 2: examples of averaged spectra (after 32 averages) at various flow levels. Greater flow is represented by a broader frequency spectrum.



Figure 3.36: Subject 3: examples of averaged spectra (after 32 averages) at various flow levels. Greater flow is represented by a broader frequency spectrum.

3.7.3 Fitting Process

The 3 sets of spectra from each subject were individually processed. For each set, the 9 models were run in turn. Parameter ranges were chosen wide enough so that the posterior functions decay to low probability at the extremes of the ranges. 32 separate parameters were used for θ_1 , 128 for θ_2 , with 10 for θ_3 . This means that a total of 40,960 function evaluations are required for each 3 parameter model and each spectrum, with 4,096 for each 2 parameter model.

The bandwidth of the fit was adjusted so that, for each spectrum, only the region well above the noise floor was included in the inference. The noise floor was calculated by examining the final 100 points in each spectrum (where no signal is present) and taking the mean. When the spectrum drops below twice the noise floor for 3 frequency bins, the cut-off frequency is set. This ensures an inference based on the Doppler shifted regions, where the majority of the power lies, and prevents results skewed by fits to noise-only bins.

3.8 Model Fit Results

For each spectrum, the posterior probability was evaluated. The results for each spectrum were normalised such that the sum of the posteriors was unity. Figures 3.37, 3.40 and 3.43 show the posterior probability for the four highest probability models as a function of the spectra index, where the spectra are sorted according to increasing flow value. A 50 point moving average filter smooths the posterior probability to highlight the trends.

The model with the largest posterior probability at each spectrum was selected and added to a cumulative total for each subject. Figures 3.38, 3.41 and 3.44 show the total number of maximum probabilities for each model.

The results from Figures 3.38, 3.41 and 3.44 illustrate not only the most likely model, but also how good the less likely models are.

Figures 3.39, 3.42 and 3.45 show samples of fitted spectra from likely models for each subject. Model 1 (the Exponential distribution) is shown by way of comparison, because it is usually assumed to describe the LDBF power spectrum. Models 3, 7 and 8 are the three most probable other models over the range of spectra and subjects so are also included.

As model 7 (the three parameter generalised Pareto) was the most likely model on the first fitting exercise, model 8 was created by taking the mean of the third parameter value. This two parameter model with fixed shape has now proven to be more probable than the three parameter model. It should be recalled that the Bayesian methodology penalises models with higher numbers of parameters which is why model 7 is less probable than the fixed shape version. This shows that fixing the third parameter at its mean value (in this distribution) does not adversely affect the fits and is a good approximation.

At lower flow values (less than around 0.01 in Figures 3.37 and



Figure 3.37: Subject 1 model fit results. The spectra are sorted by increasing flow (blue line). Only the four models with highest probability are displayed. A 50 point moving average filter has been applied to smooth the results. The posterior probability shows the relative likeliness for each model, illustrating how much better is one model than another. For example, at greater flow values, model 8 (turquoise) is approximately 5 times more likely than models 3 (green) and 9 (purple).

3.43), model 3, the Cauchy distribution, is marginally more probable, but the posteriors for the Pareto distributions are not much lower in value. As the flow gets larger, models 7 and 8 (the generalised Pareto distributions) becomes much more likely on all three subjects and model 3 moves downwards. Flow can be increased by a total power increase (i.e. greater numbers of red blood cells), or by higher frequencies (i.e. greater RBC velocities). There is, however, a positive correlation between flow and bandwidth, so that as flow increases, greater power is present in higher frequencies. This is significant as one of the aims of this thesis is to characterise the effect of frequencies which fall above the 'low bandwidth Nyquist'. So the model to choose is firstly, of course, the most likely, but also the one which is most likely at higher flow rates.



Figure 3.38: Subject 1: maximum probability results. Each bar represents the number of times that model had the greatest probability on an averaged spectrum. Model 8 has the largest number of maximum probabilities.

Model 8 is both the most likely at higher flow rates and also not unlikely at lower flow. This suggests it is the best model to use.



(b) Bar plot of posterior probabilities against model.

Figure 3.39: An example of a model fit comparing models 1, 3, 7 and 8 on one spectrum for subject 1. The posterior for all the models is shown in 3.39b. This shows, for example, that model 8 is 8.5 times more likely than model 9 on this spectrum.



Figure 3.40: Subject 2: model fit results. The spectra are sorted by increasing flow (blue line). Only the four models with highest probability are displayed (model 3 is green, model 7 is red, model 8 is turquoise and model 9 is purple). A 50 point moving average filter has been applied to smooth the results.



Figure 3.41: Subject 2: maximum probability results. Each bar represents the number of times that model had the greatest probability on an averaged spectrum. Again, model 8 has the largest number of maximum probabilities.



(b) Bar plot of posterior probabilities against model.

Figure 3.42: An example of a model fit comparing models 1, 3, 7 and 8 on one spectrum for subject 2. The posterior for all the models is shown in 3.42b. Although model 7 and 8 show a very similar fit, the posterior probability for model 8 is larger due to the Occam factor penalisation.



Figure 3.43: Subject 3: model fit results. The spectra are sorted by increasing flow (blue line). The posterior probability for model 3 is green, model 7 is red, model 8 is turquoise and model 9 is purple. Only the four models with highest probability are displayed. A 50 point moving average filter has been applied to smooth the results.



Figure 3.44: Subject 3: maximum probability results. Each bar represents the number of times that model had the greatest probability on an averaged spectrum. Again, model 8 has the largest number of maximum probabilities.



(b) Bar plot of posterior probabilities against model.

Figure 3.45: An example of a model fit comparing models 1, 3, 7 and 8 on one spectrum for subject 3. The posterior for all the models is shown in 3.45b.

3.9 Creation of Photocurrent

The best fitting model (model 8) is used to create a LDBF photocurrent, with 'flow' controllable by its parameters:

$$P(f;\theta_1,\theta_2) = \theta_1^2 \left(1 + \frac{f}{\theta_2}\right)^{-\frac{1}{1.74} - 1}$$
(3.30)

Combining P(f) (sampled at appropriate frequency resolution $\frac{f_s}{N}$) with a uniformly random phase distribution, mirroring the conjugate around the Nyquist frequency and performing an IFFT yields a photocurrent of length N. This photocurrent will form the basis of the following chapter, as algorithms for calculating blood flow are compared in full and low bandwidth conditions.



Figure 3.46: An example of the model fit on one averaged spectrum (Subject 1 spectra index 12,000 from Figure 3.37). The graph is limited from 0 to 10 kHz, illustrating the improvement the Pareto distribution has over the Exponential distribution.

3.10 Conclusion

The photocurrent of an LDBF process has been modelled by a probability distribution function. In vivo recordings under varying flow conditions were used as the data input. Bayesian Inference formed the machinery to determine the best fitting model. The spectra were binned according to their flow and concentration values, to provide groups of photocurrents with similar underlying spectral properties, and then the spectra in each bin were averaged to become suitable for the analysis. The generalised Pareto distribution, with fixed shape factor of 1.74, was inferred to be the best fitting model, particularly at higher flow values.

It has been noted that an Exponential distribution approximates the power spectrum of the LDBF photocurrent [64]. However, it has been shown that the Pareto distribution may better describe the power spectrum. This is significant because the shape of two distributions differ. The Pareto has a sharper decrease at lower frequencies and longer tail whereas the exponential is more gradual in its slope. This is illustrated in Figure 3.46. The shape of the spectrum is one of the key factors in determining the performance in low bandwidth systems, as this affects the way the high frequencies fold back onto the lower frequencies due to aliasing. Thus, a distribution which more accurately describes the spectrum will more accurately predict the low bandwidth performance.

Therefore, from now on, the Pareto distribution (model 8) will be used in this thesis as the basis for bandwidth and algorithm comparisons.

Chapter 4

A Comparison of Processing Algorithms for LDBF

4.1 Introduction

The aim of this chapter is to assess the accuracy of the LDBF algorithms described in Chapter 2 under full and reduced bandwidth conditions. In addition to this, a Laser Speckle Contrast Analysis (LASCA) algorithm will also be evaluated on the simulated data.

The recent advances in video-rate full-frame (i.e. 2-dimensional pixel array) laser Doppler imaging systems [10, 65, 108, 111, 127] call for algorithms that take less time to process and require smaller onchip silicon area (in terms of both memory and calculation space). The 'gold-standard' approach uses 1,024 point fast Fourier transforms (FFTs) [31, 82]. Performing an FFT on 4,096 pixels 25 times a second (true video rate imaging requires 25 FPS [31]) would require 102,400 FFTs a second. After the FFTs are calculated, the results must be squared, weighted and accumulated. Since the volume of calculations required per second would be inefficient and expensive to implement, cheaper and less resource intensive, but still accurate, algorithms are needed. A variety of algorithms have been proposed that offer solutions to the resource problem; however, whilst increasing the processing speed, they suffer from reduced accuracy. It is timely, therefore, to analyse the effectiveness of these algorithms.

Additionally, the algorithms may be implemented in low bandwidth systems, and the accuracy in these conditions will also be evaluated. Thus the signal processing implementation and performance of a low bandwidth system may be characterised. In this chapter, a model simulating perfect anti-aliasing will be used. In the next chapter, where a full CMOS front-end is characterised, the effect of an integrating and aliasing detector will be modelled.

Previous work in algorithm comparison has concentrated on evaluating the difference between moments of the spectrum, M_0 , M_1 , M_2 and $\frac{M_1}{M_0}$ [70, 130] rather than on variations in implementation of the first moment. In 2011, Wojtkiewicz et al. [130] analysed the effect of reducing the upper cut-off frequency to 5 kHz from 20 kHz and found that sensitivity to higher flows was reduced in all moments. Here, a greater range of cut-offs is evaluated in addition to the various methods of calculating the first moment. The effect on images is also considered, as a reduction in bandwidth or variation in algorithm implementation may or may not affect the quality of the resultant image and its subsequent interpretation. It is worth noting that in LDPI systems, colour representations of the flow results can be scaled or manipulated, either to accentuate temporal variations or to reduce noise.

The first moment of the power spectrum (proportional to mean RBC speed) is defined by:

$$\langle v^2 \rangle^{\frac{1}{2}} \propto M_1 = \int_{30Hz}^{20kHz} fP(f) \,\mathrm{d}f$$
 (4.1)

Where the bandwidth of a LDBF system is generally accepted to be between 30 Hz and 20 kHz [64].

LASCA has emerged as a parallel technique for evaluating microvascular blood flow. Therefore it seems worthwhile to use the simulated photocurrents to analyse and compare the output of the LASCA algorithm. This comparison is not possible on the experimental results because the captured photocurrent length does not provide sufficient data for the temporal LASCA analysis, but the simulated results illustrate LASCA's relationship to flow.

4.2 Theoretical Accuracy

Before assessing the algorithms on real recorded data, they were assessed theoretically using a simulated signal. This was to test how the algorithms performed using known inputs. Whilst assessment in vivo is important, it can only provide comparative results with each algorithm compared to the 'gold-standard' frequency-domain approach. This section considers the outputs of each algorithm when given a photocurrent with known properties; the results can then be used to predict in vivo performance.

4.2.1 Pareto Distribution

Chapter 3 used experimental data and Bayesian Inference model selection to suggest that the best model for the power spectrum of the LDBF photocurrent is the Generalised Pareto distribution under two variable parameters (see Equation 3.30 in Chapter 3):

$$P(f;\theta_1,\theta_2) = \theta_1^2 \left(1 + \frac{f}{\theta_2}\right)^{-\frac{1}{1.74} - 1}$$
(4.2)

Figure 4.1 plots examples of the distribution with 3 separate values for each parameter (i.e. 9 plots).

4.2.2 Methodology

There are three steps in the methodology:



Figure 4.1: Generalised Pareto distribution evaluated over ranges of θ_1 and θ_2 .

- 1. Generating LDBF photocurrents.
- 2. Processing with the algorithms.
- 3. Creating and interpreting the results.

These three steps will now be detailed in turn. MATLAB 7.12.0 (The Mathworks) was used for all simulation and processing.

4.2.2.1 Generating LDBF Data

Figure 4.2 shows the process for generating one block of photocurrent data from given input parameters. Process P1.1 firstly generates the voltage amplitude spectrum according to the square root of Equation 4.2. It evaluates the equation at discrete points, according to the length required. The IFFT algorithm (Process P1.3) generates a signal that is a summation of sinusoids at the amplitudes and phases (the phases are generated in Process P1.2) of the input phasors. These phasors are located at integer cycle lengths, and so the resultant photocurrent will be constituted of sinusoids with whole numbers of cycles. However, it is necessary to include frequencies in the photocurrent which are not

at integer bins (i.e. full cycles) of the FFT algorithm, as this would be the case in a real signal. This was performed by executing an IFFT with greater number of points than the photocurrent length required. The resultant signal from the IFFT was then truncated to the specified photocurrent length (Process P1.4). Figure 4.3 highlights the frequency resolution used for the IFFT (P1.3). For this simulation, the IFFT length was 16,384 points from 0 to 40 kHz, and the photocurrent length was 1024. A sinc function shows the frequency response of a rectangular window function for an 1,024-point FFT, along with the extra frequencies generated in the IFFT of 16,384 points.



Figure 4.2: The process for generating the photocurrent, given parameters θ_1 and θ_2 .



Figure 4.3: A sinc function, the frequency response of a rectangular window, illustrates that if only frequencies at integer bins of the FFT are included in the IFFT, then the power influence of neighbouring bins is zero. However, in this simulation, frequencies at 1/16 of the frequency resolution of the FFT are included as this more realistically represents the wide frequency content of the signal.

This is a situation similar to the white noise generation in the previous chapter, where a well sampled frequency spectrum is transformed to the time domain, and then truncated, so containing non-integer cycles. There is then a difference between the length of the amplitude spectrum generated for the inverse fast Fourier transform (IFFT) and the length of the photocurrent used in the analysis. Signal S1.3 is therefore of length 16,384, and process P1.4 reduces this to 1,024 samples.

Because the LDBF spectrum is contained in frequencies up to 20 kHz, and the Pareto distribution has infinite bandwidth, the frequencydomain function should drop to zero power above 20 kHz. To facilitate this, the spectrum is multiplied by the band-limiting function, given in Equation 4.3 and shown graphically in Figure 4.4. This function forms part of Process P1.1. A roll-off function was used rather than a sharp cut-off to simulate a real filter response. Figure 4.5 shows an example of the updated spectrum.



Figure 4.4: Process P1.1 Band-limiting function, generated according to Equation 4.3. The roll-off starts at 19 kHz.

The first moment of this updated spectrum is calculated in process P1.8 by weighting and accumulating the power spectrum samples and the output represents the real flow value. For the new amplitude spectrum, A[k] (where k is the discrete f, such that $k = \frac{fN}{f_{\text{max}}}$, N = 8192 points and $f_{\text{max}}=20$ kHz), the first moment, M_1 , is given by Equation 4.4:

$$M_1 = \frac{1}{N} \sum_{k=12}^{N=8192} k.A[k]$$
(4.4)

Because A[k], the discrete Pareto distribution, is a function of the two parameters θ_1 and θ_2 (see Equation 4.2), the real flow value is also a function of these two parameters:

$$M_1 = f(\theta_1, \theta_2) \tag{4.5}$$

In the simulations, these two parameters θ_1 and θ_2 are the controlling variables rather than the flow value.



Figure 4.5: Process P1.1 Final 'exponentially' decaying amplitude spectrum of 8192 points.

Process P1.2 generates the phase spectrum, as seen in Figure 4.6. The phase spectrum is effectively random across all frequencies, and the probability distribution uniform between π and $-\pi$.

Process P1.3 combines the amplitude and phase spectra together into complex data. The required data length is 16,384 points, so the amplitude and phase spectrum will have been generated over 8,192 points in processes P1.1 and P1.2. The first frequency sample is 0 Hz (DC) and is set to 0. The points from 8,194 to 16,384 are the mirrored (about the Nyquist point, 8193) complex conjugate of points 2 to 8,193. This will ensure the IFFT generates real data. The process then feeds the 16,384 points of real time-domain data to process P1.4, which then truncates the photocurrent to 1,024 points.

Noise is generated in process P1.5. The signal to noise ratio (SNR) specifies the ratio between the mean power in the signal, $\bar{P}(s)$, and the mean power in the noise, $\bar{P}(n)$, as in Equation 4.6. In reality, the value is variable. In the system used by Serov et al. [106], for small integration



Figure 4.6: Process P1.2 generates the phase spectrum using the random number generator 'rand()' function in MATLAB to give a uniform distribution between π and $-\pi$.



Figure 4.7: Process P1.5 generates white noise using the random number generator 'rand()' function in MATLAB, at a signal to noise ratio of 20, as indicated by Serov et al. [106].

times the SNR is 10 dB, moving up to 25 dB for larger integration times. He et al. [52] state a range of 0.2 dB to 20.25 dB for their CMOS pixel array (dependent on AC/DC ratio). In 2011, Wojtkiewicz et al. [130] calculated that at 40 dB, the error on M_1 as a result of the noise is small, at 10%. Here an SNR of 26 dB is used, small enough so as to evaluate performance in the presence of noise but not so small that the inter-algorithmic comparison is masked. A typical power spectrum of the noise data is shown in Figure 4.7.

$$SNR = \frac{\bar{P}(s)}{\bar{P}(n)} = 20 = 26 \text{dB}$$
 (4.6)

A 'DC' offset level is incorporated in process P1.7. Although it is called 'DC', in fact it refers to low frequency components at higher power than the Doppler shifted components. The offset simulates a variance in total signal amplitude, for example, if the amount of light reflected off the skin surface changes due to movement or blood volume. A sinusoidal DC drift is created (as in Figure 4.8) of 1 Hz with resolution at the sampling frequency. Therefore there is one DC value per sample. This set of numbers is multiplied by the samples in the block. The final signal (S1.5) is shown in Figure 4.9a over 0.025 s and Figure 4.9b over 2.5 s.

The whole process in Figure 4.2 was run 10,000 times over a range of θ_1 and θ_2 values. This equates to a total signal length for the simulation of:

$$T_{sim} = \frac{\text{win.length}}{f_s} \times \text{num.wins} = \frac{1,024}{40,000 \text{Hz}} \times 10,000 = 256 \text{s}$$
 (4.7)

The parameter values were created using sinusoidal functions, across the ranges observed in Chapter 3. θ_1 oscillates much faster than θ_2 because it only moves over a relatively small range, and so while θ_2 changes slowly over its large range, θ_1 oscillates quickly creating a broad range of flow values. The parameter input values are shown in Figures 4.10 and 4.11.

A bank of 10,000 photocurrents of 1024 samples in length was created and stored, along with the known flow values, ready for evaluation. The



Figure 4.8: Output of process P1.6 over 2,000 of the 10,000 windows. Each window has 1024 samples of DC data. The amplitude was selected so that the DC is approximately 10 times AC when multiplied, a value typical of LD systems [50].



Figure 4.9: Output of process P1.7 showing signal with noise and DC drift over a short time period of 0.025 s and over a longer time period of 2.5 s. The signal looks noisy, as the phase is random, but has a power spectrum related to the 'flow' of RBCs, in this case simulated by a Pareto distribution.

next section will look at the processing of simulated data using the various algorithms.



Figure 4.10: Input parameter θ_1 over 2,000 of the 10,000 windows generated. It varies sinusoidally and at high frequency compared to parameter θ_2 .



Figure 4.11: Input parameter θ_2 over 2,000 of the 10,000 windows generated. This parameter varies slowly over its large range. Together, the simultaneous slow and fast variation of θ_1 and θ_2 respectively create a broad range of distribution shapes over the 10,000 windows.

4.2.2.2 Simulation Process

A block diagram of the program flow is shown in Figure 4.12. At its simplest, the process takes a window of photocurrent with known properties (P2.1), processes it using the algorithm, stores the results and then repeats.



Figure 4.12: The simulation process for analysing the algorithms under test.

Process P2.2 removes the DC by calculating the mean of the signal, dividing it through the signal and removing the offset. This simulates the normalisation process in existing systems [50, 84].

The theoretical simulation was separated into two tests. Firstly, a full bandwidth comparison of algorithms, and secondly, a comparison under reduced bandwidth conditions. Process P2.3 sets the bandwidth of the photocurrent by performing a FFT, setting all components of the FFT above the cut-off to zero, and reversing the transform using the IFFT. This simulates a 'perfect' anti-aliasing filter. However, as discussed in Chapter 1, CMOS imagers do not have anti-aliasing circuitry built in. Therefore Chapter 5 will model the CMOS imager frequency response. Here, however, the response using 'perfect' low frequency signals will be considered.

Table 2.1 in Chapter 2 was updated to include various FFT options, such as number of averages, and that table can be referred to for authorial information and numerical implementation. Table 4.1 lists the algorithms tested.

The flow results were calculated for each window using each algorithm. Additionally, for the reduced bandwidth test, 5 upper cut-offs were assessed: 20 kHz, 10 kHz, 5 kHz, 2.5 kHz and 1.25 kHz.

Process P2.6 normalises the results because the algorithms provide

Algorithm	Abbreviation	Domain
FFT N=1,024	FFT	F
FFT (FA) $N=8$	FA8	F
FFT (FA) $N=8$ (approx. coefficients)	FA8APP	F
FFT (FA) $N=32$	FA32	F
FIR square and average	FIR S	Т
FIR absolute and average	FIR A	Т
IIR	IIR	Т
Draijer	DRAIJER	Т
Analytic/Hilbert frequency domain (new)	HILB F	F/T
Analytic/Hilbert time domain (new)	HILB T	Т
Binzoni's 2nd moment	M2	Т
Laser Speckle Contrast Analysis	LASCA	Т

Table 4.1: List of algorithms for processing in P2.4. The LASCA algorithm is tested separately in Section 4.2.4. The FFT of lengths 1,024, 32 and 8 are used, giving frequency averages (FA) of 1, 32 and 128 respectively. The Hilbert transform is performed in both frequency and time domain, with additional processing (for differentiation) performed in the time domain. The abbreviations are used on results.

an output that is proportional to real flow. The proportionality constant for each algorithm was calculated by taking the mean of the 10 largest flow values computed by the algorithm, dividing all the data by this value to normalise, then multiplying by the largest true flow value. The output of this process for the full bandwidth results is shown in Figure 4.13, and summarised in Figures 4.14 and 4.15.

There are two methods by which the reduced bandwidth results may be normalised. The first, displayed in Figure 4.17, is where the data is normalised against the full bandwidth constant. This shows the reduction in ability to determine higher flows.

The second is where the data is normalised against its own maximum values. An example of this type of normalisation is shown in Figure 4.18, and this second type is used in the results in Figures 4.19 and 4.20 as this provides a fairer comparison of the linearity and error.

Where line fits are shown, the results are fitted to a 2nd order polynomial, and linearity is assessed using the 2nd order coefficient.

The full bandwidth results are also plotted against the complexity calculated in Chapter 2.

4.2.3 First Moment Results

4.2.3.1 Full Bandwidth

The results in Figure 4.14 indicate that, as is expected and accepted, the 1,024 point FFT is an accurate method of estimating M_1 , generally taken to indicate blood flow. The frequency averaging method at N=8 (FA8) is marginally worse than its N=1,024 counterpart. The approximation of coefficients (FA8APP) seems to be a reasonable step, barely increasing the error. At N=32 (FA32), the error sits between these levels, suggesting that there is an trade-off between variance, bias and numerical integration method: at N=8, the variance is small because 128 averages have occurred but the bias is larger because the frequency bin is wider (see Chapter 3). The variance, bias and numerical integration may be more balanced at N=32 because the error is smaller. At N=1,024, the numerical integration and bias are more accurate because the frequency resolution is small, despite the variance being larger, and this algorithm is the most accurate. Therefore it seems the numerical integration and bias have a greater effect on the error than the variance.

In the time domain, because the IIR filter has a frequency response close to the ideal $f^{\frac{1}{2}}$ (see Figure 2.4), it is also accurate. The FIR filters (FIR S and FIR A) do not show nearly as good performance. The absolute and average method is only slightly less accurate than square and average, proving that the approximation of coefficients is a good one. The performance of the time domain filters is therefore a consequence of the shape of the filter, and how far it resembles the ideal. The FIR filters tend to overestimate the flow (Figures 4.13e and 4.13f); this is probably because the frequency weighted filter (response in Figure 2.3 in Chapter 2) has a higher gain at larger frequencies (>10 kHz) than the ideal $f^{\frac{1}{2}}$ weighted filter.



Figure 4.13: Full bandwidth raw data results showing 10,000 points for each algorithm. The real flow (x-axis) is a function of the two parameters θ_1 and θ_2 and calculated according to Equation 4.4. The algorithm flow (y-axis) was calculated according to the algorithms specified in Table 4.1 and detailed in Chapter 2.



Figure 4.14: The root mean square (RMS) error of the calculated flow from the real flow values under full bandwidth conditions over 10,000 windows. The full 1,024 point FFT is the most accurate algorithm, but it still has an error due to spectral leakage effects.

The analytic signal based algorithms (HILB F and HILB T) seem to perform reasonably, as in theory the equations are exact. However, the calculation of the differentials by finite difference adds some error, which is reflected in their results. The calculation of the Hilbert transform by time domain filter performs similarly to the same calculation performed in the frequency domain (as would be expected).

As expected, the approximate methods, Draijer and the second moment (M2), perform at a level between FIR filter and the others. However as Figure 4.16 reveals, these algorithms are extremely efficient in terms of resource usage and therefore are still worth considering in system implementations. The analysis on recorded experimental data will reveal how much this error impacts on a resultant flow map. All the algorithms responses are generally linear, as Figure 4.15 testifies.

Figure 4.16 illustrates the RMS error (from Figure 4.14) as a function of the complexity. As expected, the two quantities exhibit



Figure 4.15: A 3rd order line was fitted to the data for each algorithm, to visualise the linearity of the algorithms. As can be seen, they are approximately linear over the flow range with the FIR filter algorithm showing the greatest non-linearity, followed by the second moment and Draijer algorithms. The deviation, or RMS error, from the fitted line (4.15b) is in accordance with Figure 4.14.

inverse proportionality. Draijer and the second moment are extremely resource efficient but the IIR, Hilbert time domain (HILB T) and 8 point frequency average (FA8) approximately halve the RMS error and are still resource efficient. These three therefore may be the optimum algorithms in terms of balance between error and resource.



Figure 4.16: The complexity of the algorithms measured by the number of MACs required to produce one output blood flow sample from 1,024 input samples plotted against the RMS error from Figure 4.14. The Hilbert time domain (HILB T), IIR and 8 point frequency average (FA8) algorithms are both low in complexity and RMS error.

4.2.3.2 Reduced Bandwidth

Generally a reduction in the bandwidth reduces the sensitivity of larger flow values because the higher frequency content of the larger flow values is attenuated. When the data is not re-normalised, i.e. in Figure 4.17, the RMS error (i.e. the variance of points), across the bandwidths, is fairly consistent in relation to the shape of the line.

The data can be re-normalised by calculating the ratio of the mean of the largest 10 points in the full bandwidth results to the mean of the largest 10 points in each of the reduced bandwidth results, and multiplying each set of results by the ratio for that bandwidth. An example of this is shown in Figure 4.18. By performing this renormalisation, the reduction in linearity and increase in error is evident. Figure 4.19 shows the absolute value of the 2nd order coefficient versus bandwidth. All the algorithms follow a similar trend as the bandwidth is reduced, except the FIR because the RMS error is so large. The IIR and Hilbert time domain (HILB T) show consistently low RMS error and 2nd order coefficient, however the FFT algorithm, whilst it is the most accurate at full bandwidth, increases its error more sharply at lower bandwidth than the first two algorithms.

Evaluating the performance of all the algorithms under reduced bandwidth conditions, it seems that the method of calculation of M_1 does not have as big an impact on error as the reduction in bandwidth itself, as all the algorithms responses follow a similar shape. However, of the group, the analytic method by Hilbert transform and IIR both have good linearity i.e. low 2nd order coefficients (from Figure 4.19), and, from Figure 4.20, low RMS error. The under-representation of higher flow values is the major factor in determining appropriate bandwidth.






Figure 4.17: The algorithms calculating M_1 at 5 bandwidths (20, 10, 5, 2.5 and 1.25 kHz) are plotted against the real flow on the left hand side plots. A 3rd order fit to the reduced bandwidth results is plotted against the real flow on the right hand side plots.



Figure 4.18: An example of the re-normalisation process, on the 1,024 point FFT, for evaluating linearity and RMS error. The data were normalised by the ratio of the mean of the largest 10 points in the full bandwidth results to the mean of the largest 10 points in the reduced bandwidth results.



Figure 4.19: Low bandwidth linearity evaluated over 5 bandwidths. The 2nd order coefficient is the absolute, so as to measure deviation from linearity. The fit of the line is calculated after re-normalising the results. The IIR and Hilbert time domain (HILB T) show the greatest linearity at lower bandwidths.



Figure 4.20: Low bandwidth RMS Error evaluated over 5 bandwidths. The error is calculated from the fitted line. Again, the IIR and Hilbert time domain (HILB T) show the smallest RMS error at lower bandwidths. Surprisingly, the FFT algorithm is not consistently the most accurate algorithm across the bandwidth range.

4.2.4 LASCA Results

4.2.4.1 Overview

As introduced in Chapter 1, LASCA is an alternative method for estimating blood flow. Whilst the processing methodology does not aim to calculate the first moment of the photocurrent's power spectrum, it does aim to measure some quantity representing the state of the underlying velocity distribution and is becoming an accepted clinical imaging technique [16]. In 1996, Briers [15] showed that, despite the processing methodology and measured quantity being different, the physical phenomenon which underpins LDBF theory, i.e. the Doppler Effect, and the resultant interfering speckle pattern, which underpins LASCA, are the same. Because of this fact, the photocurrents generated in this section can be used to quantify the accuracy of the LASCA method in comparison to LDBF.

The LASCA algorithm calculates the ratio of the standard deviation of a time integrated speckle pattern to its mean (Equation 4.8), on, for example, a 7x7 spatial pixel region.

$$K = \frac{\sigma}{\bar{I}} \tag{4.8}$$

where K is a quantity between 0 and 1. However, it has been shown that the local spatial statistics of this speckle pattern are equivalent to the temporal statistics [44], and experimental studies have shown that a large spatial resolution can be gained by implementing a temporally based solution [1, 18]. This overcomes one of the drawbacks of spatial LASCA, i.e. reduced spatial resolution. Thus K can instead by obtained by computing the ratio of the standard deviation to the mean of 49 (for example, in comparison to the 7x7 pixel window) temporal samples, where each sample is integrated over the period τ . This is convenient for this study, where the existing photocurrents, with known properties, can be temporally averaged in a straightforward manner to provide input data for the LASCA algorithm.

Using a simple measure of K means that when the blood velocity is low, K is large as contrast is maintained for a set integration time. At higher blood velocities, the averaged photocurrent is blurred over consecutive input samples and thus K is low. This inverted situation can be rectified by taking a measure of 1/K to quantify blood flow as in [3]. It has been noted that LASCA measurements are insensitive to concentration but are sensitive to variation in velocity distribution [33]. In classical LDBF, flow is equivalent to concentration multiplied by mean velocity. To ensure that the LASCA results compute a 'flow' value, they were multiplied by the true concentration (M_0) for that window, computed in the classical way. This is a necessary step for this study as the simulated photocurrents contain modifications in both concentration and mean frequency by the parameter variation.

$$F_{LASCA} = M_0 \frac{\bar{I}}{\sigma}$$
(4.9)

where each \bar{I} and σ are computed over 49 time integrated samples.

4.2.4.2 Photocurrent Data

The previous photocurrent data was sampled at 40 kHz and generated in blocks of 1024 samples. An optimal integration time for LASCA is $\tau=5$ ms [132], therefore 200 samples of existing photocurrent data were averaged to provide one sample for the LASCA evaluation (Equation 4.10). However, in order to provide 49 samples for calculation requires $49 \times 200 = 9,800$ photocurrent samples (t=245 ms). Therefore 10,000 photocurrents were regenerated which satisfied this criterion. The first moment, M_1 , of these generated photocurrents was recorded for comparison.

LASCA[i] =
$$\frac{1}{200} \sum_{n=1}^{200} s[(i-1).200 + n]$$
 (4.10)

where i is the index of the temporal LASCA samples and s[] is the well-sampled LDBF photocurrent. The mean and standard deviation of 49 of these samples was used for the F_{LASCA} calculation.

4.2.4.3 Results

LASCA does not aim to calculate the first moment of the photocurrent's power spectrum (M_1) . However, it has been shown that M_1 scales linearly with blood flow [12] and therefore it is prudent to compare the LASCA output with this quantity.

The LASCA results from the set of photocurrents were normalised, as before, so the mean of the largest 10 LASCA flow values matched the mean of the largest 10 true flow values. The output normalised LASCA flow values were plotted against the true flow value for that photocurrent, and the 10,000 pairs are shown in Figure 4.21.



Figure 4.21: Temporal LASCA results over 10,000 windows. The RMS error was 128 (compare with Figure 4.14). The widely reported non-linearity of the LASCA output is evident.

The LASCA flow results have evident non-linearity and a high RMS error compared to the first moment calculations (see Figure 4.14: RMS error was 15 for the 1,024 point FFT versus 128 here), as is to be expected. This error can be explained by considering the time integration, standard deviation and mean calculations from a signal processing perspective. The time integration acts as a simple low pass filter, and at $\tau=5$ ms, the cut-off frequency is 200 Hz. The variance of a signal (σ^2) measures the total power at all frequencies, and the square of the mean (\bar{I}^2) measures the power at DC. Therefore K^2 is a measure of the ratio of the total AC to DC power, where the AC frequencies are attenuated above 200 Hz. K is thus related to this ratio, being the square root of the calculated quantity. Figure 4.22 illustrates how filtering the Pareto distribution under two different θ_2 parameters, corresponding to high and low flow, results in low or high K values.

The quantity being measured, K, is therefore related to the *shape* of the velocity (and frequency) distribution, and how that shape is modified as blood flow increases. The integration time sets the sensitivity of K to a particular frequency band, in this case 200 Hz. It is clear that, whilst 1/K may scale with flow, it is not and cannot be a direct replacement for a first moment calculation.



Figure 4.22: The effect of temporal averaging at $\tau=5$ ms on the LDBF spectrum. Two Pareto distributions with $\theta_2=200$ (a) and $\theta_2=1500$ (b) were generated, and plotted with the $\tau=5$ ms filter response and filter output. The first distribution shows a greater proportion of power at lower frequencies (i.e. lower blood velocities), so that when filtered, the K^2 quantity (σ^2 is area under the curve) is large. However, in (b), where power is distributed to higher frequencies, the total power in the filtered spectrum is smaller and thus K^2 , and so K, is small.

4.3 Experimental Recordings

4.3.1 Data

The Moor Instruments LDLS was used as the data source. It provides photocurrents as described in Section 3.2.3, except in this recording the laser line is moved, providing a scan of an area of tissue. The backs of the hands of three subjects (age 24 to 26, two males and one female) were recorded using the device, providing 3 images of 64x64 pixels with each pixel containing 1,024 time domain samples. Figure 4.23 shows the images from the three patients calculated by the 1,024 point FFT at full bandwidth. This is because, from the previous section, the 1,024 point FFT is the most accurate algorithm at full bandwidth.



Figure 4.23: Images recorded for comparison. The photocurrents have been converted to flow values by the 1024 point FFT algorithm. Data on row 63 was corrupted so the images only use rows 1-62.

4.3.2 Methodology

The processing of the data is similar to the simulated methodology, except the DC normalisation has already been performed by the LDLS. Figure 4.24 summarises the processing implemented. Process P3.1 retrieves the collected photocurrent and feeds it to Process P3.2. The reduction in bandwidth (process P3.2) is achieved by performing a FFT, setting all components above the cut-off to zero, and recalculating the time domain signal by IFFT. Process P3.3 then processes the data according to the algorithms listed in Table 4.1. The results are stored via Process P3.4.



Figure 4.24: Process for evaluating the images recorded by the Moor Instruments LDLS.

The normalisation value is calculated (Process P3.5) at 80% of the maximum pixel value in the image because this means that the display colour bars show the variation across the majority of the image data. The histogram in Figure 4.25 illustrates this. The RMS error is calculated from the root mean square of the differences between the full bandwidth 1024 point FFT image and the test image.

4.3.3 Results

The resultant images are displayed in Appendix B. The colour map range of the images have been adjusted, and shows that the error increases as the bandwidth is reduced. Here, the bar graphs of full bandwidth RMS errors are presented followed by line plots of bandwidth vs RMS error.

4.3.3.1 Full Bandwidth

The results from images 1 and 3 generally agree with the simulations, as the relative levels of RMS error are consistent with Figure 4.14. Image 2 shows larger errors in the frequency averaging and analytic classes than



Figure 4.25: Histogram of flow data in image 1. The normalisation value at 80% of the maximum pixel value means that the colour bars, on a scale of 0 to 1, capture the variation in the bulk of the data. The large peak around 0.1 shows 'dark' areas.

the other images. Again, the IIR and FIR have consistently small and large errors respectively.

In the simulations, the approximate algorithms had larger errors than in the experimental data. This may be because a larger range of flows were generated for the simulations than are present in the experimental images. Analysing the images in Appendix B, the Draijer and second moment (M2) algorithms show a slight loss of clarity due to increased noise, but the veins are still distinguishable. The lower flow regions between the veins reveal small over-estimations, which may be significant if very good sensitivity to flow is required e.g. in burn analysis. But if feature detection, such as the locating of perforator vessels in flap replacement surgery (where a general region of higher flow indicates the location) the fast approximations may be sufficient.



Figure 4.26: Bar plot of image 1 RMS error of the flow map compared to the 1,024-point FFT full bandwidth. The frequency average and IIR methods give the lowest error, but the fast (computationally) Hilbert and Draijer methods still only give errors of around 0.05 (i.e. 5% of the image colour range).

4.3.3.2 Reduced Bandwidth

Figures 4.29 to 4.31 show the RMS error with reduced bandwidth. The general level of the errors presented in these figures corroborates with the simulation, where at the lowest bandwidth (2.5 kHz) the RMS error is approximately 15% of the full range, except in image 2 where the errors increase to 25%.

Most features in the images in Appendix B are distinguishable down to 10 kHz bandwidth, although there is some loss of clarity at higher flow. As the bandwidth moves down to 5 kHz there is more significant distortion. At 2.5 kHz, some information about where the general regions of higher flow might be is still observable. At 1.25 kHz, information about the mere presence of blood flow is available, i.e. where there are moving scatterers, but more detail than that is indistinguishable.



Figure 4.27: Bar plot of image 2 RMS error of the flow map compared to the 1,024-point FFT full bandwidth. The results are similar to image 1, although the frequency average methods give slightly greater error here.



Figure 4.28: Bar plot of image 3 RMS error of the flow map compared to the 1,024-point FFT full bandwidth. Again, the results are consistent with image 1.



Figure 4.29: Image 1 RMS error of the flow map from the 1,024-point FFT full bandwidth. The images themselves can be seen in Appendix B. The results reveal that the IIR and Hilbert algorithms are the most accurate at lower bandwidths.



Figure 4.30: Image 2 RMS error of the flow map from the 1,024-point FFT full bandwidth, over a range of bandwidths. The images themselves can be seen in Appendix B.



Figure 4.31: Image 3 RMS error of the flow map from the 1,024-point FFT full bandwidth, over a range of bandwidths. The images themselves can be seen in Appendix B. Here, again, the IIR and Hilbert algorithms are the most accurate at lower bandwidths.

4.4 Conclusion

The performance at full and reduced bandwidth of a range of methods that compute the first moment of the LDBF power spectrum has been evaluated. The simulated photocurrents with known properties were used to test the absolute performance, before comparing the algorithms with the 1,024 point FFT on recorded images.

The simulation confirmed that the full FFT is the most accurate method of determining the first moment, as is generally expected and accepted. However, other algorithms using reduced FFT lengths with averaging and time domain methods with frequency responses close to the $f^{\frac{1}{2}}$ performed reliably, and the analysis on images showed that these faster algorithms were indistinguishable from their gold standard counterpart. The approximate class of algorithms (DRAIJER and M2) were less accurate in both the simulation and experimental analysis. However, a clinical trial would be needed in order to determine whether the error from these much faster implementations affects diagnosis. For example, a recent study that claimed that 'LDI was >90% accurate at predicting the need for excision and grafting.' [118] in burn applications. Further work could include a study which determined the effect on this figure of using the faster and/or approximate algorithms. Such a study might also overcome some of the limitations of this work, as it has not been possible to include data from subjects with skin complications or pathologies which affect microvascular blood flow because this was outside the scope and resource of the project.

In general, the reduction in bandwidth reduced the sensitivity to higher flow, the measurement of which is a requirement for accurate diagnosis of burn severity. However, for an application such as flap monitoring where a measurement of the general increase in perfusion (from zero to baseline) to the flap area is required, a lower bandwidth may be sufficient. The results suggest that a camera that samples at 10 to 20 kHz, and uses perhaps the IIR or Draijer algorithm, could be implemented and would perform adequately in this scenario. Thus a relatively cheap solution, using off the shelf components, may be produced. But if the device also needs to accurately determine the location of perforator vessels, which have marginally higher flow from baseline and so need instrumentation which is more sensitive to deeper and higher flow, it may not be adequate.

A reduction in resource requirements (see Figure 4.16) may be realised at little expense in terms of error. For very small resource usage, a set of approximate algorithms may provide adequate results. The reduction in bandwidth may be suitable for some applications, but further trials are necessary to evaluate the trade-off. In the next chapter, the effect of aliasing on the results is modelled, and whether low bandwidth flow values may be corrected for is considered.

Chapter 5

Low Bandwidth Correction

5.1 Introduction

The aim of this chapter is to characterise the error introduced into the LDBF flow value as a result of using an integrating CMOS imager, and then, using Bayesian Inference, develop a method by which the error in the flow value can be reduced.

In order to characterise the CMOS imager, the effect of the integration time on the frequency response, and how this affects the flow error, will be considered. Then, a model for aliasing in the imager will be explained and tested, before being applied to simulated photocurrents.

Because the underlying power spectrum (Pareto distribution), frequency response and aliasing effects are known, Bayesian Inference (introduced in Chapter 3) can be used to determine the most likely parameters of the distribution for an undersampled photocurrent, therefore allowing calculation of a 'real' flow value.

5.2 An Integrating CMOS Imager

5.2.1 Overview

An integrating pixel differs from a non-integrating pixel in that instead of directly and continuously converting the light intensity to current (as a conventional logarithmic photodiode will do [106]), it allows charge to collect on a storage capacitor (C_{int}) for a set period of time (the integration time, τ). The voltage on the capacitor is then read-out at the end of the period. Figure 5.1 shows a typical integrating pixel circuit.



Figure 5.1: An example of an integrating pixel. The capacitor C_{int} is reset to the voltage V_{reset} (minus the reset and shutter transistors voltage drop) by the reset signal. The shutter signal then allows the capacitor to discharge at a rate proportional to the photodiode current for the integration time, τ , at the end of which the capacitor voltage is sampled by the read-out circuitry.

The integration provides benefits in terms of signal to noise ratio (SNR) but has disadvantages in terms of sampling rates and filtering effects. The timing characteristics of the integration and sampling are illustrated in Figure 5.2.



Figure 5.2: The integration time τ must be less than the sample time T_s . τ sets the SNR and low-pass filtering characteristics of the system, and T_s controls the bandwidth.

5.2.2 Integrating Pixel

5.2.2.1 Noise

There are two types of noise that will be considered here: thermal and shot noise. Reset noise is ignored, assuming the use of correlated double sampling [28]. Thermal noise is the result of the random excitation of electrons due to thermal energy, and is 'white' i.e. has constant power at all frequencies. Shot noise is the result of the fact that the conversion of light to electrical current in the photodetector is not a continuous process. Rather, it is a discrete process, as photons arrive at the detector in 'packets' and are subsequently converted to electrons, which are also discrete charge carriers. It is also white.

The mean square thermal noise current, $\langle i_{th}^2 \rangle$, is:

$$\langle i_{th}^2 \rangle = \frac{4k_{\rm B}TB}{R} \tag{5.1}$$

where T is the temperature in degrees Kelvin, $k_{\rm B}$ is the Boltzmann constant, B is the noise equivalent bandwidth and R is the resistance of the load. The mean square of the shot noise current, $\langle i_{sh}^2 \rangle$, is:

$$\langle i_{sh}^2 \rangle = 2q \langle I_{pd} \rangle B$$
 (5.2)

where q is the elementary charge, and:

$$\langle I_{pd} \rangle = \langle I_{pc} \rangle + \langle I_{dk} \rangle \tag{5.3}$$

so that the total DC photodiode current, $\langle I_{pd} \rangle$, is the sum of the photo current, $\langle I_{pc} \rangle$, and dark current, $\langle I_{dk} \rangle$. From [109], the mean square signal, where M is the number of speckles in the photodetector area, is:

$$\langle i_s^2 \rangle = \frac{\langle I_{pc} \rangle^2}{M}$$
 (5.4)

Therefore the SNR is:

$$SNR = \frac{\langle i_s^2 \rangle}{\langle i_{th}^2 \rangle + \langle i_{sh}^2 \rangle}$$
(5.5)

The details of the experimental setup in [106] were used to calculate approximate signal to noise ratios. The authors use a 250 mW laser at 671 nm wavelength to illuminate an area with 17 cm diameter. Using a f-number of 1.2 (so numerical aperture (NA) of 0.42), the speckle size $d_{\rm sp}$ can be evaluated by Equation 5.6 [44].

$$d_{\rm sp} = 1.22 \ \lambda \ \frac{1}{2\rm NA} \tag{5.6}$$

With a pixel size of 7 μ m, M can therefore be determined. Then, using a working distance of 20 cm and focal length of 6 mm, the power incident at the pixel, when illuminating skin tissue, is 0.25 nW. With a sensitivity of 0.3 A/W, the DC photocurrent is calculated to be 74 pA. The data sheet for the sensor gives a typical dark current as 0.2 fA. The SNR at various integration times (τ) can then be calculated, assuming:

$$B = \frac{1}{2\tau} \tag{5.7}$$

For example, by using this bandwidth in Equation 5.2 at $\tau=80 \ \mu s$ and with a DC photocurrent of 74 pA, the RMS shot noise current is 0.12 pA. The results of all the SNR calculations from Equation 5.5 are given in Figure 5.3. The values are slightly larger than the experimental measurement in [106], possibly because of quantisation noise, inefficiencies in optical illumination (as fibre optics are used) or variations in the reflection coefficient of skin.



Figure 5.3: The choice of integration time is a critical factor in determining the SNR (black line), as it sets the bandwidth of the system. But the integration time also affects the maximum sampling frequency of the system (red line).

5.2.2.2 Frequency Response

Although the signal to noise calculations favour a larger integration time, the maximisation of the sample rate and frequency response of the detector favour a smaller period. The integrating detector acts as a low-pass filter. One advantage of this is that it attenuates high frequency noise components.

5.2.2.3 Sample Rate

Typical integration times range between 60 and 100 μ s [106], so that the maximum sample rates are in the range 10-16 kHz according to Figure 5.3. However, the sample rate is also a function of the capabilities



Figure 5.4: The integration time sets the -3 dB bandwidth of the system. However the cut-off is not sharp, and there is unwanted attenuation in the pass-band and poor rejection of frequencies above the -3 dB point.

of the sensor's ADCs and data throughput, and is usually determined by the size of the spatial window (resolution) required. For example, Leutenegger et al. [73] can achieve 14.9 kHz sample rate at a resolution of 480x60, or 12.4 kHz at 360x90, with other combinations possible. In order to record the whole image, the sampled regions are moved through the sensor area. For example, to achieve 480x480 pixel resolution, 8 lines of 480x60 are sampled consecutively.

The sample rate also sets the Nyquist frequency. If there is any power present in frequencies above this point, the power will be represented at frequencies up to the Nyquist frequency in the digital system. This is known as aliasing and is the subject of the next section.

5.3 Aliasing

5.3.1 Model

It is well known that in order for a digital system to fully represent an analogue signal s(t), the signal must be sampled at a frequency (f_s) greater than twice the maximum bandwidth (f_{bw}) of the analogue signal:

$$f_s > 2f_{\rm bw} \tag{5.8}$$

The digital samples are separated by $\frac{1}{f_s}$ seconds and the original signal can be perfectly reconstructed from its digital representation. If the sampling frequency, f_s , is not greater than $2f_{bw}$ then aliasing will occur: high frequencies $(f > \frac{f_s}{2})$ are 'undersampled' and manifest themselves as lower frequencies in the digital system. This can be modelled by considering a 10 Hz wave: firstly well sampled at 64 Hz (the Nyquist-Shannon criterion requiring a frequency greater than 20 Hz), as shown in Figure 5.5, and then undersampled at 16 Hz, as shown in Figure 5.6.

In Fourier transform theory, the spectral leakage effect of a known frequency can be predicted, but not reversed once the transform is complete. So also in aliasing theory, the effect of undersampling a known frequency can be predicted but not reversed. For example, it is possible to predict that 6 Hz will be the aliased frequency when a 10 Hz signal is sampled at 16 Hz, but with only the sampled information to



(a) A 10 Hz wave s(t) well sampled at 64 Hz (b) FFT frequency representation

Figure 5.5: The 10 Hz wave is accurately represented in the digital system when sampled at $f_s > 20$ Hz (64 Hz in this case).



(a) A 10 Hz wave s(t) undersampled at 16 (b) FFT frequency representation Hz

Figure 5.6: The 10 Hz wave, when sampled at 16 Hz in this case, is represented as a 6 Hz wave (i.e. aliasing occurs).

hand it is not possible to know the original frequency. It could very well have been 6 Hz originally, or 10 Hz, or 22 Hz or 26 Hz *ad infinitum*.

In general, if a frequency f with amplitude A and phase θ (i.e. $Ae^{i\theta}$) is known, the new frequency can be determined in a system, using the following conditions, where f resides between Nf_s and $(N+1)f_s$, N is an integer and f_s is the new sampling frequency.

$$f_{\text{alias}} = \begin{cases} f - Nf_s & f < Nf_s + \frac{f_s}{2} \\ (N+1)f_s - f & f > Nf_s + \frac{f_s}{2} \\ 0 & f = Nf_s \\ \frac{f_s}{2} & f = Nf_s + \frac{f_s}{2} \end{cases}$$
(5.9)

This is a folding effect, where frequencies greater than the Nyquist $(\frac{f_s}{2})$ are 'folded' back over into the lower spectrum. The frequencies greater than f_s but less than $\frac{3f_s}{2}$ are copied down to the lower spectrum directly, and the pattern continues as $N \to \infty$. The phasor $Ae^{i\theta}$ at f_{alias} is a result of the addition of the phasors of all aliased frequencies that copy onto f_{alias} , except that when the frequencies are folded (condition 2 in Equation 5.9) the conjugate of the phasor is added instead. However according to sampling theory the resultant amplitude and phase of frequencies at $\frac{f_s}{2}$, whilst being accurately predicted by this set of conditions, cannot be accurately digitally represented at the new sample spacings $\frac{1}{f_s}$ because f_s must be greater than $2f_{\text{bw}}$

This set of conditions can be applied to model the effects of aliasing on a known frequency spectrum.

5.3.2 Model Test

To test whether this set of conditions accurately describes a true undersampled photocurrent, a simulation was carried out. A set of 512 frequencies in the range 0 to 20 kHz with uniformly random amplitude and phase were generated in MATLAB (Figure 5.7). The complex conjugate of the phasors of the frequencies up to the Nyquist frequency (20 kHz) was mirrored around this frequency. An IFFT then produced a signal of 1,024 points.

The aliased spectrum may be produced in two ways: firstly, by undersampling the signal (Figure 5.8), performing a FFT and viewing the results directly, and secondly, by using the conditions in Equation 5.9 on the generated phasors.



Figure 5.7: The generated spectra for the model test. These spectra were used to generate a signal of 1024 points.



Figure 5.8: The signal, originally sampled at 40 kHz (green), was undersampled by selecting every 8th sample (red). The total signal length was 0.0256 s.

The new sampling frequency was set as 5 kHz. The results in Figure 5.9 show that the model conditions predict the aliasing effect. This model can now be used to predict aliasing in a LDBF situation.



Figure 5.9: Aliasing is accurately predicted by the model. The exception is the new Nyquist frequency (2.5 kHz), which cannot be represented properly. The original spectrum extends to 20 kHz; only 2.5 kHz is shown here as that is the range of the new spectrum. Each amplitude point of the aliased spectrum is the result of 8 summed phasors. There is little resemblance between original and aliased amplitudes because the phasors are random in amplitude and phase.

5.3.3 Anti-Aliasing Comparison

5.3.3.1 Theoretical Data

The aliasing model can, along with the frequency response of the integrating detector, be used to assess the performance of a low bandwidth camera. The same photocurrent generation method as in Chapter 4 (Figure 4.2) was used, and the simulation process was modified to include the additional integrating pixel effects. The new process is illustrated in Figure 5.10.

An integration time of 60 μ s was used as this is typical of LDPI systems [106]. The filter frequency response corresponding to this



Figure 5.10: The simulation process for evaluating the performance of a low bandwidth camera.

integration time can be seen in Figure 5.4a. This will be referred to as the detector response, and it serves to limit, but not eliminate, the effect of aliasing. Five bandwidths were tested: 20 kHz, 10 kHz, 5 kHz, 2.5 kHz and 1.25 kHz. The results were first normalised against the 20 kHz bandwidth for display in Figure 5.11, thus revealing the sensitivity to higher flow values. 2nd order polynomials were fitted to the results, and these can be viewed in Figure 5.12. Then the results were renormalised, so that the mean of the highest output flow values sat at approximately the same level for every bandwidth (as in Figure 4.18 in Chapter 4). This is so the RMS error and linearity can be directly compared.

The flow results with perfect anti-aliasing and no attenuation caused by the integrating detector were included for comparison.

Figure 5.13 shows a surprising result: as the bandwidth decreases the flow response of the system with aliasing and τ =60 µs detector response has higher linearity than the anti-aliased response. One reason for this may be to do with the shape of the spectrum. If the integration time is reduced to 50 µs, with results shown in Figure 5.14, the point at which the 2nd order coefficient line for the aliased response crosses the line for the anti-aliased response increases from approximately 7.5 kHz (Figure 5.13) to 9 kHz (Figure 5.14). This means that as more frequencies are permitted to alias, the error due to non-linearity decreases. Before these results are discussed further, the same test will be performed on experimental data.



Figure 5.11: Flow results for simulation of CMOS integrating detector response at 5 bandwidths. The aliased response has greater sensitivity to higher flow values than the anti-aliased response.



Figure 5.12: The results in Figure 5.11 were fitted with 2nd order polynomials. The linearity may be assessed by renormalising these results so that the mean of the 10 largest flow values are equivalent.

5.3.3.2 Experimental Data

The data from image 1 in Chapter 4 was processed again using the aliasing rules and integrating detector response. The full bandwidth



Figure 5.13: The 2nd order coefficient gives a measure of linearity, and the RMS error was calculated from the fitted line. As the bandwidth decreases, the aliased and τ =60 µs detector response is more linear and has slightly lower RMS error than the anti-aliased response.

FFT of the photocurrent in each pixel was taken, and then squared to calculate the power. This was passed through the detector response and the aliasing rules were applied. However, in order to see how the detector



Figure 5.14: A 50 μ s integration time confirms that as further higher frequencies are aliased to lower frequencies (i.e. not attenuated by detector), the point where the anti-aliased and aliased lines cross moves further right.

filter response increases the error in the aliased results (because it acts as an anti-aliasing filter, albeit not one with sharp cut-off), images with only aliasing and no detector response were also evaluated. The results can be seen in Figure 5.15, and the RMS error for the three scenarios (evaluated against the image at full bandwidth with no detector response) is shown in Figure 5.16.

The images confirm that aliasing (i.e. under sampling) in LDBF can actually recover some of the sensitivity to higher flow values. This can mainly be seen in terms of the location and clarity of the veins (i.e. regions of higher flow). For example, the phenomenon is particularly marked as the bandwidth moves down to 5 kHz (e.g. Figures 5.15m and 5.15n), where the ability to distinguish the veins is greater in the 'aliased' image than under anti-aliased and detector response conditions. In Figures 5.15c and 5.15h the veins are harder to distinguish and the clarity of the higher flow regions is more obscured by noise.



The full FFT was used as the M_1 calculation method.



Figure 5.16: The RMS error calculated from the full bandwidth image with no detector response. The purely aliased results show lower error as the bandwidth decreases. 5 kHz seems an optimum bandwidth as the lines start to converge slightly at lower bandwidths.

This reduction in error may be because the spectrum decays towards higher frequencies, and thus when the aliasing causes those frequencies to fold back, the effect is to add greater weight to spectral frequencies towards the Nyquist frequency. This is illustrated in Figure 5.17, which shows the Pareto distribution up to 10 kHz and the aliased spectrum with and without the detector response. It is the effect of folding this shape which helps to 'recover' the linearity.

It must be noted that the effect of aliasing is also dependent on the relative phase of the original and aliased frequencies. The aliased shape in Figure 5.17 will only be that shape if the phases in the aliased frequencies are either exactly 180° apart (for folded frequencies) or in phase (for directly copied frequencies). LDBF signals have random phase and so the quantity of aliased power is dependent on this random relative phase. However, if the spectra from multiple undersampled photocurrents are averaged, the aliased shape can be seen. The global
effect when frequency weighting and accumulating the aliased spectrum, despite its reliance on phase, is to add a 'useful error' back in and therefore recover some linearity.



Figure 5.17: The Pareto distribution alongside aliased spectra. The shape of the aliased spectrum may explain why sensitivity to higher flow values is recovered.

Having shown that the aliased flow response can, in fact, produce lower error, there are two options. The first is to accept the low bandwidth error; further work is then required to ascertain its effect on clinical interpretation as a result of the error. This will be discussed in the conclusion. The second is to try and correct the error using Bayesian Inference. This is the subject of the next section.

5.4 Correction using Bayesian Inference

5.4.1 Bayes' Theorem

If, after filtering by the detector, the shape of the aliased spectrum is known, then that information can be used to recover the full bandwidth spectrum and correct for undersampling. Again, the machinery used is that of Bayesian Inference, essentially determining what the most probable parameter values are given the power spectrum of the undersampled photocurrent and knowledge of its probable true shape. This is the Bayesian parameter selection problem [46]. Bayes' theorem can thus be cast as:

$$p(\theta_1, \theta_2 | P(f), U) = \frac{p(P(f) | \theta_1, \theta_2, U) p(\theta_1, \theta_2 | U)}{p(P(f) | U}$$
(5.10)

This is the posterior and reads as the probability of θ_1 and θ_2 given a power spectrum P(f) and an updated model function U (U will be defined in Section 5.4.2). The likelihood function then inverts the posterior:

$$p(P(f)|\theta_1, \theta_2, U) = \prod_{n=1}^{N/2} \frac{1}{\sigma_n \sqrt{2\pi}} e^{-\frac{(P(f_n) - U(f_n))^2}{2\sigma_n^2}}$$
(5.11)

to give the probability of the power spectrum occurring given some parameter values of the updated model U. It has already been shown in Chapter 3, that upon averaging, the error model of the power spectrum can be approximated by a normal distribution, and that the variance is:

$$\sigma_n^2 = k \frac{U(f_n)^2}{L}$$
(5.12)

There is no reason to prefer any particular parameter values over another, although it may be that, in the future, more informative priors are deployed in LDBF. A non-informative flat prior is used here:

$$p(\theta_1, \theta_2 | U) = p(\theta_1 | U) p(\theta_2 | U) = \frac{1}{\Delta \theta_1} \frac{1}{\Delta \theta_2}$$
(5.13)

And the normalising function (constant for all parameter values for a fixed P(f)) is:

$$p(P(f)|U) = \iint_{\Delta\theta_1 \Delta\theta_2} p(P(f)|\theta_1, \theta_2, U) p(\theta_1, \theta_2|U) \, \mathrm{d}\theta_1 \mathrm{d}\theta_2$$
 (5.14)

The aim, then, is to find the values of θ_1 and θ_2 which maximise the posterior probability function, given the updated model U. This updated model is the subject of the next section.

5.4.2 Updated Model

The Pareto distribution describes the LDBF power spectrum:

$$U_p(f;\theta_1,\theta_2) = \theta_1^2 \left(1 + \frac{f}{\theta_2}\right)^{-\frac{1}{1.74} - 1}$$
(5.15)

Because the effect of the integration attenuates the spectrum before aliasing occurs, the Pareto distribution is multiplied by the integration response. Leutenegger et al. [73] use the response (Equation 5.16) as a correction factor (to reverse the effect of integration), but it is used here to attenuate the full model:

sinc
$$(f\tau) = \frac{1}{\tau} \int_{-\tau/2}^{\tau/2} \cos(2\pi f t) dt$$
 (5.16)

This sinc function (the result of a rectangular averaging window in the time domain) determines the amplitude response of a frequency fusing an integration time τ . This stage of the model becomes:

$$U_i(f) = U_p(f) \times \operatorname{sinc}(f\tau)^2$$
(5.17)

Finally, the full bandwidth spectrum $U_i(f)$ is reduced to a set bandwidth using the aliasing rules in Equation 5.9, and becomes U(f). The process is illustrated in Figure 5.18.

For this simulation, τ was set to 60 μ s.

5.4.3 Methodology

A set of photocurrents was generated as in Figure 5.19. For each set of parameters, an amplitude spectrum of length $0.5NL \times k \times \frac{\text{High}f_s}{\text{Low}f_s}$



Figure 5.18: The model U(f) is formed in three stages - generating the Pareto with 2 parameters, passing the distribution through the detector response with set τ and finally applying the aliasing rules.

was produced, where N (the DFT length) and L (number of averages) have the same meaning as in Chapter 4. This amplitude spectrum was frequency weighted and accumulated to give the real flow value. NLwas selected to be 1,024, so that as L is increased N is reduced and the performance at different numbers of averages can be ascertained. k is the number of photocurrents to evaluate for each parameter set and values of N and L, and was chosen to be 16 to provide a large enough number of results to analyse variance. k also serves to over-sample the frequency spectrum, as has been done in previous simulations. Instead of applying the aliasing rules to the generated spectra, they were undersampled by the rate $\frac{\text{High} f_s}{\text{Low} f_s}$. For example, if the high sampling frequency is 40 kHz and the low sampling frequency is 10 kHz, then the undersampling rate is 4 and every 4th sample is selected for the reduced bandwidth photocurrent.

Once k photocurrents are generated, each can be passed to the process in Figure 5.20. The signal is split into L segments of N length and the FFT calculated. The power spectra of the L FFTs are added to produce an average. A flow value is calculated directly from this averaged spectrum before any inference, to evaluate the reduction in error, if any, produced by inference.

26 values of each parameter were tested, giving a total number of



Figure 5.19: The process for generating a set of photocurrents for the correction test. The full amplitude spectrum, after calculating the real flow value, is attenuated and undersampled to create the aliasing effect.



Figure 5.20: A photocurrent of length NL is processed to determine the corrected flow.

photocurrents for each low bandwidth value and NL pair of $26 \times 26 \times 16=10,816$. Recalling from Chapter 3 that the bias and variance are a function of N and L, several pairs of these values were tested. These pairs were, for N and L respectively, 8 and 128, 16 and 64, 32 and 32, 64 and 16, 128 and 8, and 256 and 4. The low bandwidth sample rates were 10 kHz, 7.5 kHz and 5 kHz (bandwidths of 5 kHz, 3.75 kHz and 2.5 kHz respectively). For the 7.5 kHz sample rate, High $f_s=60$ kHz to give an integer undersampling rate.

Process P6.7 calculates the flow by creating a Pareto distribution

with the maximum probability parameters, frequency-weighting and accumulating the weighted samples.

5.4.4 Results

5.4.4.1 Example Fit

Before displaying the full results for each bandwidth, some example fits from the inference will be shown. Figure 5.21 is one such example of the posterior probability function. The maximum value gives the coordinates of the most likely parameters.



Figure 5.21: An example of the posterior probability $p(\theta_1, \theta_2 | P(f), U)$ evaluated over 128 values for each parameter. The low sampling frequency was 10 kHz, N=32, L=32, input $\theta_1=1$ and $\theta_2=1,000$. The most likely parameter values, at the peak of the posterior, are at the coordinates of the input parameters.

It is computationally expensive to evaluate every set of parameters of the posterior probability space. However, the images are shown here for information. To speed up the processing of all the photocurrents, a climbing search algorithm was used. This will be discussed further in



Figure 5.22: The most likely output parameter values (from Figure 5.21) provide a full bandwidth curve to compare to the original data. Also included is the averaged spectra and the fitted line. The low sampling frequency was 10 kHz, N=32, L=32, input $\theta_1=1$ and $\theta_2=1000$.

Section 5.4.5.

The ability of the inference to recover flow varied. An example of a good fit is shown in Figure 5.22, where the full bandwidth spectrum is recovered almost exactly. The posterior probability and fits for a worse case example (where the maximum probability coordinates were further from the true coordinates) are shown in Figures 5.23 and 5.24 respectively. It is worth noting that more than one combination of parameters may fit the averaged spectrum similarly, therefore comparing the input and output flow values will give the decisive verdict on the ability of the inference to recover the true flow.



Figure 5.23: The posterior of a worse result, where the coordinates of the peak are further from the input parameters. The low sampling frequency was 10 kHz, N=32, L=32, input $\theta_1=1$ and $\theta_2=1000$. This is a result of the stochastic nature of the photocurrent, and where the sizes of N and L are limited, the inference will not always detect the true parameters.



Figure 5.24: The output parameter values (of the worse result from Figure 5.23) provide a full bandwidth spectrum. The low sampling frequency was 10 kHz, N=32, L=32, input $\theta_1=1$ and $\theta_2=1000$.

5.4.4.2 Corrected Flow

An example of the corrected flow results is given in Figure 5.25. The full set of corrected flow results are given in Appendix C. The true flow, calculated from process P5.3, is the x-axis value. The direct calculation is the result of process P6.4; the corrected flow is the result of process P6.7. To calculate the RMS error and linearity, the directly calculated flow was normalised to the corrected flow (Figure 5.25b) and 2nd order polynomials fitted. The linearity is the 2nd order coefficient and the RMS error is calculated from the fitted line. These results are given in Figures 5.26 to 5.28.

Figure 5.25 illustrates how the inference process recovers the sensitivity of the system to higher flow values. Essentially, this is because it can predict the higher frequency components based on the lower frequency components.

The RMS error and linearity change with N and L pairings. An increase in N provides more points to perform the inference fit, but an increase in L reduces the variance and brings the spectral estimates closer to the mean. The optimum pairings at all sampling frequencies seems to be either N=64 and L=16, or N=128 and L=8 based on the lowest 2nd order coefficient and RMS error from Figures 5.26 to 5.28.

Now the number of data points per photocurrent is fixed at 1,024. This could be increased to yield more points or averages at higher sampling frequencies. There is no reason that the number of averages needs to be a power of 2, but it is more important to keep the DFT calculation (i.e. N) at whole powers of 2 for efficient calculation by FFT. The limitation on NL is the length of time the photocurrent can reasonably be called stationary, and the minimum output sample rate required (f_s^{bf}) .

The absolute error values decrease as f_s increases: this is to be expected because less aliasing is occurring, which means that the recovery



Figure 5.25: Flow results for $f_s=10$ kHz, with N=128 and L=8. The inference recovers sensitivity to higher flow values, because the higher frequency components of the spectrum are recovered for the first moment calculation.

of linearity is more marked for $f_s=5$ kHz than for $f_s=10$ kHz. However the inference seems to reduce the RMS error in roughly equal measure.

The variance of the error is larger at higher flow and this is also to be expected. As the greater amount of power at higher frequencies is



Figure 5.26: Flow results for $f_s=5$ kHz. The inference goes some way to reducing the RMS error, but is very effective at correcting the non-linearity introduced as a result of reducing the bandwidth.

aliased (and also lost due to the attenuation of the integrating detector) the uncertainty of the inference increases.

The results show that, by using knowledge of the underlying statistics of the photocurrent, Bayes' theorem can be applied to reduce the error in



Figure 5.27: Flow result for $f_s=7.5$ kHz. The reduction in RMS error and 2nd order coefficient is similar to 5 kHz, except that, as is to be expected, the original values are smaller.

an undersampled photocurrent. For example, applying the inference to data sampled at 5 kHz, the RMS error can be reduced by approximately 35% (from 20 to 13) and the 2nd order coefficient decreased by 88% (Figure 5.26). To test how well this inference scheme functions on real data, it will be applied to a rotating diffuser experiment in the next



Figure 5.28: Flow results for $f_s=10$ kHz. The RMS error is the lowest where N=32, but the 2nd order coefficient where N=256.

chapter. Before concluding this chapter, a possible implementation will be suggested.

5.4.5 Implementation

Evaluating the entire posterior probability space is a computationally expensive task. However, there is no need to know the probability of each pair of parameters given a power spectrum - all that is required is the coordinates of the maximum probability. Therefore there is wasted resource in evaluating the whole probability space, and in fact much of it can be ignored. It can be used as an advantage that the log of the posterior forms a smooth function (as has been the case in all observed posterior spaces, although the possibility of local maxima cannot be ruled out) because all that needs to be done is simply 'climb the posterior hill' until the top is reached.

However, the movement is blind. Because the direction of the top of the hill is unknown, trial movements in random directions are required, and these movements are only 'accepted' when a higher point is reached. An algorithm that implements this scheme has been developed by the author of this thesis, which is summarised in Figure 5.29.



Figure 5.29: The climbing algorithm for fast search of the posterior probability to find the parameters of the maximum value.

The climbing algorithm initialises a pair of coordinates (Xcurr, Ycurr) to (64, 64), and calculates the posterior POSTcurr. The coordinates are translated to parameter values (i.e. $\theta_1 = f(X)$ and $\theta_2 = g(Y)$), but for the purposes of the algorithm an index of values of 1 to 128 for X and Y is used. A random movement (R) in the range -4 to +4 for each coordinate is generated, and the new coordinates are tested to make sure they still reside in the range 1 to 128. The posterior probability POSTnew is evaluated on these coordinates: if it is greater than POSTcurr (i.e. movement uphill has been detected), then POSTcurr becomes POSTnew and Xcurr and Ycurr are also updated. If it is not greater, the random movement is rejected and the values of Xcurr and Ycurr remain unchanged. For every random movement, after an update of POSTcurr, a counter is incremented. This acts as an exit method for the algorithm. If 128 random movements result in values of POSTnew that are less than or equal to POSTcurr, so that no movement uphill is found, the algorithm assumes the top has been reached. 128 was chosen because it is twice the number of sample spacings possible (8×8) . A balance needs to be sought between the maximum size of the random movements (R) and the number of random movements required to exit the algorithm (E); larger random movements may reach the peak faster but require more evaluations before exiting to ensure the peak has really been reached.

Figure 5.30 shows the log of the posterior. The log is easier to calculate computationally because it utilises additions rather than multiplications (see Equation 3.20). The dark blue dots show the locations of Xcurr and Ycurr, but do not reveal the random movements attempted. For this particular task, the percentage of climbing algorithm function evaluations versus the full posterior function evaluations is:

$$\frac{N_{\rm climb}}{N_{\rm full}} = \frac{415}{128 \times 128} \times 100 = 2.53\%$$
(5.18)



Figure 5.30: The log of the posterior, which is simpler to evaluate computationally. Instead of evaluating the whole posterior probability space, the path of the climbing algorithm shows the maximum can be reached in much fewer posterior evaluations (in this case the total was 415, and the number of movements was 37). The low sampling frequency was 10 kHz, N=32, L=32, input $\theta_1=1$ and $\theta_2=1,000$.

Further optimisation of R and E, by using a larger set of posteriors, may reduce the average value of N_{climb} .

5.5 Conclusion

The effect on LDBF flow values as a result of using low bandwidth CMOS imagers has been described. Aliasing, undesirable in acquisition systems, acts as a 'useful' error because the folding of the decaying spectrum increases the sensitivity to higher flow values. The low-pass filtering effect of the integrating pixel reduces aliasing (thus reducing the impact of this phenomenon) and serves to limit the noise bandwidth. However the shape of the filter response does not provide a sharp cut-off and has large side lobes, distorting amplitudes in the frequencies of interest.

Therefore, there are strong reasons to keep the integration times small, as this also allows for faster sample rates. However, this strength must be balanced with the maximisation of the SNR of the system in the presence of shot and thermal noise. There is a trade-off each system must make, dependent on laser power, illumination area, magnification, aperture size etc. Moving towards practicalities, the next chapter will describe the test of the sensitivity of two CMOS imaging systems and show that the inference techniques explained in this chapter can be used to correct the unavoidable errors of low bandwidth acquisition.

Chapter 6

CMOS Camera Evaluation

6.1 Introduction

The theoretical performance of an integrating CMOS detector has been described. It has been shown that aliasing, occurring due to the limited bandwidth of a CMOS imager, can actually serve to reduce the error in the estimated M_1 values when detecting frequency distributions with shapes like those of the power spectrum of an LDBF signal. The aim of this chapter is to evaluate the sensitivity of two CMOS imaging systems when imaging samples with small to large velocity components. The first imaging system is currently being developed at the University of Nottingham. In this device, a CMOS sensor (Micron MT9M413) is connected to an FPGA controller. The second imaging system is an off-the-shelf camera system. It is constituted of an Optronis CL600x2 camera, incorporating a Cypress Semiconductor LUPA-1300-2 CMOS sensor, with Dalsa Xcelera-CL PX4 Full frame-grabber.

In previous chapters, the tests comparing full and low bandwidth systems were based on photocurrents with known underlying properties. This was either in the simulations, where the parameters producing the photocurrent, and therefore flow, were known, or in experimental data, where the full bandwidth spectra were available. The low bandwidth results were produced by reducing the quality of the known data via subsampling, and then performing a comparison. In this chapter, however, the full bandwidth properties of the photocurrent are unknown, and all that is available at the output of the imagers is the low bandwidth (undersampled) photocurrent. In order to evaluate the sensitivity of these cameras, properties of the sample under inspection are required.

In the first experiment (using the Micron imager), the known property is the speed of a rotating diffuser. The speed of the diffuser can be used to model the distribution of Doppler-shifted frequencies present at the detector. In the second experiment, human finger blood flow is modified by cold and hot water immersion and then independently evaluated by a full bandwidth single point monitor (Moor Instruments VMS-LDF2) acting as the gold-standard. These full bandwidth results then represent the comparator. In the first experiment, the low bandwidth results are also corrected using the Bayesian Inference techniques described previously, and compared to the gold-standard.

Recent tests on LDPI systems (for example, in Leutenegger et al. [73]) do not consider the response compared with the full bandwidth result. Although the pulsatile signal and regions of high and, particularly, low flow are distinguishable, a comparison with a known gold-standard result is not given. It is acknowledged that a straightforward colour map highlighting these regions may be sufficient for clinical practice. However, knowing the accuracy of the result given by the imagers will further the chances of wider clinical uptake, because understanding the meaning of the result displayed increases confidence in the system.

For that reason, blood flow is stimulated, rather than the occlusion/release test often applied. The upper limit of the system's response can then be tested.

6.2 CMOS Sensor & FPGA

6.2.1 Overview

This device has been custom built specifically for the LDBF application by Sun [122] and the author of this thesis had no input to the design or build of the system. The commercial CMOS sensor is connected to an FPGA (Xilinx Virtex 6 XC6VLX240T), which is in turn connected to the PCI slot of a PC. The functioning of the device is controlled from the PC. The FPGA itself acts as the camera controller and receives the sampled photocurrent data (raw data) from the sensor. In its normal configuration, it also processes this data into a value representing blood flow, by calculating M_1 via the FFT. However, in the configuration for this experiment, the FPGA transmits the raw data to the computer's RAM so that it can be analysed directly in MATLAB. The experimental configuration is illustrated in Figure 6.1.

A rotating diffuser is used to test the sensitivity (Figure 6.2). This is a well known experiment [38, 102] and has predictable results. The diffuser provides a distribution of frequencies with similar shape to LDBF (i.e. monotonically decaying) and so is suitable for estimating the CMOS camera's response in an LDBF environment. The expected distribution of Doppler shifts (and so mean frequency) will first be modelled. The experimental results will then be described, followed by an analysis and comparison.

6.2.2 Theoretical Response of Diffusing Surface

The Doppler shifts, which occur as a result of illuminating a rotating diffuser, can be predicted by the standard equations in Chapter 1 [16]. The diffuser consists of a 26 mm diameter disc connected to a low speed motor driver (Maxon 110048 DC brushed). A piece of white paper is



Figure 6.1: The experimental configuration for the rotating diffuser experiment. A 580 nm laser is expanded to illuminate the disc and the static diffuser provides an interfering reference beam. A lens focuses the Doppler shifted light onto the CMOS sensor, which passes data to the FPGA and then to the PC. The integration time of the sensor was 78 μ s and hence the sampling rate was 12 kHz.



Figure 6.2: A small section of the rotating disc comprising 320 horizontal and 1 vertical pixels was imaged (highlighted green). The radius, R, is 13 mm, and the velocity, V, is variable.

attached to the front of the disk and functions as the diffusing surface. A static diffusing reflector is placed in front of the rotating disc and this provides the interfering reference beam (see Figure 6.1).

The surface profile of the white paper needs to be known in order to predict the relative phase difference between the illuminating coherent light and its reference. Because the paper is rotating, the relative phase differences are a function of time, and it is the differential of the phase differences that give the Doppler shift frequency.

The starting point is a one dimensional paper surface model. Vernhes et al. [128] list a variety of parameters describing the statistics of the surface: the dispersion, the skewness and the kurtosis. The dispersion quantifies the RMS surface height variations from the mean. The skewness quantifies whether the distribution of surface height variations has a longer tail at the positive or negative side of the slope and the kurtosis quantifies the 'peaked-ness' of the surface height variations. Leising, in an application note for a profilometer, lists typical values for white paper [72]. An RMS dispersion of 5 or 6 μ m is typical of white office paper. Gaussian statistics for the skewness (0) and kurtosis (3) will be assumed as example values vary around these regions and no information on the true values for these parameters is known.

Therefore a Gaussian process with an RMS height variation of 5.7 μ m was randomly generated in MATLAB. The length of the paper was 2²² spatial samples, corresponding to 5 mm. The surface height variations u(n) were oversampled by 2¹³, i.e. the first 2⁹ spatial frequencies had unity power, with the remainder zero power. This resulted in a spatial sampling frequency, $1/F_x$, of 838,860,800 m⁻¹. Thus, the sample number n can be related to distance in meters x by:

$$x = F_x.n \tag{6.1}$$

The data were oversampled for two reasons: firstly because the speed differences are evaluated by sub-sampling the spatial data, and no aliasing is desired, and secondly because the finite difference method of differentiation (used to determine the Doppler shifts) becomes less accurate as the frequency content approaches the Nyquist frequency. Figure 6.3 gives an example of the paper surface, with the histogram in

Figure 6.4.



Figure 6.3: The zero-mean surface height variations of the simulated paper surface. This Figure shows 1 mm of a sample length of 5 mm. The spatial sampling frequency, F_x , is 838,860,800 m⁻¹.



Figure 6.4: The histogram of surface height variations. The RMS (i.e. standard deviation) of variations was 5.7 μ m.

To simulate the difference in speed of the paper, the spatial data was sub-sampled. The degree of sub-sampling, S, is linearly related to the speed of movement v(S), and was increased from 1 to 15 in unity steps. The function u(n) becomes u(S.n). However, the speed of the slowest sub-sample rate (i.e. 1) must be chosen artificially. It was set at 0.276 mm/s as this produced a maximum speed, at the sub-sample rate of 15, of 4.1 mm/s. This corresponds to the maximum speed, at the edge, of the spinning disc used in the experiment. Similar to Equation 6.1, the time in seconds t in relation to the sample number n, is:

$$t = F_t.n \tag{6.2}$$

where F_t , constant for any S, is:

$$F_t = \frac{F_x \cdot v(S)}{S} = F_x \times 0.276 \text{ mm/s}$$
 (6.3)

Now the surface height variations of the 15 sets of data (essentially the same data with varying sub-sample rates) can be treated as a phase difference by dividing by the wavelength of the incident light, λ . For this experiment, the wavelength was 580 nm. This phase difference is then differentiated with respect to time to give an instantaneous frequency:

$$\omega = \frac{\mathrm{d}\phi}{\mathrm{d}t} = \frac{1}{\lambda} \frac{\mathrm{d}u(S.n)}{\mathrm{d}n} \frac{\mathrm{d}n}{\mathrm{d}t}$$
(6.4)

The final term in this equation, a product of the chain rule, can be evaluated from Equation 6.2 by:

$$\frac{\mathrm{d}n}{\mathrm{d}t} = \frac{1}{F_t} \tag{6.5}$$

The Doppler shift is twice the frequency, because source and detector are stationary relative to each other (see Chapter 1). Converting from the angular frequency, f_D , the Doppler shift, is:

$$f_D = 2\frac{\omega}{2\pi} = \frac{\omega}{\pi} \tag{6.6}$$

 $\frac{1}{\lambda} \frac{du(S.n)}{dn}$ is evaluated by finite difference differentiation to produce a range of instantaneous frequencies. Figure 6.5 shows an example of the frequencies produced versus time. By calculating the histogram of the frequency data, the shape of the distribution of frequencies at each speed can be determined.



Figure 6.5: An example of the instantaneous frequencies produced by differentiating the phase difference, caused by the surface height variations, with respect to time.

The simulation was performed 32 times, and the distributions at each speed averaged. The averaged distributions of the first five speeds are shown in Figure 6.6. The shapes match those observed in the experimental results of Romero in 2000 [102].

The mean frequencies of the averaged distributions were calculated by $\frac{M_1}{M_0}$. This gives a representation of the true mean frequency at each speed (shown in Figure 6.7). This mean frequency response can now be compared to the experimental data.

6.2.3 Experimental Data

This section will describe the experimentation and signal processing required to evaluate the response of the CMOS sensor. The actual collection of raw photocurrent data from the camera was not performed by the author of this thesis but was performed by the designer of the



Figure 6.6: The averaged frequency distributions of the first five speeds of a simulated rotating disc.



Figure 6.7: After 32 simulated distributions of the spinning disc were averaged for each speed, the mean frequency was calculated by $\frac{M_1}{M_0}$.

CMOS camera system [122]. However all the subsequent processing, analysis and correction was performed by the author.

The capabilities of the system limit the amount of data that can be transferred and stored on the PC. In its normal configuration, the FPGA processes all the high speed data locally and hence reduces the output data bottleneck. However, because in this configuration the raw photocurrent data is required, only a small area can be captured. A 320x1 pixel horizontal section positioned at half the height of the disc and to the right side (as shown in Figure 6.2) was selected as the region of interest.

The noise floor was measured by recording the photocurrent with no rotation of the disc, calculating the power spectrum and then evaluating the mean. Figure 6.8a shows the static noise power, which is used as a soft threshold for the rotating disc spectra calculations. The soft threshold is subtracted from the calculated power spectrum. Any resultant negative components are set to zero. The DC light level is recorded as the mean of the stationary photocurrent. As can be seen in Figure 6.8b, the light level decreases at increasing radius, due to the illumination angle of the laser. As the shot noise is reduced with the DC light level, a slight decrease in overall noise power is seen.

The DC voltage supplied to the motor was increased from 250 mV to 4.75 V in steps of 250 mV. The relationship between the supply voltage and speed of the disc at increasing radius is illustrated in Figure 6.9.

The camera, sampling at 12 kHz, collected 32 windows of 1,024 samples for each pixel. The 1,024 point FFT was used to provide a power spectrum for each window. The 32 spectra then provided an average. Thus a spectrum with low variance was produced for each radius at 1 pixel intervals, and for each DC motor supply voltage.

Two operations were then performed. The first was to directly calculate the mean frequency by evaluating $\frac{M_1}{M_0}$, and the second was to correct the spectrum using the Bayesian techniques described in the previous chapter, before calculating the mean frequency for the corrected spectra.



Figure 6.8: The noise power and DC light level of the device calculated from the stationary disc. The noise power decreases slightly towards the edge of the disc. This is because the laser illuminates at an angle and the DC light level is smaller at the edge than at the centre.

6.2.4 Comparison & Correction

Examples of the averaged spectra are displayed in Figure 6.10. The mean frequency increases because the distribution of frequencies in the



Figure 6.9: The voltage sets the speed of the disc. This graph shows the relationship between radius of the disc and surface speed, for 5 voltages.

spectrum flattens, as predicted by the simulations. Figure 6.11 shows how the mean frequency increases with increasing radius and voltage. As the voltage increases, the plots become less linear. This familiar result is caused by the integration time limiting the bandwidth, producing underestimations. The assessment of the accuracy of these results can now be evaluated by comparing experimental with theoretical data.

The result for 4.75 V is shown again in Figure 6.12. The theoretical data was overlaid by replacing the speed axis with a pixel value, using the information calculated for Figure 6.9. At the maximum speed of the disc, which corresponds to pixel 320, the speed was 4.14 mm/s and the theoretical mean frequency was 3.8 kHz.

Firstly, comparing the CMOS camera response with the theoretical response, the camera underestimates the mean frequency and only has the ability to detect mean frequencies up to around 2 kHz. However, the response may be linear enough to represent blood flow increases and decreases with sufficient accuracy, particularly in a flap monitoring situation where zero to baseline flow is the normal range.



Figure 6.10: Example averaged spectra for a motor DC supply voltage of 4.75 V. The device had significant noise components around 100 Hz and 4.25 kHz. The power results from these frequencies were discarded and replaced with averages of the surrounding values. Later, the power at these frequencies will be replaced with values calculated from a Bayesian inference fit.

Secondly, because it is known that the power spectrum is approximately Gaussian [102], the inference techniques of Chapter 5 can be used to correct the undersampled spectra.

The first step is to fit an updated model U to the recorded spectra. Instead of using a Pareto distribution, an Exponential distribution, because the Gaussian is based on this function, is input to the integration filter (τ =78 µs) and then to the aliasing model. The posterior probability was then evaluated over a 2 parameter range and the most likely parameters selected. The mean frequency of the model, using the most likely parameters, was calculated and is displayed in Figure 6.12 (light green line). An example of a fitted Exponential is given in Figure 6.13.

The fitted mean frequency results agree closely with the recorded



Figure 6.11: The mean frequency of the averaged spectra was calculated for each pixel over a range of supply voltages (i.e. increasing motor speeds). One reason that the mean frequency may flatten off at lower pixel values, where it would be expected to continue decreasing towards 0 Hz, is because the true level around 100 Hz in the recorded power spectra is unknown (due to noise in the equipment), and as there is significant power at these frequencies at lower speeds, the calculation gives a skewed result.

spectra. Visually, the Exponential fits the spectrum well, despite the variance of the averaged spectrum. The most likely parameters were then used to produce a full bandwidth (up to 20 kHz) power spectrum. The mean frequency of this spectrum was calculated and is also displayed in Figure 6.12 (dark green line). As the mean frequency of the CMOS camera output is constant at low speeds (at pixels closer to the origin of the spinning disk, Figure 6.12), the fitted and recovered spectra are also constant. One reason for the levelling is that the CMOS camera might not have been focussed directly at half the height of the disc, therefore producing Doppler shifts at greater frequencies than expected. After the point where the theoretical response crosses the recovered response,



Figure 6.12: The output from the CMOS camera is compared to the theoretical response, the mean frequency of a decaying exponential (assuming aliasing and a 78 μ s integration time) fitted using Bayesian Inference, and the mean frequency recovered using the fit parameters (r=0.986 between recovered and theoretical)). The CMOS camera response was recalculated for this graph using corrected values (from the exponential fit) for 100 Hz and 4.25 kHz.

at pixel 120, there is good agreement between the corrected values and the theoretical mean frequency (r=0.986).

In this section it has been shown that the CMOS camera, sampling at 12 kHz and with an integration time of 78 μ s, underestimates the true mean frequency, but that it can be corrected for with good accuracy using the inference techniques of the previous chapter. The next section describes the evaluation of a second camera system that images blood flow *in vivo*.



Figure 6.13: The frequency spectrum at pixel 260, fitted with a decaying exponential model U, and the representative 'corrected' spectrum. The fitted spectrum can also be used to recover the lost components around 100 Hz and 4.25 kHz.

6.3 Optronis CL600

6.3.1 Overview

This system was comprised of an Optronis CL600x2 CMOS camera, connected via two CameraLink cables to a Dalsa Xcelera framegrabber PCI card mounted in a PC. The camera is monochrome and has a maximum resolution of 1,280 pixels (width) by 1,024 pixels (height). Manufacturer software on the PC controls the camera settings and receives the data to the computer's RAM. After being stored in the RAM, the raw data was analysed offline in MATLAB. In this respect, the system is simple to assemble because it is made up of off-the-shelf components. A 40 mW infrared (785 nm) laser, aligned with a 30 mW red (660 nm) laser, is used to illuminate the sample. The red laser is provided for ease of optical configuration. The beam is expanded using a square diffuser, similar to the previous experiment, to a size of approximately 1 cm by 1 cm.

In order to assess the accuracy of this LDPI system, six volunteers had their blood flow assessed in one finger under hot and cold water stimulation. It was recorded with the Optronis system and also using the Moor Instruments VMS-LDF2 single point monitor as a comparison standard.

6.3.2 Ethical Considerations

University of Nottingham Faculty of Engineering ethical approval with application ID 2012-01, as used in Chapter 3, covers all experimentation undertaken in this chapter. Data was collected and stored in accordance with Faculty requirements and no personally identifiable information was stored with the photocurrent data.

6.3.3 Experimental Method

6.3.3.1 Camera and Optical Apparatus

Figure 6.14 shows the experimental set up. The equipment was configured in order to maximise the photocurrent SNR. Firstly the laser was expanded only to the size of the imaged region to maximise the incident optical power per unit area. Secondly the aperture of the camera lens (Schneider Xenon 50 mm) was opened to F/2.2, its widest setting. This had the disadvantage of reducing the depth of field, which meant the hand position was critical as small forward or backward movement from the focal point caused the image to be out of focus. Two ADC resolutions were available, 8-bit and 10-bit. Clearly a higher resolution is more desirable to reduce quantisation noise, but 10 bits must be transferred through the low-voltage-differential-signalling (LVDS) cabling to the PC as two bytes. Therefore this acquisition resolution doubles the data volume for the same spatial resolution and sampling rate. However, a look-up table is available on the camera which assigns the 10-bit value (range 0 to 1023) number to any value of choice in the 8 bit range (0 to 255). If the signal of interest is an AC signal with DC offset, and that DC offset is roughly known, then the ADC can sample at 10 bits and send the data as 8 bits. The signals sampled from illuminated tissue were found to be consistently in the lower quartile of the 10-bit range. Therefore any value in the range 0 to 255 was transmitted as 0 to 255 respectively, and any value above this (256 to 1,023) was also transmitted as 255. In theory this means the output would saturate at higher light levels, but in practice this situation did not occur because the full-field LD signals are not high in power. The results section will reveal typical ADC sample values.

The receiving buffer, limited by the size of the PC's RAM, allowed up to 20,000 frames to be saved at the selected resolution. At 5 kHz, this produced four seconds of data. The photocurrent at each pixel was then divided up into 1,000 sample windows. Each window was processed by 1,000 point FFT to generate a power spectrum $S_{raw}[k]$. A noise threshold N[k] was measured by recording the frequency spectrum of a static diffuser and taking the average value of the resultant power spectrum. This was used as a soft threshold on $S_{raw}[k]$ to produce S[k]. As is normal in LDBF [73], the photocurrent was normalised by dividing by the square of the mean DC value:

Flow =
$$\frac{f_s}{N} \sum_{k=1}^{500} k \frac{P[k]}{DC^2}$$
 (6.7)

However, because the sensitivity setting of the camera was set high (in order to detect the photocurrent fluctuations with acceptable resolution) the DC value recorded at each pixel was very sensitive to fixed pattern noise. As a large image resolution was used, a 10 x 10 pixel spatial Gaussian low pass filter was applied to the DC image. Over a small group of pixels (10 x 10 pixels is 0.5 mm x 0.5 mm) the true imaged DC level is effectively constant. The aim of the normalisation is to divide by the average intensity, and not simply the individual pixel value, which is contaminated by fixed noise levels. By processing in this way, the noise level of the flow map is reduced. Figure 6.15 shows an example of the processed DC image.

The speckle pattern generated as a result of illuminating a rough surface with coherent light must be resolved by the imaging system, although it is not required to be strictly resolved (this would mean a pixel size less than or equal to the speckle size) [73]. The approximate diameter of a speckle (therefore allowing coherence area A_{speckle} to be evaluated) can be calculated by the wavelength and f-number (f/#) of the imaging system:

$$D_{\text{speckle}} = 2.44.\lambda f / \# = 2.44 \times 785nm \times 2.2 = 4.2\mu m$$
 (6.8)
The number of speckles per pixel can thus be evaluated by:

$$N_{\rm speckle} = \frac{A_{\rm pixel}}{A_{\rm speckle}} = \frac{14\mu m^2}{\pi \times 2.1\mu m^2} = 14.1$$
 (6.9)

Although the ratio of AC to DC current level is inversely proportional to the number of speckles per pixel [109], the aperture, the one practically variable quantity in Equation 6.8, also controls the amount of light entering the lens. Therefore to maximise the light level it is left at its widest setting (f/2.2).¹

6.3.3.2 Recording Procedure

Six participants, aged between 21 and 30, including both males and females, were recruited for the study. After a period of rest during which the study was explained and informed consent obtained, each subject had their baseline flow recorded using the VMS-LDF2, with the probe attached onto the rear of the right index finger. The baseline flow was also then recorded using the Optronis system, where 4 s of data (20,000 frames at 5 kHz) was collected. Then, each subject's right hand was immersed in a container of water at 15°C for 2 minutes. The hand was dried and black tape, acting as a reference marker, was applied to the fingertip as shown in Figure 6.14. Firstly the true blood flow value was recorded by attaching the probe of the VMS-LDF2 to the right index finger. Then the probe was quickly removed and the finger placed in the focal range of the camera system. Again, 4 seconds of data was recorded at a 5 kHz sampling rate.

To stimulate blood flow, the right hand was immersed in a container of water at 43°C for 2 minutes. Then the same recording procedure as

¹It is worth noting that for a LASCA system, the size of the speckle should approximately equal the size of the pixel detector and therefore the aperture must be reduced to F/7.2, if using this optical set up [16]. This means less optical power reaches the detector, although this is usually compensated for by increased integration times in LASCA configurations.



Figure 6.14: The experimental configuration for recording from the Optronis CL600x2. A 40 mW 785 nm laser is aligned to a 30 mW 660 nm laser, which illuminates a 1 cm² area of the finger. Black tape is used as a reference marker and also prevents glare and other unwanted effects [106] from the fingernail. Once the data is recorded, it is stored in the PC's RAM for offline analysis in MATLAB.



Figure 6.15: A DC image processed using a Gaussian low pass filter to reduce fixed noise. One of the effects of the filter is to create a transient at the boundaries of the image.

after the cold immersion was repeated. Thus three separate recordings from each instrument, a baseline, cold and hot, were obtained for each subject.

6.3.4 Results

Once the data was recorded it was processed as previously described. An example of 1,000 samples of the photocurrent from one pixel is shown in Figure 6.16, along with its power spectrum in Figure 6.17.



Figure 6.16: An example of the photocurrent obtained from one pixel of the Optronis CL600x2 after hot water immersion. The window length is 1,000 samples at 5 kHz.



Figure 6.17: The power spectrum calculated by fast Fourier transform of the photocurrent in Figure 6.16.

In order to evaluate the flow in the tissue, a small region was selected for averaging in order to compensate for the spatial variability. It is assumed the blood flow is roughly equivalent in a region of 11x11 pixels (approximately 0.5 mm by 0.5 mm). The region of interest is illustrated in Figure 6.18, which also shows a complete blood flow map of the imaged fingertip. Averaging spatially also smooths the power spectrum, as demonstrated in Figure 6.19.



Figure 6.18: A flow image, taken with the Optronis CL600x2, of the tip of the right forefinger of subject 5 after 2 minutes immersion in hot water. The location of the black tape can be resolved (compare with Figure 6.14). The black box (upper left) indicates the location of the 11x11 pixel region averaged to evaluate the temporally varying blood flow.

The flow values created from the average demonstrate the temporal variability of the blood flow. Figure 6.20 clearly shows the pulsatile signal, with the dichrotic notch (a slight deflection on the diastolic phase of the waveform) visible on several pulses. Thus the system is capable of 'imaging the pulse.'



Figure 6.19: An average of multiple spectra from a small region (11 by 11 pixels), at one window time point (i.e. not temporally averaged), using the Optronis CL600x2.



Figure 6.20: A small spatial region's flow values (11x11) was averaged at each output time point to give 20 samples at 5 Hz. The data was collected using the Optronis CL600x2. The pulsatile signal can easily be resolved.

Table 6.1 gives the raw output from the two systems. To compare the sensitivity, the percentage change was computed for the baseline and hot flow. The percentage change used the cold result as the fixed quantity as the cold water is controlled, rather than using the baseline flow value which may be affected by pre-experimental environmental factors:

$Perc.Change = \frac{Flow - Flow_{Cold}}{Flow_{Cold}} \times 100$							(6.10)
Subject	1	2	3	4	5	6	
VMS Base	31	7	8	8	7	13	
VMS Cold	15	8	9	31	17	8	
VMS Hot	230	130	270	300	510	150	
Op. Base	-	16.7	10.1	12.6	17.1	14.7	
Op. Cold	-	10.3	12.1	12.8	15.3	13.2	
Op. Hot	-	21.8	27.9	32.7	23.7	24.1	

Table 6.1: Flow results from the Moor VMS-LDF2 and the Optronis CL600x2 for six subjects. The two systems' units are arbitrary, therefore a direct comparison of values is not possible. One set of data for subject 1 was corrupted and the flow could not be determined.

6.3.5 Discussion

It is clear that the VMS system is much more sensitive to higher flow rates. This is not surprising given the bandwidth of the VMS system is six times the size of the Optronis. However the location of the 11 x 11 pixel window used for the Optronis analysis may not have been the precise location that the VMS system analysed the flow because the area covered by the probe is small (0.5 mm fibre separation at 785



Figure 6.21: A comparison of flow results from the two systems. The red box plot represents the flow recorded using the Moor VMS system, and the blue box the flow recorded using the Optronis camera. As the units of both devices are arbitrary, the percentage change from cold was used (hence the cold box plot shows zero). The edges of the box represent the 25th and 75th percentiles, and the whiskers the extremities of the data.

nm wavelength). To counter this, every effort was made to position the probe on the centre of the apex of the final joint of the finger, the location of the pixel window. It is unlikely that small variations in this location would adversely affect the results and the conclusion considering there are few large veins near the skin surface and so the microcirculation of the skin can be assumed homogenous in this region.

It is also acknowledged that hot water immersion is an extreme method of stimulating flow and that in clinical situations it is unlikely to be raised to this extent. For example, in burn analysis, first degree burns may cause the flow to be raised only up to 500% of normal values. Thus, what this data shows is that the Optronis system, despite the large flow stimulation, is only sensitive up to around twice the baseline because the bandwidth is limited.

6.4 Conclusion

Two CMOS camera systems have been evaluated for their ability to accurately represent the first moment of a Doppler-shifted photocurrent. The first system, a custom device compromised of a CMOS sensor sampling at 12 kHz and interfaced with an FPGA, underestimated the mean frequency of a rotating disk. The true values were computed using a theoretical model. It has been shown that the errors introduced by undersampling the photocurrent can be corrected for using Bayesian Inference techniques.

The second system, an Optronis CL600x2 high speed CMOS camera sampling at 5 kHz and interfaced with a PCIe framegrabber, also underestimated flow values, but the experiments were performed on human skin rather than rotating diffuser. A Moor Instruments VMS-LDF2 provided the gold-standard flow values. A cold and hot water immersion technique was used as the stimulator. Whilst the VMS device showed the hot water stimulation increased flow by a median value of 17 times, the Optronis system at 5 kHz only increased it by twice the cold water values.

Thus simpler off-the-shelf CMOS systems can detect the Dopplershifts off moving scatterers with acceptable accuracy at 'normal' flow values, and so can detect the pulsatile signal in this range, but are not sensitive to larger frequency shifts in the way that conventional systems are.

Chapter 7

Conclusion

7.1 Introduction

Accurately identifying the condition of skin blood perfusion is crucial in a variety of medical and surgical situations. For example, by diagnosing the severity of a skin burn through examination of blood flow levels, correct remedial surgery may be provided. Additionally, by monitoring the re-perfusion of a skin flap after transplant, appropriate action may be taken to prevent flap necrosis, or partial or total flap loss. NICE has published guidelines [87] that describe how using LDBF technology reduces the cost and improves patient outcomes in the treatment of skin burns.

LDBF technology has had a presence in the marketplace for several decades in its single point monitoring configuration (LDPM), but recently a market has opened for LDPI, the imaging configuration. Reports such as the NICE guidelines have further facilitated the expansion of this market, however there is much ongoing work to try and gain widespread clinical acceptance. This work includes technological development, where much of the current focus is on providing full-field video-rate imaging devices. Clinicians need fast, reliable, repeatable and understandable diagnostic tools. Working towards this, whilst being constrained by current technological limitations, has been the subject of this thesis.

LDBF is a non-invasive optical technique for measuring blood flow. It works by measuring the frequency changes (Doppler shifts) that occur when light interacts with moving RBCs. However, there are many millions of RBCs per μ l of blood, and with a distribution of scattering angles and velocity vectors, it is not just one single frequency shift but a distribution of shifts that is produced in a sampled region. The bandwidth of these shifts is typically in the range 30 Hz to 20 kHz. LDPM devices have therefore, with appropriate anti-aliasing filters, set an ADC sampling rate of 40 kHz. However, such sampling rates have proved difficult to achieve in imaging devices. This is for several reasons; firstly because many thousands of pixels (one frame) need to be sampled at once and the data transferred off-chip, which represents a large data conversion, transfer and processing task, and secondly because the integration times for the CMOS sensors, required to provide adequate SNR, limit the maximum sampling rate. LDPI sampling rates have therefore typically been in the 5 to 15 kHz range. Since LDPI devices are entering commercial production, it is appropriate to ask how the acquisition and processing method impacts upon the results, and whether the reduction in accuracy, if any, can be corrected for post-acquisition.

7.2 Summary of Results

From conversion of light into electrical current to final blood flow image, there are several processing decisions a system designer must make. For example, the ADC sampling rate or method of calculation of M_1 . Different chapters of this thesis have dealt with these different decisions, the errors they introduce and how the errors may be compensated for. Chapter 2 detailed current methods of evaluating M_1 from a timedomain signal, and suggested a fast and novel method utilising the Hilbert transform. The computational complexity of these algorithms was quantified.

In order to compare these processing methods and how they perform in low-bandwidth scenarios, the photocurrent must be modelled. By the machinery of Bayesian Inference, Chapter 3 has shown that the power spectrum of a LDBF photocurrent can be modelled using the twoparameter Pareto distribution with fixed shape factor. Chapter 4 used this model to compare the processing algorithms and determined that, as expected, the full FFT method was the most accurate but that an IIR filter, which approximates a frequency-weighted response, performs reliably. The approximate class of algorithms performed adequately, meaning that very fast implementations may give 'good enough' results in certain scenarios. The simulated photocurrents were used to evaluate a temporal LASCA processing scheme and this revealed high error and non-linearity. The first moment analysis was repeated on real recorded images, and encouragingly these results matched the simulation. Additionally, it is confirmed that the 'Draijer' algorithm is an effective and efficient way of calculating M_1 from the noise-like photocurrent.

Chapter 5 turned to CMOS imagers, the main imaging technology used in LDPI, and modelled the response of an integrating pixel and the aliasing of high frequencies that occurs as a result of undersampling. It was then shown that, by understanding how the aliased frequencies fold or copy back into the power spectrum, the error introduced because the LDBF photocurrent is undersampled may firstly be characterised and then corrected for using Bayesian Inference. It has been shown that allowing high frequencies to alias introduces a 'useful error' that can reduce the non-linearity of the output M_1 values. This reduction is limited by the size of integration time. However the error can be further reduced using Bayesian techniques.

Chapter 6 then applied this methodology to a CMOS camera system imaging a rotating diffuser, confirming that the true mean frequency, underestimated due to low bandwidth acquisition, can be recovered using the inference technique. A second commercial CMOS camera system was then evaluated on its detection sensitivity to high blood flow using a water stimulation technique. This showed that at baseline values, the CMOS system sampling at 5 kHz was able to detect the pulsatile signal. However, when hot water stimulated the blood flow to 15 times the cold water stimulation values, the CMOS camera system was only sensitive to around twice the cold values.

7.3 Future Work

Although the Bayesian Inference error reduction technique has been applied to a rotating diffuser successfully, it has yet to be applied to a CMOS camera imaging perfused tissue. This is because the shapes of the spectra collected from the Optronis system did not match those collected from the Moor Instruments LDLS and therefore the inference could not be performed. Ideally the inference needs to be performed on a system which can output the full bandwidth photocurrent continuously (the LDLS did not do this) whilst flow is modified (for example, by the water stimulation technique). The inference system could then be tested on a sub-sampled version of the output photocurrent and compared to the full bandwidth result.

One drawback of the Bayesian method is that the computational resources required to perform the inference are not insignificant, in that several hundred function evaluations may be required per pixel per output frame. A method of reducing the computational load by utilising a climbing algorithm has been presented, but the number of function evaluations is still high. This number could be further reduced by giving the start position of the climbing algorithm at the last known peak location; this is a guess that the next blood flow value will be similar to the last. Further optimisation is required to minimise resource usage whilst still maintaining accuracy of results.

7.4 Conclusion

Bayesian Inference has been applied in a novel way to facilitate the use of low bandwidth CMOS camera systems to calculate the true first moment. It has been successfully demonstrated in a practical CMOS system imaging a rotating diffuser and been applied theoretically to blood flow data. In addition to this, aliasing CMOS systems have been characterised and a decision can be made on which applications low bandwidth systems may be suitable for.

In the introduction, two clinical applications were discussed. One, post-operative flap monitoring, does not require sensitivity to high flow values (the exception being when locating perforator vessels). The system requires accuracy of flow values between zero and normal baseline. Therefore a relatively low cost off-the-shelf CMOS camera system, processing the photocurrent by, for example, the Hilbert time domain algorithm on an FPGA or graphics card, may be used in this situation to give adequate results. There is even the possibility that such a device could be built into a small handheld unit, increasing usability, although some attention would need to be given to the safety and efficacy of the illumination system.

However in the case of burn imaging, the second application, sensitivity up to fives times baseline may be required in order to distinguish superficial burns from more severe tissue damage. In this case, a low bandwidth CMOS imager may be used, but the results need correction in order to give correct clinical diagnosis. By using Bayes' theorem, the most likely true flow value given an undersampled photocurrent, may be elicited. Further work is required to build such a system and optimise the signal processing to enable accurate real-time laser Doppler blood flow imaging.

7.5 Summary of Findings

- It has been confirmed that the Draijer algorithm is an effective method of calculating the first moment of the photocurrent's power spectrum.
- A novel Hilbert transform based processing method for LDBF is both low in resource usage and high in accuracy.
- Bayesian Inference model selection has been used to determine that the power spectrum of the LDBF photocurrents collected can be modelled more appropriately using a 2-parameter Pareto rather than the usual Exponential distribution.
- Processing methods for LDBF have been compared in full and reduced bandwidth conditions. The low resource IIR filter method is almost as accurate as the gold standard FFT.
- A reduction in bandwidth reduces the sensitivity to high blood flow, but aliasing introduces a 'useful error' that can serve to recover some sensitivity to higher flows.
- The error resulting from undersampling the LDBF photocurrent using an integrating CMOS sensor can be compensated for using Bayesian Inference. A fast implementation of the most likely parameter search has been proposed.
- This undersampling compensation method has been successfully demonstrated on data collected from a CMOS sensor imaging a rotating diffuser.
- An off-the-shelf CMOS imager undersampling the LDBF photocurrent can be used to image blood flow, but its sensitivity to high flow is reduced compared to a full bandwidth single point monitor.

Appendix A

Derivation of Absolute and Average

The following derivation shows why the absolute and average of a LDBF signal is proportional to the square and average. The variance of a signal with samples x_1 to x_M is defined as

$$\sigma^{2} = \frac{\sum_{i=1}^{M} (x_{i} - \mu)^{2}}{M}$$
(A.1)

It can be expanded and simplified to become

$$\sigma^{2} = \frac{\sum_{i=1}^{M} x_{i}^{2}}{M} - \mu^{2}$$
 (A.2)

The signal has Gaussian distribution with zero mean, so taking the square and average calculates the variance:

$$\frac{\sum_{n=0}^{M-1} (x[n])^2}{M} = \sigma^2$$
 (A.3)

Suppose the absolute value is taken of the signal. The distribution is changed from

$$p(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} \quad (-\infty \le x \le \infty)$$
(A.4)

$$p'(x) = \frac{2}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}(\frac{x-\mu}{\sigma})^2} \quad (0 \le x \le \infty)$$
 (A.5)

The mean of the distribution after the absolute process is

$$\mu' = \int_0^\infty x p'(x) dx = \int_0^\infty x \frac{2}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}(\frac{x-\mu}{\sigma})^2} dx \qquad (A.6)$$

Letting $\frac{x-\mu}{\sigma} = X$ and $dX = \frac{1}{\sigma} dx$,

$$\mu' = \int_0^\infty (\sigma X + \mu) \frac{2}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}X^2} \sigma dX$$

$$\mu' = \frac{2\sigma}{\sqrt{2\pi}} \int_0^\infty X e^{-\frac{1}{2}X^2} dX + \frac{2\mu}{\sqrt{2\pi}} \int_0^\infty e^{-\frac{1}{2}X^2} dX$$

As $\mu = 0$, then

$$\mu' = \frac{2\sigma}{\sqrt{2\pi}} \int_0^\infty X e^{-\frac{1}{2}X^2} dX$$

Letting $U = \frac{X}{\sqrt{2}}$, then $dU = \frac{1}{\sqrt{2}}dX$,

$$\mu' = \frac{2\sigma}{\sqrt{2\pi}} \int_0^\infty \sqrt{2}U e^{-U^2} \sqrt{2}dU$$
$$\mu' = \frac{2\sqrt{2}\sigma}{\sqrt{\pi}} \int_0^\infty U e^{-U^2} \sqrt{2}dU$$
$$\mu' = \frac{-\sqrt{2}\sigma}{\sqrt{\pi}} e^{-U^2} |_0^\infty$$
$$\mu' = \frac{\sqrt{2}}{\sqrt{\pi}} \sigma$$
(A.7)

Squaring equation A.7 gives

$$\mu'^2 = \frac{2}{\pi}\sigma^2 \tag{A.8}$$

Equation A.8 shows us that the result of the absolute and average with a final square is proportional to the variance, or square and average, by a factor of $\frac{2}{\pi}$. It is important to note, however, that this is only valid for Gaussian data with zero mean.

223

 to

Appendix B

Reduced Bandwidth Images

The raw data for these images was captured using the Moor Instruments Laser Doppler Line Scanner. The 3 images were processed using a variety of processing algorithms and at a range of bandwidths. The details of these algorithms and the comparison methodology can be found in Chapter 4.



Figure B.1: Image 1 evaluated over 6 bandwidths using FFT.



Figure B.2: Image 1 evaluated over 6 bandwidths using FA8.



Figure B.3: Image 1 evaluated over 6 bandwidths using FA8APP.



Figure B.4: Image 1 evaluated over 6 bandwidths using FA32.



Figure B.5: Image 1 evaluated over 6 bandwidths using FIR S.



Figure B.6: Image 1 evaluated over 6 bandwidths using FIR A.



Figure B.7: Image 1 evaluated over 6 bandwidths using IIR.



Figure B.8: Image 1 evaluated over 6 bandwidths using DRAIJER.



Figure B.9: Image 1 evaluated over 6 bandwidths using HILB F.



Figure B.10: Image 1 evaluated over 6 bandwidths using HILB T.



Figure B.11: Image 1 evaluated over 6 bandwidths using M2.



Figure B.12: Image 2 evaluated over 6 bandwidths using FFT.



Figure B.13: Image 2 evaluated over 6 bandwidths using FA8.



Figure B.14: Image 2 evaluated over 6 bandwidths using FA8APP.



Figure B.15: Image 2 evaluated over 6 bandwidths using FA32.



Figure B.16: Image 2 evaluated over 6 bandwidths using FIR S.



Figure B.17: Image 2 evaluated over 6 bandwidths using FIR A.



Figure B.18: Image 2 evaluated over 6 bandwidths using IIR.



Figure B.19: Image 2 evaluated over 6 bandwidths using DRAIJER.



Figure B.20: Image 2 evaluated over 6 bandwidths using HILB F.



Figure B.21: Image 2 evaluated over 6 bandwidths using HILB T.



Figure B.22: Image 2 evaluated over 6 bandwidths using M2.



Figure B.23: Image 3 evaluated over 6 bandwidths using FFT.



Figure B.24: Image 3 evaluated over 6 bandwidths using FA8.



Figure B.25: Image 3 evaluated over 6 bandwidths using FA8APP.



Figure B.26: Image 3 evaluated over 6 bandwidths using FA32.



Figure B.27: Image 3 evaluated over 6 bandwidths using FIR S.



Figure B.28: Image 3 evaluated over 6 bandwidths using FIR A.



Figure B.29: Image 3 evaluated over 6 bandwidths using IIR.



Figure B.30: Image 3 evaluated over 6 bandwidths using DRAIJER.



Figure B.31: Image 3 evaluated over 6 bandwidths using HILB F.



Figure B.32: Image 3 evaluated over 6 bandwidths using HILB T.



Figure B.33: Image 3 evaluated over 6 bandwidths using M2.

Appendix C

Corrected Flow Values

The following images illustrate the input (real) flow versus the output calculated flow value for simulated photocurrents at a range of bandwidths, where the output is either directly calculated (no correction) or corrected using Bayesian Inference techniques to attempt to move towards gold standard values. These techniques are described fully in Chapter 5. It can be seen that the correction recovers the sensitivity to higher flow values, however the corrected outputs are still relatively noisy. This noise level and the linearity are quantified in Figures 5.26, 5.27 and 5.28.



Figure C.1: Recovered flow values for $f_s = 5$ kHz.



Figure C.2: Recovered flow values for f_s =7.5 kHz.


Figure C.3: Recovered flow values for $f_s = 10$ kHz.

References

- K. Basak, M. Manjunatha, and P. K. Dutta. Review of laser speckle-based analysis in medical imaging. *Medical & Biological Engineering & Computing*, 50(6):547–558, Jun 2012.
- [2] G. Belcaro. Laser Doppler. Med-Orion, London, 1994.
- [3] T. Binzoni, A. Humeau-Heurtier, P. Abraham, and G. Mahe. Blood perfusion values of laser speckle contrast imaging and laser Doppler flowmetry: Is a direct comparison possible? *Biomedical Engineering, IEEE Transactions on*, 60(5):1259–1265, 2013.
- [4] T. Binzoni, T. S. Leung, A. H. Gandjbakhche, D. Ruefenacht, and D. T. Delpy. The use of the henyey-greenstein phase function in monte carlo simulations in biomedical optics. *Physics in Medicine* and Biology, 51(17):N313–N322, Sep 7 2006.
- [5] T. Binzoni, T. S. Leung, A. H. Gandjbakhche, D. Rufenacht, and D. T. Delpy. Comment on 'the use of the henyey - greenstein phase function in monte carlo simulations in biomedical optics'. *Physics in Medicine and Biology*, 51(22):L39–L41, Nov 21 2006.
- [6] T. Binzoni, T. S. Leung, D. Rufenacht, and D. T. Delpy. Absorption and scattering coefficient dependence of laser-Doppler flowmetry models for large tissue volumes. *Physics in Medicine* and Biology, 51(2):311–333, Jan 21 2006.
- [7] T. Binzoni, T. S. Leung, M. L. Seghier, and D. T. Delpy. Translational and Brownian motion in laser-Doppler flowmetry of

large tissue volumes. *Physics in Medicine and Biology*, 49(24):5445–5458, Dec 21 2004.

- [8] T. Binzoni, T. S. Leung, and D. Van De Ville. The photo-electric current in laser-Doppler flowmetry by monte carlo simulations. *Physics in Medicine and Biology*, 54(14):N303–N318, Jul 2009.
- [9] T. Binzoni, C. S. Seelamantula, and D. Van De Ville. A fast time-domain algorithm for the assessment of tissue blood flow in laser-Doppler flowmetry. *Physics in Medicine and Biology*, 55(13):N383–N394, Jul 7 2010.
- [10] T. Binzoni and D. Van de Ville. Full-field laser-Doppler imaging and its physiological significance for tissue blood perfusion. *Physics* in Medicine and Biology, 53(23):6673–6694, Dec 7 2008.
- [11] D. A. Boas and A. K. Dunn. Laser speckle contrast imaging in biomedical optics. *Journal of Biomedical Optics*, 15(1), Jan-Feb 2010.
- [12] R. Bonner and R. Nossal. Model for laser Doppler measurements of blood-flow in tissue. Applied Optics, 20(12):2097–2107, 1981.
- [13] D. I. Booi, I. B. J. G. Debats, W. D. Boeckx, and R. R. W. J. van der Hulst. A study of perfusion of the distal free-tram flap using laser Doppler flowmetry. *Journal of Plastic Reconstructive* and Aesthetic Surgery, 61(3):282–288, 2008.
- [14] G. Larry Bretthorst. Bayesian spectrum analysis and parameter estimation / G. Larry Bretthorst. Lecture notes in statistics; 48. Springer-Verlag, 1988.
- [15] J. D. Briers. Laser Doppler and time-varying speckle: A reconciliation. Journal of the Optical Society of America A-Optics Image Science and Vision, 13(2):345–350, Feb 1996.

- [16] J. D. Briers. Laser Doppler, speckle and related techniques for blood perfusion mapping and imaging. *Physiological Measurement*, 22(4):R35–R66, Nov 2001.
- [17] Alexandra Brugler, Shaun Thompson, Scott Turner, Binh Ngo, and Marc Rendell. Skin blood flow abnormalities in diabetic dermopathy. *Journal of the American Academy of Dermatology*, 65(3):559–563, 2011.
- [18] Haiying Cheng, Qingming Luo, Shaoqun Zeng, Shangbin Chen, Jian Cen, and Hui Gong. Modified laser speckle imaging method with improved spatial resolution. *Journal of Biomedical Optics*, 8(3):559–564, 2003.
- [19] K. H. Chon, S. Dash, and Ju Kihwan. Estimation of respiratory rate from photoplethysmogram data using time-frequency spectral estimation. *Biomedical Engineering*, *IEEE Transactions on*, 56(8):2054–2063, 2009.
- [20] V. Cizek. Discrete hilbert transform. Audio and Electroacoustics, IEEE Transactions on, 18(4):340–343, 1970.
- [21] Leon Cohen. *Time-frequency analysis*. Prentice Hall PTR, Upper Saddle River, N.J, 1995.
- [22] J. W. Cooley and J. W. Tukey. An algorithm for machine calculation of complex Fourier series. *Mathematics of Computation*, 19(90):297–&, 1965.
- [23] J. A. Crowe, J. Carpenter, and K. Hopcraft. An explanation for the effectiveness of the 'Draijer' algorithm for high speed laser Doppler perfusion imaging. *Med Biol Eng Comput*, 50(3):211–4, Mar 2012.
- [24] H. A. Daanen. Finger cold-induced vasodilation: a review. Eur J Appl Physiol, 89(5):411–26, Jun 2003.

- [25] D. De Backer, J. Creteur, J. C. Preiser, M. J. Dubois, and J. L. Vincent. Microvascular blood flow is altered in patients with sepsis. *Am J Respir Crit Care Med*, 166(1):98–104, Jul 1 2002.
- [26] F. F. M. de Mul, M. H. Koelink, M. L. Kok, P. J. Harmsma, J. Greve, R. Graaff, and J. G. Aarnoudse. Laser-Doppler velocimetry and monte-carlo simulations on models for blood perfusion in tissue. *Applied Optics*, 34(28):6595–6611, Oct 1 1995.
- [27] F. F. M. de Mul, J. Vanspijker, D. Vanderplas, J. Greve, J. G. Aarnoudse, and T. M. Smits. Mini laser-Doppler (blood) flow monitor with diode-laser source and detection integrated in the probe. *Applied Optics*, 23(17):2970–2973, 1984.
- [28] Y. Degerli, F. Lavernhe, P. Magnan, and J. A. Farre. Analysis and reduction of signal readout circuitry temporal noise in cmos image sensors for low-light levels. *IEEE Transactions on Electron Devices*, 47(5):949–962, May 2000.
- [29] Arjen M. Dondorp, Piet A. Kager, Johan Vreeken, and Nicholas J. White. Abnormal blood flow and red blood cell deformability in severe malaria. *Parasitology Today*, 16(6):228–232, 2000.
- [30] Klaus Dorschel and Gerhard Muller. Velocity resolved laser Doppler blood flow measurements in skin. Flow Measurement and Instrumentation, 7(3-4):257–264, 2006.
- [31] M. Draijer, E. Hondebrink, T. van Leeuwen, and W. Steenbergen. Time domain algorithm for accelerated determination of the first order moment of photo current fluctuations in high speed laser Doppler perfusion imaging. *Medical & Biological Engineering & Computing*, 47(10):1103–1109, Oct 2009.
- [32] M. Draijer, E. Hondebrink, T. van Leeuwen, and W. Steenbergen. Twente optical perfusion camera: system overview and

performance for video rate laser Doppler perfusion imaging. *Optics* Express, 17(5):3211–3225, Mar 2 2009.

- [33] Matthijs Draijer, Erwin Hondebrink, Ton van Leeuwen, and Wiendelt Steenbergen. Review of laser speckle contrast techniques for visualizing tissue perfusion. *Lasers in Medical Science*, 24(4):639–651, 2009.
- [34] L. E. Drain. The laser Doppler technique. J. Wiley, Chichester, 1980.
- [35] E. J. Droog, W. Steenbergen, and F. Sjoberg. Measurement of depth of burns by laser Doppler perfusion imaging. *Burns*, 27(6):561–568, Sep 2001.
- [36] Shlomo Engelberg. Random signals and noise : a mathematical introduction. CRC Press, Boca Raton ; London, 2007.
- [37] P. Erba, D. Espinoza, N. Koch, T. Christen, A. Serov, and W. Raffoul. Fluxexplorer: a new high-speed laser Doppler imaging system for the assessment of burn injuries. *Skin Res Technol*, Feb 15 2012.
- [38] Lee E. Estes, Lorenzo M. Narducci, and Richard A. Tuft. Scattering of light from a rotating ground glass. J. Opt. Soc. Am., 61(10):1301–1306, 1971.
- [39] I. Fredriksson, M. Larsson, and T. Stromberg. Measurement depth and volume in laser Doppler flowmetry. *Microvascular Research*, 78(1):4–13, Jun 2009.
- [40] I. Fredriksson, M. Larsson, and T. Stromberg. Model-based quantitative laser Doppler flowmetry. *Journal of Biomedical Optics*, 15(5):12, 07/09/2010 2010.

- [41] M. Fries, W. Tang, Y. T. Chang, J. Wang, C. Castillo, and M. H. Weil. Microvascular blood flow during cardiopulmonary resuscitation is predictive of outcome. *Resuscitation*, 71(2):248–53, Nov 2006.
- [42] Hitoshi Fujii, Kunihiko Nohira, Yuhei Yamamoto, Hiroharu Ikawa, and Takehiko Ohura. Evaluation of blood flow by laser speckle image sensing. part 1. Appl. Opt., 26(24):5321–5325, 1987.
- [43] P Gebuhr, JP Jorgensen, B Vollmer-Larsen, SL Nielsen, and B Alsbjorn. Estimation of amputation level with a laser Doppler flowmeter. Journal of Bone & Joint Surgery, British Volume, 71-B(3):514-517, May 1, 1989 1989.
- [44] J.W. Goodman. Speckle Phenomena in Optics, Theory and Applications. Roberts and Company, Englewood, Colorado, 2006.
- [45] V. A. Grechikhin and B. S. Rinkevichius. Hilbert transform for processing of laser Doppler vibrometer signals. Optics and Lasers in Engineering, 30(2):151–161, 1998.
- [46] P. C. Gregory. Bayesian logical data analysis for the physical sciences: a comparative approach with Mathematica support. Cambridge University Press, Cambridge, 2005.
- [47] Q. Gu. An Analogue Integrated Circuit Design for Laser Doppler Blood Flow Measurement. PhD thesis, University of Nottingham, Nottingham, 2007.
- [48] Q. Gu, B. R. Hayes-Gill, and S. P. Morgan. Laser Doppler blood flow complementary metal oxide semiconductor imaging sensor with analog on-chip processing. *Applied Optics*, 47(12):2061–2069, Apr 20 2008.
- [49] E. Haus. Chronobiology in oncology. Int J Radiat Oncol Biol Phys, 73(1):3–5, Jan 1 2009.

- [50] D. He. Full Field Laser Doppler Blood Flow Sensor. PhD thesis, University of Nottingham, Nottingham, 2008.
- [51] D. He, H. C. Nguyen, B. R. Hayes-Gill, Y. Zhu, J. A. Crowe, G. F. Clough, C. A. Gill, and S. P. Morgan. 6464 pixel smart sensor array for laser Doppler blood flow imaging. *Opt. Lett.*, 37(15):3060–3062, 2012.
- [52] Diwei He, Chayut Kongsavatsak, Barrie R. Hayes-Gill, John A. Crowe, and Stephen P. Morgan. 32 x 32 pixel array complementary metal-oxide semiconductor imaging sensor for laser Doppler bloodflow measurement. *Optical Engineering*, 50(5):054403–11, 2011.
- [53] E. Higurashi, R. Sawada, and T. Ito. An integrated laser blood flowmeter. *Journal of Lightwave Technology*, 21(3):591–595, Mar 2003.
- [54] F. Holzle, D. J. Loeffelbein, D. Nolte, and K. D. Wolff. Free flap monitoring using simultaneous non-invasive laser Doppler flowmetry and tissue spectrophotometry. *Journal of Cranio-Maxillofacial Surgery*, 34(1):25–33, Jan 2006.
- [55] N. E. Huang, Z. Shen, S. R. Long, M. L. C. Wu, H. H. Shih, Q. N. Zheng, N. C. Yen, C. C. Tung, and H. H. Liu. The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London Series a-Mathematical Physical and Engineering Sciences*, 454(1971):903–995, Mar 8 1998.
- [56] A. Humeau, W. Steenbergen, H. Nilsson, and T. Stromberg. Laser Doppler perfusion monitoring and imaging: novel approaches. *Medical & Biological Engineering & Computing*, 45(5):421–435, May 2007.

- [57] Harold Jeffreys. An invariant form for the prior probability in estimation problems. Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, 186(1007):453-461, September 24, 1946 1946.
- [58] Gwilym M. Jenkins and Donald G. Watts. Spectral analysis and its applications. Holden-Day series in time series analysis. Holden-Day, San Francisco ; London, 1968.
- [59] H. W. Jentink, F. F. M. de Mul, R. G. A. M. Hermsen, R. Graaff, and J. Greve. Monte-carlo simulations of laser Doppler blood-flow measurements in tissue. *Applied Optics*, 29(16):2371–2381, Jun 1 1990.
- [60] Price N. B. Judy He. Hemoglobin level and red blood cell count findings in normal women. JAMA: The Journal of the American Medical Association, 167(5):563–566, 1958.
- [61] S. W. Kim, S. C. Kim, K. C. Nam, E. S. Kang, J. J. Im, and D. W. Kim. A new method of screening for diabetic neuropathy using laser Doppler and photoplethysmography. *Medical & Biological Engineering & Computing*, 46(1):61–67, Jan 2008.
- [62] Donald E Knuth. Big omicron and big omega and big theta. ACM Sigact News, 8(2):18–24, 1976.
- [63] Martin Johannes Koehler, Tanja Vogel, Peter Elsner, Karsten König, Rainer Bückle, and Martin Kaatz. In vivo measurement of the human epidermal thickness in different localizations by multiphoton laser tomography. *Skin Research and Technology*, 16(3):259–264, 2010.
- [64] M. H. Koelink, F. F. M. de Mul, B. Leerkotte, J. Greve, H. W. Jentink, R. Graaff, A. C. M. Dassel, and J. G. Aarnoudse. Signal-

processing for a laser-Doppler blood perfusion meter. *Signal Processing*, 38(2):239–252, Jul 1994.

- [65] C. Kongsavatsak, D. W. He, B. R. Hayes-Gill, J. A. Crowe, and S. P. Morgan. Complementary metal-oxide-semiconductor imaging array with laser Doppler blood flow processing. *Optical Engineering*, 47(10):-, Oct 2008.
- [66] P. Kvandal, S. A. Landsverk, A. Bernjak, A. Stefanovska, H. D. Kvernmo, and K. A. Kirkeboen. Low-frequency oscillations of the laser Doppler perfusion signal in human skin. *Microvascular Research*, 72(3):120–127, Nov 2006.
- [67] M. Larsson, W. Steenbergen, and T. Stromberg. Influence of optical properties and fiber separation on laser Doppler flowmetry. *Journal of Biomedical Optics*, 7(2):236–243, Apr 2002.
- [68] Marcus Larsson and Tomas Stromberg. Toward a velocity-resolved microvascular blood flow measure by decomposition of the laser Doppler spectrum. *Journal of Biomedical Optics*, 11(1):014024–9, 2006.
- [69] M. Leahy, editor. *Microcirculation Imaging*. Wiley-VCH, Weinheim, 1st edition, 2012.
- [70] M. Leahy, A. Liebert, and R. Maniewski. Evaluation of different signal processing algorithms in laser-Doppler perfusion measurements. Opto-Ireland 2002: Optics and Photonics Technologies and Applications, Pts 1 and 2, 4876:120–127, 2003.
- [71] Chungkeun Lee, Hang Sik Shin, and Myoungho Lee. Relations between ac-dc components and optical path length in photoplethysmography. *Journal of Biomedical Optics*, 16(7):077012–077012, 2011.

- [72] Craig Leising. Paper surface roughness with 3D profilometry http://www.nanovea.com/application notes/paperroughness.pdf, accessed 20/11/2012 2010.
- [73] Marcel Leutenegger, Erica Martin-Williams, Pascal Harbi, Tyler Thacher, Wassim Raffoul, Marc Andr, Antonio Lopez, Philippe Lasser, and Theo Lasser. Real-time full field laser Doppler imaging. *Biomed. Opt. Express*, 2(6):1470–1477, 2011.
- [74] L. G. Lindberg and P. A. Oberg. Optical-properties of blood in motion. Optical Engineering, 32(2):253–257, Feb 1993.
- [75] Robert Lohwasser and Gerald Soelkner. Experimental and theoretical laser-Doppler frequency spectra of a tissuelike model of a human head with capillaries. *Appl. Opt.*, 38(10):2128–2137, 1999.
- [76] P. J. Loughlin and K. L. Davidson. Modified Cohen-Lee time-frequency distributions and instantaneous bandwidth of multicomponent signals. *IEEE Transactions on Signal Processing*, 49(6):1153–1165, Jun 2001.
- [77] Leonard Mandel and Emil Wolf. Optical coherence and quantum optics. Cambridge University Press, Cambridge, 1995.
- [78] S. M. Monstrey, H. Hoeksema, R. D. Baker, J. Jeng, R. S. Spence, D. Wilson, and S. A. Pape. Burn wound healing time assessed by laser Doppler imaging. part 2: validation of a dedicated colour code for image interpretation. *Burns*, 37(2):249–56, Mar 2011.
- [79] Stan Monstrey, Henk Hoeksema, Jos Verbelen, Ali Pirayesh, and Phillip Blondeel. Assessment of burn depth and burn wound healing potential. *Burns*, 34(6):761–769, 2008.

- [80] MoorInstruments. Laser Doppler Imager 2 Brochure Issue 2
 http://gb.moor.co.uk/product/moorldi2-laser-Doppler-imager/8, accessed 19/06/2012 2012.
- [81] MoorInstruments. Laser Doppler Line Scanner 2 Brochure Issue
 1 http://gb.moor.co.uk/product/moorldls2-ld-line-imager/10, accessed 19/06/2012 2012.
- [82] S. P. Morgan, H. C. Nguyen, B. R. Hayes-Gill, Y. Q. Zhu, J. A. Crowe, and D. W. He. Low resource processing algorithms for laser Doppler blood flow imaging. *Medical Engineering & Physics*, 33(6):720–729, Jul 2011.
- [83] G. S. Mudholkar and E. O. George. Remark on shape of logistic distribution. *Biometrika*, 65(3):667–668, 1978.
- [84] H. C. Nguyen. High Speed Processing For Laser Doppler Blood Flow Imaging. PhD thesis, University of Nottingham, Nottingham, 2009.
- [85] H. C. Nguyen, B. R. Hayes-Gill, S. P. Morgan, Y. Zhu, D. Boggett, X. Huang, and M. Potter. A field-programmable gate array based system for high frame rate laser Doppler blood flow imaging. J Med Eng Technol, 35(5-6):306–15, Jul 2010.
- [86] H. C. Nguyen, B. R. Hayes-Gill, Y. Q. Zhu, S. P. Morgan, and IEEE. A high frame rate and high accuracy implementation using an fpga for calculating laser Doppler blood flow. In 2008 IEEE Instrumentation and Measurement Technology Conference, Vols 1-5, IEEE Instrumentation & Measurement Technology Conference, Proceedings, pages 212–217. IEEE, New York, 2008.
- [87] NICE. New nice guidance on scanner to help improve burns treatment -

http://www.nice.org.uk/newsroom/pressreleases/scannerimproveburnstreatment.jsp 2011.

- [88] G. E. Nilsson, E. G. Salerud, T. Strömberg, and K. Wärdell. Laser Doppler perfusion monitoring and imaging. In Tuan Vo-Dinh, editor, *Biomedical photonics handbook*. CRC Press, Boca Raton, Fla.; London, 2003.
- [89] A. Obeid. In vitro comparison of different signal processing algorithms used in laser Doppler flowmetry. *Medical and Biological Engineering and Computing*, 31(1):43–52, 1993.
- [90] A. N. Obeid, N. J. Barnett, G. Dougherty, and G. Ward. A critical review of laser Doppler flowmetry. J Med Eng Technol, 14(5):178–81, Sep-Oct 1990.
- [91] P. A. Oberg. Laser-Doppler flowmetry. Critical Reviews in Biomedical Engineering, 18(2):125–163, 1990.
- [92] J. O'Doherty, P. McNamara, N. T. Clancy, J. G. Enfield, and M. J. Leahy. Comparison of instruments for investigation of microcirculatory blood flow and red blood cell concentration. *Journal of Biomedical Optics*, 14(3), May-Jun 2009.
- [93] Ricardo Oliveira, Sonia Semedo, Edite Figueiras, Luis F. Requicha Ferreira, and Anne Humeau. Laser Doppler flowmeters for microcirculation measurements. In *Bioengineering (ENBENG)*, 2011. ENBENG 2011. 1st Portuguese Meeting in, pages 1–4, 1-4 March 2011 2011.
- [94] J. T. Pohlmann and D. W. Leitner. A comparison of ordinary least squares and logistic regression. *Ohio Journal of Science*, 103(5):118–125, Dec 2003.
- [95] Elena Punskaya. Fast Fourier transform (fft), 2010.

- [96] Natalie Purcell and Philip J. Beeby. The influence of skin temperature and skin perfusion on the cephalocaudal progression of jaundice in newborns. *Journal of Paediatrics and Child Health*, 45(10):582–586, 2009.
- [97] V. Rajan, B. Varghese, T. G. van Leeuwen, and W. Steenbergen. Review of methodological developments in laser Doppler flowmetry. *Lasers in Medical Science*, 24(2):269–283, Mar 2009.
- [98] D. Rallan and C. C. Harland. Skin imaging: is it clinically useful? Clinical and Experimental Dermatology, 29(5):453–459, Sep 2004.
- [99] M. S. Rendell, M. F. Finnegan, J. C. Healy, A. Lind, B. K. Milliken, D. E. Finney, and R. F. Bonner. The relationship of laser-Doppler skin blood flow measurements to the cutaneous microvascular anatomy. *Microvasc Res*, 55(1):3–13, Jan 1998.
- [100] M. S. Rendell, S. S. Green, A. Catania, J. Oliveto, J. Wells, E. J. Banset, and H. Wang. Post-exercise cutaneous hyperaemia resulting from local exercise of an extremity. *Clinical Physiology*, 17(3):213–224, May 1997.
- [101] C. Riva, G. B. Benedek, and B. Ross. Laser Doppler measurements of blood-flow in capillary tubes and retinal arteries. *Investigative Ophthalmology*, 11(11):936–&, 1972.
- [102] G. G. Romero, E. E. Alanis, and H. J. Rabal. Statistics of the dynamic speckle produced by a rotating diffuser and its application to the assessment of paint drying. *Optical Engineering*, 39(6):1652– 1658, Jun 2000.
- [103] Yasser Sakr, Marc-Jacques Dubois, Daniel de Backer, Jacques Creteur, and Jean-Louis Vincent. Persistent microcirculatory alterations are associated with organ failure and death in patients

with septic shock*. Critical Care Medicine, 32(9):1825–1831 10.1097/01.CCM.0000138558.16257.3F, 2004.

- [104] L. E. Scheving, L. A. Scheving, R. J. Feuers, T. H. Tsai, and F. O. Cope. Chronobiology as it relates to toxicology, pharmacology, and chemotherapy. *Regul Toxicol Pharmacol*, 17(2 Pt 1):209–18, Apr 1993.
- [105] S. Schlosser, R. Wirth, J. A. Plock, A. Serov, A. Banic, and D. Erni. Application of a new laser Doppler imaging system in planning and monitoring of surgical flaps. *Journal of Biomedical Optics*, 15(3):-, May-Jun 2010.
- [106] A. Serov and T. Lasser. High-speed laser Doppler perfusion imaging using an integrating cmos image sensor. Optics Express, 13(17):6416-6428, Aug 22 2005.
- [107] A. Serov and T. Lasser. Combined laser Doppler and laser speckle imaging for real-time blood blow measurements - art. no. 609406. *Optical Diagnostics and Sensing VI*, 6094:9406–9406 155, 2006.
- [108] A. Serov and T. Lasser. Full-field high-speed laser Doppler imaging system for blood-flow measurements - art. no. 608004. Advanced Biomedical and Clinical Diagnostic Systems IV, 6080:8004–8004 230, 2006.
- [109] A. Serov, W. Steenbergen, and F. de Mul. Prediction of the photodetector signal generated by Doppler-induced speckle fluctuations: theory and some validations. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 18(3):622– 630, Mar 2001.
- [110] A. Serov, W. Steenbergen, and F. de Mul. Laser Doppler perfusion imaging with a complimentary metal oxide semiconductor image sensor. *Optics Letters*, 27(5):300–302, Mar 1 2002.

- [111] A. Serov, B. Steinacher, and T. Lasser. Full-field laser Doppler perfusion imaging and monitoring with an intelligent cmos camera. *Optics Express*, 13(10):3681–3689, May 16 2005.
- [112] A. N. Serov, J. Nieland, S. Oosterbaan, F. F. M. de Mul, H. van Kranenburg, H. H. P. T. Bekman, and W. Steenbergen. Integrated optoelectronic probe including a vertical cavity surface emitting laser for laser Doppler perfusion monitoring. *IEEE Transactions* on Biomedical Engineering, 53(10):2067–2074, Oct 2006.
- [113] R. Seto, R. Sawada, E. Higurashi, T. Itoh, K. Suzuki, K. Tsukamoto, T. Ikehara, R. Maeda, T. Masuda, and T. Fujitsu. Development and use of a micro-optical blood flow sensor based on system in package (sip) technology that fuses optical mems and integrated circuit to detect avian influenza. *Transducers '07 & Eurosensors Xxi, Digest of Technical Papers, Vols 1 and 2*, pages U986–U987 2616, 2007.
- [114] A.P. Shepherd and P.A. Öberg. Laser-Doppler Blood Flowmetry. Kluwer Academic Publishers, 1990.
- [115] M. R. Srinivasan. Physics for engineers. Anshan, Tunbridge Wells, 2011.
- [116] A. Stefanovska, M. Bracic, and H. D. Kvernmo. Wavelet analysis of oscillations in the peripheral blood circulation measured by laser Doppler technique. *IEEE Transactions on Biomedical Engineering*, 46(10):1230–1239, Oct 1999.
- [117] M. D. Stern. In vivo evaluation of microcirculation by coherent light scattering. *Nature*, pages 56–58, 1975.
- [118] Tara Lynn Stewart, Brandon Ball, Paul J. Schembri, Keijiro Hori, Jie Ding, Heather A. Shankowsky, Edward E. Tredget, and Wound Healing Research Group. The use of laser Doppler imaging

as a predictor of burn depth and hypertrophic scar postburn injury. *Journal of Burn Care & Research*, Publish Ahead of Print:10.1097/BCR.0b013e318257db36, 2012.

- [119] Y. Sugii, S. Nishio, and K. Okamoto. In vivo piv measurement of red blood cell velocity field in microvessels considering mesentery motion. *Physiological Measurement*, 23(2):403–416, May 2002.
- [120] H. E. Suichies, J. G. Aarnoudse, A. Okken, H. W. Jentink, F. F. M. de Mul, and J. Greve. Forehead skin blood-flow in normal neonates during active and quiet sleep, measured with a diode-laser Doppler instrument. Acta Paediatrica Scandinavica, 77(2):220–225, Mar 1988.
- [121] H. E. Suichies, C. Brouwer, J. G. Aarnoudse, H. W. Jentink, F. F. M. de Mul, and J. Greve. Skin blood-flow changes, measured by laser Doppler flowmetry, in the 1st week after birth. *Early Human Development*, 23(1):1–8, Jun 1990.
- [122] S. Sun. Laser Doppler imaging and laser speckle contrast imaging for blood flow measurement (In Preparation). PhD thesis, University of Nottingham, Nottingham, 2013.
- [123] Garry A. Tew, Markos Klonizakis, Helen Crank, J. David Briers, and Gary J. Hodges. Comparison of laser speckle contrast imaging with laser Doppler for assessing microvascular function. *Microvascular Research*, 2011.
- [124] O. Thompson, J. Bakker, C. Kloeze, E. Hondebrink, and W. Steenbergen. Experimental comparison of perfusion imaging systems using multi-exposure laser speckle, single exposure laser speckle and full-field laser Doppler. *Dynamics and Fluctuations* in Biomedical Photonics Ix, 8222, 2012.

- [125] O. B. Thompson and M. K. Andrews. Tissue perfusion measurements: multiple-exposure laser speckle analysis generates laser Doppler-like spectra. *Journal of Biomedical Optics*, 15(2), Mar-Apr 2010.
- [126] J. Umetani, C. Yuzawa, E. Okada, E. Sekizuka, C. Oshio, and H. Minamitani. Microvascular red-blood-cell velocity-measurement using the image gradient-method. *Images of the Twenty-First Century, Pts 1-6*, 11:624–625, 1989.
- [127] H. van Herpt, M. Draijer, E. Hondebrink, M. Nieuwenhuis, G. Beerthuizen, T. van Leeuwen, and W. Steenbergen. Burn imaging with a whole field laser Doppler perfusion imager based on a cmos imaging array. *Burns*, 36(3):389–396, May 2010.
- [128] Pierre Vernhes, Jean-Francis Bloch, Christophe Mercier, Anne Blayo, and Bernard Pineaux. Statistical analysis of paper surface microstructure: A multi-scale approach. *Applied Surface Science*, 254(22):7431–7437, 2008.
- [129] K. Wärdell, A. Jakobsson, and G. E. Nilsson. Laser Doppler perfusion imaging by dynamic light scattering. *Biomedical Engineering, IEEE Transactions on*, 40(4):309–316, 1993.
- [130] S. Wojtkiewicz, A. Liebert, H. Rix, and R. Maniewski. Evaluation of algorithms for microperfusion assessment by fast simulations of laser Doppler power spectral density. *Physics in Medicine and Biology*, 56(24):7709, 2011.
- [131] S Wojtkiewicz, E Wojcik-Sosnowska, M Jasik, R Maniewski, W Karnafel, and A Liebert. Assessment of speed distribution of red blood cells in the microvascular network in healthy volunteers and type 1 diabetes using laser Doppler spectra decomposition. *Physiological Measurement*, 35(2):283, 2014.

[132] S. Yuan, A. Devor, D. A. Boas, and A. K. Dunn. Determination of optimal exposure time for imaging of blood flow changes with laser speckle contrast imaging. *Applied Optics*, 44(10):1823–1830, Apr 1 2005.