Development of a Portable Time-Domain System for Diffuse Optical Tomography of the Newborn Infant Brain

Pablo Pérez Tirador

A dissertation submitted in partial fulfillment of the requirements for the degree of **Doctor of Philosophy**

of

University College London.

Department of Medical Physics & Biomedical Engineering University College London

February 19, 2021

I, Pablo Pérez Tirador, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Conditions such as hypoxic-ischaemic encephalopathy (HIE) and perinatal arterial ischaemic stroke (PAIS) are causes of lifelong neurodisability in a few hundred infants born in the UK each year. Early diagnosis and treatment are key, but no effective bedside detection and monitoring technology is available. Non-invasive, near-infrared techniques have been explored for several decades, but progress has been inhibited by the lack of a portable technology, and intensity measurements, which are strongly sensitive to uncertain and variable coupling of light sources and detector to the scalp.

A technique known as time domain diffuse optical tomography (TD-DOT) uses measurements of photon flight times between sources and detectors placed on the scalp. Mean flight time is largely insensitive to the coupling and variation in mean flight time can reveal spatial variation in blood volume and oxygenation in regions of brain sampled by the measurements. While the cost, size and high power consumption of such technology have hitherto prevented development of a portable imaging system, recent advances in silicon technology are enabling portable and low-power TD-DOT devices to be built.

A prototype TD-DOT system is proposed and demonstrated, with the longterm aim to design a portable system based on independent modules, each supporting a time-of-flight detector and a pulsed source. The operation is demonstrated of components that can be integrated in a portable system: silicon photodetectors, integrated circuit-based signal conditioning and time detection – built using a combination of off-the-shelf components and reconfigurable hardware, standard computer interfaces, and data acquisition and calibration software. The only external

Abstract

elements are a PC and a pulsed laser source. This thesis describes the design process, and results are reported on the performance of a 2-channel system with online histogram generation, used for phantom imaging. Possible future development of the hardware is also discussed.

Impact Statement

UCL is one of the pioneering institutions in the field of medical optics, having contributed major advances in near-infrared spectroscopy and imaging technologies for non-invasive monitoring of brain development and injury in neonates and older infants. For example, UCL has been a leader in the development of diffuse optical tomography (DOT) systems and methods for generating images of blood volume and oxygenation in the infant brain. These have been applied to the assessment of infants with hypoxic-ischaemic encephalopathy (HIE) and perinatal arterial ischaemic stroke (PAIS), conditions that affect the distribution of cerebral oxygen around the time of birth, and are severely underdiagnosed.

Time domain diffuse optical tomography (TD-DOT), which involves measuring the flight-times of near-infrared photons through tissue, enables quantitative maps of tissue optical properties to be produced, potentially enabling clinicians to understand these conditions better and propose methods for monitoring and treatment. Furthermore, time-domain datatypes (such as mean flight-time) are largely immune to the uncertainty and variability in surface coupling (e.g. due to hair) that contaminates simple intensity measurements. A new generation of silicon-based TD-DOT devices is currently being developed to achieve more portability, affordability, energy efficiency and replicability of the designs. This project contributes to this effort by proposing and documenting two prototypes of time domain systems, capable of measuring the times of flight of photons with a high temporal resolution (of the order of picoseconds) exploiting state-of-the-art technology and hardware design.

In this thesis, I explore the requirements and technical characteristics of a

Abstract

portable TD-DOT technology, and present the design methodology for a complete electronic signal conditioning chain that detects individual photon events, obtains their time of flight from the laser source, and builds time-of-flight histograms that can be used for DOT image reconstruction. The focus of my design methodology has been the use of inexpensive, commercially available components, programmable hardware, and easily portable data acquisition software to produce a modular design that does not depend on custom silicon. A thorough documentation ensures that the design can be easily replicated or adapted. The complete system has been used to obtain topographical maps of tissue equivalent phantoms, and is readily adapted as a tool for scanning human subjects. The final desk-top prototype presented and evaluated in the thesis can be further integrated into a more portable form factor suitable for use in clinical environments, such as the cot-side in a hospital or in an ambulance. I also discuss how the system can be adapted into a distributed system form factor, consisting of multiple probes that house individual sources, detectors and electronics for signal processing. This constitutes a step towards a wearable system, a long term objective of this research. A wearable system will allow TD-DOT to be used in more varied environments and reach new patients, likely transforming the use of optical methods for neuroimaging.

Acknowledgements

This project has been made possible with the funding provided by the Engineering and Physical Sciences Research Council to UCL's Centre for Doctoral Training in Medical Imaging. I would also like to thank the CDT's staff for their support throughout my PhD.

I would like to thank my primary supervisor Professor Jem Hebden for his time and dedication and for his invaluable help with the organisation of the project. I extend my thanks to my second supervisors Dr Samuel Powell, for his technical and academic input, and Dr Danial Chitnis for his direction and technical expertise during the early stages of the project.

I would like to give recognition to Dr Konstantinos Papadimitriou for his collaboration and perspective during a very challenging period of the project and to Dr Dimitrios Airantzis for his personal advice.

I am grateful to the people I met during my time at UCL – all the people at Gowerlabs, with whom I shared space and equipment for a long time, and the colleagues at BORL who were there during the good and bad times (theirs and mine). Thanks also to all the people I met while living at Goodenough College. They were a true second family.

Finally, this could not have been possible without the support of my parents and my sister, who have lived through my PhD by my side and helped me keep my confidence during this challenge.

List of Acronyms

- ADC Analog-to-Digital Converter. 78
- AOTF Acousto-Optic Tuneable Filter. 92
- ASIC Application Specific Integrated Circuit. 42

BRAM Block Random Access Memory. 81

- CFD Constant Fraction Discriminator. 83
- CW Continuous Wave. 46
- DAC Digital-to-Analog Converter. 103
- **DNL** Differential Non-Linearity. 78
- **DOI** Diffuse Optical Imaging. 65
- **DOT** Diffuse Optical Tomography. 65
- **DPF** Differential Path-length Function. 68
- EEG Electroencephalogram. 38
- FPGA Field Programmable Gate Array. 80
- FWHM Full Width at Half Maximum. 93
- GUI Graphical User Interface. 81

- HDL Hardware Description Language. 81
- HIE Hypoxic Ischaemic Encephalopathy. 29
- IC Integrated Circuit. 133
- **IDE** Integrated Development Environment. 95
- INL Integral Non-Linearity. 79
- **IRF** Impulse Response Function. 85
- laser Light Amplification by Stimulated Emission of Radiation. 73
- LC Logic Cell. 80
- LED Light Emitting Diode. 48
- LSB Least Significant Bit. 78
- LV-PECL Low Voltage, Positive Emitter Coupled Logic. 99
- **MONSTIR** Multichannel Optoelectronic Near-infrared System for Time-domain Image Reconstruction. 40
- mutex Mutual Exclusion (lock). 164
- ND Neutral Density. 94
- NECL Negative Emitter Coupled Logic. 102
- NIM Nuclear Instrumentation Module. 91
- NIR Near Infrared. 64
- NIRS Near Infrared Spectroscopy. 30
- **OD** Optical Density. 94
- p-p Peak-to-peak. 97

- PAIS Perinatal Arterial Ischaemic Stroke. 29
- PCB Printed Circuit Board. 93
- PLL Phase Locked Loop. 75
- PMDF Photon Measurement Density Function. 67
- **PMOD** Peripheral Module. 137
- **PMT** Photomultiplier Tube. 69
- **RF** Radio Frequency. 95
- **RMS** Root Mean Square. 54
- SiPM Silicon Photo Multiplier. 72
- SMA Sub-Miniature connector version A. 93
- **SNR** Signal to Noise Ratio. 147
- SPAD Single Photon Array Detector. 71
- SPI Serial Peripheral Interface. 95
- TCSPC Time Correlated Single Photon Counting. 95
- TD-DOT Time Domain Diffuse Optical Topography. 67
- TDC Time-to-Digital Converter. 51
- **TPSF** Temporal Point Spread Function. 68
- TTL Transistor-Transistor Logic. 93
- UART Universal Asynchronous Receiver-Transmitter. 112
- VCSEL Vertical Cavity Surface-Emitting Laser. 74
- VHDL Very high speed integrated circuit Hardware Description Language. 81
- XOR Exclusive Or. 112

Contents

Li	List of Acronyms		11
1	Intr	duction	29
	1.1	Aims and objectives	29
	1.2	Organisation of the thesis	32
2	Proj	ect Background	35
	2.1	Diffuse optical imaging systems for HIE and PAIS diagnosis	35
		2.1.1 Commercial TD-DOT systems	39
		2.1.2 TD-DOT systems and progress towards miniaturisation	39
		2.1.3 Integrated systems	43
		2.1.4 Other miniaturised systems	46
	2.2	Time to digital converters	51
	2.3	Conclusions	57
3	The	retical Background	63
	3.1	Fundamentals of biomedical optics	63
	3.2	Diffuse optical imaging and time domain	65
	3.3	Detectors	68
	3.4	Light sources – laser pulse generation	72
	3.5	Time to digital converters	75
	3.6	Digital converter characteristics and error correction	78
	3.7	FPGA and hardware design	80
	3.8	Thresholding and fraction discrimination	83

Contents

	3.9	Time b	bases and calibration	. 85
	3.10	High f	requency design	. 87
4	Desi	gn of O	Pptical and Electronic Signal Conditioning	91
	4.1	Comp	onents and prototyping platforms	91
		4.1.1	Light source	91
		4.1.2	Light detectors	. 92
		4.1.3	Event detection circuits	. 94
		4.1.4	Development boards	. 95
	4.2	Design	n of the signal conditioning system	. 96
		4.2.1	Characteristics of the input signals and design requirements	97
		4.2.2	Design iterations and refinement	. 98
	4.3	Analy	sis of system performance	. 104
		4.3.1	Pulse shape and optimisation	. 104
		4.3.2	Stability and jitter	105
5	Deve	lopme	nt of an FPGA-based Time-to-Digital Converter	109
	5.1	Initial	direct-to-histogram design	. 109
	5.2	Non-li	nearity characterisation and correction	. 113
		5.2.1	Simulation of the delay line and code bubbles	. 117
		5.2.2	Estimation of the DNL	. 117
	5.3	Initial	results	. 118
	5.4	Limita	ations of the design and reprogramming	. 123
	5.5	Tests ı	using the signal conditioning system	. 127
	5.6	Conclu	usions	. 129
6	Deve	lopme	nt of a Time-of-Flight System Based on a Commercial Time	è-
			^N onvortor	133
	to-D	igital C		155
	to-D 6.1	Charae	cteristics of the Texas Instruments TDC7201	133
	to-D 6.1 6.2	Charad Systen	cteristics of the Texas Instruments TDC7201	133 133 135
	to-D 6.1 6.2 6.3	Charae Charae Systen One-cl	cteristics of the Texas Instruments TDC7201	133 133 135 137

Contents

		6.3.2	Characterisation tests using the custom conditioning system	142
		6.3.3	Single channel system – evaluation on tissue-like phantoms	151
		6.3.4	Single channel system – tests on human volunteers	163
	6.4	Two-cl	hannel system	164
		6.4.1	Design differences for the two channel system	164
		6.4.2	Tests using the custom conditioning system	170
		6.4.3	Two channel system – test on phantoms	173
	6.5	Conclu	isions	178
7	Con	clusions	s and Future Work	183
	7.1	Conclu	isions	183
	7.2	Wearal	bility and power distribution	185
	7.3	Compa	arison to other work	187
	7.4	Future	work	188
Ap	pend	ices		192
A	Circ	uits and	d Schematics	193
B	Cod	e		203
	B .1	Simula	ation of the direct-to-histogram TDC	203
	B.2	Simula	ation of the TDC7201 system's count rate	205
С	Data	asheets		209
D	Soft	ware To	ools	211
Bi	bliogı	aphy		212

2.1	MONSTIR II and its individual components, taken from Cooper et	
	<i>al</i> [63]	41
2.2	A TD-NIRS system designed by Buttafava et al [93]	43
2.3	Photomicrograph of the line sensor from Erdogan et al [107]. Blue	
	and red SPADs for five pixels are zoomed in	45
2.4	The multimodal NIRx Nirsport 2, shown with the cap and the	
	portable central unit. Taken from NIRx's website [134]	50
2.5	The LUMO NIRS system from Gowerlabs Ltd, with cap and central	
	unit. Taken from Gowerlabs's website [135]	50
3.1	Absorption spectra in the optical wavelength range for some tissue	
	chromophores [174]	64
3.2	PMDFs for a continuous wave beam for different source-detector	
	separations [185]. The white lines indicate the boundaries between	
	the scalp and skull, skull and CSF, and CSF and grey matter, from	
	top to bottom	67
3.3	Timing diagram of a Vernier oscillator TDC, adapted from [155].	
	The events are represented by the low to high transition in the level	
	of the signal. The red lines mark the moments the status of the delay	
	line is sampled.	76
3.4	Timing diagram of a 4-element delay line. The events are repre-	
	sented by the low to high transition in the level of the signal. \ldots	77
3.5	Workflow for hardware design in FPGAs (adapted from [201])	82

3.6	Time walk in signal thresholding (left) and constant fraction thresh-
	olding (right)
3.7	Diagram of the base structure of a CFD, showing the signal shape
	on each branch
3.8	Diagrams of types of transmission types, taken from Ulaby and Ra-
	vaioli [210]
4.1	Boxes used for housing and shielding the components of the imag-
	ing system
4.2	Pulse shaping and amplifier circuits used for the first stage of the
	system development
4.3	Circuit of a constant fraction discriminator based on the design of
	Hansang Lim [222]
4.4	A monostable multivibrator design using two comparators, repro-
	duced from [223]
4.5	Circuit including amplifiers used with the custom comparator board. 104
4.6	Mean time of the IRF from the signal conditioning system plotted
	against the period of stabilisation time. The coloured bands repre-
	sent a range of ± 10 ps (blue) and ± 20 ps (red) over the mean time
	after 200 min
4.7	Mean time of the IRF from the signal conditioning system plotted
	against the period of stabilisation time after the system is stable.
	The coloured bands represent a range of ± 10 ps (blue) and ± 20 ps
	(red) over the average mean time
5.1	Block diagram and truth table of an adder with carry. If the input
	is fixed to one of the highlighted combinations, the output carry
	reproduces the input carry
5.2	Schematic of a delay line with two layers of registers
5.3	A toggling register. The output's transitions are caused by the rising
	edges of the input (marked in red)

5.4	Timing diagram comparison of an event detection in the delay line with thermometer code and applying the XOR function between
	registers. The red signals represent the state of the registers
5.5	Histograms of the variation of the frequency of a periodic signal passing through the FPGA's I/O banks
5.6	Input-to-output delay measured for a 40-element delay line (upper) and variation of the delay between delay elements (lower) 116
5.7	Timing diagram showing the effect of <i>code bubbles</i> on a histogram. 118
5.8	Histograms generated by the simulation code running with a mono- tonically increasing delay line and with code bubbles
5.9	Plot of a noise histogram, DNL and correction factors for the direct- to-histogram TDC
5.10	Histogram obtained with a photodiode, as collected (blue) and after applying the DNL correction factor (red). The mean is indicated by a vertical line
5.11	Position of the histogram peak and mean of the histogram collected with a photodiode for different free space source-detector separations.121
5.12	Comparison of the histogram response of the SPAD at high intensity levels in the SPC-130 card and the FPGA. Due to the saturation, the timing gets distorted, narrowing the main peak and reducing the count rate at some time points
5.13	Comparison of the histogram response of the SPAD at lower inten- sity levels (4.5 OD) in the SPC-130 card and the FPGA
5.14	Position of the peak and mean time of the histogram as delay in- creases for a SPAD test
5.15	Block diagram of the TDC system with a block memory-based his- togram generator. Data buses are represented on a thicker line 125

5.16	A 50 MHz clock with 30% duty cycle being sampled at different
	frequencies. To be conservative, the coincidence of a rising edge
	with a rising edge is considered an event miss. Each time division
	is 1.25 ns
5.17	A TPSF collected using the FPGA TDC with timestamp histogram
	generation, before (blue) and after (red) merging bins together. The
	amplitude of the filtered TPSF has been divided by 2
5.18	Mean time of the IRF from FPGA system plotted against the period
	of stabilisation time. The coloured bands represent a range of ± 10
	ps(blue), ± 20 ps (red) and ± 50 ps (magenta) over the mean time
	after 200 min
5.19	Mean time of the IRF from FPGA system plotted against the period
	of stabilisation after the system is stable. The coloured bands rep-
	resent a range of ± 10 ps(blue), ± 20 ps (red) and ± 50 ps (magenta)
	over the average mean time
6.1	A block diagram of the main components and signals for the FPGA
	for a one-channel system. The numbers in brackets represent the
	bus width in number of bits.
6.2	I/O register map for the one-channel system
6.2	
6.3	Flow diagram of the software for the single channel system 141
6.4	Noise profile, DNL and correction factor for an experiment using
	the MONSTIR II laser and the single channel TDC system 143
6.5	A TPSF collected from a phantom using the single channel system
	before (blue) and after (red) applying DNL correction
6.6	Mean time of the IRF from the TDC7201 system plotted against
6.6	Mean time of the IRF from the TDC7201 system plotted against time during the period of stabilisation. The coloured bands repre-
6.6	Mean time of the IRF from the TDC7201 system plotted against time during the period of stabilisation. The coloured bands repre- sent a range of ± 10 ps (blue), ± 20 ps (red) and ± 50 ps (magenta)

6.7	Mean time of the IRF from the TDC7201 system plotted against
	time when the system is stable. The coloured bands represent a
	range of ± 10 ps(blue) and ± 20 ps (red) with respect to the average
	mean time
6.8	Plot of the calculated distance (mean, maximum and minimum)
	against the true source-detector distance for a free-space experiment
	in steps of 2 mm. The dashed lines represent a slope of 1 with dif-
	ferent offsets
6.9	Plot of the intensity (log scale) and mean time of flight (linear scale)
	with max-to-min variation against attenuation (OD) for a free-space
	experiment. The range where the mean time is stable is highlighted
	in blue
6.10	Plot of the output count rate of the TDC7201 system against the
	input start frequency for two different stop frequencies
6.11	Plot of the output count rate of the TDC7201 system obtained with
	a computer simulation compared with real count rate
6.12	Schematic of the setup for phantom scans using MONSTIR II 152
6.13	Imaging setup used for experiments during the project. Top row
	(left to right): two DC sources and a PC. Bottom row (left to right):
	a phantom with a SiPM probe, RF amplifier boxes, comparator box
	and FPGA box
6.14	Schematic and picture of a single channel probe
6.15	Schematic of the rod phantom (from [122])
6.16	TPSFs obtained on the rod phantom, with the target in two different
	positions: between the source and detector (red) and at the edge of
	the phantom (blue)
6.17	Schematic diagrams of the eight-target phantom. All dimensions
	are given in mm
6.18	Schematic of a line scan of the first row of the eight-target phantom
	showing the positions of the probe and the points scanned 157

6.19	Light intensity and mean time for one row of the eight-target phan-	
	tom, scanned in steps of 10 mm. The positions of the four targets	
	are indicated by the dashed vertical lines	58
6.20	Schematic of a raster scan of the eight-target phantom showing the	
	positions of the probe and the area scanned	59
6.21	2D images generated from intensity and mean time measurements	
	of a surface scan of the phantom, raster scanned in steps of 10 mm.	
	The positions of the targets are indicated by red dots	60
6.22	Intensity and mean time values for five rows of the eight-target	
	phantom, scanned in steps of 10 mm. The positions of the targets	
	are indicated by vertical lines. The red dot in the top row represents	
	interpolated data	61
6.23	Illustration of the most likely photon paths when scanning a target	
	at different depths with a 40 mm source-detector separation 1	63
6.24	Plot of the mean time of flight against the absorption coefficient for	
	different target depths, from a Monte Carlo simulation preformed	
	by Diana Sakaan.	63
6.25	Flow diagram of the mutex that controls the access to the SPI lines 1	65
6.26	Block diagram of a two-channel system with two microcontrollers	
	working in parallel. The dashed lines represent signals that are con-	
	nected to the microcontroller appearing on the head of the arrow 1	66
6.27	Flow diagrams for the programs in each of the two channels of the	
	parallelised system. The steps that require access to the SPI mutex	
	are indicated in blue	67
6.28	Block diagram of a two-channel system with one microcontroller	
	coordinating two memory banks	68
6.29	Flow diagram for the two-channel system operated through one mi-	
	crocontroller and the modified register reading subroutine 1	69

6.30	Plot of the output count rate of the two-channel TDC7201 system
	against the input start frequency, for several combinations of start
	frequency. For each channel, the coloured lines represent the fre-
	quency on the other channel
6.31	Two alternative topologies to connect the sync signal from the com-
	parator board to two stop inputs
6.32	Comparison of the sync signal for two channels using two alterna-
	tive circuit topologies
6.33	Comparison of the free-space IRFs obtained from one channel using
	two connection topologies for the sync signal
6.34	Two alternative topologies to connect the synchronisation signal to
	two channels without introducing an impedance mismatch 174
6.35	Rendering (courtesy of Ioana Albu) and dimension diagrams of the
	dual channel probe. Dimensions are in mm
6.36	A scan of one row of the eight-target phantom with the two-channel
	system
6.37	Comparison of the evolution of the intensity and mean time of flight
	of all source-detector combinations in a row scan of the eight-target
	phantom. The positions of the targets are marked by blue strips.
	Position marks in mm
6.38	Evolution of the DPF of all source-detector combinations with po-
	sition for a row of the eight-target phantom. The positions of the
	targets are marked by blue strips
6.39	Plot of the jitter of the mean time of flight for each position for the
	two channel row scan experiment
A.1	Schematics for the non-working designs of the main comparator
	CFD and sync processing circuits. (Main comparator and CFD) 194
A.1	Schematics for the non-working designs of the main comparator.
	CFD and sync processing circuits. (Main comparator and CFD) 195

A.1	Schematics for the non-working designs of the main comparator,
	CFD and sync processing circuits. (Sync circuit)
A.2	PCB layout for the non-working designs of the main comparator,
	CFD and sync processing circuits. (Scaled to 50%)
A.3	Schematics for the functional comparator board
A.4	PCB layout for the functional board
A.4	PCB layout for the comparator board (cont.)

List of Tables

2.1	Comparison of time domain transmittance systems and characteris-	
	tics – sources	59
2.2	Comparison of time domain transmittance systems and characteris-	
	tics – detectors	60
2.3	Comparison of time domain transmittance systems and characteris-	
	tics – timing circuit	61
4.1	Performance requirements for the detector and TDC	98
4.2	Performance requirements for the source	98
4.3	Performance requirements for the probe system	98

Chapter 1

Introduction

1.1 Aims and objectives

Neonatal Hypoxic-Ischaemic Encephalopathy (HIE) is a term for a series of conditions caused by asphyxia during or immediately prior to birth. The lack of oxygen or the alterations of blood flow are known to present a risk to the child during and right after birth and in some cases cause later complications in normal brain development. PAIS (Perinatal Arterial Ischaemic Stroke) or neonatal stroke is a condition caused by the thrombotic occlusion of cerebral arteries during delivery that commonly causes cerebral palsy and seizures during infancy. It is a significant cause of later difficulties in brain development leading to slow development of speech, behavioural and cognitive impairment, and motor complications. The incidence rate of HIE has been estimated as 1.3:1000 births – with some hospitals estimating it to be as high as 8:1000 births [1], depending on the exact definition of HIE they employed. The specific incidence rate for PAIS has been estimated as 1:2300 -1:5000 births, which compares with a 1:5000 incidence of large artery stroke in adults. However, the medical knowledge of these conditions, and their diagnosis and treatments are not well developed [2]. It is known that immediate action, like cooling, improves the outcome in the long term and that these conditions develop in the first 72 hours of life, so correct and early identification is essential [2, 3, 4].

One of the reasons why the clinical developments have been inhibited is the

Chapter 1. Introduction

lack of diagnostic tools to obtain information on the infants' condition. Currently, the only means of confirmation for PAIS is neuroimaging, which is usually performed using magnetic resonance imaging (MRI). This requires the infant to be transferred to a specialised MRI unit, which delays the treatment, and is often not an option for the most critically ill infants. Ultrasound scanning has been explored as an alternative diagnostic, but in general it is insufficiently sensitive or has inadequate spatial resolution [5, 6] – although performance has improved since the 1990s and ultrasound is now considered a general practice tool for routine monitoring, reports still find that it under-diagnoses cases and MRI has to be used for more conclusive assessment [7, 8, 9]. Other related conditions, such as perinatal haemorrhagic stroke (PHS) or sinovenous thrombosis share the same diagnostic requirements.

Near-infrared spectroscopy (NIRS) and diffuse optical tomography (DOT) have already demonstrated some promise in the diagnosis of PAIS [10, 11, 12, 13, 14]. These modalities rely on low power light sources that shine non-ionising NIR light through the patient's tissue, typically to reveal changes in blood volume and blood oxygenation. DOT is safe for use on newborn infants, and involves portable instrumentation which can be operated closer to the cot side in an intensive care environment. DOT can be performed by measuring either the intensity of the light (called 'continuous wave' imaging), the time of flight of individual photons ('time domain') or the phase of modulated light ('frequency domain'). Time and frequency domain devices offer the advantages of determining photon pathlengths, and thus enable more quantitative measurements of physiological parameters, and of providing measurements which are independent of unknown and variable surface coupling. However, despite its potential advantages, the time domain methodology has not yet been fully validated in the clinical environment, and the associated instrumentation is often bulky, expensive, and requires high voltage, power consuming electronics, such as photomultiplier tubes (PMTs). The lack of portability and the high power consumption have so far inhibited use of the technology in environments where it could potentially make the greatest impact, such as in the back of an

1.1. Aims and objectives

ambulance, which has in turn hindered the validation and potential interest in the technique itself. In recent years, advances in technology have produced new optical detectors, sources, signal processing electronics and fast timing electronics in silicon which have enabled the production of new time-domain devices that require less power, are safer to use and fit in smaller, potentially even wearable form factors. Commercial wearable continuous wave NIRS systems have already appeared in the market, and miniaturized time domain systems are currently under development [15, 16].

The primary aim of this PhD project is the development and evaluation of a prototype portable time-domain DOT system (TD-DOT system) based on state-ofthe-art silicon technology which can measure the flight times of photons diffusely transmitted through tissue. One of the objectives for this project is to design a system using silicon based technology, commercially available components and standard communication interfaces, comparing different alternatives where possible. The system is to be affordable, easy to reproduce and maintain, and easy to modify or reconfigure during deployment or use. The sizes of the detector and other parts of the system will be kept small, ensuring the system is as portable as possible, and the voltages and power consumption will remain low to avoid any hazard to the patient. To achieve this, silicon photomultipliers were evaluated as detectors and small commercial integrated circuits were utilised for the time detection and data collection circuits. Another objective is to design a system which is modular, i.e. consisting of an arbitrary number of individual 'tiles' which can be networked together. Ideally, each tile would support an independent source, detector, and timing electronics. Networking multiple tiles into an array will require a means of synchronizing them, so that arrival times of photons emitted by a source on one tile can be measured by a detector on a different tile. A further objective was then to exam the requirements and components to make a scalable and interconnected system possible.

1.2 Organisation of the thesis

Chapter 2 is dedicated to a review of the existing literature about near infrared systems. Special focus will be given to the technical developments in recent time domain imaging systems, and to the technological trends in miniaturisation of the components that integrate the device. Advances in wearable continuous wave systems and in integrated circuits that incorporate sources, detectors and timing circuits is also reviewed, as they inform the way time domain portable systems can be shaped in future. The developments in implementation of time-to-digital converters in FPGAs are also examined, including the proposals that served as basis for the time domain system integrated in the FPGA.

Chapter 3 describes the theoretical background behind this project, explaining the concepts used in the development of the system and the way they have informed the design. The chapter will detail the fundamentals of biomedical optics, electronic design and time-of-flight detection and calibration that are key to time domain optical imaging.

In chapter 4, the optical and electronic components employed in the experiments and in the system design are reported, with a justification for their selection. Additionally, this chapter reports on the design of a signal conditioning system which registers the pulses received from the optical detectors, filters the noise, and generates electronic pulses readable by the time converter. The system development and testing processes are described, and results of experiments to evaluate the stability of the system are presented.

The following two chapters, 5 and 6, describe the design and testing of two alternative time converter and histogramming systems. First, a system fully integrated in a reconfigurable integrated circuit, an FPGA, with connection to a PC is discussed. For this system, two versions of a one-channel system are described: one with direct generation of histograms from the photon events and another based on timestamps. Experiments are presented which characterise the linearity of the system and its stability, and which indicate that the performance is inadequate for imaging experiments. Second, a mixed system using a commercial integrated circuit

is described. This system uses an external TDC7201 integrated circuit, which provides two channels to measure the time of flight, and an FPGA to build histograms and communicate with a PC. The read out, histogram building and communication schemes are described for a single-channel and dual-channel system, explaining the calibration method of the time base for a multi-channel system. The stability, linearity and count rate results are discussed, giving a range of conditions where imaging experiments can be performed. Chapter 6 presents results of imaging experiments on tissue-equivalent phantoms.

The final chapter provides a summary of the work and a discussion of the results, comparing them with other existing systems, and lays out possible future lines of work.

The appendices provide complementary material for the thesis. Appendix A contains schematics for circuits explained in chapter 4. Appendix B contains code used for simulations of the systems in chapters 5 and 6. Due to its length, the full code for the systems described in the thesis cannot be fully reproduced in print, but it is freely available at UCL's research data catalog in https://figshare.com/s/4c754a417f61a899021b. Finally, appendix C contains datasheets for the main components used in the design.

Chapter 2

Project Background

2.1 Diffuse optical imaging systems for HIE and PAIS diagnosis

The large majority of DOI systems used for brain imaging research in general have been based on continuous wave NIRS [17], although there exist a number of published reports of imaging of neonates using time-domain NIRS [18].

The first NIRS measurements in newborn infants, following experiments in laboratory animals, were obtained in 1985 by Brazy, Darrell, Lewis, Mitnick and Jöbsis [19], who observed changes in haemoglobin and cytochrome aa₃ absorption. Reynolds, Edwards, Wyatt and colleagues [20, 21] later obtained the first quantitative measurements of oxygenation and haemodynamic parameters on sick newborn infants, including changes in haemoglobin concentrations, blood volume and blood flow in the brain. Cope and Delpy [22] used a new 4-wavelength photon counting system to measure changes in absorption caused by changes of haemodynamics and the positioning of the infant. Also in 1988, Delpy and colleagues [23] published the first demonstration in biological tissue of a device that recorded the time of flight of individual photons in order to estimate their average pathlength and quantitatively calculate the concentration of cromophores. Chance *et al* [24] would follow this in 1990 with a frequency domain system, which measures the phase of the transmitted

Chapter 2. Project Background

light intensity. These advances opened the opportunity to many studies during the 1990s and early 2000s to obtain the values of the mean photon pathlengths and several other optical properties of the adult and infant brains, as well as clinical studies of haemodynamics, metabolism and functional response to visual and motor stimuli and monitoring of the fetal brain during labour.

The team led by Dr Benaron at Stanford University recorded the first tomographic image of the newborn brain in the year 2000 [25], with a system that measured the time of flight of photons between points of the head circumference of the infant, proposed by Benaron and Hintz's team [26]. Their system used a flexible headband holding the source and detector fibres and employed a low power laser (100 µW) and single detector, taking 2 to 6 hours to record a single static image and limiting the source-detector separation to 50 mm. They employed a simple reconstruction algorithm that back-projected the optical properties calculated from individual TPSFs across statistically predicted photon paths, which ignores the three-dimensional nature of the problem and the heterogeneity of the structure of the infant head. Despite this, this system was capable of correctly identifying intracranial haemorrhage and regions of low oxygenation after acute stroke. The technical limitations were later overcome by UCL with the MONSTIR system, which used a higher laser intensity and simultaneous acquisition for all the detector fibres, allowing a full head scan to be performed within 10 minutes [27, 28]. The introduction of more complex reconstruction algorithms based on iterative model fitting by Arridge et al [29] and Bluestone et al [30], among others, opened the path to more accurate 3D reconstruction of the images. This summary of early developments does not intend to cover the full technical and clinical development of NIRS. More details on the early developments of NIRS and time resolved imaging can be found in more detail in the following reviews by Hebden and Austin [31, 32] and Schmidt [28].

More recently, NIRS has been used in clinical studies to measure predictive biomarkers of hypoxia or brain damage during different phases of the child development or during intervention (surgery, positioning in the cot for observation) and
to evaluate cerebral autoregulation. Reviews of the advances in this field have been given by Greisen *et al* in 2011 [33], da Costa *et al*. in 2015 [34], or by Dix *et al*. in 2017 [35]. By measuring cerebral blood saturation and comparing it to mean blood pressure and the tissue oxygenation index (HbO₂/HbT) with coherent function analysis, Wong *et al* [36, 37] first observed the lack of cerebral autoregulation in sick infants, while Gilmore *et al* [38] employed time domain analysis to describe a new correlation index (the cerebral oximetry index) that can identify failure in autoregulation during the critical first 3 days of life. Another correlation index, this time between arterial pressure and heart rate, TOHRx, has been defined by Mitra *et al* [39], which can be used to define optimal values of mean arterial blood pressure, to distinguish the outcome of patients, in retrospective analysis, and to potentially obtain a real time index to inform clinicians of the status of the neonates. Other studies have focused on the technology used for infant NIRS such as Barker *et al* [40] and Emberson *et al* [41], which describe the prevention and correction of motion artefacts.

The ability of NIRS to serve as a monitor of cerebral oxygenation in infants, and to alert clinicians of changes or of clinical deterioration, are still being assessed. Although no large scale investigation has been published yet, the SafeBoosC consortium is dedicating efforts to characterise the contributions and limitations of NIRS to monitor preterm infants during the first days of life [42]. Using continuous wave brain oximeters, they have conducted a phase II randomised study with 166 neonates, which successfully evaluated the use of a cerebral oximeter to measure oxygenation and reduce the burden of out-of-normal-range blood flow [42, 43, 44]. The next stage, the SafeBoosC-III study, currently in progress, is aimed at assessing the efficiency of brain oximetry combined with clinical guidelines to reduce the risk of death or developmental problems in preterm infants up to the 36th week, and includes over 370 babies. Another set of randomised trials, based on a continuous wave system, was published by Pichler *et al* [45]. Despite discrepancies in the values in normoxia with the SafeBoosC-II publications, this also demonstrated the feasibility of using cerebral regional oxygen saturation to guide intervention in infants.

Meanwhile neonatal seizures, which correspond to episodes of low oxygenation, have been imaged by neoLAB, a group jointly formed by researchers based at the Rosie Hospital in Cambridge and at UCL (Lee *et al* [46]).

Time-domain technology has also been used for studies of the infant brain. Ishii *et al* [47] and Fujioka *et al* [48] used commercial systems developed by Hamamatsu to investigate the differences in perfusion and oxygenation between normalterm infants and those who are either preterm or small for their gestational age, during the first three days of life. They conclude that time resolved measurements can reveal significant differences in blood oxygenation and other markers in children depending on the stage of their neurodevelopment. Nakamura *et al* [49] observed an increase in several biomarkers (oxygen saturation or StO₂ and HbT) in infants with HIE during the first 24 h after birth. Ijichi *et al* [50], Spinelli *et al* [51] – both using TD-NIRS – and Pagliazzi, Giovanella *et al* [52] – with a combined TD-NIRS and DCS instrument – have conducted experiments in neonates to assess the optical properties of the newborn brain, yielding several databases of absorption and scattering coefficients, and values of DPF, blood volume and concentration and StO₂. These are in good agreement with simulations but they conclude that standard reference values have not been obtained yet.

Recent DOT research at UCL concerning infant brain imaging includes: a breakthrough paper by Cooper *et al.* [53] that identified transient haemodynamic phenomena during seizures using combined EEG and DOI; a paper by Brigadoi *et al* [54] which builds a 4D atlas combining high density DOT and MRI information; an article by Singh *et al* presenting the first images of an infant seizure with whole-scalp coverage combining EEG and DOT [55], and a paper by Chalia *et al* [56] that investigates the response of EEG and DOI to infant seizures.

An evaluation of the application of time domain technology for diffuse optical imaging of the newborn brain is described in the PhD thesis of Laura Dempsey at UCL [57], which details the use of novel data processing and probe geometries for more efficient image acquisition, and reports on new findings on the hemispheric symmetry of the blood volume associated with PAIS.

Meanwhile, the BabyLux project [58], involving a network of companies, research centres and universities from Spain, Italy, Germany and Denmark has produced a new device combining reflectance TD-NIRS and diffuse correlated spectroscopy (DCS), with the aim of obtaining simultaneous measurement of oxygen concentrations in blood, microperfusion and metabolism on infants. This device has been demonstrated in preliminary studies on adults and has begun to be tested in real-life hospital settings [52, 59].

2.1.1 Commercial TD-DOT systems

Several time domain systems have been developed and commercialised for medical imaging research purposes. Systems available from commercial companies include those built by PicoQuant [60] or AUREA Technology [15], and those produced by Hamamatsu for use in mammography and brain monitoring (the TRS-10 and TRS-20, based on PMTs, and tNIRS-1, based on cooled SiPMs) [61]. From these, the Hamamatsu systems are the most portable, being presented in compact designs fitting in one case, but no commercial wearable or lightweight time resolved device has been released yet.

Despite these commercial systems, a significant amount of research in the development of TD-DOT and TD-fNIRS systems is primarily based at universities, mainly in Europe and the United States, with almost half of the equipment dedicated to research being built in-house, as estimated by Lange and Tachtsidis [18]. The most recent publications focus on establishing clinical significance and on technological progress in order to make optical imaging more accessible.

2.1.2 TD-DOT systems and progress towards miniaturisation

For time domain imaging, most of the systems currently in use for research share a common structure, being built around a picosecond-pulsed source (e.g. laser diodes, fibre lasers, Ti:sapphire lasers, etc.) with light delivered via fibres to the patient. The detectors used for earlier systems, such as streak cameras and PMTs, are now being replaced by silicon based SiPMs and SPADs. The time detection has been tradi-

tionally employed by PTAs, ICCDs and TCSPC devices, and the data processing and image reconstruction is performed by computers that are powerful enough to accommodate complex forward models of light propagation in the brain and solve the corresponding inverse problems [27, 62].

University College London has pioneered the development of time domain systems for both breast imaging (having demonstrated the clinical potential of TD-DOT for diagnosis and assessment) and imaging of the newborn infant brain, where research is still active. UCL has produced two generations of TD-DOT devices, known as MONSTIR (multichannel optoelectronic near-infrared system for timedomain image reconstruction) [27] and MONSTIR II [63]. Construction of the first generation MONSTIR system was completed in 1999 and was later employed to perform studies of breast cancer detection and response to therapy [64, 65] and on infant brain oxygenation [66, 67, 68]. This first system consisted of 32 sourcedetector pairs, with an external laser source. The light was time-multiplexed, illuminating one source at a time and detecting using multichannel PMTs. The timing for each detector was registered using a CFD and a picosecond time analyser. The second generation MONSTIR II (2014), appearing in Figure 2.1, replaced the laser source with a supercontinuum laser and the timing and histogram building was performed by a TCSPC card, allowing for a more compact device. This system has been used to study infants at the Rosie Hospital in Cambridge as well as for phantom studies at UCL [57].

Researchers at UCL have also demonstrated a time-domain system that uses spread-spectrum modulation as a faster alternative to traditional single-photon counting, based on VCSELs [69]. Papadimitriou *et al* validated the device on phantoms and arterial occlusion experiments. The use of VCSELs significantly improves portability thanks to the reduced footprint of the light transceiver.

Some recently published DOT systems include: a series of non-contact systems developed at Physikalisch-Technische Bundesanstalt (PTB) in Berlin [70, 71], which acquires diffuse reflectance measurements rather than transmittance measurements; a 3D imager also built at the Politecnico di Milano with a photodiode Image removed due to copyright restrictions

Figure 2.1: MONSTIR II and its individual components, taken from Cooper et al [63].

array that sacrifices time resolution (97 ps) in order to achieve a 5 mm spatial resolution [72] using new Fourier domain algorithms to reconstruct the image; the CCD camera-based system built at the Martinos Center for Biomedical Imaging, Boston [73]; or the random bit sequence system reported by Mo and Chen [74], that replaced the laser by a modulated diode and the TCSPC module by a digital-toanalog converter. These articles show experimental support for the use of SPADs, alternative timing systems and novel acquisition schemes as important components for the development of compact systems.

Researchers from the departments of Physics and Electronics, Informatics and Bioengineering in Politecnico di Milano, Italy, have also produced a range of time of flight systems dedicated to medical research. Using systems which measure diffuse reflectance [75, 76], their main focus of research has been characterising the optical properties of phantoms [77] and tissues, applied to breast imaging [78] and the study of brain hemodynamics [79]. Politecnico di Milano has released two TDfNIRS systems known as fOXY (2006) [80] and fOXY2 (2013) [81], also based on PMTs and TCSPC timing devices with diode lasers as sources, as well as other multichannel systems [82]. The emphasis of their current work is miniaturisation [83]. They have published numerous articles on the design of a new time domain

spectroscopy system, analysing the feasibility of SPADs and SiPMs [84, 85, 86] for single photon counting. Recently they have described a prototype single photon counting device based on a SPAD [87], an 8 channel reflectance imaging device based on a SiPM [88] and a DOT probe [89, 90], and evaluated a miniaturised laser source [91]. One of their latest and most advanced prototypes is a SiPM based imaging device (2017) [92]. It employs a pulsed diode laser source with a fixed threshold, a PC controlled time-to-digital converter module (TDC) and pulse conditioning circuits, and has been demonstrated on tissue-simulating phantoms and for simple in vivo measurements. Another of their recent lines of research, and the most successful in terms of achieving a portable and integrated system, began with a publication of the characteristics of a compact, fine-resolution one channel TD-NIRS system, designed as first step towards a portable time domain system [93]. This design integrates the control timing system to trigger a laser photodiode with circuitry to pick up pulses from a SiPM, shape the pulses and feed them to a timeto-digital converter that fits within a $200 \times 160 \times 50$ mm portable metal box, as shown in Figure 2.2. The TDC was fabricated on an ASIC, while the control and PC communication was implemented on an FPGA, meaning that not all the system is reconfigurable. As described in a paper by Renna et al in 2019 [94], this idea was expanded to produce a 2-channel, 8-wavelength system. This system employs an optical switch to multiplex the output of 8 photodiodes at different wavelengths and build histograms from two SiPMs, giving two simultaneous channels with a userselectable integration time. This system was tested on a phantom, but only results for a single channel were reported. The approach used for the design of the device shares common features with the work presented in this thesis; the similarities and differences will be addressed in the Conclusions chapter. A recent thesis by Anurag Behera [95] suggests that development and in vivo testing of small devices with SiPMs and sensors in direct contact with tissue is in active development.

Tables 2.1, 2.2 and 2.3 provide a summary of the characteristics of the sources and detectors used for transmittance time-domain imaging that are most relevant to the specification of the TD system proposed in this thesis in Chapter 4.



Figure 2.2: A TD-NIRS system designed by Buttafava et al [93]

2.1.3 Integrated systems

As an alternative to the module-based systems discussed above, efforts are being made to develop systems that integrate photon sensing, through small silicon detectors, and time of flight readout, often with time-to-digital converters or TDC (discussed in Section 2.2), into a single chip or in a unified, small form factor.

In a 2017 paper, Burri, Bruschini and Charbon [98] report on LinoSPAD, a breakthrough CMOS SPAD and FPGA-based TDC combination with a 50 ps resolution. This TDC does not include a direct-to-histogram architecture, and instead builds the histograms from a thermometer-to-binary encoder, and has not yet been validated using phantoms or *in vivo* experiments.

Later, in 2019, two papers by Saha *et al.* [99, 100], in a team including the aforementioned researchers, report on a 10 mm² chip integrating a VCSEL and a SiPM for imaging in direct contact with the sample. The chip includes all signal processing and photon identification with time gating done in a custom CMOS process, with the additional option to connect the output of the SiPM to a more advanced TCSPC. The advantage of this design is that the power consumption of the amplifiers integrated in the CMOS chip is very low and reduces the overall footprint. So far this design has only been validated with scattering experiments using phantoms.

Another trend in design is the integration of TDCs with SPADs or other pho-

todetectors to produce cameras and optical probes with automatic histogram generation for time-of flight dependent applications, like DOT and lifetime spectroscopy. For example, in a 2018 paper where they propose a probe for reflectance imaging, Alayed *et al.* report that CMOS SPADs, that is SPADs that can be manufactured in a silicon CMOS process usually with a small footprint, are reducing their price compared with SiPMs [101]. Bruschini, Homulle and Charbon [102] point out in their 2017 review that there is an increasing trend for the integration of imaging technology with the photosensors, and they foresee further developments in this technology.

One recent example of a development in this area is Nolet *et al*'s readout ASIC for SPADs [103]. They propose a pixelated readout circuit that can be connected to an array of 256 SPADs, effectively turning it into a 1.1 mm \times 1.1 mm SiPM with a per-pixel timestamp generator with a timing resolution of 10 ps. External circuitry is then used to manage the throughput of information, which the authors explain is a significant bottleneck. The main application for this circuit is the timing of PET events, but similar technology could be employed for near-infrared imaging.

In 2015 Pavia *et al* presented a chip for NIROT (near infrared tomography) [104], cited in review papers as an important step towards integration of CMOS imagers. In this type of technology, the SPAD and the TDC are integrated into different layers in the same 3D fabrication process to obtain a single IC with the full functionality. In their proposal, clusters of SPADs are connected together to a shared time-to-digital converter with 50 ps resolution, with an arbitration circuit that detects collisions of events with a 'winner-takes-all' protocol. This can output a stream of data at 1040 Mbps. This chip was tested using tissue-simulating phantoms, and exhibited good response to variations in intensity and mean time of flight. This architecture was later overtaken by systems like the Ocelot, described in 2019 [105] by a collaborative team from Zurich, Lausanne and Delft, that employs parallel structures, similar bus architectures and collision detection but improves the number of detectors and is more readily demonstrated in practical applications (in this case, LIDAR at 30 frames/s).

Image removed due to copyright restrictions

Figure 2.3: Photomicrograph of the line sensor from Erdogan *et al* [107]. Blue and red SPADs for five pixels are zoomed in.

Another instance of this type of detector technology are the line sensors developed by researchers at the Universities of Edinburgh, Dundee and Durham, including a 1024×8 sensor (described by Erdogan *et al* in 2017 [106]) and a 512×16 imager (described by Erdogan *et al* in 2019 [107]). A photomicrograph of one of the sensors is shown in Figure 2.3. Their technology consists of an array where each pixel incorporates a SPAD, a TDC and a 32-bin, 11-bit histogram generator with a configurable or 'zoomable' resolution of 51.20 ps to 6.55 ns per bin, and several clock networks for the TDCs. The design allows three readout modes – SPC, TCSPC and on-chip histogramming – with data rates of the order of millions or billions of events per second (depending on the mode). This allows fluorescence lifetime imaging with SPADs centered on different wavelengths to be performed on-chip.

Yet another alternative proposal is to change the patterns of activation of the array of detectors to read out the signal with one timing circuit based on compressed sensing, as is reported in a 2019 paper by Farina *et al* [108], where a single timing circuit receives the events from a high density array of detectors. Instead of reading out element-by-element, the pixels are activated in groups with patterns of activation designed to scan the sample in the spatial frequency domain. This method, although less competitive in terms of performance – they have not yet acquired sub-second images – allows for existing equipment to be used and can drive down the cost of manufacture. The paper demonstrates its use for time-domain fluorescence and

transmittance imaging.

Some commercial systems include: Philips' dSiPM for PET imaging; ICs from small companies such as PhotonForce and MPD [102]; and modules from STMicroelectronics, that integrate light emitters and SPADs for time-of-flight ranging applications with customised round-trip time measuring algorithms [109]. One such module (the VL6180X) was demonstrated in an article by Hebden, Shah and Chitnis [110], which describes a low cost testing probe to characterise the properties of liquid phantoms. While the results were not a true measurement of the photon path length, the IC was able to follow the changes in absorption and scattering with good correlation. A later version of the module, the VL53L1X, was evaluated at UCL as part of an undergraduate project [111]. While power efficient, affordable, easy to interface and promising for some time domain imaging situations, the readout scheme was slow and the time measurements exhibited dependency on the type of reflecting surfaces and the illumination.

In contrast with modular systems, systems integrating the detector or the readout circuitry in custom processes utilise a lower footprint and can achieve high timing resolutions of tens of picoseconds, making them a potential technology for portable time domain imaging systems. However, they typically lack flexibility, as the design is more monolithic and system components are already set in the silicon and as such are more difficult to manufacture independently. The current weak points of the technology, as identified in Bruschini *et al*'s review [102] are the immaturity of the technology, the efficiency of the readout schemes – that nonetheless have improved since the publication of the review, the established competition from other types of single-photon systems, and the mostly experimental nature, rather than commercial, of many of the designs. Thus while these detectors could be integrated in the future in DOT or NIRS systems, they are not readily available now.

2.1.4 Other miniaturised systems

The fastest development in miniaturisation of diffuse optical tomography systems has been produced for continuous wave (CW) imaging because of the simplicity

of the sampling and the availability of components when compared to time domain imaging. In the last decade multiple CW wearable systems and prototypes have been produced [16].

Some early wearable designs were based on a support headset, which houses the sources and detectors, linked by cables to a control box that contains the sampling technology and communicates with a PC. Examples of this technology include the systems produced by Atsumori et al [112] and Kiguchi et al [113], which were only tested on a few adult volunteers as the components were cumbersome to use and electrically unsafe. Another system reported by Piper et al in 2013 [114] successfully reduced the size of the control box and the bulk of the cables. Later developments of this concept led to the system reported by Lareau et al [115] and Safaie et al [116], which combines EEG probes with optodes and transmits the data via a wireless connection to a central PC. These technologies have mostly been superseded by other formats, more suitable for wearable systems, and some designs are still emerging such as that presented by Bergonzi et al in 2018 [117], which uses 200 µm diameter fibres (24 sources, 28 detectors) that are lightweight enough to be worn by the patient even while walking, as the full cap weighs just 1 lb, to obtain superpixel images, in which the signal-to-noise ratio and the detectivity threshold are enhanced via a noise subtracting algorithm.

More portable designs have been achieved by placing the sources and detectors on a probe that then connects to a control unit, making the design fibreless. In contrast with the previous approach, the control of the system can be distributed, with on-probe sampling (providing a better SNR) and a reduced central unit. Thus the sources and detectors do not need cumbersome cabling and can be arranged more comfortably on the scalp with a higher density. The first designs with this approach used a mixture of rigid and flexible PCBs, with the flexible PCB serving as interconnection between the rigid parts. Examples of these were produced by Muehlemann *et al* in 2008 [118], who published the first prototype with flexible PCBs that connected to a central PC via Ethernet. Later, in 2016, Hallacoglu *et al* [119] achieved one of the first fiberless high-density systems by placing VCSELs

and detector diodes at very short separations, employing a combination of flexible PCBs and interconnecting cables, which was accepted for commercial development for clinical applications by Cephalogics. Again, the high density of electrical cables and the distribution of power among the probes can limit the scalability due to power consumption for this type of device.

More recent systems employ a completely modular design to address scalability and reduce the cabling. Each module integrates a limited number of sources, detectors and dedicated control electronics all designed to cover the patient's scalp completely and adapt to its shape. The control electronics enable each individual module to sample the signal from the desired source with different multiplexing schemes, creating a large array of sources and detectors from small individual units. Examples of this type of system include that proposed by Von Lühmann et al in 2015 [120], which included most of the source and detector control electronics into distributed modules, with a central unit carrying out the most demanding parts of the data acquisition. This system was later improved by increasing the digitalisation resolution and adding accelerometry for motion correction [121]. In 2016, Chitnis et al [122, 123] demonstrated a fibreless, 8-wavelength distributed system (called micro-NTS) which integrated 2 LED sources and 4 photodiodes into a small PCB 'tile', with all the amplification and digital sampling done on the 'tile', achieving a low (fW) power consumption. Four tiles could be interconnected with a single ribbon cable, making the placement on the scalp quite easy, at least for a small number of modules. Zimmermann et al [124] described another portable system in 2013 which incorporated the sources and detectors and the control electronics in separate stacked PCBs, and was the first demonstration of SiPMs for fNIRS applications, This design was later improved by Wyser et al in 2017 [125] by adding a new communications protocol similar to that of Chitnis et al. Finally, Strangman et al [126] reported in 2018 on the applications of a custom made device, known as the NINscan, which integrates 64 channels (NIRS, ECG, respiration) with accelerometer based motion correction and which can operate for 24 h while powered on AA batteries; this system has been evaluated for applications related to sports and high

altitude mountaineering [127].

Newer DOT systems are being developed which either attempt to minimise the occurrence of motion artifacts, or which integrate components to detect and compensate for motion. Piper et al [114] provided the first demonstration of fNIRS during an outdoors sports activity in 2014. Algorithms like MARA or wavelet analysis are being applied to the data itself to correct for motion artefacts, as explained in Pinti et al's 2018 review on wearable fNIRS [128], and by following Di Lorenzo's recommendations in 2019 [129]. Chiarelli et al [130] published in 2019 a specification for a system which employs EEG-style sensors to reduce the movement of the optodes on the scalp and uses commercial EEG caps for fNIRS, adding lock-in techniques to improve the SNR in the photodetector signal. Meanwhile, motion sensor data from accelerometers, gyroscopes and magnetometers to correct fNIRS data (e.g. recorded on a walking subject) has been reported by Siddiquee et al in 2018 at Florida International University [131] and by Brigadoi et al at UCL in 2019 [132]. Also, new devices are being designed to produce high-density DOT images, like the work by Bergonzi *et al* cited above or the prototype system by Maira *et al* published in 2020 [133], which integrates 12 LEDs and 13 SiPMs in a 80 mm \times 80 mm grid.

Two of the latest commercial systems for portable NIRS are the Nirsport 2, produced by NIRx [134], and LUMO, produced by Gowerlabs Ltd [135]. Nirsport 2, shown in Figure 2.4 is a 200 channel multimodal helmet system compatible with NIRS and EEG, suitable for on-device storage or wireless transmission. LUMO, shown in Figure 2.5 was launched by Gowerlabs Ltd in 2019 and is a high-density system, derived from the micro-NTS, consisting of an array of hexagonal 'tiles' housing three sources and four detectors that can be inserted into or removed from a wearable cap that provides the interconnection between them, making the number of source-detector pairs and the area covered more flexible. A wireless wearable hub is in development. LUMO has been validated in adults and is now being examined for use on infants as well (see the article by Zhao *et al* [136]).

These latest developments in wearable CW-DOT reveal the directions that TD-



Figure 2.4: The multimodal NIRx Nirsport 2, shown with the cap and the portable central unit. Taken from NIRx's website [134].



Figure 2.5: The LUMO NIRS system from Gowerlabs Ltd, with cap and central unit. Taken from Gowerlabs's website [135].

DOT can follow, with wearable distributed probes and scalable modular designs, which can facilitate important new areas of application, such as new functional and psychological studies and industrial applications [15]. The efforts dedicated to re-

duce the footprint of the individual components of time domain devices, together with the research in new architectures and sampling techniques that parallelise the sampling and information processing among distributed elements, are advancing the field towards low power systems based on individual probes that can be used for research in a clinical setting. It has become apparent that the 'tiled' form factor is advantageous for power distribution, to reduce noise and to make the imaging setup more flexible. On a technical level, the accessibility of the individual components required to assemble a system, the coordination schemes between channels in a wearable or low-power system, and the efficiency of the image reconstruction, are all topics that are less discussed in the published literature and require further research. *In vivo* experiments and more complete phantom and *ex vivo* imaging experiments should also be performed to thoroughly validate portable systems.

2.2 Time to digital converters

The time of flight measuring device is one of the key components on which TD-DOT relies, and as such it has received the most attention during the first stages of this PhD project.

When work on fine-precision sub-microsecond time interval measurement devices began to be reported, in the second half of the 20th century [137, 138, 139], the main technique used to calculate the time between two events was time-to-pulse height conversion, also known as start-stop measurement, where the start signal (an event) triggers a voltage ramp that is halted by the stop signal [140]. The voltage level is then proportional to the time of arrival, which can be digitised if required, turning the converter into a TDC (time-to-digital converter), in which each time of arrival is assigned a digital code. Other techniques developed were based on oscillators, employing interpolation circuits or counters, based on the Vernier principle, but these did not achieve as fine a timing resolution with the same level of linearity.

Nowadays [141, 142], although commercial time-to-amplitude converters are still available [143], CMOS fabrication processes have facilitated the development

of digital delay line and oscillator TDCs where all the components can be implemented in silicon. This has resulted in a larger number of commercial TDCs based on this technology, and in the availability of designs targeted at reconfigurable hardware. All of these designs are sensitive to the variability in industrial processes, that will render two systems to be non-identical even if manufactured at the same time, and to changes in temperature, that affect the frequency drift of the oscillators and the delay of the delay lines. Calibration and on-the-fly compensation is often implemented on modern TDCs.

Several commercial TDCs are available, such as the modules sold by Surface Concept [144] and Laser Components [145]. These are typically too large to be suitable for a wearable or highly portable device. TDCs in an integrated format are less common and their usual purpose is range detection and distance measurement, which require a lower sample rate than TD-DOT and often have higher jitter levels. The availability of standalone integrated circuits is considerably lower than that of PC cards or small boards from TDC manufacturers. Texas Instruments, AMS and Maxim Integrated have the largest range of available TDCs on integrated circuits. Some of these models are not suitable for time domain imaging, as they do not include separate event and reference inputs, relying instead on an oscillator as the time base. From available integrated circuit options, the Texas Instruments TDC7200/7201 [146] and the AMS TDC-GP22 [147] and GPX2 [148] models were found to be particularly suitable for this project. The TDC7200 and TDC-GP22 chips are oscillator based and work on CMOS logic levels, making them easy to interface and operate. The oscillators impose a minimum and a maximum on the measurable time of flight but their timing resolution of around 50 ps make them potential choices for the prototype system. The AMS chip reports a 500 kS/s rate, whereas the Texas Instruments chip does not quote a rate. The GPX2 is a high-end chip that operates at LVDS levels for fast interfacing (20 MS/s) and fine resolution (20 ps). The Texas Instruments chip was chosen for experimentation because of its lower price, faster availability and the option to use two separate 'start' event channels. Tests were later conducted which found that the sampling rate is 10 times less

than that of the otherwise similar AMS chip.

Alternatively, TDCs have been implemented in custom processes [141, 149] and some recent designs have successfully been coupled with SPADs for photon counting [91]. Many designs, based on both oscillators and delay lines, have been produced and reported for ASIC and FPGA, as they reduce the cost of prototyping and manufacture and give more control to the designer over the characteristics of the converter and how it is integrated with other components. Two types of TDC designs can be implemented this way, oscillator converters and delay lines.

To implement oscillator TDCs in an FPGA, a circuit that delivers a reliable oscillator needs to be built. This has been achieved by using phase locked loops (PLL) or by joining logical gates in a loop, to form a simple metastable circuit [150, 151, 152]. The oscillators can then be coupled to counters to obtain the timestamp of the events. Oscillator TDCs are very area efficient; in its simplest form, the ring oscillators can be made to occupy a few logic cells, and fine adjustment of the frequency easily delivers picosecond resolution. On the other hand, as these designs rely on metastability and the complex behaviour of electronic components, software simulation of the TDC has proven to be challenging but not impossible. These factors also hinder in-chip debugging and the design of test suites in general. Comprehensive information on a sound test methodology does not appear to be available in the published literature. For these reasons, this architecture was discarded as an option.

Delay line TDC designs are abundant in literature, with recent papers reporting resolutions of the order of 10 ps. The base design of the delay line, making use of intrinsic FPGA components, is simple enough to implement in few lines of code, but comes at the cost of area and efficiency [153]. With process delays of the order of tens of picoseconds, implementing a delay line that covers time spans of several nanoseconds quickly consumes the FPGA fabric. In addition, the combination of the manufacturing process and the place and route allocation does not guarantee the linearity of the delay line.

The main trend in TDC design couples an encoder with several systems to

correct for non-linearities. When the signal is digitised and propagated through the delay line, it produces chains of '1's and '0's that can be interpreted to signify an event arrival and how many elements it has gone through. These chains are known as a thermometer code, because the length of the chain resembles the measurement of temperature in a mercury thermometer. The most common approach is using fast, clock-asynchronous counters to add the number of significant '1's (or '0's). This is known as the fine encoding, that can be subtracted from a lower frequency coarse counter to obtain an absolute timestamp or used on its own to work with a timestamp relative to the stop (or clock) signal.

Variations on this design include a second delay chain on the clock with a slight difference in the delay, forming a Vernier delay line TDC [154, 155]. While the principle of operation is the same, the resolution can be increased from the mere delay of the delay elements to the difference in delays between both delay lines. This approach was considered difficult to implement and test, because now the process differences of two delay lines have to be characterised.

Many articles describe the successful implementation of other types of delay line TDCs within FPGAs, motivated by PET and ultrasound applications. In 2008, Wu and Shi [156] presented an approach for digital calibration and correction of a delay line TDC called the *wave union*. In their design, a train of pulses is generated for every event, yielding multiple hits in the thermometer code that can be interpolated to obtain a measurement with finer resolution than that provided by the delay line. This also reduces the sensitivity to non-linearities. An implementation by Wang *et al* in 2011 [157] achieved a 9 ps RMS resolution in a Virtex 4, an FPGA architecture two versions older than the test platforms employed in other articles of that time.

In 2009, another scheme was proposed by Favi and Charbon [158]. Often in TDC designs, the delay line needs to be clear before encoding another event, as the chains of events can overlap or mask each other. By encoding the events with a level transition, rather than exclusively with a '1' or a '0', and adding several delay lines that operate simultaneously, more events can be identified and counted within

one clock cycle, reducing the dead time of the TDC to less than 500 ps. This paper also included calibration and interpolation schemes to deal with irregularities in the delay line that cause so-called code bubbles. If the delay in the delay line is not monotonically increasing, small chains of 1s or 0s can be found in the middle of larger chains of the opposite level, causing problems when encoding. They propose that dithering the histogram can mask this artefact for small bubbles.

Menninga, Favi, Fishburn and Charbon built on this design in a 2011 paper [159], where they describe a design fully implemented TDC in an FPGA with a 10 ps timing resolution and low non-linearity, and in another paper in 2013 [160] focused on open source PET applications of the technology. Their thorough testing approach, comparing different FPGA models and mapping the delays inside the FPGA, represents a point of reference for this work. Part of this initiative resulted in a simplified version being released as open source hardware by the University of Delft [161].

Also in 2011, Sébastien Bourdeauducq published a TDC design based on the paper by Favi and Charbon to be used in CERN's White Rabbit project [162]. The design is publicly available for download [163] and takes a different approach for some of the issues described in the original paper. It does not use a toggling signal register, preferring a pulse widening circuit instead, and includes a temperature compensation system. It was designed with the intention of adding wave union interpolation in the future. This design was used as a primary reference in my project for the design of an FPGA-based TDC.

In 2014, Dutton *et al* [164] proposed an alternative to thermometer encoding. The traditional encoding has an inherent dead time, caused by the overlapping of events and the speed of the thermometer-to-binary converters, that could only be solved by using multiple delay lines. If the toggling signal from Favi and Charbon is compared with a series of XOR gates, the exact position of the event transitions can be detected immediately and asynchronously. In addition, this series of event localisations, called one-hot encoding in the paper, can be used to feed counters and build a histogram at the same time that the signal arrives – the paper calls this

'direct-to-histogram' conversion. This avoids the need for storing the timestamps values and all the processing can be done on the histogram itself, limiting the sampling speed to that of the input signal. Events can still go undetected if they are too close together, as either the delay of the toggling register or the delay elements can cause them to register as only one event. Rather than improving the resolution, Dutton *et al*'s proposal improves the device's efficiency. This idea has been tested for the FPGA TDC presented in this work, with the direct-to-histogram technique being explored in more detail in Section 5.1.

Chen, Zhang and Li later addressed the non-linearity problems of this design in 2017 [165]. In their paper, they claim a 10.5 ps resolution TDC with non-linearities of less than 1 least significant bit, a unique achievement, by combining the direct-to-histogram architecture with the tuned delay line proposed by Won and Lee in 2016 [166]. The tuned delay line is a particular implementation based on the one-bit adder, the component typically used to create delay lines in FPGAs. Traditionally only the carry output of these components is used to tap the line, but Won and Lee proposed that alternating between the carry and sum outputs for a particular FPGA model had the potential of linearising the delay. In addition to the tuned delay line and the XOR gate system, the paper by Chen, Zhang and Li also discussed using different clock phases to improve the event collection. None of these changes have been implemented in this project, but they remain ideas to potentially experiment with.

Other TDC designs that have been explored include the designs of Aloisio *et al* [167] and Torres *et al* [168] and those described in the theses of Amiri [169], Yuan [170] and Daigneault [155]. In his thesis, Daigneault develops a TDC aided by software planning that solves non-linearities using an automated design. This approach provided valuable ideas on how to manage resource allocation and design for FP-GAs, but the specific program requires deep knowledge of the FPGA architecture and reduces the portability of the code to other FPGA models.

2.3 Conclusions

During the past 10 years, there has been increased focus on the clinical translation of optical techniques for infant brain monitoring as a result of greater miniaturisation and portability of opto-electronic components. Technology which is small and portable is particularly important when attempting to apply such techniques to patients in a critical care environment, at the bedside, or outside the hospital. Research into continuous wave imaging devices for both infants and adults has made outstanding progress towards portable and wearable systems. More portable time domain systems have been slower to appear, but recently there have been major advances in silicon photomultiplier technology, custom time of flight circuitry, and low-cost ultrashort-pulse light sources.

The published academic literature has not yet documented any portable time domain system with numbers of discrete channels or wavelengths equivalent to those achieved for continuous wave devices. Neither is there published literature addressing the issue of coordinating the time references across channels in a multichannel time-domain system, and no significant *In vivo* results have yet been reported using such a system.

The systems described in the following chapters were designed to be applied ultimately to topographic and tomographic image reconstruction, and consequently particular attention is paid to the interaction between two or more channels and the feasibility of scaling up the system to an arbitrary number of channels.

The system design is also focused on wearability, the availability of the components, and the flexibility to employ the components for different configurations. Thus the components used during this project were chosen to potentially fit in all inclusive probes, and were readily available from commercial suppliers. This contrasts with other published proposals that utilise custom ASICs and/or employ bulkier components if their performance is better. This thesis explores the advantages, limitations, and replaceability of the different components in a modular design. This thesis also documents some of the less discussed aspects of circuit design, the integration of time-of-flight measuring techniques with optical detectors, system stability, sampling limitations, and system sensitivity to external interference.

	Table 2.1: Comparise	on of time domain tran	smittance syst	ems and characteri	stics – sources
System	Type	No. <i>J</i>	P max	Rep. rate	Pulse width
	Supercontinuum laser	Λ 1011			
Monstir II [63]	NKT SC400	AIIy 410 7400 am	5 mW/nm	40 MHz	6 ps
	fibre coupled	410-2400 1111			
Dt.f 2017 [02]	TO-8 laser diodes	2	0/Mm 2 C	40 MHz	
[CE] / IUZ BURIAND	fibre-coupled	670, 830 nm	V/MIII C.7	(max 80 MHz)	sd nzz
Dame 2010 F041	TO-8 laser diodes	8			75 175 22
K ellitä 2019 [94]	fibre-coupled	635–1050 nm	4./ III W / A		sd cc1-c/
Hamamatsu	laser diodes	3			
tNIRS-1 [96]	fibre-coupled	755, 816, 850 nm			

Table 2.1: Comparison of time domain transmittance systems and characteristics -

)				
			Not published	MPPC (SiPM)	Hamamatsu tNIRS-1 [96]
2 K	15 V (syst 50 V? (Sil	300–1000 kHz	PDE 5-10%	SiPM Hamamatsu*	Renna 2019 [94]
	95 V	300 kHz	PDE 7-10%	SiPM Excelitas C30742-11-050-T1	Buttafava 2017 [93]
	11.5–15.5 (module)	125–375 Hz	80-90 mA/W	PMT Hamamatsu H8224P-50MOD	Monstir II [63]
	V	Dark ct.	NIRS sensitivity	Туре	System

Table 2.2:
Comparison
of time
e domain
n transmittanc
ce systems and
l charact
eristics
 detectors

circuit
- timing
haracteristics -
s and c
transmittance system
domain
of time
Comparison
ë
ıble 2,
E

serial to FPGA 400 kS/s (system)

Not published

Custom Direct estimation of StO2, ...

Hamamatsu tNIRS-1 [96]

3 MS/s (max)

CMOS

10 ps

Custom ASIC [97]

Buttafava 2017 [93] Renna 2019 [94]

SPC-150

Chapter 3

Theoretical Background

3.1 Fundamentals of biomedical optics

The term *optical imaging* refers to a family of imaging techniques that use light as the interrogating source to derive structural and functional information. In this context, light is defined as electromagnetic radiation with wavelengths in the range from 100 nm to 30,000 nm, including ultraviolet, visible, and infrared radiation. It can be considered to propagate through space in packets of energy called photons, and the signal differences essential for imaging arise from the different interactions between photons and matter. Over this range of wavelengths, the primary interactions in biological tissue are absorption, reflection and refraction (where reflection and refraction at a microscopic scale give rise to the interaction we call scattering). [171, 172, 173]

Absorption is the loss of light photons due to various molecule-dependent mechanisms, such as the excitation of molecules to higher energy states. For a given material, a characteristic absorption coefficient, represented by μ_a can be defined. This coefficient represents the likelihood of a photon being absorbed per unit length or, equivalently, the reciprocal of the average distance a photon travels before being absorbed. Absorption in biological tissue is highly wavelength dependent. This is a determining effect in optical imaging of tissue, where optical absorption by the principal light absorbers or *chromophores* – water, melanin and

haemoglobin – is highly dominant except for a *diagnostic or therapeutic window* around the near-infrared (NIR) region, shown in Figure 3.1, where absorption is low enough to allow light penetration of several centimetres at safe levels of exposure.

Image removed due to copyright restrictions

Figure 3.1: Absorption spectra in the optical wavelength range for some tissue chromophores [174].

The interactions known as *reflection* and *refraction* are due to internal variation in refractive index (i.e. change in local speed of light). When light reaches a boundary between media with different refractive index, some of the energy is reflected and some is transmitted. If a beam of light arrives at the boundary at an oblique angle of incidence, the transmitted component is refracted (i.e. the direction of the beam changes). Multiple reflections and refractions occur at the scale of individual cells, producing what is known as *scatter*. This would cause a pencil beam of light incident on tissue to diverge into a diffuse distribution. Scattering can be characterised by a scattering coefficient, μ_s , equal to the probability of scatter per unit length. Its reciprocal is the mean free-path.

While μ_s influences the proportion of photons that are scattered away from

their original direction, the scattering phase function, $p(\mathbf{s}, \mathbf{s}')$, describes the new direction the photons will take.

For scattering media p is a continuous probability density function describing the probability of a photon changing its direction from its initial direction \mathbf{s} to its final direction \mathbf{s}' . For homogeneous tissues, such as blood, it can be assumed that the phase function is only dependent on the angle θ between the two direction vectors, $p(\mathbf{s}, \mathbf{s}') = p(\cos(\theta))$ or $p(\mathbf{s}, \mathbf{s}') = p(\theta)$. The scattering probability for different directions will depend on the medium, with some media scattering the light equally in all directions (isotropically) and others scattering in a preferential direction (anisotropically). The anisotropy coefficient $g = \int_{4\pi} p(\mathbf{s}, \mathbf{s}')(\mathbf{s} \cdot \mathbf{s}') d\Omega$ is equal to the mean cosine of the scattering angle, such that g is a number between -1 and 1, with g = 0 corresponding to istotropic scattering. This allows a new scattering coefficient to be defined, the *reduced* or *transport* scattering coefficient, as $\mu'_s = \mu_s(1-g)$. By introducing the effects of direction into the coefficient, scattering can be re-scaled as a diffuse phenomenon replacing multiple small-angle scatters with a single isotropic scatter.

Scattering is the dominant effect when NIR radiation travels through tissue, with typical mean paths of a few tens of µm. However, NIR photons can be scattered many thousands of times before being absorbed, and thus scattered light can be detected over distances of several tens of millimetres. *Diffuse* optical imaging (DOI) is the family of techniques that take advantage of this scale of distances to image large volumes of tissue.

3.2 Diffuse optical imaging and time domain

Diffuse optical tomography (DOT) techniques work by shining light of known intensity and wavelength (and often modulation frequency or pulse duration) at known multiple locations into tissue and measuring the light that is diffusely reflected (i.e. scattered within the tissue and then re-emitted within a few centimetres from the source) or transmitted (i.e. across large thicknesses of tissue). The imaging pro-

Chapter 3. Theoretical Background

cess then consists of determining the internal distribution of optical properties (i.e. μ_a and μ'_s) that is consistent with the measurements. This process is both illconditioned (the number of unknowns far exceeds the number of measurements) and ill-posed (the solution is not unique) [175, 176]. It requires an accurate model of light transport through the tissue (known as the forward problem) in order to generate a matrix describing the sensitivity of the external measurements to perturbations in optical properties at discrete locations within the imaged volume, and a means of inverting the matrix to reconstruct the internal properties for a given set of measurements (known as the inverse problem). A sophisticated software package for DOT image reconstruction has been developed at UCL, known as TOAST [177, 178].

During the past twenty years, DOT has been widely used as a research tool for imaging soft tissue, and specifically to map functional and metabolic activity. The most common applications of DOT have been breast and brain imaging [15]. Breast imaging has been explored as a safer and more effective alternative to x-ray mammography for cancer screening, and also to assess tissue response following surgery and/or therapy [65, 179, 180, 181]. However, brain imaging has become the dominant application, with DOT systems now widely used by researchers to study the functional behaviour of the healthy brain, particularly on patient groups and in environments where functional MRI (fMRI) is inappropriate or impossible. Such groups include developmental studies on babies and young children, and newborn infants in intensive care with known or suspected neurological injury, such as due to complications resulting from extreme prematurity. The penetration depth of NIR radiation in the infant head is sufficiently high to achieve 3D optical imaging of the entire brain [67, 182, 68]. Other DOT studies have included investigations of skin tumours, muscle imaging, and molecular imaging. [183, 184]

There are two basic approaches to DOT image reconstruction, known as 'absolute imaging' and 'difference imaging'. Absolute imaging involves generating a 3D image of the absolute optical properties of the medium from a single set of measurements, whereas difference imaging involves acquiring an image which reveals the changes in the internal distribution of properties which occur between the acquisition of two discrete sets of data. For detection and characterisation of PAIS in infants, we are unlikely to have data acquired before the onset of the condition, and therefore it will necessarily rely on absolute imaging.

Most DOT instruments measure the intensity of light using continuous-wave (CW) sources. Consequently, they are highly sensitive to variation in the coupling efficiency of light into and out of the tissue surface, which results in unpredictable variations in intensity. Because the coupling is unknown, the instrument can only reliably measure intensity *ratios* where the coupling (assumed constant between successive measurements) cancels out. The resulting *differences* in the logarithms of the intensity can then be used to reveal changes in brain optical properties. However, TD-DOT measures photon times-of-flight which are largely insensitive to coupling, and thus offers the facility to perform absolute imaging of optical properties.

Time Domain DOT or TD-DOT involves illuminating tissue using a source of very short pulses of NIR light with a duration of a few picoseconds, and acquiring histograms of the times of flights of photons which emerge at points on the surface at a separation d from the point of illumination. For a given source and detector location, the sensitivity of the measurement to the volume explored by the detected photons is known as the photon measurement density function (PMDF).

Image removed due to copyright restrictions

Figure 3.2: PMDFs for a continuous wave beam for different source-detector separations [185]. The white lines indicate the boundaries between the scalp and skull, skull and CSF, and CSF and grey matter, from top to bottom.

As a rough rule-of-thumb, the average depth of the PMDF in soft tissue is around half the source-detector distance d. Thus, to probe superficial tissues, the source-detector separation is kept low, but is increased when requiring good sensitivity to deeper organs. For example, to probe the infant cortex, a source-detector separation of d = 3 - 4 cm is sufficient. However, to generate an optical image of the entire brain, sources and detectors must be distributed around the head, including separations of up to 12 cm. This requires very sensitive large-area detectors, and has been achieved by the UCL TD-DOT system, MONSTIR [27, 63].

The measured histogram of the times of flight of photons is known as either the *photon distribution of time of flight* (DTOF) or the *temporal point spread function* (TPSF). This distribution becomes broader when tissue scattering increases and when the source-detector separation increases, and become narrower when tissue absorption increases. It is commonly used to estimate the mean time of flight $\langle t \rangle$ and a unitless parameter known as the differential path length factor (DPF):

$$DPF = \frac{\langle t \rangle c}{d}.$$
(3.1)

The DPF is dependent on both μ_a and μ'_s , and for the neonatal brain the DPF is typically around 5 [186].

The technology required for TD-DOT is necessarily complex, requiring accurate time discrimination at a sub-nanosecond level and a means of acquiring a histogram. Until recently, this could only be achieved using expensive and bulky electronics, such as PMTs, high voltage time discriminators or, in the best scenario, PC-based timing cards, which severely limits their portability. Now, advances in silicon processes have enabled the construction of sensors, timing circuits and measuring electronics in smaller form factors that require lower voltages to operate.

The following sections are devoted to explaining the principal components of a TD-DOT system. A review of the current state of the art and prior developments in time domain instrumentation were presented in Chapter 2, and the next chapters describe the technical aspects of the design of TD systems and how their performance are assessed.

3.3 Detectors

The measurements employed in diffuse optical imaging, such as light intensity and photon times of flight, derive from the electrical signal obtained from a suitable light detector. Quantum detectors, where the photoelectric effect is exploited, are commonly used for imaging because of their efficiency and fast signal response. There are two main kinds that will be described later: tube based and semiconductor based.

The photoelectric effect, first described by Heinrich Rudolf Hertz and later formalised by Albert Einstein in his theory on the corpuscular nature of light is a

[...] phenomenon in which electrically charged particles are released from or within a material when it absorbs electromagnetic radiation. [187]

When light is incident upon a material, the photons can be absorbed. If the photon energy (proportional to the frequency of the light, v, as E = hv) is above a certain threshold, free electrons or ions can be generated. Each material possesses a characteristic work function, W, so the maximum kinetic energy of a released particle will be $E_{max} = hv - W$. [188]

For the photodetectors used in optical imaging, the particles released are electrons, which produce a variable current. The sensitivity of these devices is characterised in terms of the voltage or current generated per unit of incident light power, while their quantum efficiency is the probability of generating an electron per incident photon. The other main defining characteristics are the signal rise time and the dark count (equal to the number of electrons produced per unit time when no light is incident).

Photomultiplier tubes (PMT) [189] have traditionally been the most commonly used type of light detector for diffuse optical imaging. They consist of: an evacuated glass tube where the cathode is made from, or coated with, a metal with a known sensitivity to light; a dynode chain (an electron amplifier) that amplifies the incident photocurrent by a large factor (10^5 to 10^7); and a coupling to a current amplifier. When a low flux of photons is incident, the PMT output signal consists of discrete nanosecond-width voltage pulses of varying heights. A multilevel detector is then used to discriminate events independently of height in photon counting applications. Because metals have large work functions, PMTs can require bias voltages

Chapter 3. Theoretical Background

of hundreds of volts or sometimes kilovolts to excite a sufficiently large number of electrons. Despite that, the dark current for PMTs can be significant: in the range of several nanoamperes to microamperes. The sensitivity of the photocathode to thermally emitted electrons may cause the PMT's output to emit a weak pulse even when there is no light. The intensity of this pulse depends on several factors, such as the temperature of the system and the bias voltage. The impact of these dark pulses on the count rate of the system will depend on the pulse discrimination technique – if the intensities of the dark pulses are lower than those of the light pulses, they can be eliminated using a thresholded pulse discriminator.

Other photodetectors used for communications and imaging are based on the properties of semiconductors [190, 191]. Semiconductor materials have a crystalline structure where the energy gaps between electronic levels are lower than in insulating materials. When a semiconductor is *doped* with impurities, compound lattices can be obtained such that they readily release (*N-type*, doped with *donors*) or accept (*P-type*, doped with *acceptors*) electrons in their higher energy bands when stimulated. The lack of an electron in a layer, able to accept a new electron, is called a hole and is treated as a virtual positively-charged carrier in semiconductor physics.

Electronic components like diodes or transistors consist of alternating layers of donor and acceptor doped semiconductors. The junction between P-type and Ntype layers (a PN junction) only allows current to flow in one direction from N to P unless the applied voltage is large (known as the *breakdown* voltage) when flow from P to N is possible.

Light can act as a stimulus for the electron interchange. If the energy of the incident photons is greater than the electronic gap, some electrons can be promoted to the conductance band, generating an electron-hole pair that gives rise to a *photocurrent*. Devices made with one layer of semiconductor with ohmic contacts are called *photoconductors*, while those based on donor and acceptor junctions are called *photovoltaic detectors*. The photodetectors used for imaging, as described below, are all photovoltaic.

3.3. Detectors

PIN (P-intrinsic-N) photodiodes are constructed from a layer of undoped intrinsic semiconductor (I-type) sandwiched between P-type and N-type layers. Light generates electron-hole pairs, with higher probability in the I layer. The electrons then flow towards the N layer while the holes move towards the P layer. If the diode is reverse biased, this allows electrons to flow from the N to the P layer, generating a current. The bias voltage determines the gain, i.e. the magnitude of the current per unit of photon flux. Photodiodes are fast and cheap, although not very sensitive to low light levels. For single photon counting applications, such as TD-DOT, a multilevel detection circuit like a constant fraction discriminator (CFD) has to be employed, similar to a PMT.

If the voltage surpasses the breakdown voltage, the high electric field can accelerate a carrier with energies greater than the characteristic ionisation energy of the semiconductor. The accelerated carriers then interact with the lattice causing more electrons to be liberated, leading to further interactions and so on. This is known as the avalanche effect which can ultimately destroy the junction if the structure is not designed for it. Avalanche Photodiodes (APD) are designed to exploit the avalanche effect. Operating at a voltage above the breakdown limit improves sensitivity to light as one photon event can release a very large response from the detector. The avalanche effect continues as long as the sensor is biased, so an additional circuit to stop the avalanche, a so-called *quenching circuit*, is required for single photon counting. The quenching circuit lowers the bias voltage and enables the APD to respond to the next photon. The biasing voltages needed for APDs are of the order of hundreds of volts, and the large electric fields produce high temperatures in the device. This is the technology behind SPADs (single photon avalanche diodes), which employ APDs designed for photon counting, a quenching circuit, and thermoelectric cooling in an encapsulation. The thermoelectric cooler actively regulates the amount of heat transferred out of the APD to a heat sink to maintain a constant temperature, guaranteeing a constant responsivity. For TD-DOT applications, SPADs have a faster response to individual photons than regular photodiodes and their sensitivity is higher.

Chapter 3. Theoretical Background

Improving the sensitivity and the active area of photodiodes and APDs leads to an increase in the capacitance. As the area increases more photons can be detected, but the capacitance is then larger and the detection of one photon keeps the device inactive for longer. This problem has been solved in the design of *silicon* photomultipliers (SiPM), a photodetector intended to operate like a PMT but with a smaller form factor and lower bias voltage. SiPMs are arrays of small APDs (up to thousands) joined together in a network that share the same bias voltage and the same output. When a photon arrives, it excites one of the photodiodes, leaving the rest in a responsive state. The signal from the many APDs is summed in the single output, giving a multilevel signal similar to the one generated by PMTs. The greatest advantage of SiPMs when compared to SPADs is that while the active area of the device is equivalent to that of an APD, the internal capacitance of the individual photodiodes is smaller and each one operates independently. This gives SiPMs greater speed with lower bias voltages - of the order of tens of volts. Manufacturing costs of SiPMs have been reduced, making them currently less expensive than SPADs. The biggest disadvantage of SiPMs for photon counting, as required for TD-DOT, is that a CFD is needed as is the case for PMTs.

During this project photodiodes, SPADs and SiPMs were all evaluated as candidate technologies for the detector in a portable time domain imaging system. The selection of components is described in the next chapter.

3.4 Light sources – laser pulse generation

For optical imaging, the wavelength and power of the chosen light source determine important criteria such as tissue contrast, penetration depth, volume of interrogated tissue, acquisition time, and the safety of the technique. For most diffuse optical imaging applications, tissue is illuminated at two or more discrete wavelengths within the NIR diagnostic window. For time domain imaging, a source of short pulses of NIR light are required, ideally with a pulse width no greater than a few picoseconds, and with an average power of several milliwatts consistent with max-
imum permissible exposure to the skin [171, 15].

A *laser* (originally an acronym of Light Amplification by Stimulated Emission of Radiation) is a system capable of generating a coherent field of light at a high intensity. The first patent for a laser was filed in 1957 with development starting in the 1960s and its basic structural elements have remained the same. A laser consists of three main components: a *medium*, that can be gaseous, solid or liquid, where the light amplification will take place, a *pump* to excite the molecules in the medium and start the lasing process and a structure of *optical feedback* that determines the number of times that the light passes through the cavity before being emitted. This structure is analogous to that of an electronic oscillator [192, 193].

For the lasing effect to take place, the pumping process must excite the medium to the point in which more of its atoms or molecules are in an excited state (in a higher quantum energy level) than in a base state. This is known as *population inversion*. Once inversion is achieved, the radiation that passes through the medium and matches a range of frequencies characteristic of the material is amplified. From this point the amplitude and timing of the output beam will follow the variations of the input pumping source. By adding an optical feedback loop using mirrors and lenses the interference inside the medium cavity can make the system oscillate. The output of this system will be a highly directional, monochromatic beam.

There are two types of laser, depending on the nature of the oscillation: continuous wave lasers, where the emission is continuous, and pulsed lasers, where the modulation in the input pump or the light interference inside the cavity (phenomena like *mode locking* and *chirping*) cause the output to be formed of narrow packets of light. For TR-DOI applications pulsed lasers based on solid state technology, specially semiconductor lasers, are often used. As explained in the previous section, in a PN junction the electron-hole dynamics can be stimulated by light. The reverse phenomenon can also occur: the recombination of electrons and holes generates the emission of a photon with the wavelength depending on the distance between energy levels in the material. For laser emission to be efficient, the crystalline structure of the diode must allow for a precise placement of the electrons and holes. The diode crystals are usually *grown* in layers of silicon under a specific structure and then cut to size.

The first time resolved systems used Ti:sapphire lasers, which are bulky systems consisting on a titanium-doped sapphire crystal pumped by another laser that provided a high average power and MHz pulse frequencies. Portability requirements stimulated a move toward fibre lasers and diode lasers (side-emitting diodes and VCSELs) [15].

In a fibre laser the non-linear effects introduced by an optical fibre, like dispersion and scattering are exploited to modify the amplitude and the frequency components of the light, generating laser pulses at the desired frequency bands from a pre-existing pump. One special type of fibre laser is the supercontinuum laser [194], that utilises the interference phenomenon known as *four wave mixing* to generate short pulses of light (*solitons*). Over time, the solitons will cover a series of wavelengths spanning a wide range of the optical spectrum. These laser systems are more portable than Ti:sapphire lasers, and allow easy wavelength tunability by means of a filter at high average power. Their main drawbacks are the poor uniformity of the spectrum and the robustness and maintenance of the gain medium.

Diode lasers are generally side-emitting semiconductor crystals. They can be electronically driven to produce a train of pulses at a high frequency of the order of megahertz. Because of their small size and low cost they have been part of time domain systems either as pumps for fibre lasers or applied directly to the sample. However, they cannot be driven at a high power (tens of milliwatts or more) without broadening the pulses and their size is often too large for high density systems.

A VCSEL (vertical cavity source-emitting laser) is a particular type of diode laser that presents an alternative to side emitting diodes. Their structure enables the light emission to occur in the top layer of the semiconductor crystal [193]. The resulting beam has a narrow spectral width and is more focused, with a profile (circular, rather than elliptical) which can be more easily coupled to another device. Although the average output power of a VCSEL is limited (up to a few milliwatts), their structure lends itself to be incorporated in arrays or high density systems and is currently regarded as a key component for future portable time domain instrumentation.

The experiments presented in this thesis employed two different types of laser. A single-wavelength fibre laser was used for initial system validation and characterisation, while a supercontinuum laser was used for phantom imaging experiments, which benefited from additional output power and the facility to select the wavelength. Although the design of a miniaturised laser source was not an objective of this project, it is important to note that the use of laser diodes would greatly improve the portability of an imaging system and they are already being employed by some experimental time domain systems. Thus, future work to achieve a wearable system will inevitably involve the introduction of laser diode sources.

3.5 Time to digital converters

A Time-to-Digital Converter (TDC) is a circuit that assigns a digital code to the time at which an event is detected [195, 155]. This type of circuit is used to measure the time interval elapsed between two signals, usually known as *start* and *stop*, with very fine resolution, such as the times of arrival of photons with respect to the source laser reference. When the measured interval of time is very short (less than 1 ns) or needs a fine, sub-nanosecond precision, sampling the number of cycles of a simple oscillator is not feasible or efficient, as the frequency would need to be of the order of one terahertz. Several solutions have been developed to address this problem. The first were based on a two-step conversion with a time-to-analogue converter that translated a time interval into a charging voltage, followed by an analogue-to-digital converter. More recently, advances in microelectronics design have given rise to two faster TDC structures: one based on phase-locked loops (PLL) and another one based on delay chains, valued for their ability to be implemented in fixed silicon processes and in reconfigurable hardware.

The *Vernier oscillator* architecture uses two free-running oscillators with different frequencies in a mechanism similar to a chronometer (see Figure 3.3). The start and stop signals activate one oscillator each. A phase locking circuit is used to detect when the phase of the two oscillators has matched. By measuring the number of oscillator cycles for both frequencies, a fine and discretised time measurement can be obtained as the difference between the two measurements $(n_1t_1 - n_2t_2)$ in the figure). This architecture needs just a small number of components and the interpolation circuits are simple. However, the oscillators need to be reliable, stable, and well characterised. When based on reconfigurable hardware, the oscillators and PLLs are not trivial to implement or simulate in design software and a standard test strategy does not exist [150, 3].



Figure 3.3: Timing diagram of a Vernier oscillator TDC, adapted from [155]. The events are represented by the low to high transition in the level of the signal. The red lines mark the moments the status of the delay line is sampled.

For the delay chain approach, called a *tapped delay line*, the event or start signal is replicated through a chain of delaying elements (see Figure 3.4). The stop signal is used as a clock to sample the delay elements. This way, the event signal is sampled at many discrete points in time, with a resolution that depends on the delay of the individual elements rather than the sampling clock. If each delay element delays the signal by a constant time t_d , the delay line discretises the time of arrival as $n_{detect} = \lfloor t_{event}/t_d \rfloor$, where n_{detect} is the number of measured delay elements. If a periodic stop signal is used, the total delay of the delay elements will cover a fraction of the clock period, which means that by employing the number of elements with the equivalent delay of one period is enough to cover all possible start-to-stop

times of arrival. This architecture is simple to implement and test, especially in reconfigurable hardware, because the delay can be intrinsically provided by any silicon component's propagation delay and the registered signals provide debugging and test information. However, the delay provided by the individual elements varies due to the manufacturing process, and is also usually dependent on temperature and on the internal wiring. Thus, delay line TDCs are more prone to non-linearities. Many strategies have been developed to calibrate for the timing variability, to correct the non-linearities, and to improve the timing resolution. The most common are periodically measuring the resolution to adjust the digital codes, applying correction factors, as explained in Section 3.6, and adding elements to the base structure, like fixed routes, logic gates between outputs or extra stages with registers.



Figure 3.4: Timing diagram of a 4-element delay line. The events are represented by the low to high transition in the level of the signal.

A variation on the previous scheme, the *delay locked loop* was proposed in the 1990s [196], which uses a bias voltage regulated by a PLL to control the delay line. However, this architecture is not usually seen in modern applications.

Chapter 5 presents some implementations of a tapped delay line TDC on an FPGA with different architectures to detect and store photon events, and a description of experiments to evaluate their performance. Meanwhile, chapter 6 describes

an imaging system based upon a commercial oscillator TDC.

3.6 Digital converter characteristics and error correction

The behaviour of the TDC is analogous to that of an analogue-to-digital converter (ADC) in that a continuous physical quantity is discretised and then encoded. The minimum variation of the input that produces a variation in the output adds one unit (1 bit) to the code, so the precision and the conversion errors are often expressed in terms of the least significant bit (LSB). The ideal TDC works under two main assumptions:

- 1. The codes or bins are monotonically increasing in time. This means that a longer time should always be encoded after a shorter time and that there are no crossings or changes of order.
- Codes or bins are of the same width, or span the same range of arrival times.
 For a uniform source of arrival times, all bins should add up to the same value.

The performance of the digital conversion for this kind of device can be characterised in terms of three types of error [197, 198, 199]:

- 1. Offset errors: a static error that shifts the origin of the conversion curve. It is uniform for all the codes.
- 2. Differential non-linearity (DNL) error: an error related to the step size for each code, that is the variation of the input that causes each code to be generated in the output. It is defined as the difference between the actual step size and the ideal size (1 LSB/step). For each step or bin *k* starting at time t_{k-1} and ending at time t_k , the bin width is $\Delta t_k = t_k t_{k-1}$ and the DNL is calculated as

$$DNL(k) = \frac{\Delta t_k - \Delta t_{\text{ideal}}}{\Delta t_{\text{ideal}}} = \frac{t_k - t_{k-1}}{\Delta t_{\text{ideal}} - 1}.$$
(3.2)

3. Integral non-linearity (INL) error: an error related to the DNL, caused by the cumulative effect of uneven step sizes, and which describe the deviation of the input/output curve from a straight line. It measures how monotonical is the increase in bin time and how steep it is relative to the starting time, t_0 . It is calculated as

$$INL(k) = t_k - \frac{t_0 + k\Delta t_{\text{ideal}}}{k\Delta t_{\text{ideal}}} = \sum_{r=0}^k \frac{DNL(r)}{k}.$$
(3.3)

The errors related to the fulfillment of the second condition of ideal operation, the DNL and INL errors, derive from differences in the fabrication process or imprecisions in the oscillators, that cause the bins obtained from the TDC to be of uneven widths or intermediate detection steps to be absent (*missing codes*). Some of these variations can also be caused by changes with time, temperature or electrical conditions such as the input voltages.

The designers of TDCs often incorporate methods to monitor such variations while measuring signals and actively adjust the code generation to reduce the errors. In the case of the more static causes of DNL, an offline correction method may be applied. Once the DNL is measured for each bin, k, some investigators (e.g. Dutton et al [164] and Chen et al [165]) suggest that a correction factor can be calculated per bin such as

$$C(k) = \frac{1}{1 + DNL(k)}.$$
(3.4)

When applied, the differences in bin size are smoothed out. A *code density test* can be used to measure the DNL. In this test, a series of input signals with a random uniform distribution are applied to the converter [197]. For a TDC this distribution can be obtained from two signals with uncorrelated frequencies or from a source of noise, such as stray light on a photodetector. Ideally, and with a sufficiently large sample size, this results in a uniform histogram at the output. In reality, however, the resulting histogram will show peaks oscillating around a mean (approximately the gain of the ideal converter) due to the nonlinearities. This cumulative histogram can be normalised by the mean to isolate the nonlinear effects and ensure that the calculation is independent of the time range and of the LSB size. Ideal bins now should have a value of 1, while non-ideal ones fall above or below. By subtracting 1, the difference becomes apparent and the DNL falls between -1 and infinity [200]:

$$DNL_{\text{normalised}} = \frac{hist}{av(hist)} - 1.$$
 (3.5)

This correction method was implemented in the data acquisition software for all the systems described in the following chapters.

3.7 FPGA and hardware design

An FPGA (Field Programmable Gate Array) [201, 202] is a re-configurable integrated circuit made up of many discrete hardware blocks, each one capable of performing one of many different logic functions, that can be configured and interconnected to generate a more complex hardware design. This results in a system that, while being outperformed by customised silicon integrated circuits, provides access to parallel, fast and flexible electronics for development and testing of architecture designs.

An FPGA's internal structure or *fabric* is composed of hardware blocks called *logic cells* (LCs) by Xilinx or *logic elements* (LEs) by Intel/Altera, grouped together in a hierarchy. In each of these logic cells the following elements can be present:

- A look-up table (LUT) that stores the output values expected for any combination of inputs. This is the most basic block in the FPGA design logic, and can be used to act as simple logic gates or similar stateless components. For some FPGAs, LUTs can also serve as small memories or shift registers.
- A flip-flop (FF), or a chain of flip-flops, to store data or the state of a circuit.
- Small combinatorial circuits. In Xilinx FPGAs, for instance, LCs include a small chain of four one-bit adders with carry called CARRY4.

By configuring the values in the LUTs and the connections between the elements inside an LC, a small logic block is created. Then, for more complex structures LCs can be routed together. For fabrication, LCs are grouped in structures called slices and CLBs in Xilinx terminology. This is meant to improve the performance of cells that work together. Additionally, inside the FPGA fabric, manufacturers include other ready-to-use hardware circuits called *macros* such as arithmetic units, communication and DSP blocks, clock PLLs, and fast block memory (block RAMs or BRAMs).

Hardware design involves creating a logic scheme that solves a particular problem and then identifying a physical layout of components that behave like the logic scheme. For complex designs, this becomes a time-consuming optimisation process that is usually handled by a software package. The workflow from problem capture to physical implementation is standard across manufacturers and only the terminology and optimisation algorithms change depending on the choice of FPGA.

The behaviour of the system is described either with a hardware description language (HDL), like VHDL (used for this project) or Verilog, or using GUI circuit capture. The first approach is more flexible and commonplace, as a detailed, device-specific design can be achieved without needing a complex editing program. For HDLs, the behaviour of the system is described using logic interactions between binary signals (or groups of signals), processes and state machines. Thus, the hardware designer can specify both stateless behaviours – that is, functions that only depend on the input values – and behaviours with state, where the logic function changes with time or succession of events. Behavioural blocks, called components or modules depending on the HDL, can be used as black boxes hierarchically to facilitate easier design and debugging. This contrasts with software design in that every behavioural block is supposed to be responsive at the same time. The final design will be parallel, instead of sequential (although the internal behaviour of an individual component could be sequential). HDLs also allow the designer to introduce device specific constraints, so manipulation of LCs and some interconnections is possible at this stage.

The high-level code is then analysed by a *synthesis* program that divides the behavioural structure into individual logic gates and functions. This yields a still

high-level design which consists of many interconnected logic blocks (such as multiplexers, memories, logic gates, etc.).

Thereafter, in a third stage called *implementation* or *placement and routing*, the synthesised circuit is converted into an equivalent combination of logic cells and macro cells and the software algorithms determine the optimum locations within the FPGA fabric where the routing between elements is consistent with the timing and performance constraints. From this result, a bitmap mapping the configuration can be generated and programmed into a flash memory, from where the FPGA will be programmed when turned on.

The workflow described above, illustrated in Figure 3.5, is supported by simulation tools that can be used at any stage in the process to verify/test the design with various degrees of realism.



Figure 3.5: Workflow for hardware design in FPGAs (adapted from [201]).

The flexibility of the FPGA allows for the implementation of cores capable of executing arbitrary code, sometimes based on the behaviour of silicon microprocessors, often referred to as *soft processors* or *soft cores*. A combination of sequential code execution and dedicated hardware can reduce the complexity of a design and ease the maintenance and debugging of the *glue* components, those which provide coordination between functional units, while retaining control over the parts of the design that require processing efficiency. Some manufactures offer soft microprocessors optimised for their own architectures, such as Intel/Altera's Nios II [203]

and Xilinx's MicroBlaze [204] (more complex, C compatible but more resource consuming) and PicoBlaze [205] (simpler, less powerful, but more lightweight) to-gether with compilers, assemblers and detailed programming guides.

3.8 Thresholding and fraction discrimination

When attempting to determine the time at which an electronic pulse occurs, it is often inappropriate just to measure the time at which the signal crosses (i.e. exceeds or falls below) a single threshold. As exemplified in Figure 3.6, if the pulse amplitude varies over a wide dynamic range, the time at which a threshold is crossed could correspond to any point within the pulse profile, or might not detect the pulse at all. This phenomenon is known as *time walk* and is characteristic of some light detectors like PMTs and SiPMs. However, if the threshold is set to a fixed fraction of the pulse amplitude, as illustrated on the right diagram in Figure 3.6, the dependency of the timing on the variation in pulse amplitude is almost eliminated.



Figure 3.6: Time walk in signal thresholding (left) and constant fraction thresholding (right)

Designing a circuit around this concept is therefore more complicated than a fixed threshold discriminator and a compromise should be made between histogram spread due to time walk and circuit development time.

A constant fraction discriminator (CFD) is a circuit that adaptively thresholds a signal and emits a pulse when the signal has crossed a threshold determined in relation to its maximum value. The basic structure of a CFD, as shown in Figure 3.7, consists of a comparator that detects the zero crossing between an inverted and attenuated version of the signal and a delayed input signal. The comparator removes the dependency on the signal amplitude, such that the time of crossing depends only on the ratio between the two branches, the constant fraction, f and the delay, t_d .



Figure 3.7: Diagram of the base structure of a CFD, showing the signal shape on each branch.

Let us suppose that the leading edge of the input signal is a linear ramp, $V_{in}(t) = At$. The inverted and attenuated signal is therefore given by $V_a(t) = -fAt$, (where f < 1) while the delayed signal is given by $V_d(t) = A(t-t_d)$. Thus the zero crossing time, when the sum of both signals is zero, will be given by:

$$V_a(t_{\text{crossing}}) + V_d(t_{\text{crossing}}) = -fAt_{\text{crossing}} + A(t_{\text{crossing}} - t_d) = 0, \qquad (3.6)$$

$$t_{\text{crossing}} = \frac{t_d}{1 - f},\tag{3.7}$$

which is independent of the amplitude.

The zero-crossing that activates the output of the comparator occurs when the attenuated signal reaches its maximum level. For the desired fraction to be obtained, the delay time has to be proportional to the maximum. Thus, for a signal with a given characteristic rise time t_{rise} , the optimal delay time to achieve a fraction f is given by $t_{d,opt} = t_{rise}(1 - f)$. This means that the design of a CFD depends on the shape of the input signal and will work optimally only for such signals. Time walk or timing errors can occur if the input signal has a shape different from that expected, and also if there is significant noise in the pulse train or DC shifts on one of the branches. A broad dynamic range and the choice of t_d have been

demonstrated to influence the jitter in the signal in different ways. CFD design considerations will be explored in Section 4.2.2.

3.9 Time bases and calibration

The output TPSF obtained by any TD system can be modelled as consisting of an ideal TPSF, which contains information related to the optical properties of the medium, modified by the impulse response function of the system, which includes the temporal characteristics of the source, the detector, and all intermediate components of the system, such as the optical fibres [206]. For a system consisting of multiple sources and detectors, a measured TPSF $M_{n,m}(t)$ using source *n* and detector *m* can be represented by:

$$M_{n,m}(t) = D_{n,m}(t) * I_{n,m}(t) + f_{n,m}(t), \qquad (3.8)$$

where $D_{n,m}(t)$ is the object TPSF that we intend to measure, $I_{n,m}(t)$ is the impulse response function or IRF of the system, $f_{n,m}(t)$ is the background stochastic noise, and * represents a convolution. The IRF of the system includes all the effects mentioned earlier and is different for each source-detector pair. This means that the mean times derived for different sources and/or detectors cannot be compared until an appropriate calibration is performed. The effect of the additive background noise can be minimised with techniques that are easy to implement, first by reducing its influence (shielding the device, eliminating background and stray light) and second by simple data processing techniques, such as windowing the TPSF. Removing the influence of the convolution is more challenging. In principle, if the value of $I_{n,m}(t)$ is known, the measurement can be deconvolved to obtained the desired TPSF. A method was developed for performing this calibration for the data acquired using MONSTIR [207]. It requires additional measurements to be acquired as follows;

• A so-called source calibration, S_n , which accounts for the length and dispersion differences between sources fibres. To obtain it, the full set of source fibres are placed in a ring around a transparent cylindrical block, in the centre

Chapter 3. Theoretical Background

of which is a small target coupled to a single detector. All the source fibres are illuminated in sequence, and the detector measures a TPSF in each case. This calibration measurement needs to be recorded just once, unless any changes are made to the source fibres or laser optics.

• A so-called absolute calibration, *A_n*, which corresponds to a measurement acquired when source *n* is coupled directly to detector *n*. This can be achieved by placing the ends of the fibres in contact. For the 32-channel MONSTIR system, each optode consists of a co-axial source fibre surrounded by a detector fibre bundle, and thus these measurements were approximated by measuring the specular reflection from the surface of the object. Because the temporal characteristics of detectors change over time due to drift, etc., it is necessary to acquire this absolute calibration measurement at regular intervals, and as close in time to the object measurement as possible.

The calibrated TPSF, $D_{n,m}$, is obtained from these measurement using

$$D_{n,m}(t) = (M_{n,m}(t) * S_m(t)) \circledast (A_m(t) * S_n(t)),$$
(3.9)

where \circledast denotes deconvolution.

This method was inspired by an approach reported by Eda, *et al* [208], and has been used and refined extensively for measurements using both MONSTIR systems at UCL [57].

As observed by Elizabeth Hillman in a paper [209] and later her PhD thesis, when calculating calibrated measurements of moments (e.g. mean time, variance, skew, etc.) it is not necessary to perform a full deconvolution. The convolution between two functions results in the sum of their corresponding moments. Therefore for mean time:

$$\langle t \rangle_{M(t)} = \langle t \rangle_{D(t)} + \langle t \rangle_{I(t)}.$$
 (3.10)

Thus, to obtain the mean time of the deconvolved signal, it is sufficient to subtract

the mean time of the calibration IRF from the that of the measured TPSF:

$$\langle t \rangle_{D(t)} = \langle t \rangle_{M(t)} - \langle t \rangle_{I(t)}. \tag{3.11}$$

For the experiments reported in this thesis, the IRF was measured for each source-detector pair by physically placing the output fibre in contact with the detector fibre, with a thin, highly-attenuating ND filter sandwiched between them. The filter ensured that the detector was not saturated. However, it is noted that this method only works efficiently when the number of pairs is low, and it would not be appropriate for a system with a large number of sources and detectors.

3.10 High frequency design

When designing circuits aimed at operating at high frequencies (of the order of a megahertz and above) certain effects that are typically ignored in low frequency design need to be taken into consideration to guarantee that the output signals are not distorted and will preserve their stability, ensuring their compatibility with other circuits and their timing is preserved. Several design techniques exist to ensure these requirements are met. The following summary of possible sources of distortion draws information from personal experience and that of people I have collaborated with to produce the designs for this project, books on electromagnetism and high frequency analysis, such as Ulaby and Ravaioli [210] and Pozar [211] and design websites with the personal and technical experience of the authors, including EDN [212], Ian Poole's electronics notes [213] and Microwaves 101 [214].

One of the sources of distortion stems from the characteristics of the signal routes. Mismatches in impedance along a transmission line, in this case a circuit trace, cause the power delivered to a load to be below the optimum and cause fractions of the signal to reflect and travel in the opposite direction, producing distortions. These effects become significant when the length of the transmission line and the wavelength of the signal being processed are of the same order. As a rule of thumb, transmission line effects should be considered if the line length between source and load is 1/10 of the signal wavelength [212]. For instance, for a 40 MHz signal, $\lambda = c_{\text{copper}}/f = (2 \cdot 10^8 \text{ m/s})/(40 \cdot 10^6 1/\text{s}) = 5 \text{ m}$, so lengths of the order of 0.5 m, a length easily reached by the combination of PCB traces and cables, can already be considered transmission lines. The rise and fall times of the pulses expected in the current application are in the order of one nanosecond, bringing the maximum frequencies to handle close to the gigahertz.

Three components determine the properties of a transmission line and its characteristic impedance: a conductor plane, a ground or return plane, and a dielectric, which is non-conductive. Depending on the relative position of these elements, different topologies exist, as shown in Figure 3.8. The most common types occurring in PCB design are those involving one strip of metal against a back plane, mainly the microstrip. Other factors that affect the characteristic impedance are the width and thickness of the planes and the material of which they are made. The value of the characteristic impedance of a line can be estimated using one of many numerical models of the electromagnetic transmission. These are often available as software calculation tools. The calculation tools consulted for this project were found to differ up to 1 Ω in their estimates, which was considered an acceptable error.



Figure 3.8: Diagrams of types of transmission types, taken from Ulaby and Ravaioli [210].

A second source of noise is inductive and capacitive coupling between elements of the circuit like pins, pads, traces and the ground plane. These parasitic effects can increase the transition time of the signals or introduce extra pulses or ringing. To avoid them, all the pads should be soldered flat on their traces, the circuit should be cleaned of any soldering flux, and the ground plane should be as continuous as possible, without large interruptions by any other trace. In cases where there are more than one layer acting as a power plane, the vias connecting the layers should be as close to the target component as possible.

A third source of noise which must be considered is electromagnetic interference. Depending on the size of the traces and the range of frequencies it is expected that parts of the circuits will act as antennae, either picking up signals from external sources or emitting radiofrequency signals. For this reason, during this project all of the circuits were isolated using aluminium boxes, acting as Faraday cages, and the coaxial connections were kept as short as physically allowed.

Chapter 4

Design of Optical and Electronic Signal Conditioning

4.1 Components and prototyping platforms

In this section, the characteristics and rationales behind the choice of components and test elements will be explored.

4.1.1 Light source

During this project, the focus of the design was on the light detection, signal processing and TPSF building, using reliable and stable light sources with known characteristics. Three laser sources were employed, each providing picosecond pulses at near-infrared wavelengths to illuminate the sample and a reference synchronisation output at the laser's pulse repetition frequency.

Initial tests were conducted using a bespoke IMRA A-70 pulsed fibre laser, emitting 780 nm pulses of picosecond duration at a repetition rate of 40 MHz. The laser module provided access to an optical fibre output with maximum power 70 mW and a NIM synchronisation output that generated a narrow pulse for each optical pulse. This unit was used for the first tests using photodiodes, but it ultimately failed during tests and had to be replaced with other laser sources.

The first of these was an IMRA Femtolite A-15 (780 nm) fibre laser, which

92 Chapter 4. Design of Optical and Electronic Signal Conditioning

provides 180 fs-wide pulses at a specified mean power of 15 mW (experimentally, closer to 20 mW) and with a repetition rate of 48 MHz. To obtain a synchronisation signal, the output from the laser was split using a reflective neutral density filter with a high attenuation ratio $(10^3 \text{ or } 3 \text{ OD})$. The transmitted component provided an eye-safe laser beam for free-space tests, while the reflected component is fed into an external photodiode (Thorlabs DET025). An enclosure around the laser constructed with blackout cardboard ensured that any stray or reflected beams from the laser are absorbed and only the eye-safe output exits via a small aperture. The enclosure was later extended around the detector.

For later fibre-coupled testing on phantoms, the laser integrated within MON-STIR II was used. This system provides switchable fibre couple outputs and a synchronisation signal, all controllable via software installed on MONSTIR II's console. The laser source is a NKT SC400 supercontinuum laser, which operates with a central wavelength of 1100 nm and a spectral bandwidth from 410 to 2400 nm. The laser produces pulses of approximately 6 ps duration at a 40 MHz repetition rate. The output from the laser is coupled via a polarizing beamsplitter into two acousto-optic tuneable filters (AOTFs, model 2986-01, Gooch and Housego), each acting on one polarization state. The AOTFs have a tuneable wavelength range of 640-1100 nm, and a transmission bandwidth of approximately 5 nm. After filtering, the beams from the two AOTFs are recombined. Light lost via the higher order Bragg diffractions from one AOTF is used to illuminate a photodiode which provides a synchronisation signal. The output laser beam is optically switched into any one of 32 fibres, configurable with the console's software. MONSTIR II also includes PMT photodetectors and photon counting units, which were not used for this project.

4.1.2 Light detectors

Three types of sensor have been investigated for the prototype systems: photodiodes, SPADs and SiPMs.

A Roithner Lasertechnik QL78F6S laser diode [215] was used to test the initial

timing systems. This is a package that includes a laser diode and a monitoring photodiode operating in the infrared, designed to check the operation of the laser diode. If the laser pin is ignored and the monitoring photodiode is reverse biased, it can also be used as a standalone detector with sensitivity in the infrared. This way the QL78F6S can be used alternatively as a detector or as a laser source for quick prototypes or early stage experiments.

An Excelitas Technologies SPCM-AQRH-10-FC SPAD [216] was chosen to test the histogram functions in the FPGA system during the initial functionality tests with the IMRA A-70 system. It is a SPAD module incorporating a detector, an internal power supply and a TTL output, that simplifies interfacing with other electronic components. It generates a pulse every time a photon is detected. The SPAD then needs to recover its base state, resulting in a dead time of 29.3 ns (as reported in the calibration sheet) when no photon can be detected. Under normal use for imaging, this should not present a problem, but for free-space testing this means that the usual pulse rate of the laser (above 40 MHz, at least 25 ns) will most likely saturate the SPAD and events will be missed, so a proper attenuation factor should be added. The SPAD module was later replaced by a SiPM because the SPAD showed an unexpectedly large variation in the output timing (FWHM of the IRF much larger than 1 ns) that made time sensitive experiments more difficult to perform. A SiPM was selected for the most complex timing experiments, because it combined a good sensitivity to light intensity with low jitter and the potential to be used in PCB designs.

The selected SiPM was an On Semiconductor/SensL MicroFC-10010 [217] from the C-Series, available on a test board fitted with the necessary passive components and SMA connectors. The test board provides two outputs: a fast output (used for experiments) which provides unamplified fast pulses (rise time of 0.3 ns and pulse with of 0.6 ns) through a capacitor and a standard output connected to the anode. At the time of purchase, it was one of the newest additions to SensL's catalogue and was later carried over to On Semiconductor's catalogue after it acquired the company, and thus provides a state of the art performance. Although its

central wavelength lies in the blue region, it maintains an adequate sensitivity in the near-infrared range of around 5%, competitive with PMTs. Its maximum dark count rate is 30–96 kHz at 2.5 V overvoltage and 21°C.

All detectors tend to saturate at high light intensities, when too many photons arrive at the same time. If more than one photon is detected in one duty cycle, the pile-up will cause a distortion in the resulting histogram. This is easily solved by reducing the light intensity using optical filters. A rule of thumb, explained in Florian Schmidt's doctoral thesis [28], is to adjust the intensity so that only one photon event is detected for every 1% of the laser pulses. This makes the probability of detecting two photons almost negligible, $(1\%)^2 = 0.01\%$. Other limits may be imposed by the rest of the components. The safe limit in counts per second for the lasers used in this project is 400,000 (for the 40 MHz laser) and 480,000 (for the 48 MHz laser). A set of neutral density (ND) filters with fixed attenuation and a rotating variable ND filter were used to regulate the intensity during experiments. The attenuation in these filters is expressed in a logarithmic scale as $A = \log_{10}(\Phi_i/\Phi_o)$, where Φ_i and Φ_o are, respectively, the input and output flux. The attenuation is sometimes referred to as *optical density* or OD as well.

4.1.3 Event detection circuits

To achieve pulse discrimination for sensors with multi-level outputs, such as photodiodes and SiPMs, a constant fraction discriminator (CFD) is desirable (although a fixed threshold comparator was designed later in the project when the range of variability in the SiPM output was known to be constrained). For the initial setup, a commercial Ortec 9307 pico-Timing Discriminator [218] was used. This CFD includes an adjustable threshold to distinguish event detections from noise, and provides both TTL and NIM outputs. However, the TTL output is slow, so the NIM output had to be manually converted into a signal compatible with the other electronic components.

For my final experiments, a custom circuit board was built including both the comparator and the voltage level correction to interface the synchronisation and detector to the TTL components.

Throughout this project, a TCSPC (time-correlated single photon counting) card from Becker & Hickl GmbH, model SPC-130 [219], was used as a reference to validate the timing characteristics of the measured histograms recorded with the prototype devices. The card is connected to a PC via a PCI slot, and bundled software enables time of flight histograms to be acquired with picosecond resolution.

4.1.4 Development boards

The TDC and the histogram generation were implemented on a Nexys 4 DDR board from Digilent [220]. This entry-level board includes a Xilinx Artix-7 FPGA with access to features useful for educational purposes and testing, such as LEDs, switches, a serial USB connection and generic input/output ports. It provides a middle-end FPGA in an environment that is suitable for debugging and experimentation. The board is programmed using the Xilinx Vivado IDE [221], which is available free of charge for Artix-7 targets.

In addition to the FPGA-based TDC implementation, an off-the-shelf TDC was also evaluated in the final system design. The Texas Instruments TDC7200 [146] was chosen for this purpose. It operates at standard TTL levels for the inputs and outputs, so any design compatible with the FPGA will be valid for this chip as well. The TDC is managed via an SPI interface, which allows the master system to initiate measurements, configure parameters, and read out the measured times. The manufacturer also provides a test board on which the chip is mounted, which includes an oscillator to provide a time base for the TDC, impedance-matched SMA connections for the inputs, and access to all pins from the chip using pin headers. An adaptor PCB with PMOD connectors was printed to connect this board to the Nexys 4 inputs.

Finally, the signal conditioning was implemented using a combination of Minicircuits RF blocks and a bespoke comparator PCB. The comparator design is described in detail in Section 4.2.2.

The system including the initial conditioning RF blocks, the comparator and



(a) Connected boxes. From left to right: RF amplifier boxes, comparator box and FPGA box.





(c) FPGA and TDC7201 (open box).

Figure 4.1: Boxes used for housing and shielding the components of the imaging system.

the TDC can fit in a set of four boxes. The RF amplifiers are housed in two $110 \times 190 \times 60$ mm boxes (one for the sync and one for the detectors), the comparator board fits in a $120 \times 120 \times 58$ mm box and the FPGA (with the TDC7201 evaluation board if it is used) is housed in a $150 \times 260 \times 70$ mm box. A picture including the boxes is shown in Figure 4.1.

4.2 Design of the signal conditioning system

Signal conditioning circuitry is required between the sensors and the time-to-digital converter, serving as a bridge between the optics and the time-resolved logic. The main requirement is to pick up the narrow spikes generated by the two (probe

and synchronisation) detectors, and generate pulses that the logic components can recognise without missing events or generating excessive noise events, while maintaining a faithful time relationship between the two signals. During the project, two types of conditioning circuit were designed, one using only commercial block amplifier and bias circuits and later a custom design on a printed circuit board. The characteristics of these systems are described below.

4.2.1 Characteristics of the input signals and design requirements

The characteristics of the signals from the various photodetectors and the laser synchronisation output were as follows:

- The signals emitted by the detector photodiode, the SiPM, and the photodiodes used to provide a sync signal were short pulses (width < 5 ns) with variable voltage level (around 1 to 2 mV for the SiPM and around 10 mV for the photodiodes). All required a 50 Ω termination load and amplification to facilitate thresholding of the signal.
- The synchronisation signal output from the original IMRA laser and the output from the CFD followed the NIM standard: 0 V to -1 V.

To be compatible with the inputs of the FPGA and Texas Instruments TDC7201 chip, the signals needed to be converted to 3 V peak-to peak (p-p), with a minimum of 2.1 V to generate a low-to-high transition, on a 50 Ω load.

With respect to the form factor, the objective of the first prototype circuit was merely to provide a functional signal to test the hardware; for the second system, my aim was to build a circuit out of easy to find and affordable integrated circuits that could fit on a single board, and ideally on a small 'tile', that could be externally adjusted for variations in the noise floor and the threshold level.

In terms of performance, a set of ideal and minimum requirements was captured for the proposed system, as summarised below in tables 4.1, 4.2 and 4.3. The design of a source was given lower priority than the detector, as the system prototypes could be tested with an external source and the development status on laser diodes and VCSELs did not guarantee that the design requirements for the system would be met within the available time.

	Ideal	Minimum	
Captured data	Temporal histograms	Mean time $\langle t \rangle$	
Bin size	<1/2 temporal resolution		
Temporal resolution	<100 ps	350 ps	
Max. photon count rate	>10 ⁵ photons/s	10 ⁴ photons/s	

Table 4.1: Performance requirements for the detector and TDC

Table 4.2: Performance requirements for the source

	Ideal	Minimum
Source location	Directly mounted on tile	External, fibre-coupled
Wavelengths	Multiple (760 – 860 nm)	Two
Average power	Several mW	0.5 mW
Pulse width	<100 ps	300 ps
Pulse repetition rate	80 MHz	10 MHz
Source switch rate	10 Hz	Any

 Table 4.3: Performance requirements for the probe system

	Ideal	Minimum
Measurement geometry	On tile and between tiles	On tile only
Data acquisition speed	10 Hz	Any
Number of tiles	8	2
On-tile S-D separation	30 mm	25 – 35 mm
Imaging	Real time	Post-processed

4.2.2 Design iterations and refinement

Initially, a circuit based on several operational amplifiers, bloc RF amplifiers and bias-tees in combination with a CFD (Ortec 9307 pico-Timing Discriminator) were used in the circuit layout shown in Figure 4.2. The CFD was used to apply an adaptive threshold to the output of the photodiodes and the SiPM, reducing the time walk effect, as the FPGA and other circuits do not admit inputs with a variable

amplitude. Except for the laser sync signal, which is already in the NIM range, each branch uses an amplifier to set the amplitude of the signal into an acceptable range – 50 to 100 mV p-p – and a bias tee shifts the NIM signal into the LV-PECL range (1.6 V to 2.4 V) prior to a pulse shaping stage. In this stage, which has a common design on the sync and detector lines, the input pulses are stretched from the original duty cycle of about 16 %, which is not detected by most TTL and CMOS logic. The input signal triggers the clock of a flip-flop with a feedback loop that connects the positive output to reset; after a clock pulse the output will be active while the feedback output has not yet triggered the reset, thus the delay of the feedback loop – effectively the length of the cable ℓ , as $t_{delay} = \ell/(\frac{2}{3}c_0)$ – fixes the pulse width. Finally, the signal is amplified and shifted to the desired range. The main disadvantage of this approach is that the design is not portable as it depends on bulky components with many different input voltages, bias voltages and cables, but it provided an initial setup for experiments and a base for later designs.

To produce a more portable thresholding and pulse shaping circuit, two possible designs were considered:

- A custom CFD. The output of the SiPM is expected to vary in amplitude for different input intensities, even in the single-photon counting regime. With this design the possibility of time walk is reduced, but the design process for signals with a low amplitude and strict time constraints is more complicated.
- 2. A comparator circuit. This design is much simpler to implement and introduces less stages where the timing of the signals can be compromised due to noise, but is more prone to time walk issues. This design would only be considered after testing it with the target signals and making sure that the variations in the single-photon counting regime would not cause distortion in the histograms.

First, construction of a CFD was attempted, based on a circuit proposed by Hansang Lim [222]; a schematic is shown in Figure 4.3. It corresponds to the basic structure of a CFD, with an attenuation branch created using a voltage divider and



Figure 4.2: Pulse shaping and amplifier circuits used for the first stage of the system development.

a delayed branch which is inverted by connecting to the inverting input of a comparator. There were two issues that proved difficult to overcome with this design:

- The delay was provided by a coaxial cable, but this would not be feasible for a small format PCB. Alternatively, a short delay could be provided by a PCB trace (but this could still be too large for a small 'tile') or an RC network.
- 2. The baseline of the attenuated branch must be controlled (for example with the V_{offset} input in Figure 4.3). In practice, the CFD implemented on a PCB exhibited a level of time walk on clean inputs generated with a function generator that was observable on an oscilloscope as a constant phase shift between the input and the output of up to 5 ns, which was unacceptable. This could have been due to the poor quality of the traces or the unreliability of the offset source.



Figure 4.3: Circuit of a constant fraction discriminator based on the design of Hansang Lim [222]

The second alternative, a comparator, was then attempted. This proved to introduce no significant distortions to the TPSF when compared with the thresholding performed by the TCSPC card, so it was considered adequate. Other recent probe designs using a SiPM have obtained stable results using a comparator with a single threshold [93].

The test cycle for the next iterations of the design was done principally on PCBs, printing one board for each iteration. By testing individual integrated circuits and sub-circuits using separate test boards, it became apparent that the interfacing between them made the circuit more sensitive to noise and unstable, at times not matching the functionality of the full circuit once integrated into a single board. This is important, not only because it conditioned the way that non functional designs were tested and the time employed, but also because it affected the performance of the those that were functional. This is less common during the development of larger systems that can rely on block circuits, which are already insulated and designed to work as separate 'design pieces'. It would be important, when further developing the latest version of the design to include as many elements that function together on a single board as possible, to assess the real functionality of the components and to ensure a good timing performance.

The first designs of a comparator were based on a structure similar to that of the circuit in Figure 4.2, with the MAX9600 previously used for the CFD design,

replacing the CFD block. This type of comparator operates in the NECL voltage range (-1.7 V to -0.9 V) and meets the ideal performance requirements as it provides fast response with sub-nanosecond propagation delay, sub-picosecond jitter, and tracking frequency of the order of gigahertz while being usable with a low number of power and control voltages. For the pulse shaping, an attempt was made to utilise the same flip-flop structure used in the past, substituting the last amplification and bias stage by an NECL to 3V-TTL level translator. The models of level translator available on the market cannot provide enough current to 50 Ω lines, so the last stage including an amplifier and a bias-tee was also evaluated. Tests also showed that the pulse broadening circuit's output was not a full range low-to-high signal; it generated a noisy final output signal with a slow rising edge that did not match the required voltage range. It was hypothesised that the comparator's output, while attempting to follow a rapidly changing, low amplitude input, was not producing enough current to drive the PCB trace and the flip-flop's input. Therefore, after introducing several intermediate stages to attempt to buffer or amplify some of the intermediate outputs, it was decided that this circuit configuration would not be able to provide the desired behaviour. The last design using this philosophy is reproduced for reference in the appendix, in Figures A.1 and A.2.

Following this, an alternative design was developed which addressed the problems with the voltage levels, by designing a new pulse shaping circuit, and reduced the reliance on external cables. The design was produced in collaboration with Dr Konstantinos Papadimitriou (currently Senior Electronic Systems Engineer at Envisics) and the complete schematic is reproduced in Appendix A, Figures A.3 and A.4.

The first modification made to the design was substituting the flip-flop with a chain of LTC6752 comparators in a one shot multivibrator configuration. As suggested in the application notes published by Analog Devices [223], two comparators and an RC network can be used together in the configuration shown in Figure 4.4 to produce a 3V-TTL pulse with a fixed width proportional to the value of C_{timing} . With the aid of software simulation, a value of 10 pF was chosen for a pulse width

of approximately 10 ns. Comparator 1 acts as a threshold comparator. In this case, it is triggered by the pulses coming from the fast NECL comparator, which is more efficient at registering the fast pulses from the input lines. Its output triggers comparator 2 and begins to charge C_{timing} . When it is charged, comparator 2's output goes low and deactivates comparator 1. The comparators have a small footprint and are affordable, so this presents an alternative to the flip-flop design that does not utilise external delays and can be integrated into a small PCB. As this comparator model is not designed to drive a 50 Ω transmission line and possible variations in amplitude were expected based on simulations and previous experiments, an amplification by a factor of 2 was introduced into the last stage with an operational amplifier.

A second modification was the addition of a potentiometer to the hysteresis pin of the NECL comparators. This input affects the hysteresis cycle of the output, making it more or less sensitive to changes in the input and can be adjusted by changing the value of the input resistance, effectively modifying an internal voltage. This way the triggering on background noise and to small pulses can be regulated.

The final change introduced to the design was the addition of an 8-channel digital-to-analog converter (DAC), AD5362, that can be configured externally through an SPI interface. This component controls the individual thresholds for the input signals, the logic levels to activate the comparators, and the internal thresholds used for the monostable (the '-' input of comparator 1 in Figure 4.4). An Arduino MKR ZERO board was used as an interface between the DAC and a PC through a USB serial interface, so all the voltages can be reconfigured manually when needed.

This new circuit configuration still relies on external amplifiers for the SiPM and synchronisation lines to make the input amplitude to the comparator circuit large enough. The combination of circuits is shown in Figure 4.5. For the SiPM branch, a signal with amplitudes between 60 mV and 100 mV is produced, and for the synchronisation signal, the amplitude of the amplified signal is 100 mV. If this circuit design is to be translated into a 'tile' form, the circuits may be substituted with surface mount alternatives (e.g. in the Minicircuits catalog) or by surface



Figure 4.4: A monostable multivibrator design using two comparators, reproduced from [223].

mount RF operational amplifiers, such as the Texas Instruments LMH3401.



Figure 4.5: Circuit including amplifiers used with the custom comparator board.

4.3 Analysis of system performance

The circuits described above were tested using the IMRA laser with the SiPM as detector and a photodiode for the synchronisation signal, monitoring the output pulse shape through an oscilloscope and the time distribution using the Becker & Hickl TCSPC card.

4.3.1 Pulse shape and optimisation

There are four user-adjustable inputs that affect the pulse shape of the output:

- The input voltage threshold: this voltage should be set higher than the noise floor but low enough that a large proportion of the input pulses are detected. If the fraction of the input above the threshold does not give the comparator enough transition time, the output will not produce a complete pulse (low level to high level) and at times it will not produce a pulse at all.
- The hysteresis level: this should be set so the noise is filtered but without letting more than one input pulse to be registered as the same pulse.
- The monostable threshold: this voltage is used to threshold the NECL output of the comparator after removing the DC component. Modifying it can also change the width of the pulse or its amplitude.
- The final DC filter: after the amplifier, an optional DC filter can be applied to set the baseline level of the pulses to zero. This may not be needed, depending on the signal.

The values for each input were set according to the expected voltages, and later tweaked by hand. In the case of the synchronisation branch and additional 1 dB attenuator was inserted to reduce the maximum voltage to a safe level below 3 V.

During testing a variation and stabilisation trend was observed on the outputs, especially on the synchronisation branch. After an input threshold is set, specially after powering on the circuit, the amplitude of the output increases, and then stabilises, sometimes with more abrupt changes the first few minutes after powering up. This significantly influences the stability of the system and the behaviour reported in the next section.

4.3.2 Stability and jitter

The full signal conditioning chain, consisting of the IMRA laser, amplifiers and comparator is expected to suffer fluctuations in the signal due to variation in temperature and external noise. These factors affect the average time of flight and the shape of the impulse response function (IRF). In order to characterise this behaviour, the mean system response was measured when the signal conditioning system was connected to a SiPM illuminated with a free-space beam of attenuated laser pulses on the input and to a Becker & Hickl TCSPC card on the output. The card software was used to collect one histogram for 1 second every 2 minutes for a period of 5 hours, totalling 150 free-space IRFs.

Figure 4.6 shows a plot of the delay time, that is, the evolution of the mean time-of-flight of the IRFs. There is an arbitrary initial offset of 34750 ps because of the time delay caused by the cables and other circuit components. For approximately 50 minutes the mean time shifts dramatically. During this period the laser is warming up, the temperature of the components in the rest of the system is increasing, and the stabilisation effect in the comparator board, described in the previous section, has not fully settled producing a jump in the mean time. Thereafter the time shift is much slower, and after approximately 200 minutes (3 h 20 min) the mean time varies within a range of 40 ps. After this warm-up period, the mean time varies mostly within a range of 20 ps, which we can consider the jitter of the system without the TDC. Figure 4.7 shows the variation of the mean time over a period of 300 s (5 min), acquiring an IRF every 2 s.



Figure 4.6: Mean time of the IRF from the signal conditioning system plotted against the period of stabilisation time. The coloured bands represent a range of ± 10 ps (blue) and ± 20 ps (red) over the mean time after 200 min.



Figure 4.7: Mean time of the IRF from the signal conditioning system plotted against the period of stabilisation time after the system is stable. The coloured bands represent a range of ± 10 ps (blue) and ± 20 ps (red) over the average mean time.

I conclude that the initial warm-up period for the IMRA laser used for these measurements is around 50 minutes, which is similar to the warm-up time recorded for the Femtolite laser used in MONSTIR II, as reported by the researchers that regularly use it. To avoid an excessive time shift recorded in the first 3 h it was decided to warm up the system for at least that amount of time before each experiment, as part of the experimental protocol.

For wearable systems and more critical ambulant setups, the warm-up time should be reduced in order to make the system more responsive. The laser source and the comparator board were the two components that affected the warm-up time more significantly and where improvements can be made. A different laser source, like a photodiode, can have a different warm-up time, speeding up the stabilisation of the system. The stabilisation process for the comparator can depend on the placement of the components on the board, the jumper pins and cabled connections. Reducing the amount of connectors can reduce capacitive effects, making changes in the signal shape faster.
Chapter 5

Development of an FPGA-based Time-to-Digital Converter

5.1 Initial direct-to-histogram design

The principal criteria for the design of a reconfigurable TDC on an FPGA which meets the project requirements were as follows:

- The design should allow for several channels on the same device. This not only saves components, but also enables multiple detectors to be integrated onto the same 'tile'.
- The communication protocol with the rest of the components should be simple and fast. Sending data should not disturb the collection process, to avoid photon events being missed.
- The FPGA should process the data to some degree (e.g. build the histograms or calculate statistics). This parallelises the processing load among 'tiles', potentially enabling real time processing.

The initial designs for the FPGA TDC were based on the work by Sébastien Bourdeauducq [163] and the notes and code made available by the Delft University of Technology [161]. Both present a delay line-based TDC using a VHDL CARRY4 primitive to implement the delay elements.

110 Chapter 5. Development of an FPGA-based Time-to-Digital Converter

A CARRY4 gets implemented in hardware as a block of four one-bit adders with carry. The logic of the individual adder is shown in Figure 5.1. If the inputs are fixed to 1 and 0 the output carry will propagate the input carry with a small delay. The length of this delay depends on the fabrication process (that is, on the physical characteristics of the transistors that make up the adder), on the irregularities in the manufacturing process, and on the temperature of the device. For modern FPGAs the delay is typically a few tens of picoseconds, which is suitable for TD-DOT applications.



Figure 5.1: Block diagram and truth table of an adder with carry. If the input is fixed to one of the highlighted combinations, the output carry reproduces the input carry.

The input event and synchronisation signals' low and high levels (logical 0 and 1 respectively) are discriminated in the input buffers of the FPGA according to TTL logic thresholds. At this stage the signals have not been sampled, so they retain their timing relationship. The event arrivals can now be represented within the logic as a transition from 0 to 1 (a rising edge of the pulse train), and the event signal then propagates through the CARRY4 delay line. The synchronisation signal serves as the sampling clock for a series of flip-flop registers that records the status of the delay line when a synchronisation pulse is received. The first implemented version of the system, illustrated in Figure 5.2, incorporated a double flip-flop layer, as suggested by Bourdeauducq. This structure is intended to compensate for possible metastable states caused by the internal clock routing.

During the place-and-route process, the optimisation software can locate the CARRY4s and registers far apart from each other, rendering a configuration and per-



Figure 5.2: Schematic of a delay line with two layers of registers.

formance that could introduce non-linearities and that is not easily reproduced (i.e. the optimiser converges to one of multiple solutions). Some investigators such as Daigneault [155] and Bourdeauducq [163] have recommended fixing the logic cell locations, either manually or using an automated program. Whatever the approach, the naming conventions for physical locations in the code will be tied to the particular architecture of the FPGA, and adapting the code for a different model or manufacturer will involve modifying them.

For this project the location of the three chains (delays and flip-flops) was fixed by hand, using the LOC directive (to fix the logic cell location) and the BEL property (to choose one of the physical FFs) to tie the flip-flops to the desired positions. The Xilinx Artix 7 can fit up to 700 delay elements in one 'column' of fabric. For most of the experiments, a smaller number of delay elements was chosen as it is faster to synthesise, while still covering the TPSF's time range. The output of this process is a so-called thermometer code, in analogy to how a mercury thermometer shows the temperature: it contains as many consecutive '1's as the number of delays the signal has gone through, with the rest set to '0'. It can then be translated into a fine timestamp code. If the pulse shape is too narrow or if many pulses occur simultaneously the transitions can be difficult to detect, and some authors suggest systems for asynchronous signal shaping, usually a flip-flop at the input that transitions after a set amount of time or when the event signal has gone through the delay line.

112 Chapter 5. Development of an FPGA-based Time-to-Digital Converter

The pulse shaping proposed by Dutton et al [164] was considered for use in the system. It involves a toggling register, where the rising edge of the input signal toggles the output: 0 to 1, or 1 to 0 (see Figure 5.3). Now, an event is represented by any level transition (not just a pulse) and the number of delay elements is reflected in the number of both '0's and '1's. Dutton et al [164] also propose that using this technique in combination with XOR gates enables multi-event detection and automatic histogram generation to be achieved.



Figure 5.3: A toggling register. The output's transitions are caused by the rising edges of the input (marked in red).

An XOR gate between two registers will have a high output when the signals are different (1 vs 0 or 0 vs 1). This means a transition in the signal, and thus an input event, can be identified at any point in the delay line, and the length of the delay line it crossed can be tracked without the need for counting 1s or 0s. Also, multiple events (multiple transitions) can easily be identified. Figure 5.4 shows a comparison between the thermometer code and the XOR output for one event.

This idea was applied to the design in the form of a set of XOR gates connected in turn to a set of adders with register – one per XOR gate. Each register holds a 16-bit or 32-bit value corresponding to the histogram bin in that position; 1 is added to that value for every synchronisation pulse, when the XOR gate indicates it.

To complete the design of the timing system, the test board communicates with the PC via the inbuilt FTDI UART to USB link, which appears to a host PC as a serial communications port. The host PC sends commands made up of ASCII characters to start the communication or to set up the number of cycles to sample (integration time). Components in the hardware design handle these requests, and start



Figure 5.4: Timing diagram comparison of an event detection in the delay line with thermometer code and applying the XOR function between registers. The red signals represent the state of the registers.

and stop the operation of the delay line accordingly. After generating a histogram, the FPGA sends the values recorded for each bin. Initially, a human-readable rendition of the bin contents (in hexadecimal) was sent, for debugging purposes. This was later substituted with the raw binary values to speed up the transmission.

The PC-side software was programmed in Python, using the pySerial library to communicate with the serial port. The software sends the desired configuration for the integration time, receives the data and saves it in matrix form in MAT-files, readable by Python and MATLAB.

5.2 Non-linearity characterisation and correction

When a version of the delay line TDC was first tested as part of an MRes project [224], some of the early results exhibited histograms containing artefacts in the form of narrow peaks. This phenomenon occurred with different detectors, e.g. a SPAD and a photodiode. The narrow peaks inhibited the identification of the true peak of the TPSF and thus made it difficult to assess the timing characteristics as the

114 Chapter 5. Development of an FPGA-based Time-to-Digital Converter

signal progressed along the individual delay elements in the FPGA. The origin of this problem was initially attributed to signal integrity issues, such as interference or distortion due to low bandwidth. To isolate the source of distortion, the influence of the external lines and connections was tested by measuring the distribution of jitter through the signal chain with a Rhode & Schwartz RTO2044 oscilloscope, on loan to the lab. The oscilloscope was used to record the period and frequency of the input signal over a large sampling time (several seconds, corresponding to millions of cycles) and accumulate a histogram. As reference, a 40 MHz square signal was generated using a low jitter (7.7 ps) Silicon Labs Si5324 PLL-based clock generator, whose output distribution can be seen in Figure 5.5a. The output was connected to the FPGA, which was programmed to directly route two pins without any other circuitry in between. This makes it possible to measure the effect of passing the signal through two I/O drivers. The result of this experiment, shown in Figure 5.5b, reveals that the FPGA adds jitter, as the SD of the distribution is observed to increase to 20.18 ps, and the shape of the signal is distorted. None of the features in the histogram matched the high frequency spikes seen in the histograms, which meant the I/O ports were not the source of this distortion.

After reviewing the papers on which the design is based, it was hypothesised that the most likely causes of the distortion of the histograms are non-linearities related to the bin width and/or the monotonicity of the delay chain, which become apparent after reading out the delay line. The latter, known as the source of so-called *code bubbles*, is frequently mentioned in the published literature, although solutions are not always suggested. One of the practical solutions, proposed by Bourdeauducq, is to estimate the delay of each delay element's output (using certain timing analysis tools in the design software) and then sort them manually. This solution was not easily reproduced using the version of the Xilinx tools employed for this project, and a way of replicating the procedure with the new language syntax was not found. This method also poses immediate limitations, as it cannot calculate the delay on the real hardware and it is not guaranteed to work across different software versions or FPGA manufacturers.



(a) Variation of the period and frequency of a signal generated by a Si5324 clock generator.

Meas 1	Current	+ Peak	- Peak	mu (Avg)	RMS	StdDev	Event Count	Wave Count
Period	20.027 ns	20.136 ns	19.843 ns	20 ns	20 ns	20.18 ps	2395200	4800
Frequency	49.933 MHz	50.395 MHz	49.663 MHz	50 MHz	50 MHz	50.45 kHz	2395200	4800



Figure 5.5: Histograms of the variation of the frequency of a periodic signal passing through the FPGA's I/O banks.

116 Chapter 5. Development of an FPGA-based Time-to-Digital Converter

An experiment was designed to test this hypothesis on the individual delay elements. If a signal is measured both at the input and the output of the delay line for different lengths, the evolution of the total delay can be estimated. A delay line was created within the FPGA fabric and configuration bitmaps were generated routing the outputs for all possible delays to an external pin. The last route between the delay and the output is ultimately decided by the design software and introduces some uncertainty, but which was considered sufficient to estimate the delay, even if an accurate value could not be extracted.

The delay between the signals at the input and the output of the delay line was measured with the Rhode & Schwartz oscilloscope. A plot of the average delays obtained from a 40 element-delay line (equivalent to 10 CARRY4s) is shown in Figure 5.6. It shows that the step size between delays is not constant and the input-to-output delay is not monotonically increasing. The lower graph in the figure shows the variation of delay between consecutive delay elements which in some cases is negative, that is, the signal arrived faster to the output. The mean delay estimated by this method was 12.67 ps.



Figure 5.6: Input-to-output delay measured for a 40-element delay line (upper) and variation of the delay between delay elements (lower).

5.2.1 Simulation of the delay line and code bubbles

To better understand the effects of non-linearities on the time-of-flight histogram, a computer simulation was created using Python. A summary of the simulation steps is reproduced in pseudocode in Listing 1 and the full code can be read in Appendix B.1. It takes a list of the cumulative delays for each delay element and a noise distribution, and simulates the behaviour of the delay line by tracking the state of the toggled event signal through the delays, registers, and XOR gates, and calculates the resulting histogram. By modifying the order of the cumulative delays, it is possible to simulate the situation that was observed in the experiments. The routes within the FPGA are distributed from central nodes into the different elements. It is possible for the input signal to arrive at the register associated with a longer delay slightly before the signal from a shorter delay. When this occurs, a single event arrival can be registered in two different locations, that are identified by the XOR gates as two events. In this situation, illustrated schematically in Figure 5.7, the chains of 0s and 1s from the registers will show discontinuities (that are referred to as *bubbles*). When a histogram is calculated from this, a double peak appears. And when jitter is added to the signal, the probability of some times of arrival coinciding with the discontinuity increases, giving the resulting histogram the appearance as observed experimentally. An example of the histograms generated by simulating 500 events is shown in Figure 5.8.

5.2.2 Estimation of the DNL

The procedure to estimate the DNL as described in Section 3.6 was applied to the FPGA design using random counts produced by a photodiode as the event input. The random histogram, the DNL and the correction factor are shown in Figure 5.9. The dominant effect is a pattern of positive DNL spikes occurring every fourth delay, and regularly in larger groups of 4 delays. This is consistent with the expectation that adder elements inside the CARRY4 are close together, while the FPGA groups the cells in blocks that are then separated from each other.



Figure 5.7: Timing diagram showing the effect of *code bubbles* on a histogram.



Figure 5.8: Histograms generated by the simulation code running with a monotonically increasing delay line and with code bubbles.

5.3 Initial results

To test the FPGA system, two types of experiment were carried out: one to test the linearity of the system and another one to test the individual delays and assess the effect of code bubbles.

For the linearity tests, a photodiode and a SPAD were aligned with the IMRA A-70 fibre laser beam in free space. The laser output fibre was mounted on a trans-

Algorithm 1 Pseudocode for the delay line simulation
function GETHISTOGRAM(signalTimes, clockTimes, delays)
histogramBins \leftarrow zeros(# delays)
for $I \leftarrow 0$, # clockTimes do
$delayedSignal_J \leftarrow 0$
$delayedSignal_{J-1} \leftarrow 0$
xorOutput \leftarrow zeros(# delays – 1)
for $J \leftarrow 0, \#$ delays do
$delayedSignal_{J-1} \leftarrow delayedSignal_J$
signalTimes _{delayed} \leftarrow signalTimes + delays[J]
$K \leftarrow FINDLAST(signalTimes_{delayed} < clockTimes[I])$
Find the index of the last event just before the clock
if exists(K) then
delayedSignal _J \leftarrow MOD(K + 1, 2)
else
delayedSignal _J $\leftarrow 0$
end if
▷ The value of the toggled signal depends on the num. of the cycle
if J>0 then
$xorOutput[J-1] \leftarrow delayedSignal_J \oplus delayedSignal_{J-1}$
end if
end for
histogramBins
end for
return histogramBins
end function

lation stage to allow displacement up to 50 mm and neutral density (ND) filters were used to regulate the intensity of the fibre output and the intensity reaching the SPAD. The filters exhibited slight discrepancies from their nominal value, so they were calibrated using a Thorlabs power meter. Additionally, external delays were employed to adjust the timing. The histograms were recorded using the FPGA system and a Becker & Hickl SPC-130 timing card for comparison.

Figure 5.10 shows a histogram collected from the photodiode's output before and after applying the DNL correction factor. The shape of the histogram is smoother after correction, although some spikes still remain. The main peak of the distribution, however, seems to be more easily identifiable. Figure 5.11 shows the result of increasing the distance between the source fibre and the detector using the translation stage: the positions of the peak and the means of the acquired his-



Figure 5.9: Plot of a noise histogram, DNL and correction factors for the direct-tohistogram TDC.

tograms are both observed to increase. The result suggests that the TDC is linear and that the time delay per FPGA logic block is approximately 8–11 ps (although this estimate is improved in later experiments). It was also concluded that the DNL correction coefficient was useful to recover information from the acquired data, and therefore it should be applied to subsequent measurements. The translation stage results were difficult to replicate and were not repeated with other detectors until additional hardware was available to physically secure the elements in place.

Using the SPAD as detector and a Hamamatsu delay unit to shift the sync signal in time, TPSFs were obtained at different delays and with different ND attenuation levels in order to capture different time ranges of the histogram in a short delay line. Figure 5.12 shows the TPSF without a filter (at around 25 ns), with the SPAD saturating. On saturation, the TPSF becomes narrower and timing artefacts appear (possibly due to the dead time of the SPAD not allowing the output signal to trigger or due to pile-up in the detectors), making the count at some bins drop towards 0. The TPSF in Figure 5.13 (also at around 25 ns) was collected when the beam was attenuated by 4.5 OD. The histograms obtained with the FPGA and a 400-element



Figure 5.10: Histogram obtained with a photodiode, as collected (blue) and after applying the DNL correction factor (red). The mean is indicated by a vertical line.



Figure 5.11: Position of the histogram peak and mean of the histogram collected with a photodiode for different free space source-detector separations.

delay line exhibited features observed in the results obtained with the TCSPC card. If the average of the histogram is computed, the effects of the external delay can be calculated, with the progression seen in Figure 5.14. By doing this, the individual delay per logic block can be estimated to be between 10 ps and 17 ps, which is larger than that observed with the photodiode.

From previous experiments performed at UCL, the FWHM of the SPAD's freespace IRF was expected to be a half of what was observed using the TCSPC card





(b) Histogram at high light intensity levels collected with the FPGA

Figure 5.12: Comparison of the histogram response of the SPAD at high intensity levels in the SPC-130 card and the FPGA. Due to the saturation, the timing gets distorted, narrowing the main peak and reducing the count rate at some time points

(1.5 ns), and this was attributed to a malfunction of the SPAD. Although the FPGA was shown to be sensitive to the timing changes, the Excelitas SPAD was considered unsuitable to evaluate the TDC. Therefore the SensL SiPM was used for subsequent testing. Meanwhile other limitations explored in the next section prompted the redesign of various components of the TDC chain.



(a) Histogram with a 4.5 OD attenuation factor collected with the SPC-130 card.



Figure 5.13: Comparison of the histogram response of the SPAD at lower intensity levels (4.5 OD) in the SPC-130 card and the FPGA.

5.4 Limitations of the design and reprogramming

The main limitation of this design is the loss of control over which events trigger the histogram generation. Further tests revealed that the count rate produced by the system could be up to twice the true input rate of events. This means that the uncorrected *bubbles* in the delay line were having a larger effect than initially estimated, and they were not easy to mask with regular correction techniques. Careful inspection of the TPSFs recorded during the experiments also revealed that the



Figure 5.14: Position of the peak and mean time of the histogram as delay increases for a SPAD test.

spike artefacts were occurring in spite of the DNL correction, which meant that the calculation of mean flight time would also be affected by the code bubbles.

A further investigation of a way to extract the state of the delay line for every delay element could lead to a means of reordering the delay outputs which is independent of the FPGA version and of the features available in the design software. Equipment such as a digital delay generator, which was not available at that stage of the project, could have helped perform the necessary measurements.

The XOR gate direct-to-histogram system was then substituted by a timestamp-based histogram generator. The thermometer code output from the delay line is reduced to a number by an adder that counts the number of 1s and this is passed to a new component that stores an updated histogram, as illustrated in Figure 5.15. The histogram generation now needs two sub-components to avoid false counts and to avoid missing events.

The first sub-component implements a method to identify the events. The histogram generator is no longer synchronous with the event signal, but still has to communicate with the delay line, which means that two systems with different update rates (the sync frequency and the internal oscillator frequency) must be coordi-



Figure 5.15: Block diagram of the TDC system with a block memory-based histogram generator. Data buses are represented on a thicker line

nated. To keep the main system running on a reliable, oscillator-generated clock, a status signal needs to be generated whenever an event arrives and is sampled by the synchronisation pulse and while the state of the delay line registers is stable. The system clock must be fast enough to sample the synchronisation signal in order to find when the registers are stable; this signal can be as fast as 48 MHz and with a duty cycle as low as 30%. Figure 5.16 illustrates different sampling clock frequencies. It shows that a system clock of 400 MHz is adequate, although this is a fast clock that can generate *race conditions* in some sections of the design, as reported by the timing analysis software. A race condition is a situation that occurs when the propagation delay of a group of signals and their frequency are of a similar order – then, small variations in the routing may cause two events to arrive in an unexpected order, provoking an unexpected system behaviour. A 100 MHz clock is a safer option in this respect and is more than sufficient to keep a high speed link with the central PC, but can miss many events. While maintaining several clock domains is generally a bad practice that can lead to timing errors in the signals, a compromise was made by allowing the event signal generated with a 400 MHz clock to run for several cycles until the 100 MHz components acknowledge receiving it.

The second sub-component is the memory that stores the histogram bins. It works by assigning each timestamp to a memory address and then adding 1 to that



126 Chapter 5. Development of an FPGA-based Time-to-Digital Converter

Figure 5.16: A 50 MHz clock with 30% duty cycle being sampled at different frequencies. To be conservative, the coincidence of a rising edge with a rising edge is considered an event miss. Each time division is 1.25 ns.

address for each event. Retrieving the stored data, adding 1 to the value and storing the new value can take several clock cycles. The component was programmed to perform this process in 2 clock cycles: one to retrieve the old value of the bin from memory (a 1024-level memory buffer) and another to store the new value, ensuring that all events are stored before processing. The histogram component reads timestamps from the buffer while it is not empty, relieving the need for precise timing while new events arrive – the component with the quick reaction to events is the buffer, which has a simpler logic. The histogram memory is also accessed for reading by the component that sends the data to the PC, which is still tasked with collecting the data and correcting it. In order to avoid any ringing in the TPSFs due to small differences in timing between histogram bins, a new histogram is generated in the PC by merging bins in pairs – each new bin contains the counts from two of the old bins. This makes the effective time step become twice the delay of the

CARRY4 elements.

Before any further tests, the timing resolution of the new system was measured with a function generator. Skipping the histogram generation, the raw timestamp from the delay line was sent to a PC. Two square signals were generated with the function generator, applying a small phase difference between them. This phase was slowly increased until a change was observed in the delay line, accounting for a delay of approximately 17 ps, which confirms previous estimations. After applying bin filtering the effective resolution of the system is approximately 34 ps per bin.

An example of an IRF collected with this system, before and after applying the bin correction, appears in Figure 5.17. The amplitude of the corrected IRF has been reduced to a half for the sake of comparing the shapes.



Figure 5.17: A TPSF collected using the FPGA TDC with timestamp histogram generation, before (blue) and after (red) merging bins together. The amplitude of the filtered TPSF has been divided by 2.

5.5 Tests using the signal conditioning system

Two SMA-to-PMOD connectors with a 50 Ω termination were built to connect the comparator board to the Nexys 4 DDR board. A stability test was performed under similar conditions to the Becker & Hickl test described in Section 4.3.2. The stability of TPSFs collected using the MONSTIR II supercontinuum laser was also assessed.

A plot of the mean time-of-flight of the system response calculated over a 5 h period is shown in Figure 5.18. The mean time stability over a shorter 5 minute period is shown in Figure 5.19. Although the drift in mean time largely disappears after a period of 3 h as seen before, there is still a significant short-term variability. The mean time varies between 100 ps and 200 ps, even after isolating the development boards and taking care that the connection cables are not damaged. This amount of jitter, equivalent to path distances of 30–60 mm, is too large to distinguish the time differences produced by small changes in oxygenation or blood volume. The reduction in quality of the TPSFs can be attributed to the SMA-to-PMOD converison. The PMOD interfaces are not necessarily designed for high speed signals or for complex pulse shapes and the I/O banks have been demonstrated to add significant amounts of jitter to a clean signal. For this reason, I decided to focus efforts on developing an alternative system based on the Texas Instruments TDC7201 integrated circuit.



Figure 5.18: Mean time of the IRF from FPGA system plotted against the period of stabilisation time. The coloured bands represent a range of ± 10 ps(blue), ± 20 ps (red) and ± 50 ps (magenta) over the mean time after 200 min.



Figure 5.19: Mean time of the IRF from FPGA system plotted against the period of stabilisation after the system is stable. The coloured bands represent a range of ± 10 ps(blue), ± 20 ps (red) and ± 50 ps (magenta) over the average mean time.

5.6 Conclusions

This chapter has presented the design and testing process of delay line time-todigital converters on an entry-level FPGA. Starting from the same basic configuration, a chain of CARRY4 blocks acting as delay elements, two readout schemes were explored: a direct-to-histogram architecture, which builds a histogram as events arrive looking for differences in the contents of the delay line, and a timestamp-based architecture with independent histogram generation. The optimisation process for both architectures has been described. Timing systems incorporating the two architectures were built and tested by connecting them to different sensors and measuring their linearity and stability.

The direct-to-histogram readout architecture can identify all events with a dead time of a few picoseconds (conditioned by the delay of the toggling register at the start of the delay line), but has proven to be difficult to debug, as the state of the routes between the delay line and the next components are hard to monitor, and too sensitive to non-linearities with the XOR comparators introducing a noisy pattern in the histogram and possibly masking event arrivals.

The alternative timestamp architecture reduces the risk of masking events and

130 Chapter 5. Development of an FPGA-based Time-to-Digital Converter

the impact of non-linearities by only processing one event at a time and adding the state of the whole delay line. The main disadvantage is that the sampling frequency is limited to that of the system clock (hundreds of megahertz). By adding a buffer before the histogram registers, the event histogram can be calculated taking three clock cycles without accidentally missing events for low count rates that do not cause the buffer to overflow.

For both TDC types an algorithm to correct non-linearities was successfully employed to improve the shape of TPSFs and reduce the errors made when calculating statistics.

Lastly, a complete chain consisting of a light source, a detector and a TDC connected to a computer was demonstrated. For both TDC architectures, the integration time can easily be set by the user by modifying the value of an independent counter. The computer software can be written to collect the TPSFs, apply correction and filtering algorithms and extract features of interest, like the mean time of flight, the FWHM and the total intensity.

The main limitations of the tested setups were the sensitivity to signal integrity errors and to non-linearities. The connections from the comparator board to the digital I/O ports of the FPGA evaluation board were prone to voltage mismatches, resulting in distortions of the histogram and increased levels of jitter that made the more demanding imaging experiments impossible. To mitigate this problem, a different evaluation board with properly matched SMA connections should be considered. If the FPGA is used on a custom design, the inputs can also be designed to match the impedance and output levels from the previous stages. Furthermore, the delay line TDCs are intrinsically more susceptible to the effect of non-linearities on the overall time resolution. Some techniques to improve the resolution were not tested for this project, but can be implemented in the future, such as analysing and re-routing the delay elements to ensure the delay of the signal paths is monotonically increasing, or implementing a *wave union* architecture.

Considering further miniaturisation of the imaging system into a distributed 'tile' style, the role of the FPGA should be reconsidered. The size of the integrated circuit is large, requiring too much space for individual probes. Using one FPGA per probe would also be too expensive. As the FPGA can serve more than one channel simultaneously, the integrated circuit could be better employed within a hub between several sensors, collecting the information from more than one probe.

The following chapter presents a more complete and functional imaging system with components that are easier to incorporate within a single 'tile' format.

Chapter 6

Development of a Time-of-Flight System Based on a Commercial Time-to-Digital Converter

6.1 Characteristics of the Texas Instruments TDC7201

The Texas Instruments TDC7201 is an integrated circuit (IC) introduced in 2016 designed to measure the delay times between start and stop pulses in two independent channels, using TTL input signals. Texas Instruments sells it at \$1.35 per unit for an order of 1000 units, with distributors in the UK selling a single unit for about £4. It is also available with a test board, the TDC7201-ZAX-EVM, which is available for \$199 from the manufacturer and for about £180 in the UK. It has been designed for ultrasonic, LIDAR and radar applications, and its data sheet suggests it can also be used for collision detection and flow measurement. Articles have been published about the use of this chip (or its single-channel version, the TDC7200) for LIDAR [225], for the measurement of wind speed [226], and to monitor the status of Ethernet cables [227], but in this chapter its potential for use in time-domain diffuse optical imaging will be explored.

The TDC7201 internal structure is based on a ring oscillator paired with a multicycle calibration system that accounts for variations with time and tempera-

134 Chapter 6. Development of a TOF System Based on a Commercial TDC

ture. The device can provide measurements down to a resolution of approximately 55 ps with a measurement uncertainty of 35 ps, as per the data sheet, on a variable time range, depending on a mode setting. The IC communicates with other devices using a standard SPI interface and interrupt line, thus it can be seen as a black box compatible with any device that can perform the read out. The IC pins provide one input channel to the device (1 MOSI line) and two separate outputs, one for each channel (2 MISO lines).

The configuration and read out schemes are driven by instructions sent to the IC via SPI to an internal set of register addresses (8 bits per register). The most relevant ones hold the measurement and calibration data (24 bits each) as well as the configuration for the number of internal calibration cycles to be performed each measurement, the triggering edge for each channel, the number of stop pulses to be read, and the number of measurement cycles to average for each measurement (for the calibration of the resolution and/or for up to 128 measurements). The configuration and read out speeds are limited by the external clock period and by the number of registers needed to be read for each transaction. The limit for this (effectively the count rate) is not specified within the data sheet and has to be determined experimentally.

The measurement scheme follows these steps, recommended in the data sheet:

- 1. Power up the device and set the appropriate pins to reset its status.
- 2. Send the configuration via SPI.
- 3. Start a new measurement. The IC then waits for a start signal and then for a corresponding stop.
- After the last stop is received, the interrupt line is driven low (active) and the IC remains on hold until the interrupt status is cleared externally. No new measurements are performed.
- 5. Externally, via the SPI line, the results of the measurements are retrieved and the interrupt status is cleared. The system is ready for a new measurement (back to step 3).

To calculate the time of flight, data from three registers are required: TIMEn, CALIBRATION1 and CALIBRATION2. The values of CALIBRATION1 and CALIBRATION2 are used to calculate the calibration value of the least significant bit of the TDC code –that is, the resolution. The value of TIMEn is the raw time measurement for the n-th stop pulse. Given the external system clock period T_{clk} , the time of flight is calculated as follows:

$$TOF_n = \texttt{TIMEn} \cdot LSB_{norm} \tag{6.1}$$

where:
$$LSB_{norm} = T_{clk} \cdot \frac{\# \text{ calibration periods} - 1}{\text{CALIBRATION2} - \text{CALIBRATION1}}$$
 (6.2)

6.2 System design and development

The principal requirements for the system are equivalent to those given for the FPGA-based system (see Section 5.1). They translate to the TDC7201 as follows:

- The design should allow for several channels on the same device. The TDC7201 offers two independent channels that can be used for two simul-taneous detectors. Multiple ICs can be connected to the same microcontroller to expand the number of channels.
- The communication protocol with the rest of the components should be simple and fast. Sending data should not disturb the collection process, to avoid missing photon events. The TDC7201 offers a standard SPI interface that can work at high transmission rates (up to 8 16 MHz). However, the overhead imposed by the register-based protocol can affect the count rate, as I will discuss in a later section.
- The 'tile' should process the data. The TDC7201 itself does not include any facility to build histograms or average the data, other than averaging several stop pulses for the same start pulse, so the processing tasks should be performed by a microcontroller or an equivalent device.

The data collection was evaluated first on an ARM Cortex M4 based STM32

136 Chapter 6. Development of a TOF System Based on a Commercial TDC

Discovery test board, where development is made easy by C programming and the manufacturer's libraries, but ultimately, the design was completed on the Xilinx Artix 7 FPGA, as it allows for a faster readout rate and histogram building and includes more memory resources.

The central component of the FPGA architecture is the picoBlaze soft processor [205], written originally by Ken Chapman. It consists of a set of design files for a small 8-bit microcontroller and template files to connect it to memory blocks and external peripherals to be written by the user. The licence for this processor core allows it to be used freely only on Xilinx devices. Several soft processors, such as PacoBlaze and PauloBlaze, are behaviourally equivalent to picoBlaze – but not necessarily optimised for any particular FPGA – and are distributable under free licences, so they can serve as replacements should the system be exported to a different architecture. picoBlaze is programmed through an assembly instruction set called KCPSM. The assembler program accompanying the core can also be substituted by Opbasm, an open source assembler for the same instruction set.

The picoBlaze code provides only a core capable of executing sequential code (that can only be written in assembly), a set of registers, a data memory (up to 256 bytes) and a program memory (up to 4000 instructions). The rest of the peripherals needed for an application are to be provided by the system programmer and interfaced to the microcontroller through 8-bit buses. For applications that require more complex data paths, such as this project, this type of simple microcontroller can act as a coordination component that does not process data, but holds together the functionality of the key components, like the SPI and UART communication blocks, and the memory. For a system consisting of multiple steps, a microcontroller can be easier to debug and re-program than a complex state machine while keeping the resource use low -26 logical slices and 1 BRAM (for the Artix 7A100T used in this project, that represents 26 out of 15850 logic cells -0.16% – and 1 BRAM out of 135).

On the hardware side, the TDC7201-ZAX-EVM evaluation board is already fitted with SMA inputs, an 8 MHz oscillator and interface pins compatible with

Texas Instruments' Launchpad development boards. An adaptor was made on a PCB to connect the relevant pins to the Nexys 4 DDR PMOD connectors. The TDC7201-ZAX-EVM also provides an external port to modify the oscillator frequency, but this made the SPI connection work less reliably, so the base 8 MHz oscillator was retained.

6.3 One-channel system

The structure of a single channel system is based on a three-step process: configure the TDC7201 and start a time counter; retrieve timestamps and request new measurements during the integration time; and send the information to the central node or PC. The time spent on each step should be optimised to maximise the count rate of the system.

The TDC7201 is configured when the system is reset to operate with only one stop pulse and with the minimum number of averaging cycles for calibration = 2. Thus the numerator for equation 6.2 is, # calibration periods -1 = 2 - 1 = 1.

The primary challenge with timing comes with handling a high rate of incoming data points. If data is sent to the computer straight away (and stored on its slow memory) the whole process is slowed down. On the other hand, storing all the data before sending requires large amounts of fast storage. It was decided to store the data in the form of a histogram, which is conventional for time-of-flight data, and send the values of the histogram bins after the integration time is finished. Because the number of bins is fixed, this solution can handle different integration times and amounts of data points and its behaviour is very predictable – the same amount of data is stored and sent every time.

The TDC7201 chip handles 24-bit data points. If one bin is kept per timestamp, that makes $2^{24} \approx 1.7 \times 10^7$ possible bins for the histogram, which is larger than the memory capacity of the FPGA. In theory, if the stop signal has a period of 25 ns (40 MHz), the timestamps will be restricted to that period. Given that each bin's width is approximately 55 ps, 25 ns is equivalent to 455 bins, a number

138 Chapter 6. Development of a TOF System Based on a Commercial TDC

representable with 9 bits. In the case of this project, the timestamps were found to be constrained to a range of approximately 10 bits, with a leading gap of blank bins. With 16 bits (65536 bins), it is possible to collect histograms (TPSFs) from the optical system and also test the system with function generators while keeping the memory resources low. Both the one-channel and two-channel systems reported here keep a block RAM memory (one per channel) of 65536 bins with 24 bits per bin. A hardware switch is provided to reduce the number of bins sent to the PC to 768 (300 in hexadecimal) during imaging experiments to speed up the sending process to 300 ms per histogram.

A block diagram of the hardware structure for the FPGA is shown in Figure 6.1, which highlights the most relevant components and signals. A full map of the signals assigned to each input and output port is shown in Figure 6.2. The picoBlaze microcontroller is the central block that coordinates the actions of the other components: starting the operation of the communication blocks and the integration time counter, signalling the arrival of an event to the histogram, and requesting the data that have to be sent to the PC. The most time-sensitive operations are assigned to *constant output* ports. With a dedicated instruction, an immediate value can be written to these ports in one clock cycle, saving the extra cycle that is usually needed to store an immediate value in a register. This makes operations where the status of the output port is a well known, fixed value, such as starting the time counter or sending a signal to the SPI, much faster.

The microcontroller does not handle the bulky data paths for the histogram directly. Moving blocks of 24 bits with an 8-bit processor is slow, because it needs 3 separate input ports to access the data and 3 instruction cycles for each memory write or memory read operation. The data buses connecting the incoming data to the histogram block and the histogram block to the UART module are wired in the VHDL code so the microcontroller does not need full access to the data. Because the UART can only send 8 bits of data at a time, some switching is still needed, but this function is left to a VHDL block. Finally, the data are sent back to the PC through a 921600 bps link, the maximum standard rate that worked reliably with



Figure 6.1: A block diagram of the main components and signals for the FPGA for a onechannel system. The numbers in brackets represent the bus width in number of bits.

the FPGA-to-PC USB connection.

Two other steps were taken to speed up the system. First, the TDC calibration value is read only once per experiment. After measuring this value's variation over the course of 1 minute, it was found to change by a maximum of 300 fs, and that it was consistent across different experimental sessions. There remains a bug in the code by which an impossible calibration value is read for certain acquisitions. Although this is not a problem if the same TDC chip is used in a stable environment (given the small variation in the value), if the system is integrated into a portable device it could cause long-term shifts in the data. The calibration was correctly read more consistently with lower frequency inputs produced by a function generator; hence a possible solution could be to add an initial step where the TDC reads signals not from the SiPM but from oscillators with a fixed frequency. The second step taken to speed up the system is to read the histogram memory in reverse. This way the condition to stop the software counter can be expressed with only one instruction (one comparison with zero). This also renders the TPSF time scale the expected way - displaying a steep rise followed by a more slowly decaying tail. The program flow for the single channel system is shown in Figure 6.3.

The operation of the laser source has to be controlled manually, since there is no interface available to connect the FPGA board, hence the coordination between the data acquisition and the laser is purely manual. To obtain a reference IRF or

			I	nput ports					
Bit no.	7	6	5	4	3	2	1	0	
Port A Status	0	0	Timer done	Memory ready	SPI interrupt	SPI ready	UART ready	PC data arrived	
Port B UART	UART data (in)								
Port C SPI	SPI data[70] (in)								
Port D SPI	SPI data[158] (in)								
Port E SPI	SPI data[2316] (in)								
Port F Switches	0	0	0	0	Input switches				
			0	utput port	s				
Bit no.	7	6	5	4	3	2	1	0	
Port O UART	UART data (out)								
Port P SPI	SPI data (out)								
Port Q SPI	SPI address (out)								
Port R Debug	Debug								
Port T Histogram	Memory address[70]								

140	Chapter 6	Development of	f a TOF System	Based on a	Commercial TDC
140	Chapter 0.		a ror bystem	Dasca on a	

Figure 6.2: I/O register map for the one-channel system.

0

Constant output ports

4

0

Memory address[16..8]

Byte no.

3

0

ТΙ

reset

2

0

Data

switch

ТΙ

OSCEN

1

UART

read

Event

SPI length

8/24

0

UART

start

Reset

memory

SPI

start

Port U Histogram

Bit no.

Port c

UART

Port k

Status

Port x

SPI

7

0

0

0

6

0

Enable

memory

0

5

0

Reset

time

0

background noise measurement during the experiments, the source fibre and detector components are adjusted by hand. However, to implement the system into a 'tile' form it would be necessary to establish a connection protocol that controls the laser operation and/or the FPGA acquisition program accordingly.



Figure 6.3: Flow diagram of the software for the single channel system.

6.3.1 Non-linearity characterisation and correction

For the TDC7201 version of the system it was decided to apply an error correction protocol similar to the one used with the FPGA-based TDC, involving collection of data from the SiPM without laser illumination after each measurement. This serves as an indication of the background noise due to dark counts and stray light that can be subtracted from the measurements if necessary. This sampling of a random distribution also enables DNL correction coefficients to be calculated, to compensate for bin size non-uniformities coming from the TDC and other components in the system. Collecting background noise data for each experiment guarantees that the bin range of the noise will be the same as the range of the data, since they will have been collected using the same sync frequency.

Figure 6.4 shows a typical example of the background noise recorded during a phantom experiment, and the corresponding correction factor applied to the data. The wavy pattern originates in the components that come before the TDC, as it can be observed in other data captures from the Becker & Hickl TCSPC and the FPGA, while the higher frequency spikes are due to the DNL of the TDC itself. Figure 6.5 shows an example of a TPSF collected from a phantom with the experimental setup described in later sections before and after applying the correction; the smoothing effect of the DNL correction is evident.

6.3.2 Characterisation tests using the custom conditioning system

The temporal drift and level of jitter of the system were evaluated by reproducing the same experiment that was performed using the Becker & Hickl system and the FPGA: a series of free-space IRFs were collected for a period of 300 min (5 h), with 2 minutes between each acquisition; then a series of IRFs was collected every 2 seconds for a period of 5 minutes to assess the short term jitter. A plot of the long term drift of the mean during warm-up is shown in Figure 6.6, while the jitter following warm-up is shown in Figure 6.7. The TI TDC system exhibited



Figure 6.4: Noise profile, DNL and correction factor for an experiment using the MON-STIR II laser and the single channel TDC system.



Figure 6.5: A TPSF collected from a phantom using the single channel system before (blue) and after (red) applying DNL correction.

a stronger time drift than the previous results obtained with the Becker & Hickl TCSPC during the 50 min – 120 min period but thereafter settled to a stable value. The figure also shows the effect of unexpected external noise in the system. For instance, at approximately 180 min, an external source of interference (probably an air conditioning system) caused the mean time to change significantly and then

144 Chapter 6. Development of a TOF System Based on a Commercial TDC

return to its original value. The chance of this happening is reduced by using proper insulation. The jitter exhibits a reasonably satisfactory variation of about 20 ps for this experiment, although this was later observed to increase to 60 ps during some experiments, especially when the (electrically noisy) MONSTIR II equipment was switched on.



Figure 6.6: Mean time of the IRF from the TDC7201 system plotted against time during the period of stabilisation. The coloured bands represent a range of ± 10 ps (blue), ± 20 ps (red) and ± 50 ps (magenta) with respect to the average mean time after 200 min.

To assess the stability and linearity of mean time measurements, the SiPM was mounted on a translation stage so that the pulse arrival time could be adjusted manually. The translation stage was moved in steps of 2 mm, increasing the distance between the output end of the laser fibre and the detector. Twenty IRFs were acquired for each step with the Becker & Hickl TCSPC and the TDC7201. For jitter between 20 ps and 40 ps the minimum distinguishable step size would be expected to be between 6 mm and 12 mm. Figure 6.8 compares the distances calculated from the IRF mean times for data obtained with both timing systems. The vertical blue lines represent the range of values for each position and the red crosses represent the average value. The average range (maximum - minimum) was 2.67 mm for the Becker & Hickl measurement and 3.34 mm for the TDC7201, less than what was expected from the stability experiment but indeed greater than the step size. While


Figure 6.7: Mean time of the IRF from the TDC7201 system plotted against time when the system is stable. The coloured bands represent a range of ± 10 ps(blue) and ± 20 ps (red) with respect to the average mean time.

the evolution of the distance appears linear, some of the distances overlap, and it is difficult to distinguish a difference for some of the steps. There is also a noticeable leap in the slope for the TI data (6.8b). Due to the closeness of the source and the detector, it is possible that a slight change in laser intensity or beam alignment altered the pulse shape or the timing. The Becker & Hickl TCSPC is more robust to these changes thanks to its internal CFD. It is also possible that the TDC7201 chip introduced some uncertainty in the timestamps that was not registered as a change in the mean time.

Another conclusion that can be drawn from this experiment is that the sum of IRFs can mask the effects of jitter on the individual acquisitions, as the mean of the data is very stable. If the stability of the system needs to be evaluated before obtaining real world measurements, it should be measured on individual IRFs with a short acquisition time.

The stability of the IRF mean time during changes in incident intensity was also evaluated. A combination of fixed and (rotating) variable ND filters was inserted in the free-space beam between the source fibre and the detector. The light intensity was decreased in steps of 0.5 OD, with blackout material placed around the detector to avoid any stray light (from the reflective filters) reaching the detector. For each



Figure 6.8: Plot of the calculated distance (mean, maximum and minimum) against the true source-detector distance for a free-space experiment in steps of 2 mm. The dashed lines represent a slope of 1 with different offsets.

fixed filter combination (4 OD, 5 OD and 6 OD) the variable filter was rotated to obtain the intermediate steps. The resulting measurements of IRF mean time of flight and intensity are shown in Figure 6.9, the three colours representing the three fixed filter combinations. The mean time appears reasonably stable for both systems above a certain attenuation threshold, although estimates at the highest attenuations have large uncertainties because of the low signal to noise ratio (SNR). At the lowest attenuations, the probability of multiple photons being detected during each sync cycle becomes significant, and the TPSFs suffer 'pile-up' distortion which decreases the apparent IRF mean time. The shape of the output pulses from the comparator will also be altered, as the threshold setup was meant for a different input. This does not affect the Becker & Hickl TCSPC's CFD inputs, which are prepared for increases in intensity or slight shape variations, but it does affect the input of the TDC7201, with a threshold-based input, making the range of variation (max to min) of the mean time much larger in this situation. The range of attenuations where mean time is most stable is highlighted in blue in the figures.

Although the log intensity decreases with attenuation for the TDC7201 system, the intensity values did not match those obtained with the Becker & Hickl TCSPC and for a range of attenuation settings (4.5 OD – 5.5 OD) it was not linear. This implies that the count rate of the input signal was not being correctly sampled by the TDC and further characterisation of this behaviour was needed; as will be explained in the following paragraphs, the TDC7201 cannot match the input rate at count rates close to 10^5 events per second. Where it is linear, the slope is -0.85 (a decrease of 8.5 per decade of attenuation) for the TDC7201, close to the -0.88 of the Becker & Hickl. These slopes are not equal to 1, but their similarity suggests it may be due to the calibration of the absorption values in the filter wheel. This could also explain the discontinuities seen in the intensity graphs in Figure 6.9, where the last point in every coloured segment is at a higher intensity than the first point of the next, suggesting the attenuation was smaller.

Another series of experiments were performed to characterise the maximum count rate achievable by the TDC7201 chip. Neither the data sheet nor the informa-



(b) Plots for the TDC7201
Figure 6.9: Plot of the intensity (log scale) and mean time of flight (linear scale) with maxto-min variation against attenuation (OD) for a free-space experiment. The range where the mean time is stable is highlighted in blue.

8.0

7.5

7.0

24350

4.0

4.5

5.0

5.5

6.0

attenuation (OD)

6.5

7.0

7.5

8.0

 10^{3}

4.5

5.0

5.5

6.0

attenuation (OD)

6.5

4.0

tion published online by the manufacturer provide a value of the maximum count rate or the minimum time between acquisitions. Experimentally, this value can be determined by comparing the event rate of an input signal (e.g. derived from a function generator) to the output count rate of the generated histograms. The maximum input frequency where the count rates match will determine the maximum count rate and the dead time of the system. A plot of the measured count rate against the true input count rate, for two different stop frequencies of 6 MHz and 15 MHz, is shown in Figure 6.10. This experiment establishes that the maximum count rate is 81 kHz, and thus the maximum sampling frequency is 81 kS/s (corresponding to a sampling period of approximately 12.35 μ s). After an event is detected, the system cannot register another until this *dead time* period has elapsed. Thus when the

input rate exceeds 81 kHz, two or more events occur during the sampling period, yet only the first is counted. This results in the observed distortion. As the input rate increases, the output rate is not capped at 81 kS/s, but instead is still a function of the input rate – after the maximum is reached the slope of the count decreases depending on the input frequency as

$$slope = \frac{1}{n+1}$$
, where $n = \left\lfloor \frac{\text{count rate}[\text{kHz}]}{81} \right\rfloor$. (6.3)



Figure 6.10: Plot of the output count rate of the TDC7201 system against the input start frequency for two different stop frequencies.

To understand this process better, a simulation was created using Python. Given an input frequency, a sequence of times are calculated at which an input event would be generated. These are then counted, but any events which occur within a 12.35 μ s window following a counted event are ignored. Later, to make it a more realistic simulation, the dead time was also applied after the stop pulse. The code for this simulation can be consulted in Appendix B.2. With this second condition, the sections where the count rate decreases also appear in the simulation – thus the loss of counts is also conditioned to the stop frequency. Figure 6.11 compares the simulation results with the real data for the two types of simulation.

If the data acquired for this experiment are compared with the results of the





(b) Count rate simulated by adding the dead time to the stop pulse time.

Figure 6.11: Plot of the output count rate of the TDC7201 system obtained with a computer simulation compared with real count rate.

intensity experiment above, it can be observed that the TDC7201 has measured count rates of 84 kS/s, higher than 81 kS/s. As the stop frequency has some observed effect on the count rate, it is possible that the faster frequency of 48 MHz can slightly increase the count rate. This situation could not be simulated with the function

generator as it could not generate this frequency. Another factor not explored here is the effects of the TDC's oscillator. If its frequency is tuned to 16 MHz, the response rate of the IC should be faster, making the count rate larger.

The consequence of this investigation is that experimental measurements with the TDC7201 must operate in the linear range (approximately 0 - 81 kS/s) to ensure that the recorded intensity is accurate and that the TPSF shape is not distorted. Thus, to guarantee stable and linear measurements during any acquisition, the light intensity incident on the detector should be decreased until the TDC produces changes in measured count rate proportional to the changes in incident intensity. It will then be necessary to record data over long integration times (e.g. around 30 - 60 s in the case of the experiments reported here) to achieve adequate signal-to-noise ratio. This type of compromise is typical of time resolved systems – the count rates and integration time of the proposed system are similar to those of MONSTIR II.

6.3.3 Single channel system – evaluation on tissue-like phantoms

To evaluate the functionality and performance of the system, and to explore whether localised variations in absorption within a tissue-like medium can be identified using measurements of mean-flight-time, a series of experiments was performed on phantoms containing targets of different absorption coefficients. For these phantom experiments, the following set up conditions were applied:

- All the circuits were shielded within aluminium boxes to reduce external electromagnetic interference and to provide a way for heat to dissipate.
- Cables connecting all the circuits were as short as possible, and crossing of other cables was avoided, to reduce attenuation and crosstalk respectively.
- When required, time delays were generated by introducing longer cables (as commercial delay units were found to introduce noise).
- Blackout cloth was wrapped around the phantom and the connecting fibres to



Figure 6.12: Schematic of the setup for phantom scans using MONSTIR II.



Figure 6.13: Imaging setup used for experiments during the project. Top row (left to right): two DC sources and a PC. Bottom row (left to right): a phantom with a SiPM probe, RF amplifier boxes, comparator box and FPGA box.

reduce the detection of background light.

- The laser and the comparator board were allowed to warm up for at least 3 h before the experiments began.
- Following the warm-up period, several TPSFs were acquired at the start of the session to ensure that the detected intensity was within the linear regime and that the mean time was stable.

The experimental setup is illustrated in Figure 6.12. A photograph of the main elements of the setup and the connection to a PC is shown in Figure 6.13.

With the help of an undergraduate student (Ioana Albu), a series of probes were

designed and 3D-printed to house the SiPM evaluation board and secure a source fibre at a fixed distance (either 30 mm or 40 mm) from the SiPM (see Figure 6.14). The end of the fibre was held a few millimetres above the surface of the phantom to enable the divergence of the beam to illuminate across the full 4 mm diameter aperture of the probe. The printed material was optically opaque, but a 10 mm layer of black foam was added to the lower surface of the probe to prevent light reflected off the surface of the phantom reaching the SiPM directly.

The first experiment on phantoms was used to confirm that the system was stable and verify that time of flight differences could be observed in response to an absorption change within the phantom. The 30 mm-separation probe was attached to a slab phantom with a movable rod containing a target. This phantom was designed to evaluate portable NIRS systems (see D. Chitnis, *et al.* [122]) and consists of a 95 × 175 × 60 mm slab of epoxy resin with a cylindrical cavity to fit a rod of the same material ($\emptyset = 10 \text{ mm}$, $\ell = 130 \text{ mm}$) with a target at its centre. By mixing the epoxy resin with known concentrations of titanium dioxide particles and near-infrared dye, the phantom is given a transport scattering coefficient of $\mu_s = 1.0 \text{ mm}^{-1}$ and an absorption coefficient of $\mu_a = 0.01 \text{ mm}^{-1}$, except for the target which has an absorption of $\mu_a = 0.1 \text{ mm}^{-1}$. The phantom is illustrated in Figure 6.15.

The probe was positioned so that the rod was aligned perpendicular to the source-detector axis, and directly below the mid-point between the source and detector, as shown in Figure 6.15. TPSFs were acquired for the target in two different locations: first, located directly below the midpoint between the source and detector; second, located at the extreme edge of the phantom, furthest from the detector. 20 TPSFs were acquired for each rod position (equivalent to a 20 s integration time). These TPSFs were then corrected for non-linearities and summed to obtain a single TPSF for each position. In this way a difference measurement can be obtained comparing the mean time with a target against the baseline. Figure 6.16 shows both TPSFs, where a difference in mean time is apparent, which is found to correspond to 66.8 ps. Although the intensities of the TPSFs appear similar, the overall count



(a) Schematic and dimensions of a single channel probe (file courtesy of Ioana Albu).



(b) Picture of the assembled single channel probe.

Figure 6.14: Schematic and picture of a single channel probe.

rate ($\sum TPSF/(20 s)$) of the TPSF with a target (62.4 kPh/s) is slightly lower to the number of photon counts without a target (65 kPh/s). The two being similar in intensity suggested that the laser intensity was not properly calibrated and the TDC7201 operated outside of the linear regime. After observing this and performing a test to characterise the count rate, more care was taken during subsequent experiments to



Figure 6.15: Schematic of the rod phantom (from [122]).





Figure 6.16: TPSFs obtained on the rod phantom, with the target in two different positions: between the source and detector (red) and at the edge of the phantom (blue).

Another, more complex phantom was designed and built specifically for this project with the help of undergraduate student Arihant Bijjala. Like the rod phantom, this was fabricated from optically clear polyester resin to which was added known concentrations of TiO₂ particles to provide scatter and near-infrared-absorbing dyes to provide absorption. The phantom, illustrated in figure 6.17, is a slab of dimensions $129 \times 190 \times 40$ mm with optical properties (at a wavelength of 800 nm) corresponding to $\mu'_s = 1.00 \text{ mm}^{-1}$ and $\mu_a = 0.01 \text{ mm}^{-1}$. The slab contains two sets of cylindrical targets (length 10 mm and diameter 10 mm) with absorption

156 Chapter 6. Development of a TOF System Based on a Commercial TDC

coefficients $5\times$, $10\times$, $20\times$ and $50\times$ that of the surrounding slab. The two sets of targets were embedded at two different depths from the surface: 10 mm and 20 mm, measured to the center of each cylinder.



Figure 6.17: Schematic diagrams of the eight-target phantom. All dimensions are given in mm.

The first experiment performed using this phantom involved scanning a single line by manually translating the 30-mm separation probe along the phantom surface in steps of 10 mm. The probe was oriented perpendicular to the row of four targets (at the depth of 10 mm), such that the the mid point between the source and detector was directly above a line through the centre of the four targets. A diagram showing the positions of the probe and the scanned points is shown in Figure 6.18.

A total of 40 TPSFs were recorded at each of 16 discrete positions, where data were collected for 1 second for each TPSF (equivalent to a 40 s integration time). The supercontinuum laser system was tuned to a wavelength of 800 nm, and an output power of 5 mW (confirmed using a power meter). An internal 2 OD filter within the MONSTIR II optics was selected, reducing the output power of the source fibre to approximately 0.05 mW, with the aim of ensuring that the detected count rate remained sufficiently low to avoid distortion of the TPSFs.

To calibrate the time of flight, a reference IRF of the system was also acquired by placing the fibre in contact with the detector with a high ND filter (6 OD) be-



(a) Alignment positions of the probe. The upper right corner was aligned with each mark. The starting position is highlighted in blue.



(b) Scanned points with position numbering

Figure 6.18: Schematic of a line scan of the first row of the eight-target phantom showing the positions of the probe and the points scanned.

tween them to adjust the intensity. The mean time of this IRF serves as the time origin for the TPSFs.

After applying the DNL correction and adding the 40 TPSFs together, the mean time and intensity were computed for each position of the probe. A plot of the mean time and intensity values as a function of probe position is shown in Figure 6.19. The locations of the four intensity minima and mean time maxima correspond with the known positions of the targets, marked by vertical lines in the figure. It can be concluded that the system is able to pick up the absorption changes produced by these targets.

As seen in the experiment characterising the behaviour of the system with increasing attenuation, the intensity surpasses the 81 kS/s limit at certain points, suggesting that the dead time does indeed depend on the stop frequency and perhaps on the distribution of the arriving events (the function simulator was generating a

158 Chapter 6. Development of a TOF System Based on a Commercial TDC

uniform distribution, different from a TPSF). As the system was very close to the count rate limit, it is possible that distortion was introduced in the measurements. This is apparent in the difference between peaks in the intensity graph, where the change in intensity does not fully match the absorption of the targets and in the mean time graph, where the background's average time of flight is lower than the targets (suggesting more absorption), where the intensity is closer to the upper limit of the system.



Figure 6.19: Light intensity and mean time for one row of the eight-target phantom, scanned in steps of 10 mm. The positions of the four targets are indicated by the dashed vertical lines.

Finally, a full raster scan of the phantom was performed. For this experiment

the slab surface was divided into 5 rows, 16 mm apart, each sampled at 15 positions, 10 mm apart. A probe with a 40 mm source-detector separation was used, to increase the sensitivity to the row of deeper targets, with the source-detector axis perpendicular to the row direction. A schematic showing the positions of the probe and the area scanned in the phantom is shown in Figure 6.20.



(a) Alignment positions of the probe. The upper right corner was aligned with each mark. The starting position is highlighted in blue.



(b) Scanned area with row and column numbering

Figure 6.20: Schematic of a raster scan of the eight-target phantom showing the positions of the probe and the area scanned.

A total of 20 TPSFs were recorded at each position, with an integration time of 1 second for each TPSF (equivalent to a 20 s integration time). While keeping the 2 OD filter, the power of the laser was reduced to 4 mW to reduce the risk of exceeding the 81 kS/s count limit. The file for one acquisition (row 1 position 6) was unfortunately lost, so the missing data point was simulated by interpolation, i.e. by averaging the data recorded at neighbouring locations – (1,5), (1,7) and (2,6).

2D greyscale images of the calculated mean time and intensity values are



160 Chapter 6. Development of a TOF System Based on a Commercial TDC

Figure 6.21: 2D images generated from intensity and mean time measurements of a surface scan of the phantom, raster scanned in steps of 10 mm. The positions of the targets are indicated by red dots.

shown in Figure 6.21, and the intensity and mean time values for each row are plotted in Figure 6.22. Again, the locations of the intensity minima and mean time maxima correspond with the known positions of the targets, marked by vertical lines. Towards the lower edge of the phantom, targets become harder to distinguish. It is possible that this effect was due to the proximity of the detector to the edge of the phantom, where more stray light is allowed to enter.

As expected, the targets produce a localised decrease in intensity. For the top



Figure 6.22: Intensity and mean time values for five rows of the eight-target phantom, scanned in steps of 10 mm. The positions of the targets are indicated by vertical lines. The red dot in the top row represents interpolated data.

162 Chapter 6. Development of a TOF System Based on a Commercial TDC

row of targets, where the target depth is 10 mm, the intensity minimum is greater as the target absorption increases (from left to right). For the bottom row of targets, at a depth of 20 mm, the intensity minima are are harder to differentiate. This is because the 3D probability distribution of detected light (the so-called photonmeasurement density function or PMDF) is broader at greater depths, resulting in lower spatial resolution. As is well known, the PMDF for mean time is narrower than than for intensity, and the mean time values differentiate the deeper targets more clearly [228]. Curiously, however, the influence of the absorbing targets on the measured mean time values is different for the two different depths. For the top row of targets, the mean time clearly increases with increasing target absorption (from left to right). For the bottom row of targets, located at a larger depth of 20 mm, the mean time decreases with increasing target absorption.

This phenomenon can be explained by considering the paths traversed by photons between the source and detector, as illustrated in Figure 6.23 for different target depths. When targets are closer to the surface, detected photons with longer flight-times are more likely to have been able to traverse beneath the targets without interacting with it, and thus the target preferentially absorbs shorter flight-time photons (leading to an increase in mean time). For the deeper targets, the detected photons with shorter flight-time are more likely to avoid being absorbed (leading to a decrease in mean time).

This observation has been confirmed using Monte Carlo simulations performed by masters student Diana Sakaan. She simulated a slab with optical properties $\mu_a = 0.01 \text{ mm}^{-1}$ and $\mu'_s = 1.0 \text{ mm}^{-1}$, illuminated with a source-detector separation of 40 mm. A spherical target with the same volume as the cylindrical targets in the real phantom is simulated for depths of 10 mm and 20 mm and absorption coefficients of 0.01, 0.05, 0.1, 0.2 and 0.5 mm⁻¹, injecting 10,000,000 photons for each simulation. The mean time of flight calculated with the simulation for both target depths is plotted in Figure 6.24. While the mean time for the 20 mm-deep target follows a monotonically decreasing trend, as the absorption tends to remove longer flight-time photons, the curve for the 10 mm-deep target has a rising trend.



Figure 6.23: Illustration of the most likely photon paths when scanning a target at different depths with a 40 mm source-detector separation.



Figure 6.24: Plot of the mean time of flight against the absorption coefficient for different target depths, from a Monte Carlo simulation preformed by Diana Sakaan.

6.3.4 Single channel system – tests on human volunteers

In vivo experiments on a human volunteer were planned for March 2020. Unfortunately, due to UCL's shutdown during the COVID-19 pandemic, these tests could not be performed.

6.4 Two-channel system

6.4.1 Design differences for the two channel system

To exploit the availability of two parallel channels on the TDC7201 chip, two issues must be addressed. First, the IC has two independent output lines, but only one input where all commands, including the read requests, are sent; the TDC channel receiving commands is selected via an independent chip select line, but clashes will occur when requesting information from both channels at the same time. Second, neither the IC (via commands) nor the test board offer the possibility of connecting the same signal to two stop inputs; the same synchronisation signal must be externally connected to both TDCs without harming the signal integrity or introducing additional jitter to the system.

In order to solve the first problem, two different strategies were tested: either attempt to read both channels at the same time if two events have been received or prioritise one channel over the other.

To implement the simultaneous read strategy, the architecture of the controller in the FPGA needs to be doubled and parallelised – two independent microcontrollers sort out the events belonging to their respective channel into a separate histogram memory and exchange coordination signals to start the acquisition and to request a simultaneous read operation if an event is registered on both channels. The main risk of this strategy is insufficient coordination between the microcontrollers. If the coordination signals are not issued at the right moment, one or both channels can get blocked; but if too much signaling is used, the additional overhead caused by it can reduce the performance of the system. The method I implemented to evaluate this strategy employs a simple two-way handshake between the channels to change between the steps of the acquisition process and then relies on an external VHDL *mutex* to coordinate the access to the SPI component. This mutex works as a 'locked' signal that keeps a channel waiting while the SPI is busy, as illustrated in the flow diagram shown in Figure 6.25.

If a channel makes a request to the SPI, it will not be acknowledged until the

component is in the 'idle' state, and the microcontroller issuing the request will remain waiting. The SPI component was also modified to allow the case where one microcontroller sends data to two channels simultaneously – to simplify the program, only microcontroller number 1 can request a simultaneous read. Figure 6.26 shows a block diagram of the structure of the parallel system, while Figure 6.27 compares the program flow for each microcontroller.



Figure 6.25: Flow diagram of the mutex that controls the access to the SPI lines.



Figure 6.26: Block diagram of a two-channel system with two microcontrollers working in parallel. The dashed lines represent signals that are connected to the micro-controller appearing on the head of the arrow.



(*) Does not control the UART operation. Just waits until the PC request has arrived.

Figure 6.27: Flow diagrams for the programs in each of the two channels of the parallelised system. The steps that require access to the SPI mutex are indicated in blue.

168 Chapter 6. Development of a TOF System Based on a Commercial TDC

The second strategy – giving priority to a channel – can be implemented using just one microcontroller that has access to both chip select lines and to two histogram memories. As the program is sequential, the status of one channel is checked before the other, giving it greater priority. If this blocks the controller from reading the non-prioritised channel, the count rate in the two channels will be not match – the maximum count limit of one channel will be significantly greater than the other. Alternatively, if the status of both channels is checked at the same time, before either channel is read, the channel with the greater input event rate can *starve* the other, that is, cause the channel receiving fewer events to be even less likely to be read. To minimise this shortcoming, in my implementation one channel (channel 1) is read before the other (channel 2), but the status of channel 2 is checked immediately after channel 1's, regardless of whether there was an event or not. To enable the register reading subroutine to work generically with two different channels, a memory address is passed to it, and its contents allow the subroutine to identify the TDC channel. A block diagram of the system is shown in Figure 6.28 and a flow diagram illustrating the operation of the register reading subroutine is shown in Figure 6.29.



Figure 6.28: Block diagram of a two-channel system with one microcontroller coordinating two memory banks.



Figure 6.29: Flow diagram for the two-channel system operated through one microcontroller and the modified register reading subroutine.

170 Chapter 6. Development of a TOF System Based on a Commercial TDC

For both strategies, the calibration value is requested once per integration period for each channel.

In practice, the alternative strategies exhibit the same behaviour, which is explored in more detail in the next section: the maximum count rate is 'divided' between the channels, and the linear regime cuts off at approximately half the maximum count rate of the one-channel system. Unfortunately the extra time spent on coordination slows the process down, such that the gains from the parallel architecture become negligible. The parallel solution was also more prone to freezing after the last acknowledgement step, likely because of a mishandled coordination signal, and therefore the sequential solution was implemented for phantom experiments. Incidentally, this approach also requires slightly less resources from the FPGA, as it does not need to implement two sets of microcontrollers. The final version of the system utilised 65.95% of the FPGA memory blocks, but only 3% of the slice logic and 20% of the clocking networks.

Using 16-bit histograms, only two channels can fit within the 7A100T FPGA, but the resource count can be decreased if the memory is restricted to the 768 bin range used for experiments. Reducing the number of bins to $2^{10} = 1024$ (the smallest power of 2 > 768) decreases the number of resources needed in the FPGA to 2.22% of the memory blocks. The Nexys 4 DDR board has four generic 8-pin I/O ports available and uses one port and two additional pins to communicate with the TDC7201 evaluation board, meaning that, with the appropriate hardware adaptations, three boards (6 channels) could be controlled simultaneously with the same FPGA.

6.4.2 Tests using the custom conditioning system

To calibrate the experimental data and to investigate the behaviour of the twochannel system, two tests were carried out – one to determine the maximum count rate of each channel and another one to assess the signal integrity of the system. During initial experiments, one of the channels on the TDC7201 evaluation board stopped working and it was replaced with a new one. The LSB calibration values

171

(i.e. the temporal sampling intervals) were observed to be different between channels and also between boards: for the first evaluation board, the values were 54.3 ps and 52 ps, while for the second board, they were 59.8 ps and 60.6 ps. The cause of any differences in the calibration values for two apparently identical boards from the same manufacturer is unknown.

An equivalent to the one-channel count limit test was performed on the twochannel system. On this occasion, three signals are used: a common 15 MHz stop, generated with a function generator, and two start frequencies, one provided by a function generator and the other one provided by an oscilloscope. A range of frequency combinations were tested by leaving one frequency constant, and incrementing the other frequency in 10 kHz steps, while the total count rate is derived from the output histograms. Figure 6.30 shows a plot of the calculated output count rate against the input rate for a range of start frequencies, for different fixed values of the other start frequency. Although an asymmetry in count rate between channels is noticeable – channel 1 is able to register higher count rates than channel 2 – in general, count rates below 40 kS/s are correctly registered in both channels. For two-channel acquisitions, the total event rate per channel should stay below this threshold to avoid non-linear distortion of the output histograms.



Figure 6.30: Plot of the output count rate of the two-channel TDC7201 system against the input start frequency, for several combinations of start frequency. For each channel, the coloured lines represent the frequency on the other channel.

It should be noted that 40 kHz is about half the count rate maximum observed for the single channel system, and that the combined count rate for both channels –

that is the maximum number of rate of events that the TDC7201 chip could respond to - peaks at 93.7 kHz, a number which is close to the maximum observed during the phantom experiments using the single channel system (see Section 6.3.3).

The second series of tests were designed to characterise the effect of two different connection topologies for the sync signal. Each input to the TDC evaluation board has a 50 Ω load. Ideally each should receive a clean signal free from noise or distortion, and both signals should arrive at the same time (or with a fixed and known delay) to make the time bases of both channels comparable. Two possible connection topologies using the comparator board, coaxial cables and the TDC were tested. The circuits for each topology are shown in Figure 6.31. Each has its advantages and drawbacks.



Figure 6.31: Two alternative topologies to connect the sync signal from the comparator board to two stop inputs.

For configuration (*a*) the sync signal is derived from the same IC, providing a fixed and constant delay (ideally zero if both branches are the exact same length). However, the load impedance of 25Ω (due to two 50Ω resistors in parallel) does not match the impedance of the circuit, and creates distortions (an oscilloscope capture of the two sync signals is shown in Figure 6.32a).

Configuration (b) which employs two different thresholding circuits avoids this distortion (see Figure 6.32b). However, as each sync is then independent, the timing difference between the channels has some added uncertainty.

When comparing the results of the two alternatives, configuration (a) did introduce additional jitter – it was as high as 60 ps in some of the experiments – but



Figure 6.32: Comparison of the sync signal for two channels using two alternative circuit topologies.

combination (*b*) produced distortions in the histograms. A comparison of the resulting IRFs, recorded using the Becker & Hickl TCSPC card, is shown in Figure 6.33.

This distortion can be explained as timing errors introduced in the initial stage of the comparator circuit. For configuration (b) the original sync signal, which has a low voltage, has to feed two different comparator channels. The noise introduced when splitting the signal can cause the circuit to pick up pulses at erroneous times. As a consequence a decision was made to employ configuration (a) for phantom experiments. Extra precautions should be taken in practice, because sudden changes in impedance cause the comparator output to spike momentarily, which can damage the circuits connected to the sync. To avoid this, the sync lead should not be connected to the two TDC inputs in sequence.

If the comparator circuit and the TDC are to be inserted into the same circuit board, Figure 6.34 illustrates alternatives topologies which avoid the impedance mismatch. This should reduce the distortion and provide better results.

6.4.3 Two channel system – test on phantoms

The two-channel system was evaluated by performing a scan of one row of the eighttarget phantom, using a new two-detector probe, also 3D printed with the help of



Figure 6.33: Comparison of the free-space IRFs obtained from one channel using two connection topologies for the sync signal.



Figure 6.34: Two alternative topologies to connect the synchronisation signal to two channels without introducing an impedance mismatch.

Ioana Albu. The scanned row corresponded to the targets at a depth of 20 mm. This probe, shown schematically in Figure 6.35, can accommodate two SiPM boards, with a vertical space between them of 19 mm, and two fibres in the fibre holders, allowing diffuse reflection measurements using two sources and two detectors. The additional vertical height of one of the detectors will produce a very small delay in the flight times registered by the TDC, which can be calibrated away using the IRF, and a possible loss of signal intensity which should be negligible. With this setup four separate source-detector measurements can be obtained – two measurements at 35 mm separation and two at 49.5 mm (using diagonally opposite sources/de-



Figure 6.35: Rendering (courtesy of Ioana Albu) and dimension diagrams of the dual channel probe. Dimensions are in mm.

tectors). For each source illumination, two TPSFs are collected simultaneously. In addition, four different IRFs are now necessary to calibrate the data – one for each source-detector combination. For this experiment, each fibre and detector were connected directly together, with a very high OD filter between them, and then the IRF was manually collected. The computer software can then treat each source-detector combination as a different virtual channel.

A schematic diagram of the experiment setup and the area covered by the probe is shown in Figure 6.36. Soft black foam was placed on the lower surface of the probe (with holes for each source and detector) to reduce stray light, and the whole probe and phantom were wrapped in black cloth. For each position 20 TPSFs were collected, integrating for 1 s in each TPSF.

The resulting evolution of the intensity and the mean time of flight measured for all four channels as the probe was scanned across the phantom is shown in



(a) The probe positions on the phantom. The upper right corner was aligned with a mark on the phantom surface. The starting position of the probe is highlighted in blue.



(b) The source and detector positions with the direct measurement and cross-measurement paths. Each colour represents one probe position.

Figure 6.36: A scan of one row of the eight-target phantom with the two-channel system.

Figure 6.37. The general trend is a decrease in mean time as the probe scans across targets of increasing absorption (except for the value in position 11, where stray light may have influenced the measurement to the probe being so close to the edge of the phantom).

The two different source-detector separations produce measurements with noncomparable time offsets, and relate to different scanned volumes. The DPF, introduced in equation 3.1, serves as a distance-normalised value that facilitates the comparison between the measurements. In this case, for each data point, the DPF is calculated as $DPF = (c_0/n) \cdot \langle t \rangle / d_{SD}$, where $c_0 = 3 \times 10^8$ m/s is the speed of light in vacuum; n = 1.55 is the refraction index of the phantom's resin; $\langle t \rangle$ is the mean



Figure 6.37: Comparison of the evolution of the intensity and mean time of flight of all source-detector combinations in a row scan of the eight-target phantom. The positions of the targets are marked by blue strips. Position marks in mm.

flight time obtained from the TPSF calibrated using the IRF; and d_{SD} is the sourcedetector separation (35 mm or 49.49 mm). Figure 6.38 shows the DPF values for all four channels plotted against position, defined as the horizontal location of the midpoint between the corresponding source and detector for that channel. The vertical blue bands indicate the locations of the four targets. Unfortunately the data does not exhibit the expected correlation between DPF (i.e. mean time) and target location, as was observed for the single channel system. Some channels have recorded a very unexpected variation. Excessive noise in the cables for channel number 2 or a slight difference in the comparator settings (either the threshold voltage or the hysteresis potentiometer) might explain the poorer stability compared with channel number 1. Noisy mean time values could be attributed to a poor signal to noise ratio combined



Figure 6.38: Evolution of the DPF of all source-detector combinations with position for a row of the eight-target phantom. The positions of the targets are marked by blue strips.

with the lower count capacity of the two-channel system, that could be improved by increasing the intensity of the laser. However, the intensity of the laser can also affect the sync signal, since it also feeds the sync photodetector and thresholding circuit, potentially leading to errors in timing. The jitter of the mean time (maximum – minimum) for most measurements was around 15 ps for channel 1 and 25 ps for channel 2, but some measurements past position 100 mm varied by over 80 ps, suggesting that external noise and signal degradation contributed significantly to the errors. A plot of the jitter for each position is shown in Figure 6.39.

Unfortunately, these hypotheses could not be tested and the experiment could not be repeated due to the unforeseen closure of UCL during the COVID-19 pandemic.

6.5 Conclusions

In this chapter I have described the design of the component architecture and readout for a time domain imaging system based on an affordable commercial TDC integrated circuit, and demonstrated its operation in a single-channel and a dualchannel mode. The linearity, resolution and stability of the system were tested and



Figure 6.39: Plot of the jitter of the mean time of flight for each position for the two channel row scan experiment.

1D and 2D imaging experiments on phantoms were presented.

The main constraint on the design was the need for fast and continuous acquisition. For any number of channels, an FPGA collects only the minimum number of IC registers to calculate a timestamp of the events and then creates a histogram, which can be sent over to a PC. If two channels are being polled at the same time, the event interrupt lines are read in sequence causing a slight imbalance in count rate. Compared to the FPGA-based systems introduced in the previous chapter, the main differences are the sequential nature of the data acquisition, which can lead to more events being missed, the linearity, which is superior due to the ring oscillator implementation, and the stability of the TDC, which is guaranteed by the manufacturing and calibration processes.

The characteristics of the system were measured experimentally. After the laser and signal conditioning have completed their stabilisation period, the jitter of the mean time can go down to 20 ps, or up to 60 ps when interference from external equipment is higher. Using metal boxes for isolation reduces the jitter and the occurrence of spurious spikes. The actual bin size, which is nominally 55 ps, depends on the individual evaluation board used, but remains stable with a variation of less than a picosecond over a period of 1 minute. For experiments, it is advisable

180 Chapter 6. Development of a TOF System Based on a Commercial TDC

to let the system warm up and then measure the bin resolution at the start of the experiment, and again during the experiment if it takes more than 1 h. The system was shown to behave linearly, with a sensitivity to detect changes in mean flight time of tens of picosecondst, dependent on the jitter.

The count rate of the system was also found to be variable, depending on the clock frequencies of the TDC, the event signal, and the sync signal. The maximum count rate for the single channel system was estimated to be 81 kS/s with a function generator and around 90 kS/s using a laser's sync signal. For the dual-channel system, the count rate is halved. The system should be operated below this count rate, by reducing the intensity of the light output, to avoid distortion in the TPSFs and erroneous estimations of the mean time of flight.

Other experiments involved scanning the probe across the surface of phantoms. The system could detect differences in time of flight due to targets at different depths and with different absorption coefficients. Results also explored the relationship between mean-time changes and the depth of targets, which was then confirmed using Monte Carlo simulations.

To enable 3D imaging, data from two simultaneous channels were collected. However, the system was not able to identify changes in the mean time for stability reasons that could not be fully explored.

These experiments have shown that the Texas Instruments TDC7201, or an integrated circuit with similar characteristics, can be successfully employed to measure photon times of flight differences for diffuse optical tomography, an application which is more demanding than those for which it was originally designed. This enables the possibility of building a portable imaging system with affordable, easy to obtain, easy to configure, low footprint components.

Time constraints and the unexpected closure of facilities due to the Covid-19 pandemic left some experiments to be done in the future. To demonstrate the full capabilities of the system, a complete 3D scan should be performed and image reconstructed using the dual-channel system. Eventually, this system should be also tested on human volunteers to assess the sensitivity to biological changes
(e.g. oxygenation, blood volume or flow changes) and the variability of the data, discriminating between the differences introduced by the system (due to movement, variable surface coupling, and noise) and those due to the changing physiology of the subject

To integrate the TDC7201 into a more portable or distributed device, the IC can be incorporated into the end of the signal conditioning chain. This reduces the size of the device, and also decreases the noise introduced by interference and connectors. By taking advantage of the two channel configuration, and using an FPGA as a hub for multiple chips, a distributed system with multiple detectors per probe can also be designed.

Chapter 7

Conclusions and Future Work

7.1 Conclusions

In this thesis a pickup and signal conditioning circuit and two types of system to measure the time-of-flight of photons have been proposed. These systems have subsequently been tested with different experimental setups: a function generator for functionality tests; a free space setup with a pulsed laser and modifiable light attenuation and source-detector distance; and a fibre-coupled setup to conduct transmittance experiments on tissue-equivalent phantoms. In addition to the final system several alternatives for their design have been proposed.

The conditioning circuit has been designed utilising low footprint off-the-shelf components and is capable of picking up amplified pulses from photodiodes, in order to obtain a synchronisation signal, and from SiPMs, to register single photon events. The threshold voltages and noise sensitivity can be externally adjusted. The circuit provides 3 V pulses suitable for use with many commercial circuits with a jitter on the output of around 10 ps. The possibility of converting the circuit into a constant fraction discriminator has also been discussed, and with sufficient development, could be used in other time of flight related projects. The CFD provides more precise identification of the time of arrival of events and will be necessary if the detector's output has a larger range of variation in voltage than in the single-photon regime used during the experiments in this thesis.

Chapter 7. Conclusions and Future Work

The first timing circuit is fully integrated within an FPGA. A delay line made from internal elements detects the individual times of arrival of events, while a histogram of the times of arrival is simultaneously built in the internal memory. Two different versions have been explored, one generating timestamps as an intermediate step, capable of operating at 100 MS/s and one with a direct-to-histogram architecture, with virtually no dead time. This systems can achieve temporal resolutions of 34 ps/bin once non-linearity corrections are applied. The reconfigurable design allows the user to redefine the number of bins produced by the system or introduce new elements in the design without the need to replace the physical circuits or, although this has not been tested, to program multiple channels on the FPGA. The experiments have shown that, while the time-to-digital converter programmed in the FPGA behaves linearly, the individual delay produced by the delay line elements is highly non-uniform and requires software correction, and the histograms obtained with the full signal conditioning system exhibit a large jitter.

The second timing circuit employs a commercial time-to-digital converter integrated circuit to provide the timing of events and an FPGA to read out the data and build histograms of the times of flight. This system can achieve up to 90 kS/s in single-channel measurements and up to 40 kS/s in dual-channel measurements with a jitter as low as 20 ps or as high as 60 ps depending on the external noise levels. This mixed system is still reconfigurable to an extent, as the readout program and histogram building are performed inside the FPGA. Its sampling rate is comparable to that of the UCL MONSTIR II system, requiring several seconds to record at one position with adequate SNR, i.e with a number of counts for the TPSF about two orders of magnitude above the noise level.

For both systems, the data are sent through a high speed serial-over-USB connection to a central PC, which stores the histograms and is then used for further data processing. The data acquisition chain, which uses a widely available and inexpensive communication protocol and software written in Python, can be easily exported to other computers or integrated in low power systems and portable single board computers, provided that reliable data storage is available. If this system is

developed into a wearable format, the central control unit can be easily built to deliver relevant statistics, such as mean time, jitter of the mean, intensity, or full width at half maximum of the histogram.

Details have been provided throughout the thesis on the technical implementation and on possible alternatives to the final design of the working system, which can be used to replicate the design or applied to other optical systems with different requirements.

Imaging results have been obtained for setups consisting of one source and one detector (one TDC channel), and two sources with two detectors (two TDC channels, yielding four measurements per location), with the aid of 3D printed probes that support the SiPMs and source fibres. These probes are not suitable for brain imaging measurements on infants, but they could be used for simple *in vivo* experiments. The single-channel setup was able to observe differences in average times of flight on tissue-equivalent resin-based phantoms containing targets of various absorptions set at two different depths. The dual-channel system produced mixed results. Unfortunately no results could be presented for *in vivo* experiments.

The systems described in this thesis have two important limitations. First, they rely on fibre coupling for the probes and on coaxial cables between some block circuits, so they are not immediately portable. Second, the timing results exhibited variable amounts of jitter and sensitivity to changes in the input. The systems are very dependent on the fine tuning of thresholds, cable lengths and positions, and quality of the contacts. Integrating the detector and signal processing within the same circuit, as would be achieved in a tile-based system, as discussed in Section 2.1.4, can improve the signal quality, reduce the jitter, and make measurements more consistent.

7.2 Wearability and power distribution

The primary aim of this project was to develop a time domain imaging system that is easy to translate to a portable, distributed and inexpensive form factor, for use in

Chapter 7. Conclusions and Future Work

the back of an ambulance or by the cot side. In its current form, the system is not easily portable, as it relies on a fibre-coupled laser, and is not suitable for imaging studies on infants, where a safe and comfortable means of coupling the system to the scalp is essential. However, most of the elements can be easily transferred to a single, small-sized PCB, facilitating the design of a wearable system.

For the current version of the system presented in this thesis, a 100×100 mm PCB was used in order to accommodate many signal routing combinations, and test points and means of power generation. This is too large to integrate on a wearable probe. For a higher integration level and a reduction in size, a PCB with more than 4 layers can be used. This can accommodate the power distribution, ground planes and signal distribution between stages of the signal conditioning circuit in the extra layers with the configurations that are known to work well, as the footprints of the integrated circuits themselves are not large. The voltage converters should be moved away from the main PCB as they are the main source of heat and of current consumption, and thus should not be placed close to human skin.

Moreover, more elements should be fitted on the same PCB, taking advantage of the new power routing, to achieve better signal transmission and ultimately improve the SNR, incorporating the TDC and the signal amplification. The latter has been performed using large profile circuit blocks, but there exist lower profile alternatives suitable for this application, such as equivalent or similar circuits from the same manufacturer, Minicircuits, with a lower profile, or high frequency operational amplifiers as demonstrated by Buttafava *et al* [93]. With a smaller probe/PCB, the SiPM and possibly a laser can be included with the rest of the circuitry – one of the challenges as explained by Zhao and Cooper [16] is isolating the circuit appropriately, as the bias needed for operation is very high. The last element used to build the system, the FPGA, can be placed outside of the main 'tile' into a control unit and connected to the communication lines provided by multiple detectors, saving space on the PCB while fully using the FPGA resources.

The cost of production also benefits from using a distributed system. The size of the PCB, rather than the number of units, is usually the factor that drives the price

in PCB manufacturing. For example, sets of up to 5 identical circuits can be printed by the company PCBtrain, which assembled the circuits used during the project, at a cost of around £400.

Finally, a wearable system will require a suitable integrated light source. The source should small in size (adequate for a portable PCB), have a low power consumption, yield an average output power of the order of milliwatts, and be able to provide short (ideally picosecond) pulses at a high (MHz) repetition rate. The two most likely candidates for use as sources are laser diodes and VCSELs. Traditional photodiodes are typically larger and more power consuming than VCSELs, making the latter the most adequate choice in this regard. However, the output power and choice of wavelengths are currently limited. Both types of source would need to be connected to a pulsed driver that is able to produce picosecond pulses at high frequencies with a very low jitter. The solutions available on the market for a low profile driver are scarce and not yet fully developed. VCSEL technology is an active and rapidly advancing field, so new options for a portable source are expected in the next few years. An alternative to an in-tile source would be to retain an external laser fibre-coupled to each 'tile'. This would greatly reduce the portability of the device, making it non wearable, but still suitable for use at the cot-side or in the back of the ambulance.

7.3 Comparison to other work

With respect to other recent small or low power time domain systems, the two proposals described in this thesis compare differently. The fully FPGA-based system achieves resolutions comparable to other FPGA TDCs based on a similar architecture and, with a more stable signal chain or an improved direct-to-histogram design, could compete in terms of count rate with imagers made from ASIC or custom silicon. On the other hand, the TDC7201-based system, while operational, possesses a coarser time resolution and a slower count rate.

In terms of portability and accessibility of the design, the TDC7201-based sys-

tem has significant unique advantages, providing a way to assemble a cost effective time domain imaging system from off-the-shelf components that could be replaced based on availability.

The system described by researchers at Politecnico di Milano [93, 94] is similar in architecture to the one proposed in this thesis. Their time domain system includes two time-to-digital channels, a similar circuit for the pulse shaping and thresholding, and an FPGA-based central control unit. It also employs SiPMs used as detectors and experiments are performed using a fibre-coupled laser. The choice of elements and connection architecture come from similar design principles and therefore it is not unexpected to see similarities.

Some key differences between the two approaches nevertheless exist. The main difference is that while the Milan group's work relies on custom made integrated circuits for the time-to-digital converter, implemented on ASICs, this thesis proposes and documents ways to implement them in reconfigurable hardware, providing more flexibility to the system implementation. Another notable difference is the discussion of the time base coordination in multiple channel systems and about the scalability of the design, which has had a stronger focus in this thesis than in the papers published by the Milan group. In contrast, their work has focused more on multiple wavelength imaging, which has not been reported in my thesis. One of the ultimate aims of this research project, wearability, is not an expressed aim of the work published by the Milan group. For example, the integration with portable probes is not fully discussed and their system remains fibre coupled. The use of cer-tain circuit components, such as ASICs or diode lasers does not necessarily translate into a small 'tile'-like probe.

7.4 Future work

Some validation experiments and development challenges remain to be addressed with the systems developed during this project. As a short term goal, the unfinished experiments mentioned in previous chapters should be completed to validate the

multiple channel imaging system using phantoms and for simple *in vivo* measurements on human volunteers. The resulting data can be used to generate μ_a and μ'_s surface maps (or 3D reconstructions of the volume). A comparison of these results with the expected values, or with simulated models, can provide another validation step for the system and be used to calculate calibration values, if needed.

Another immediate aim is to write up an account of the design and performance of the system based on the Texas Instruments TDC7201, suitable for publication in a scientific journal.

A longer term objective of this work is further miniaturisation to achieve a fully integrated probe. Given the developments of the TDC7201-based system, a smaller PCB can be produced containing (at least) the signal processing and TDC, with optical components coupled to the head via optical fibres. At this stage, the power distribution to the PCB and the data communication lines can already be tested by building an external power supply and using flexible cabling (e.g. ribbon cable) to connect the PCB to the control FPGA. This way, subsequent iterations of the design can focus more on the in-tile distribution of components and the scalability of the system. For example, as the number of tiles grows, the amount of data can be overwhelming to the communication interface, so faster communication systems, data compression, on-tile calculation of statistics and saving data in an on-tile file system would have to be considered. The form factor of the probe casing will have to be redesigned to fit the new PCB, being careful to make it light and, with the addition of a suitable layer of insulation, comfortable for the patient.

Following this, to produce a fully wearable 'tile', it would need to incorporate detectors (SiPMs) and sources (VCSELs). The first challenge posed by this is that the power distribution issues, and associated safety concerns when worn, become more complex and require stricter insulation and possibly heat dissipation. As mentioned above, finding a ready to use, suitable light source is currently very challenging.

With any choice of detector or source, the impact of the motion of the probe relative to the head on the acquired measurements should be studied. Although time-of-flight measurements should be robust to changes in photon count rate (intensity), alterations due to movement of the sources and detectors and electrical noise should not be discounted.

Another major issue to address when developing a distributed system is the coordination of the time base between 'tiles'. In this setup each detector or pair of detectors will have a different and independent oscillator, with its own inherent variability in frequency. This will cause drifts in the time bases for each 'tile', making them impossible to compare, and preventing reliable measurements when sources and detectors are mounted on different tiles. The system needs a way of coordinating these time bases or at least determining the differences between them relative to a common time. This issue has not been discussed extensively in the published literature related to miniaturised or portable TD-NIRS systems, but has been addressed with respect to other devices, where two possible approaches are evident. The first approach consists of adding a time 'marker' to each collected histogram which is linked to the time at which the source is pulsed and that does not depend on the tile's time base. For MONSTIR II, for example, this is achieved optically during the 'absolute calibration' step (see Section 3.9), where the timing for each detector is calibrated against the laser repetition cycle. To replicate this approach for a distributed system (where sources are mounted on every tile) would require additional optical fibres between tiles, so that a very small fraction of the source power would be fed directly to neighbouring tiles, providing a time stamp for the histograms acquired on those tiles. A more integrated and elegant solution would be to coordinate the time bases purely by electronic signals. The same pre-pulse could be generated electronically and distributed to the TDCs with a controlled delay. A disadvantage of this being that electric transmission lines are typically noisier than optical lines and therefore introduce more jitter. The second approach involves implementing a network of clock coordination information. Numerous algorithms exist for sharing timing information over a distributed system, mainly to coordinate the real time in networked computers. The main disadvantage of this approach is that the sharing of information and the computing of the common clock base add a significant overhead to the system, which depends on having real time access to data. The use of an FPGA for coordination can ease the overhead by using parallelised components for each task.

Finally, the FPGA system should be validated with other components. With the proper coaxial connections existing for other evaluation boards, a stable time domain system completely integrated in the same circuit can be already tested. The footprint of the FPGA can make it more challenging to design 'tile'-like probes with one FPGA per probe. Using one FPGA as the central node between various sources and detectors is a more cost- and space-efficient solution. Development of this kind of system can also lead to exploration of the online processing of timestamps – methods for the selection and correction of timestamps or histograms have been used in the past to improve the quality of the acquisition and the application of existing or new correction methods can compensate for the limitations in monotonicity inherent to the FPGA.

Appendix A

Circuits and Schematics















(a) Main comparator and CFD



Figure A.2: PCB layout for the non-working designs of the main comparator, CFD and sync processing circuits. (Scaled to 50%)











(a) Top layer



(**b**) Bottom layer (mirrored)

Figure A.4: PCB layout for the functional board.



(c) Ground plane



Figure A.4: PCB layout for the comparator board (cont.).

Appendix B

Code

B.1 Simulation of the direct-to-histogram TDC

```
1 import numpy as np
2 from matplotlib import pyplot as plt
3 import operator
4
5 ## Function definitions ##
6 def nullfunction(size):
7
       return np.array([0]*size)
8
9
   def generateClkPoints(period, numPoints, offset, jitterDistribution):
       realOffset = int(np.round(offset % period))
10
       numCycles = int(np.floor_divide(numPoints, period))
11
       idxBase = np.array(range(numCycles))*period + realOffset
12
       jitter = (period/22 * jitterDistribution(size=numCycles)).astype(int)
13
       idxJitter = idxBase + jitter
14
15
16
       return idxJitter.astype(int)
17
18
   def getPointByPointHistogram(signalPts, clk, dels):
19
       histogramBins = np.zeros(len(dels)-1)
20
21
       for I in range(len(clk)):
22
           delayedSignal = 0
           delayedSignal_1 = 0
23
24
           xorSignal = np.zeros(len(dels)-1)
25
           for J in range(len(dels)):
26
                signalDelayedPts = (signalPts + dels[J]).astype(int)
27
```

```
idxBeforeClk = np.where(signalDelayedPts<clk[I])[0]</pre>
28
29
                delayedSignal_1 = delayedSignal
                if len(idxBeforeClk)>0:
30
                    delayedSignal = ( idxBeforeClk[-1] + 1 ) % 2
31
32
                else:
33
                    delayedSignal = 0
34
                if J>0:
35
                    xorSignal[J-1] = delayedSignal^delayedSignal_1
36
37
            histogramBins += xorSignal
38
        return histogramBins
39
40
   ## Problem conditions ##
41
   clkLen = 100000
42
   clkPeriod = 200
43
44
   signalJitter = np.random.normal # Add random jitter
45
46
   # Axis
47
48
   delAx = np.array([ 3, 6, 9, 12, 15, 18, 21, 24, 27, 30,
                      33, 36, 39, 42, 45, 48, 51, 54, 57,
49
50
                      60, 63, 66, 69, 72, 75, 78, 81, 84,
51
                      87, 90, 93, 96, 99, 102, 105, 108, 111,
52
                      114, 117, 120, 123, 126, 129, 132])
53
54
      # Monotonically increasing
   dels = np.array([ 0, 3, 6, 9, 12, 15, 18, 21, 24, 27, 30,
55
                     33, 36, 39, 42, 45, 48, 51, 54, 57, 60,
56
57
                     63, 66, 69, 72, 75, 78, 81, 84, 87, 90,
                     93, 96, 99, 102, 105, 108, 111, 114, 117,
58
                     120, 123, 126, 129, 132])
59
60
      # Two swappes delays
61
62
    dels2 = np.array([ 0, 3, 6, 9, 12, 15, 18, 21, 24, 27, 30,
                      33, 36, 39, 42, 45, 48, 51, 54, 57, 60,
63
                      63, 66, 69, 75, 78, 81, 72, 84, 87, 90,
64
                      93, 96, 99, 102, 105, 108, 111, 114, 117,
65
                      120, 123, 126, 129, 132])
66
67
   # Signals
68
   clk = generateClkPoints(clkPeriod, clkLen, 2, nullfunction)
69
    sigBase = generateClkPoints(clkPeriod, clkLen, 0, signalJitter)
70
71
72
73
```

```
74 ## Simulate for a range of input delays ##
75 for I in range(115,160,3):
76
       print I
77
       sigIdx = sigBase + I
       monoHist = getPointByPointHistogram(sigIdx, clk, dels)
78
79
       swapHist = getPointByPointHistogram(sigIdx, clk, dels2)
80
81
       plt.figure()
       plt.plot(delAx, monoHist, 'b*-')
82
       plt.plot(delAx, swapHist, 'r*-')
83
84
       plt.xlabel('Bin no.')
       plt.ylabel('Count')
85
86
87 plt.show()
```

B.2 Simulation of the TDC7201 system's count rate

```
1 import numpy as np
2 import matplotlib.pyplot as plt
3
4 ## Function definitions ##
5 def nearestT(Tclk, tevent):
6
       ee = np.finfo(np.float32).eps # An epsilon to tip np.floor
                                      # in case of precision errors
7
       return (np.floor(tevent/Tclk + ee)+1) * Tclk
8
9
10 ## Problem conditions ##
11 # True data
12 c1K = np.array([ 1001, 2001, 3001, 4001, 5001, 6001, 7001,
     8001, 9001, 10001, 11001, 12001, 13001, 14001, 15001, 16001,
13
    17001, 18001, 19001, 20001, 21001, 22001, 23001, 24001, 25001,
14
    26001, 27001, 28001, 29001, 30001, 31001, 32001, 33001, 34001,
15
    35001, 36001, 37001, 38001, 39001, 40001, 41001, 42001, 43001,
16
    44001, 45001, 46001, 47000, 48000, 49000, 50000, 51000, 52000,
17
    53000, 54000, 55000, 56000, 57000, 58000, 59000, 60000, 61000,
18
    62000, 63000, 64000, 65000, 66000, 67000, 68000, 69000, 70000,
19
    71000, 72000, 73000, 74000, 75000, 76000, 77000, 78000, 79000,
20
21
    80000, 81000, 79099, 57087, 43320, 42511, 43001, 43501, 44001,
    44501, 45001])
22
23 c10K = np.array([45001, 50001, 55001, 60001, 65001, 70000,
24
    75000, 80000, 56677, 60001, 63334, 66667, 70001, 73334,
    76667, 80000, 70077, 65001, 67501, 70001, 72501, 75001,
25
    77501, 80000, 80500, 68001, 70001, 72001, 74001, 76001,
26
    78001, 80001, 80743, 70237, 71667, 73334, 75001, 76667,
27
```

```
78334, 80001, 81667, 73235, 72886, 74286, 75715, 77144,
28
    78572, 80001, 81429, 78642, 73792, 75001, 76251, 77501,
29
    78751, 80001, 81251, 80592, 74940, 75604, 76667, 77778,
30
    78889, 80001, 81112, 80560, 83161, 79024, 77001, 78001,
31
    79001, 80001, 81001, 80725, 80358, 78998, 77273, 78182,
32
33
    79092, 80001, 80910, 81788, 80309, 79956, 77504, 78334,
34
    79167, 80001, 80834, 81667, 80817, 79440])
35
   ##
36
   T_fmax = 1e6/np.max(c1K) # Period taken from max count rate
37
38
   T_{sampling} = 1e3/15000
39
40
   num_freq = 1000
41
   Tsim = 1e6
   counts1 = np.zeros(num_freq)
42
   counts2 = np.zeros(num_freq)
43
44
   ## Simulations ##
45
   # Simulation type 1: after each event the dead time is T_fmax
46
   for ff in range(num_freq):
47
48
        print('1. {}k'.format(ff+1))
        T_us = 1e3/(ff+1)
49
50
        t_current = 0
51
52
       while t_current < Tsim:</pre>
53
            nn = nearestT(T_us, t_current)
54
            counts1[ff] += 1
55
            t_current = nn+T_fmax
56
57
   # Simulation type 2: the dead time applies after a "stop"
   for ff in range(num_freq):
58
        print('2. {}k'.format(ff+1))
59
        T_{us} = 1e3/(ff+1)
60
        t_current = 0
61
62
63
       while t_current < Tsim:</pre>
64
            nn = nearestT(T_us, t_current)
65
            nsampling = nearestT(T_sampling, nn)
            counts2[ff] += 1
66
67
            t_current = nsampling+T_fmax
68
69
   ## Plots ##
70
   fAx = np.arange(num_freq)+1
71
72 plt.figure()
73 plt.plot(fAx, counts1/1000, 'bx', label="simulated")
```

Appendix B. Code

```
74 plt.plot(fAx[0:90], c1K/1000, 'c--', label="real")
75 plt.plot(fAx[89::10], c10K/1000, 'c--')
76 plt.xlabel('true count rate (kHz)')
77 plt.ylabel('measured count rate (kHz)')
78 plt.figure()
79 plt.plot(fAx, counts2/1000, 'bx', label="simulated")
80 plt.plot(fAx[0:90], c1K/1000, 'c--', label="real")
81 plt.plot(fAx[89::10], c10K/1000, 'c--')
82 plt.xlabel('true count rate (kHz)')
83 plt.ylabel('measured count rate (kHz)')
84 plt.show()
```

Appendix C

Datasheets

Datasheets removed due to copyright restrictions:

- IMRA-A1
- SC400
- On Semiconductor C Series
- MAX9600
- LTC6752
- AD5362/5363
- Texas Instruments TDC7201
- Digilent Nexys 4 DDR

Appendix D

Software Tools

For the development of this project, the following software tools were used:

- Python 2.7 and IPython through Anaconda, for data collection and processing and for the generation of figures with the following libraries:
 - pySerial by Chris Liechti.
 - Matplotlib, Numpy and Scipy.
- MATLAB from Mathworks, for some of the data processing.
- Xilinx Vivado for FPGA programming and debugging.
- EAGLE, from Autodesk, and Altium Designer for PCB design.
- LTSpice from Analog Devices and TINA-TI from Design Soft and Texas Instruments for circuit simulation.

For the typesetting of this document the following tools were used:

- LATEX and BibTEX, as provided by Overleaf. The document is typeset using the Times Roman, Helvetica and Computer Modern typefaces and the following additional packages: listings, algorithm, algpseudocode, rotating, multicol, outlines, glossaries, notoccite, includepdf.
- LibreOffice Draw for diagrams and schematics.
- Xcircuit for circuit schematics.

Bibliography

- Jennifer J Kurinczuk, Melanie White-Koning, and Nadia Badawi. Epidemiology of neonatal encephalopathy and hypoxic-ischaemic encephalopathy. *Early human development*, 86(6):329–338, 2010.
- [2] T. N.K. Raju, K. B. Nelson, D. Ferriero, J. K. Lynch, and the NICHD-NINDS Perinatal Stroke Workshop Participants. Ischemic perinatal stroke: Summary of a workshop sponsored by the national institute of child health and human development and the national institute of neurological disorders and stroke. *PEDIATRICS*, 120(3):609–616, Aug 2007.
- [3] Antoine Giraud, Clémence Guiraut, Mathilde Chevin, Stéphane Chabrier, and Guillaume Sébire. Role of perinatal inflammation in neonatal arterial ischemic stroke. *Frontiers in Neurology*, 8, Nov 2017.
- [4] M. F. Rafay, A.-M. Pontigon, J. Chiang, M. Adams, D. A. Jarvis, F. Silver,
 D. MacGregor, and G. A. deVeber. Delay to diagnosis in acute pediatric arterial ischemic stroke. *Stroke*, 40(1):58–64, Sep 2008.
- [5] F Cowan. Does cranial ultrasound imaging identify arterial cerebral infarction in term neonates? Archives of Disease in Childhood - Fetal and Neonatal Edition, 90(3):F252–f256, May 2005.
- [6] Jeffrey M. Perlman, Nancy K. Rollins, and Denise Evans. Neonatal stroke: Clinical characteristics and cerebral blood flow velocity measurements. *Pediatric Neurology*, 11(4):281–284, Nov 1994.

- [7] Gerda van Wezel-Meijler, Sylke J. Steggerda, and Lara M. Leijser. Cranial ultrasonography in neonates: Role and limitations. *Seminars in Perinatology*, 34(1):28–38, Feb 2010.
- [8] Tristan T. Sands and Tiffani L. McDonough. Recent advances in neonatal seizures. *Current Neurology and Neuroscience Reports*, 16(10), Sep 2016.
- [9] Jacqueline Salas, Aylin Tekes, Misun Hwang, Frances J. Northington, and Thierry A.G.M. Huisman. Head ultrasound in neonatal hypoxic-ischemic injury and its mimickers for clinicians: A review of the patterns of injury and the evolution of findings over time. *Neonatology*, 114(3):185–197, 2018.
- [10] Mateo Calderon-Arnulphi, Ali Alaraj, Sepideh Amin-Hanjani, William W. Mantulin, Chiara M. Polzonetti, Enrico Gratton, and Fady T. Charbel. Detection of cerebral ischemia in neurovascular surgery using quantitative frequency-domain near-infrared spectroscopy. *Journal of Neurosurgery*, 106(2):283–290, Feb 2007.
- [11] Joseph P Culver, Turgut Durduran, Daisuke Furuya, Cecil Cheung, Joel H Greenberg, and A G Yodh. Diffuse optical tomography of cerebral blood flow, oxygenation, and metabolism in rat during focal ischemia. *Journal of Cerebral Blood Flow & Metabolism*, 23(8):911–924, Aug 2003.
- [12] Zengyong Li, Ming Zhang, Qing Xin, Guoqiang Chen, Feifei Liu, and Jianping Li. Spectral analysis of near-infrared spectroscopy signals measured from prefrontal lobe in subjects at risk for stroke. *Medical Physics*, 39(4):2179–2185, Mar 2012.
- [13] Susanne Muehlschlegel, J. Selb, M. Patel, S. G. Diamond, M. A. Franceschini, A. G. Sorensen, D. A. Boas, and L. H. Schwamm. Feasibility of NIRS in the neurointensive care unit: A pilot study in stroke using physiological oscillations. *Neurocritical Care*, 11(2):288–295, Aug 2009.

- [14] Brian R. White, Steve M. Liao, Silvina L. Ferradal, Terrie E. Inder, and Joseph P. Culver. Bedside optical imaging of occipital resting-state functional connectivity in neonates. *NeuroImage*, 59(3):2529–2538, Feb 2012.
- [15] Antonio Pifferi, Davide Contini, Alberto Dalla Mora, Andrea Farina, Lorenzo Spinelli, and Alessandro Torricelli. New frontiers in time-domain diffuse optics, a review. *Journal of Biomedical Optics*, 21(9):091310, Jun 2016.
- [16] Hubin Zhao and Robert J. Cooper. Review of recent progress toward a fiberless, whole-scalp diffuse optical tomography system. *Neurophotonics*, 5(01):1, Sep 2017.
- [17] Marco Ferrari, Joseph P. Culver, Yoko Hoshi, and Heidrun Wabnitz. Special section guest editorial: Clinical near-infrared spectroscopy and imaging of the brain. *Neurophotonics*, 3(3):031401, Sep 2016.
- [18] Frédéric Lange and Ilias Tachtsidis. Clinical brain monitoring with time domain NIRS: A review and future perspectives. *Applied Sciences*, 9(8):1612, Apr 2019.
- [19] JE Brazy, DV Lewis, MH Mitnick, and vander Vliet FF Jöbsis. Noninvasive monitoring of cerebral oxygenation in preterm infants: preliminary observations. *Pediatrics*, 75(2):217, 1985.
- [20] EOR Reynolds, JS Wyatt, D Azzopardi, DT Delpy, EB Cady, M Cope, and S Wray. New non-invasive methods for assessing brain oxygenation and haemodynamics. *British Medical Bulletin*, 44(4):1052–1075, 1988.
- [21] J.S. Wyatt, D.T. Delpy, M. Cope, Susan Wray, and E.O.R. Reynolds. Quantification of cerebral oxygenation and haemodynamics in sick newborn infants by near infrared spectrophotometry. *The Lancet*, 328(8515):1063– 1066, Nov 1986.

- [22] Mark Cope and David T Delpy. System for long-term measurement of cerebral blood and tissue oxygenation on newborn infants by near infrared transillumination. *Medical and Biological Engineering and Computing*, 26(3):289–294, 1988.
- [23] David T Delpy, Mark Cope, Pieter van der Zee, Simon Arridge, Susan Wray, and JS Wyatt. Estimation of optical pathlength through tissue from direct time of flight measurement. *Physics in Medicine & Biology*, 33(12):1433, 1988.
- [24] Britton Chance, Michael B Maris, J Sorge, and MZ Zhang. Phase modulation system for dual wavelength difference spectroscopy of hemoglobin deoxygenation in tissues. In *Time-Resolved Laser Spectroscopy in Biochemistry II*, volume 1204, pages 481–491. International Society for Optics and Photonics, 1990.
- [25] David A. Benaron, Susan R. Hintz, Arno Villringer, David Boas, Andreas Kleinschmidt, Jens Frahm, Christina Hirth, Hellmuth Obrig, John C. van Houten, Eben L. Kermit, Wai-Fung Cheong, and David K. Stevenson. Noninvasive functional imaging of human brain using light. *Journal of Cerebral Blood Flow & Metabolism*, 20(3):469–477, Mar 2000.
- [26] Susan R. Hintz, David A. Benaron, John P. van Houten, Joshua L. Duckworth, Frank W. H. Liu, Stanley D. Spilman, David K. Stevenson, and Wai-Fung Cheong. Stationary headband for clinical time-of-flight optical imaging at the bedside. *Photochemistry and Photobiology*, 68(3):361–369, Sep 1998.
- [27] Florian E. W. Schmidt, Martin E. Fry, Elizabeth M. C. Hillman, Jeremy C. Hebden, and David T. Delpy. A 32-channel time-resolved instrument for medical optical tomography. *Review of Scientific Instruments*, 71(1):256–265, Jan 2000.
- [28] Florian E W Schmidt. *Development of a time-resolved optical tomography system for neonatal brain imaging*. phdthesis, University of London, 2000.
- [29] Simon R. Arridge and Martin Schweiger. Image reconstruction in optical tomography. *Philosophical Transactions of the Royal Society of London*. *Series B: Biological Sciences*, 352(1354):717–726, Jun 1997.
- [30] Avraham Y. Bluestone, Gassan Abdoulaev, Christoph H. Schmitz, Randall L. Barbour, and Andreas H. Hielscher. Three-dimensional optical tomography of hemodynamics in the human head. *Optics Express*, 9(6):272, Sep 2001.
- [31] Jeremy C. Hebden. Advances in optical imaging of the newborn infant brain. *Psychophysiology*, 40(4):501–510, Jul 2003.
- [32] Jeremy C. Hebden and Topun Austin. Optical tomography of the neonatal brain. *European Radiology*, 17(11), May 2007.
- [33] Gorm Greisen, Terence Leung, and Martin Wolf. Has the time come to use near-infrared spectroscopy as a routine clinical tool in preterm infants undergoing intensive care? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1955):4440–4451, Nov 2011.
- [34] Cristine Sortica da Costa, Gorm Greisen, and Topun Austin. Is near-infrared spectroscopy clinically useful in the preterm infant? Archives of Disease in Childhood - Fetal and Neonatal Edition, 100(6):F558–F561, Jul 2015.
- [35] Laura Marie Louise Dix, Frank van Bel, and Petra Maria Anna Lemmers. Monitoring cerebral oxygenation in neonates: An update. *Frontiers in Pediatrics*, 5, Mar 2017.
- [36] F. Y. Wong, T. S. Leung, T. Austin, M. Wilkinson, J. H. Meek, J. S. Wyatt, and A. M. Walker. Impaired autoregulation in preterm infants identified by using spatially resolved spectroscopy. *PEDIATRICS*, 121(3):e604–e611, Mar 2008.

- [37] Flora Y. Wong, Reshma Silas, Simon Hew, Thilini Samarasinghe, and Adrian M. Walker. Cerebral oxygenation is highly sensitive to blood pressure variability in sick preterm infants. *PLoS ONE*, 7(8):e43165, Aug 2012.
- [38] M M Gilmore, B S Stone, J A Shepard, M Czosnyka, R B Easley, and K M Brady. Relationship between cerebrovascular dysautoregulation and arterial blood pressure in the premature infant. *Journal of Perinatology*, 31(11):722– 729, Mar 2011.
- [39] S Mitra, M Czosnyka, P Smielewski, H O'Reilly, K Brady, and T Austin. Heart rate passivity of cerebral tissue oxygenation is associated with predictors of poor outcome in preterm infants. *Acta Paediatrica*, 103(9):e374–e382, Aug 2014.
- [40] Jeffrey W. Barker, Andrea L. Rosso, Patrick J. Sparto, and Theodore J. Huppert. Correction of motion artifacts and serial correlations for real-time functional near-infrared spectroscopy. *Neurophotonics*, 3(3):031410, May 2016.
- [41] Lauren L. Emberson, Stephen L. Crosswhite, James R. Goodwin, Andrew J. Berger, and Richard N. Aslin. Isolating the effects of surface vasculature in infant neuroimaging using short-distance optical channels: a combination of local and global effects. *Neurophotonics*, 3(3):031406, Apr 2016.
- [42] SafeBoosC. https://www.rigshospitalet.dk/english/ departments/juliane-marie-centre/department-of-neonatology/ research/safeboosc/Sider/default.aspx, 2018. Accessed 4/8/2020.
- [43] S. Hyttel-Sorensen, A. Pellicer, T. Alderliesten, T. Austin, F. van Bel, M. Benders, O. Claris, E. Dempsey, A. R. Franz, M. Fumagalli, C. Gluud, B. Grevstad, C. Hagmann, P. Lemmers, W. van Oeveren, G. Pichler, A. M. Plomgaard, J. Riera, L. Sanchez, P. Winkel, M. Wolf, and G. Greisen. Cerebral near infrared spectroscopy oximetry in extremely preterm infants: phase II randomised clinical trial. *BMJ*, 350(jan05 2):g7635–g7635, Jan 2015.

- [44] Simon Hyttel-Sorensen, Gorm Greisen, Bodil Als-Nielsen, and Christian Gluud. Cerebral near-infrared spectroscopy monitoring for prevention of brain injury in very preterm infants. *Cochrane Database of Systematic Reviews*, Sep 2017.
- [45] Gerhard Pichler, Berndt Urlesberger, Nariae Baik, Bernhard Schwaberger, Corinna Binder-Heschl, Alexander Avian, Jasmin Pansy, Po-Yin Cheung, and Georg Marcus Schmölzer. Cerebral oxygen saturation to guide oxygen delivery in preterm neonates for the immediate transition after birth: A 2-center randomized controlled pilot feasibility trial. *The Journal of Pediatrics*, 170:73–78.e4, Mar 2016.
- [46] Chuen Wai Lee, Robert J Cooper, and Topun Austin. Diffuse optical tomography to investigate the newborn brain. *Pediatric Research*, 82(3):376–386, May 2017.
- [47] Hiroki Ishii, Takeshi Takami, Tao Fujioka, Norio Mizukaki, Atsushi Kondo, Daisuke Sunohara, Akinori Hoshika, Osamu Akutagawa, and Keiichi Isaka. Comparison of changes in cerebral and systemic perfusion between appropriate- and small-for-gestational-age infants during the first three days after birth. *Brain and Development*, 36(5):380–387, May 2014.
- [48] Tao Fujioka, Takeshi Takami, Hiroki Ishii, Atsushi Kondo, Daisuke Sunohara, and Hisashi Kawashima. Difference in cerebral and peripheral hemodynamics among term and preterm infants during the first three days of life. *Neonatology*, 106(3):181–187, 2014.
- [49] Shinji Nakamura, Kosuke Koyano, Wataru Jinnai, Satoshi Hamano, Saneyuki Yasuda, Yukihiko Konishi, Toru Kuboi, Kenji Kanenishi, Tomoko Nishida, and Takashi Kusaka. Simultaneous measurement of cerebral hemoglobin oxygen saturation and blood volume in asphyxiated neonates by near-infrared time-resolved spectroscopy. *Brain and Development*, 37(10):925–932, Nov 2015.

- [50] Sonoko Ijichi, Takashi Kusaka, Kenichi Isobe, Kensuke Okubo, Kou Kawada, Masanori Namba, Hitoshi Okada, Tomoko Nishida, Tadashi Imai, and Susumu Itoh. Developmental changes of optical properties in neonates determined by near-infrared time-resolved spectroscopy. *Pediatric Research*, 58(3):568–573, Sep 2005.
- [51] Lorenzo Spinelli, Lucia Zucchelli, Davide Contini, Matteo Caffini, Jacques Mehler, Ana Fló, Alissa L. Ferry, Luca Filippin, Francesco Macagno, Luigi Cattarossi, and Alessandro Torricelli. In vivo measure of neonate brain optical properties and hemodynamic parameters by time-domain near-infrared spectroscopy. *Neurophotonics*, 4(04):1, Aug 2017.
- [52] Marco Pagliazzi, Martina Giovannella, Bjørn Andresen, Agnese De Carli, Monica Fumagalli, Gorm Greisen, Davide Contini, Antonio Pifferi, Lorenzo Spinelli, Turgut Durduran, Udo M. Weigel, and Alessandro Torricelli. The BabyLux device: Baseline hemodynamic and optical properties of the newborn brain and the reproducibility of the measurements. In *Biophotonics Congress: Biomedical Optics Congress 2018 (Microscopy/Translational/Brain/OTS)*. OSA, 2018.
- [53] R.J. Cooper, Jeremy C. Hebden, H. O'Reilly, S. Mitra, A.W. Michell, N.L. Everdell, A.P. Gibson, and T. Austin. Transient haemodynamic events in neurologically compromised infants: A simultaneous EEG and diffuse optical imaging study. *NeuroImage*, 55(4):1610–1616, Apr 2011.
- [54] Sabrina Brigadoi, Paul Aljabar, Maria Kuklisova-Murgasova, Simon R. Arridge, and Robert J. Cooper. A 4d neonatal head model for diffuse optical imaging of pre-term to term infants. *NeuroImage*, 100:385–394, Oct 2014.
- [55] Harsimrat Singh, Robert J. Cooper, Chuen Wai Lee, Laura Dempsey, Andrea Edwards, Sabrina Brigadoi, Dimitrios Airantzis, Nick Everdell, Andrew Michell, David Holder, Jeremy C. Hebden, and Topun Austin. Mapping cor-

tical haemodynamics during neonatal seizures using diffuse optical tomography: A case study. *NeuroImage: Clinical*, 5:256–265, 2014.

- [56] Maria Chalia, Chuen Wai Lee, Laura A. Dempsey, Andrea D. Edwards, Harsimrat Singh, Andrew W. Michell, Nicholas L. Everdell, Reuben W. Hill, Jeremy C. Hebden, Topun Austin, and Robert J. Cooper. Hemodynamic response to burst-suppressed and discontinuous electroencephalography activity in infants with hypoxic ischemic encephalopathy. *Neurophotonics*, 3(3):031408, May 2016.
- [57] Laura Alicia Dempsey. Development and application of diffuse optical tomography systems for diagnosis and assessment of perinatal brain injury.
 PhD thesis, Department of Medical Physics and biomedical Engineering, University College London, 2018.
- [58] BabyLux Project. https://www.babylux-project.eu/, 2020. Accessed 4/8/2020.
- [59] Martina Giovannella. BabyLux device: a diffuse optical system integrating diffuse correlation spectroscopy and time-resolved near-infrared spectroscopy for the neuromonitoring of the premature newborn brain. *Neurophotonics*, 6(02):1, May 2019.
- [60] Diffuse optical tomography and imaging. https://www. picoquant.com/applications/category/life-science/ diffuse-optical-tomography-and-imaging, 2013. Accessed 23/1/2018.
- [61] Yukio Yamada, Hiroaki Suzuki, and Yutaka Yamashita. Time-domain nearinfrared spectroscopy and imaging: A review. *Applied Sciences*, 9(6):1127, Mar 2019.
- [62] Alessandro Torricelli, Davide Contini, Antonio Pifferi, Matteo Caffini, Rebecca Re, Lucia Zucchelli, and Lorenzo Spinelli. Time domain functional NIRS imaging for human brain mapping. *NeuroImage*, 85:28–50, Jan 2014.

- [63] Robert J. Cooper, Elliott Magee, Nick Everdell, Salavat Magazov, Marta Varela, Dimitrios Airantzis, Adam P. Gibson, and Jeremy C. Hebden. MONSTIR II: A 32-channel, multispectral, time-resolved optical tomography system for neonatal brain imaging. *Review of Scientific Instruments*, 85(5):053105, May 2014.
- [64] Jeremy Hebden, T. Bland, Elizabeth M.C. Hillman, A. Gibson, N. Everdell, David T. Delpy, Simon R. Arridge, and M. Douek. Optical tomography of the breast using a 32-channel time-resolved imager. In *Biomedical Topical Meeting*, page SuE5. Optical Society of America, 2002.
- [65] Louise C. Enfield, Adam P. Gibson, Nicholas L. Everdell, David T. Delpy, Martin Schweiger, Simon R. Arridge, Caroline Richardson, Mohammad Keshtgar, Michael Douek, and Jeremy C. Hebden. Three-dimensional timeresolved optical mammography of the uncompressed breast. *Applied Optics*, 46(17):3628, May 2007.
- [66] Jeremy Hebden, Elizabeth M.C. Hillman, A. Gibson, N. Everdell, R. Yusof, David T. Delpy, Simon R. Arridge, T. Austin, and Judith H. Meek. Time resolved optical imaging of the newborn infant brain: initial clinical results. In *Biomedical Topical Meeting*, page TuE2. Optical Society of America, 2002.
- [67] Jeremy C Hebden, Adam Gibson, Rozarina Md Yusof, Nick Everdell, Elizabeth M C Hillman, David T Delpy, Simon R Arridge, Topun Austin, Judith H Meek, and John S Wyatt. Three-dimensional optical tomography of the premature infant brain. *Physics in Medicine and Biology*, 47(23):4155–4166, Nov 2002.
- [68] T. Austin, A.P. Gibson, G. Branco, R.Md. Yusof, S.R. Arridge, J.H. Meek, J.S. Wyatt, D.T. Delpy, and J.C. Hebden. Three dimensional optical imaging of blood volume and oxygenation in the neonatal brain. *NeuroImage*, 31(4):1426–1433, Jul 2006.

- [69] Konstantinos I. Papadimitriou, Laura A. Dempsey, Jeremy C. Hebden, Simon R. Arridge, and Samuel Powell. A spread spectrum approach to time-domain near-infrared diffuse optical imaging using inexpensive optical transceiver modules. *Biomedical Optics Express*, 9(6):2648, May 2018.
- [70] Heidrun Wabnitz, Dieter Richard Taubert, Mikhail Mazurenka, Oliver Steinkellner, Alexander Jelzow, Rainer Macdonald, Daniel Milej, Piotr Sawosz, Michal Kacprzak, Adam Liebert, Robert Cooper, Jeremy Hebden, Antonio Pifferi, Andrea Farina, Ilaria Bargigia, Davide Contini, Matteo Caffini, Lucia Zucchelli, Lorenzo Spinelli, Rinaldo Cubeddu, and Alessandro Torricelli. Performance assessment of time-domain optical brain imagers, part 1: basic instrumental performance protocol. *Journal of Biomedical Optics*, 19(8):086010, Aug 2014.
- [71] Laura Di Sieno, Heidrun Wabnitz, Antonio Pifferi, Mikhail Mazurenka, Yoko Hoshi, Alberto Dalla Mora, Davide Contini, Gianluca Boso, Wolfgang Becker, Fabrizio Martelli, Alberto Tosi, and Rainer Macdonald. Characterization of a time-resolved non-contact scanning diffuse optical imaging system exploiting fast-gated single-photon avalanche diode detection. *Review of Scientific Instruments*, 87(3):035118, Mar 2016.
- [72] Juan Mata Pavia, Edoardo Charbon, and Martin Wolf. 3d near-infrared imaging based on a single-photon avalanche diode array sensor. In Andreas H. Hielscher and Paola Taroni, editors, *Diffuse Optical Imaging III*. SPIE, Jun 2011.
- [73] Juliette Selb, Jonathan J. Stott, Maria Angela Franceschini, A. Gregory Sorensen, and David A. Boas. Improved sensitivity to cerebral hemodynamics during brain activation with a time-gated optical system: analytical model and experimental validation. *Journal of Biomedical Optics*, 10(1):011013, 2005.

- [74] Weirong Mo and Nanguang Chen. Fast time-domain diffuse optical tomography using pseudorandom bit sequences. *Optics Express*, 16(18):13643, Aug 2008.
- [75] Andrea L. Lacaita, Giancarlo Riparmonti, Pier Andrea Francese, and Sergio D. Cova. Single-photon optical-time-domain reflectometer at 13 μ m with 5-cm resolution and high sensitivity. *Optics Letters*, 18(13):1110, Jul 1993.
- [76] G. Ripamonti, F. Zappa, and S.D. Cova. Effects of trap levels in singlephoton optical time-domain reflectometry: evaluation and correction. *Journal of Lightwave Technology*, 10(10):1398–1402, 1992.
- [77] Antonio Pifferi, Alessandro Torricelli, Rinaldo Cubeddu, Giovanna Quarto, Rebecca Re, Sanathana K. V. Sekar, Lorenzo Spinelli, Andrea Farina, Fabrizio Martelli, and Heidrun Wabnitz. Mechanically switchable solid inhomogeneous phantom for performance tests in diffuse imaging and spectroscopy. *Journal of Biomedical Optics*, 20(12):121304, Jul 2015.
- [78] H. Rinneberg, D. Grosenick, K. Moesta, H. Wabnitz, J. Mucke, G. Wübbeler,
 R. Macdonald, and P. Schlag. Detection and characterization of breast tumours by time-domain scanning optical mammography. *Opto-Electronics Review*, 16(2), Jan 2008.
- [79] D. Comelli, A. Pifferi, A. Taroni, A. Torricelli, R. Cubeddu, F. Maritnelli, G. Zaccanti, and J.-M. Tualle. In vivo time-resolved multi-distance spectroscopy of human forehead: a step towards optical characterization of brain tissue. In *Biomedical Optics*. OSA, 2006.
- [80] D. Contini, A. Pifferi, L. Spinelli, A. Torricelli, and R. Cubeddu. Design and characterization of a two-wavelength multi-channel time-resolved system for optical topography. In *Biomedical Optics*. OSA, 2006.
- [81] D. Contini, R. Re, M. Turola, L. Spinelli, G. Romano, R. Cubeddu, and A. Torricelli. Multi-channel time-resolved functional near infrared spec-

troscopy system. In Bruce J. Tromberg, Arjun G. Yodh, and Eva M. Sevick-Muraca, editors, *Optical Tomography and Spectroscopy of Tissue X*. SPIE, Mar 2013.

- [82] Davide Contini, Alessandro Torricelli, Antonio Pifferi, Lorenzo Spinelli, Floriano Paglia, and Rinaldo Cubeddu. Multi-channel time-resolved system for functional near infrared spectroscopy. *Optics Express*, 14(12):5418, 2006.
- [83] Davide Contini, Alberto Dalla Mora, Simon Arridge, Fabrizio Martelli, Alberto Tosi, Gianluca Boso, Andrea Farina, Turgut Durduran, Edoardo Martinenghi, Alessandro Torricelli, and Antonio Pifferi. Time-domain diffuse optics: towards next generation devices. In *Diffuse Optical Imaging V*. OSA, 2015.
- [84] Alberto Dalla Mora, Edoardo Martinenghi, Davide Contini, Alberto Tosi, Gianluca Boso, Turgut Durduran, Simon Arridge, Fabrizio Martelli, Andrea Farina, Alessandro Torricelli, and Antonio Pifferi. Fast silicon photomultiplier improves signal harvesting and reduces complexity in time-domain diffuse optics. *Optics Express*, 23(11):13937, May 2015.
- [85] Sanathana Konugolu Venkata Sekar, Alberto Dalla Mora, Ilaria Bargigia, Edoardo Martinenghi, Claus Lindner, Parisa Farzam, Marco Pagliazzi, Turgut Durduran, Paola Taroni, Antonio Pifferi, and Andrea Farina. Broadband (600–1350 nm) time-resolved diffuse optical spectrometer for clinical use. *IEEE Journal of Selected Topics in Quantum Electronics*, 22(3):406– 414, May 2016.
- [86] Antonio Pifferi, Alberto Dalla Mora, Edorado Martinenghi, Davide Contini, Laura di Sieno, Rebecca Re, Simon Arridge, Fabrizio Martelli, Alberto Tosi, Andrea Farina, Turgut Durduran, and Alessandro Torricelli. Photonics advancements in time-domain diffuse imaging: Towards hand-held and wearable devices. In *Biomedical Optics 2016*. OSA, 2016.

- [87] Davide Tamborini, Mauro Buttafava, Alessandro Ruggeri, and Franco Zappa. Compact, low-power and fully reconfigurable 10 ps resolution, 160 range, time-resolved single-photon counting system. *IEEE Sensors Journal*, 16(10):3827–3833, May 2016.
- [88] M. Mazurenka, L. Di Sieno, G. Boso, D. Contini, A. Pifferi, A. Dalla Mora, A. Tosi, H. Wabnitz, and R. Macdonald. Non-contact in vivo diffuse optical imaging using a time-gated scanning system. *Biomedical Optics Express*, 4(10):2257, Sep 2013.
- [89] Rebecca Re, Edoardo Martinenghi, Alberto Dalla Mora, Davide Contini, Antonio Pifferi, and Alessandro Torricelli. Probe-hosted silicon photomultipliers for time-domain functional near-infrared spectroscopy: phantom andin vivotests. *Neurophotonics*, 3(4):045004, Oct 2016.
- [90] E. Martinenghi, L. Di Sieno, D. Contini, M. Sanzaro, A. Pifferi, and A. Dalla Mora. Time-resolved single-photon detection module based on silicon photomultiplier: A novel building block for time-correlated measurement systems. *Review of Scientific Instruments*, 87(7):073101, Jul 2016.
- [91] Laura Di Sieno, Jan Nissinen, Lauri Hallman, Edoardo Martinenghi, Davide Contini, Antonio Pifferi, Juha Kostamovaara, and Alberto Dalla Mora. Miniaturized pulsed laser source for time-domain diffuse optics routes to wearable devices. *Journal of Biomedical Optics*, 22(08):1, Aug 2017.
- [92] Andrea Farina, Susanna Tagliabue, Laura Di Sieno, Edoardo Martinenghi, Turgut Durduran, Simon Arridge, Fabrizio Martelli, Alessandro Torricelli, Antonio Pifferi, and Alberto Dalla Mora. Time-Domain Functional Diffuse Optical Tomography system based on fiber-free silicon photomultipliers. *Applied Sciences*, 7(12):1235, Nov 2017.
- [93] Mauro Buttafava, Edoardo Martinenghi, Davide Tamborini, Davide Contini, Alberto Dalla Mora, Marco Renna, Alessandro Torricelli, Antonio Pifferi, Franco Zappa, and Alberto Tosi. A compact two-wavelength time-domain

NIRS system based on SiPM and pulsed diode lasers. *IEEE Photonics Journal*, 9(1):1–14, Feb 2017.

- [94] Marco Renna, Mauro Buttafava, Anurag Behera, Marta Zanoletti, Laura Di Sieno, Alberto Dalla Mora, Davide Contini, and Alberto Tosi. Eightwavelength, dual detection channel instrument for near-infrared timeresolved diffuse optical spectroscopy. *IEEE Journal of Selected Topics in Quantum Electronics*, 25(1):1–11, Jan 2019.
- [95] Anurag Behera. Novel devices and methodologies for optical functional brain imaging and neuromonitoring. PhD thesis, Dipartimento di Fisica, Politecnico di Milano, 2020.
- [96] tnirs-1. https://www.hamamatsu.com/jp/ja/product/type/C12707/ index.html, July 2019. Accessed 10/8/2020.
- [97] Bojan Markovic, Simone Tisa, Federica A. Villa, Alberto Tosi, and Franco Zappa. A high-linearity, 17 ps precision time-to-digital converter based on a single-stage vernier delay loop fine interpolation. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(3):557–569, mar 2013.
- [98] Samuel Burri, Claudio Bruschini, and Edoardo Charbon. LinoSPAD: A compact linear SPAD camera system with 64 FPGA-based TDC modules for versatile 50 ps resolution time-resolved imaging. *Instruments*, 1(1):6, Dec 2017.
- [99] Sreenil Saha, Yuankang Lu, Frederic Lesage, and Mohamad Sawan. Wearable SiPM-based NIRS interface integrated with pulsed laser source. *IEEE Transactions on Biomedical Circuits and Systems*, 13(6):1313–1323, Dec 2019.
- [100] Sreenil Saha, Samuel Burri, Claudio Bruschini, Edoardo Charbon, Frederic Lesage, and Mohamad Sawan. Time domain NIRS optode based on null/small source-detector distance for wearable applications. In 2019 IEEE Custom Integrated Circuits Conference (CICC). IEEE, Apr 2019.

- [101] Mrwan Alayed, Darek Palubiak, and M. Deen. Characterization of a Time-Resolved Diffuse Optical Spectroscopy prototype using low-cost, compact single photon avalanche detectors for tissue optics applications. *Sensors*, 18(11):3680, Oct 2018.
- [102] Claudio Bruschini, Harald Homulle, and Edoardo Charbon. Ten years of biophotonics single-photon SPAD imager applications: retrospective and outlook. In Ammasi Periasamy, Peter T. C. So, Karsten König, and Xiaoliang S. Xie, editors, *Multiphoton Microscopy in the Biomedical Sciences XVII*. SPIE, Feb 2017.
- [103] Frédéric Nolet, William Lemaire, Frédérik Dubois, Nicolas Roy, Simon Carrier, Arnaud Samson, Serge A. Charlebois, Réjean Fontaine, and Jean-Francois Pratte. A 256 pixelated SPAD readout ASIC with in-pixel TDC and embedded digital signal processing for uniformity and skew correction. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 949:162891, Jan 2020.
- [104] Juan Mata Pavia, Mario Scandini, Scott Lindner, Martin Wolf, and Edoardo Charbon. A 1 × 400 backside-illuminated SPAD sensor with 49.7 ps resolution, 30 pJ/sample TDCs fabricated in 3d CMOS technology for near-infrared optical tomography. *IEEE Journal of Solid-State Circuits*, 50(10):2406– 2418, Oct 2015.
- [105] Chao Zhang, Scott Lindner, Ivan Michel Antolovic, Juan Mata Pavia, Martin Wolf, and Edoardo Charbon. A 30-frames/s, 252 × 144 SPAD flash LiDAR with 1728 dual-clock 48.8-ps TDCs, and pixel-wise integrated histogramming. *IEEE Journal of Solid-State Circuits*, 54(4):1137–1151, Apr 2019.
- [106] Ahmet T. Erdogan, Richard Walker, Neil Finlayson, Nikola Krstajic, Gareth O. S. Williams, and Robert K. Henderson. A 16.5 giga events/s 1024 × 8 SPAD line sensor with per-pixel zoomable 50ps-6.4ns/bin histogramming TDC. In 2017 Symposium on VLSI Circuits. IEEE, Jun 2017.

- [107] Ahmet T. Erdogan, Richard Walker, Neil Finlayson, Nikola Krstajic, Gareth Williams, John Girkin, and Robert Henderson. A CMOS SPAD line sensor with per-pixel histogramming TDC for time-resolved multispectral imaging. *IEEE Journal of Solid-State Circuits*, 54(6):1705–1719, Jun 2019.
- [108] A. Farina, A. Candeo, A. Dalla Mora, A. Bassi, R. Lussana, F. Villa, G. Valentini, S. Arridge, and C. D'Andrea. Novel time-resolved camera based on compressed sensing. *Optics Express*, 27(22):31889, Oct 2019.
- [109] STMicroelectronics. Proximity Sensors. https://www.st.com/en/ imaging-and-photonics-solutions/proximity-sensors.html. Accessed 20/7/2020.
- [110] Jeremy C. Hebden, Ruchir Shah, and Danial Chitnis. Probe for evaluating the absorbing and transport scattering properties of turbid fluids using lowcost time-of-flight technology. *Journal of Biomedial Optics*, 22(5):055009, May 2017.
- [111] Mihir Milind Rao. Evaluating a very low-cost commercial optical device for LIDAR imaging. resreport, Department of Medical Physics and biomedical Engineering, University College London, 2019.
- [112] Hirokazu Atsumori, Masashi Kiguchi, Akiko Obata, Hiroki Sato, Takusige Katura, Tsukasa Funane, and Atsushi Maki. Development of wearable optical topography system for mapping the prefrontal cortex activation. *Review* of Scientific Instruments, 80(4):043704, Apr 2009.
- [113] M. Kiguchi, H. Atsumori, I. Fukasaku, Y. Kumagai, T. Funane, A. Maki, Y. Kasai, and A. Ninomiya. Note: Wearable near-infrared spectroscopy imager for haired region. *Review of Scientific Instruments*, 83(5):056101, May 2012.
- [114] Sophie K. Piper, Arne Krueger, Stefan P. Koch, Jan Mehnert, Christina Habermehl, Jens Steinbrink, Hellmuth Obrig, and Christoph H. Schmitz. A

wearable multi-channel fNIRS system for brain imaging in freely moving subjects. *NeuroImage*, 85:64–71, Jan 2014.

- [115] Etienne Lareau, Frederic Lesage, Philippe Pouliot, Dang Nguyen, Jerome Le Lan, and Mohamad Sawan. Multichannel wearable system dedicated for simultaneous electroencephalography/near-infrared spectroscopy real-time data acquisitions. *Journal of Biomedical Optics*, 16(9):096014, 2011.
- [116] J Safaie, R Grebe, H Abrishami Moghaddam, and F Wallois. Toward a fully integrated wireless wearable EEG-NIRS bimodal acquisition system. *Journal of Neural Engineering*, 10(5):056001, Jul 2013.
- [117] Karla M. Bergonzi, Tracy M. Burns-Yocum, Jonathan R. Bumstead, Elise M. Buckley, Patrick C. Mannion, Christopher H. Tracy, Eli Mennerick, Silvina L. Ferradal, Hamid Dehghani, Adam T. Eggebrecht, and Joseph P. Culver. Lightweight sCMOS-based high-density diffuse optical tomography. *Neurophotonics*, 5(03):1, Aug 2018.
- [118] Thomas Muehlemann, Daniel Haensse, and Martin Wolf. Wireless miniaturized in-vivo near infrared imaging. *Optics Express*, 16(14):10323, Jun 2008.
- [119] Bertan Hallacoglu, Jason W. Trobaugh, Kate L. Bechtel, and Chandran V. Seshagiri. Blood phantom verification of a new compact DOT system. In *Biomedical Optics 2016*. OSA, 2016.
- [120] Alexander von Luhmann, Christian Herff, Dominic Heger, and Tanja Schultz. Toward a wireless open source instrument: Functional near-infrared spectroscopy in mobile neuroergonomics and BCI applications. *Frontiers in Human Neuroscience*, 9, Nov 2015.
- [121] Alexander von Luhmann, Heidrun Wabnitz, Tilmann Sander, and Klaus-Robert Muller. M3ba: A mobile, modular, multimodal biosignal acquisition architecture for miniaturized EEG-NIRS-based hybrid BCI and monitoring. *IEEE Transactions on Biomedical Engineering*, 64(6):1199–1210, Jun 2017.

- [122] Danial Chitnis, Dimitrios Airantzis, David Highton, Rhys Williams, Phong Phan, Vasiliki Giagka, Samuel Powell, Robert J. Cooper, Ilias Tachtsidis, Martin Smith, Clare E. Elwell, Jeremy C. Hebden, and Nicholas Everdell. Towards a wearable near infrared spectroscopic probe for monitoring concentrations of multiple chromophores in biological tissue in vivo. *Review of Scientific Instruments*, 87(6):065112, Jun 2016.
- [123] Danial Chitnis, Robert J. Cooper, Laura Dempsey, Samuel Powell, Simone Quaggia, David Highton, Clare Elwell, Jeremy C. Hebden, and Nicholas L. Everdell. Functional imaging of the human brain using a modular, fibre-less, high-density diffuse optical tomography system. *Biomedical Optics Express*, 7(10):4275, Sep 2016.
- [124] R. Zimmermann, F. Braun, T. Achtnich, O. Lambercy, R. Gassert, and M. Wolf. Silicon photomultipliers for improved detection of low light levels in miniature near-infrared spectroscopy instruments. *Biomedical Optics Express*, 4(5):659, Apr 2013.
- [125] Dominik Wyser, Olivier Lambercy, Felix Scholkmann, Martin Wolf, and Roger Gassert. Wearable and modular functional near-infrared spectroscopy instrument with multidistance measurements at four wavelengths. *Neurophotonics*, 4(04):1, Aug 2017.
- [126] Gary E. Strangman, Vladimir Ivkovic, and Quan Zhang. Wearable brain imaging with multimodal physiological monitoring. *Journal of Applied Physiology*, 124(3):564–572, Mar 2018.
- [127] Quan Zhang, Xiangguo Yan, and Gary E. Strangman. Development of motion resistant instrumentation for ambulatory near-infrared spectroscopy. *Journal of Biomedical Optics*, 16(8):087008, 2011.
- [128] Paola Pinti, Clarisse Aichelburg, Sam Gilbert, Antonia Hamilton, Joy Hirsch, Paul Burgess, and Ilias Tachtsidis. A review on the use of wearable functional

near-infrared spectroscopy in naturalistic environments. *Japanese Psychological Research*, 60(4):347–373, Jul 2018.

- [129] Renata Di Lorenzo, Laura Pirazzoli, Anna Blasi, Chiara Bulgarelli, Yoko Hakuno, Yasuyo Minagawa, and Sabrina Brigadoi. Recommendations for motion correction of infant fNIRS data applicable to multiple data sets and acquisition systems. *NeuroImage*, 200:511–527, Oct 2019.
- [130] A.M. Chiarelli, G.C. Giaconia, D. Perpetuini, G. Greco, L. Mistretta, R. Rizzo, V. Vinciguerra, M. F. Romeo, A. Merla, and P.G. Fallica. Wearable, fiber-less, multi-channel system for continuous wave functional near infrared spectroscopy based on silicon photomultipliers detectors and lockin amplification. In 2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). IEEE, Jul 2019.
- [131] Masudur R. Siddiquee, J. Sebastian Marquez, Roozbeh Atri, Rodrigo Ramon, Robin Perry Mayrand, and Ou Bai. Movement artefact removal from NIRS signal using multi-channel IMU data. *BioMedical Engineering On-Line*, 17(1), Sep 2018.
- [132] Sabrina Brigadoi, Aman Ganglani, Hubin Zhao, and Robert J. Cooper. Integrating motion sensing and wearable, modular high-density diffuse optical tomography: preliminary results. In Hamid Dehghani and Heidrun Wabnitz, editors, *Diffuse Optical Spectroscopy and Imaging VII*. SPIE, Jul 2019.
- [133] Giovanni Maira, Antonio M. Chiarelli, Stefano Brafa, Sebania Libertino, Giorgio Fallica, Arcangelo Merla, and Salvatore Lombardo. Imaging system based on silicon photomultipliers and light emitting diodes for functional near-infrared spectroscopy. *Applied Sciences*, 10(3):1068, Feb 2020.
- [134] NIRSPORT 2. https://nirx.net/nirsport. Accessed 4/8/2020.
- [135] LUMO. https://www.gowerlabs.co.uk/lumo, 2019. Accessed 4/8/2020.

- [136] Hubin Zhao, Sabrina Brigadoi, Reuben Hill, Greg Smith, Luke Dunne, Elisabetta Maria Frijia, Samuel Powell, Nick L. Everdell, and Robert Cooper. Advances in wearable high-density diffuse optical tomography: first applications of a new commercial technology and development of an infant-specific research device. In Hamid Dehghani and Heidrun Wabnitz, editors, *Diffuse Optical Spectroscopy and Imaging VII*. SPIE, Jul 2019.
- [137] R.E. Green and R.E. Bell. Notes on a fast time-to-amplitude converter. Nuclear Instruments, 3(3):127–132, Sep 1958.
- [138] G. Culligan and N. H. Lipman. Fast transistorized time-to-pulse-height converter. *Review of Scientific Instruments*, 31(11):1209–1214, Nov 1960.
- [139] W. Weber, C. W. Johnstone, and L. Cranberg. Time-to-pulse-height converter for measurement of millimicrosecond time intervals. *Review of Scientific Instruments*, 27(3):166–170, Mar 1956.
- [140] Dan I. Porat. Review of sub-nanosecond time-interval measurements. *IEEE Transactions on Nuclear Science*, 20(5):36–51, 1973.
- [141] P. Napolitano, A. Moschitta, and P. Carbone. A survey on time interval measurement techniques and testing methods. In 2010 IEEE Instrumentation & Measurement Technology Conference Proceedings. IEEE, 2010.
- [142] Józef Kalisz. Review of methods for time interval measurements with picosecond resolution. *Metrologia*, 41(1):17–32, Dec 2003.
- [143] Fermilab electronic systems engineering. eg & g/ortec directory. http:// prep.fnal.gov/catalog/hardware_info/egg_ortec/, October 2008. Accessed 9/1/2018.
- [144] Time measurement. http://www.surface-concept.com/products_ time.html, 2017. Accessed 9/1/2018.
- [145] Time-to-digital converter for photon counters. https: //www.lasercomponents.com/de-en/product/

time-to-digital-converter-for-photon-counters/, 2018. Accessed 9/1/2018.

- [146] TDC7200. http://www.ti.com/product/TDC7200, 2017. Accessed 9/1/2018.
- [147] TDC-GP22 Ultrasonic flow converter. http://ams.com/eng/Products/ Precision-Time-Measurement/Time-to-Digital-Converters/ TDC-GP22, 2018. Accessed 24/1/2018.
- [148] TDC-GPX2 Time-to-Digital Converter. http://ams.com/eng/Products/ Precision-Time-Measurement/Time-to-Digital-Converters/ TDC-GPX2, 2018. Accessed 24/1/2018.
- [149] Nahema Marino, Federico Baronti, Luca Fanucci, Sergio Saponara, Roberto Roncella, Maria G. Bisogni, and Alberto Del Guerra. A multichannel and compact time to digital converter for time of flight positron emission tomography. *IEEE Transactions on Nuclear Science*, 62(3):814–823, Jun 2015.
- [150] Mahantesh P. Mattada and Hansraj Guhilot. Area efficient vernier time to digital converter(TDC) with improved resolution using identical ring oscillators on FPGA. In *INTERNATIONAL CONFERENCE ON SMART STRUCTURES AND SYSTEMS - ICSSS'13*. IEEE, Mar 2013.
- [151] Chao Chen, Shengwei Meng, Zhenghuan Xia, Guangyou Fang, and Hejun Yin. An FPGA-integrated time-to-digital converter based on a ring oscillator for programmable delay line resolution measurement. *Journal of Electrical and Computer Engineering*, 2014:230803, 2014.
- [152] Min-Chuan Lin, Guo-Ruey Tsai, Chun-Yi Liu, and Shi-Shien Chu. FPGAbased high area efficient time-to-digital IP design. In *TENCON 2006 - 2006 IEEE Region 10 Conference*. IEEE, 2006.

234

- [153] Jinyuan Wu. Several key issues on implementing delay line based TDCs using FPGAs. *IEEE Transactions on Nuclear Science*, 57(3):1543–1548, Jun 2010.
- [154] GS Jovanovic and MK Stojcev. Vernier's delay line time-to-digital converter. Series A: Applied Mathematics, Informatics & Mechanics, Scientific Publications of the state University of Novi Pazar, 1(1):11–20, 2009.
- [155] Marc-André Daigneault. Utilisation de la reconfiguration dynamique des FPGA pour le contrôle précis et exact des délais dans les convertisseurs temps à numérique. mathesis, Départament de Génie Électrique, École Polytechnique de Montréal, 2009.
- [156] Jinyuan Wu and Zonghan Shi. The 10-ps wave union TDC: Improving FPGA TDC resolution beyond its cell delay. In 2008 IEEE Nuclear Science Symposium Conference Record. IEEE, Oct 2008.
- [157] Jinhong Wang, Shubin Liu, Lei Zhao, Xueye Hu, and Qi An. The 10-ps multitime measurements averaging TDC implemented in an FPGA. *IEEE Transactions on Nuclear Science*, 58(4):2011–2018, Aug 2011.
- [158] Claudio Favi and Edoardo Charbon. A 17ps time-to-digital converter implemented in 65nm FPGA technology. In *Proceeding of the ACM/SIGDA international symposium on Field programmable gate arrays - FPGA '09*. ACM Press, 2009.
- [159] Harmen Menninga, Claudio Favi, Matthew W. Fishburn, and Edoardo Charbon. A multi-channel, 10ps resolution, FPGA-based TDC with 300MS/s throughput for open-source PET applications. In 2011 IEEE Nuclear Science Symposium Conference Record. IEEE, Oct 2011.
- [160] MatthewW. Fishburn, L. Harmen Menninga, Claudio Favi, and Edoardo Charbon. A 19.6 ps, FPGA-based TDC with multiple channels for open source applications. *IEEE Transactions on Nuclear Science*, 60(3):2203– 2208, Jun 2013.

- [161] Harald Homulle and Edoardo Charbon. FPGA designs for reconfigurable converters. http://cas.tudelft.nl/fpga_tdc/index.html, 2015. Accessed 9/1/2018.
- [162] Maciej Lipinski, Tomasz Wlostowski, Javier Serrano, and Pablo Alvarez.
 White rabbit: a PTP application for robust sub-nanosecond synchronization.
 In 2011 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication. IEEE, Sep 2011.
- [163] Sébastien Bourdeauducq. OHWR TDC core. https://www.ohwr.org/ projects/tdc-core/wiki, 2011. Accessed 9/1/2018.
- [164] Neale Dutton, Johannes Vergote, Salvatore Gnecchi, Lindsay Grant, David Lee, Sara Pellegrini, Bruce Rae, and Robert Henderson. Multiple-event direct to histogram TDC in 65nm FPGA technology. In 2014 10th Conference on Ph.D. Research in Microelectronics and Electronics (PRIME). IEEE, Jun 2014.
- [165] Haochang Chen, Yongliang Zhang, and David Day-Uei Li. A low nonlinearity, missing-code free time-to-digital converter based on 28-nm FPGAs with embedded bin-width calibrations. *IEEE Transactions on Instrumentation and Measurement*, 66(7):1912–1921, Jul 2017.
- [166] Jun Yeon Won and Jae Sung Lee. Time-to-digital converter using a tuneddelay line evaluated in 28-, 40-, and 45-nm FPGAs. *IEEE Transactions on Instrumentation and Measurement*, 65(7):1678–1689, Jul 2016.
- [167] A Aloisio, R Lomoro, S Loffredo, V Izzo, P Branchini, R Cicalese, and R Giordano. High-resolution time-to-digital converter in field programmable gate array. CEN, 2008.
- [168] J. Torres, A. Aguilar, R. Garcia-Olcina, P. A. Martinez, J. Martos, J. Soret, J. M. Benlloch, P. Conde, A. J. Gonzalez, and F. Sanchez. Time-to-digital converter based on FPGA with multiple channel capability. *IEEE Transactions on Nuclear Science*, 61(1):107–114, Feb 2014.

- [169] Amir Mohammad Amiri. Design and implementation of a high resolution, multi-hit time-to-digital converter (TDC) on FPGA. phdthesis, Département de Génie Électrique, École polytechnique de Montréal, 2007.
- [170] Qiuchen Yuan. A High Resolution Time-to-Digital Converter on FPGA for Time-Correlated Single Photon Counting. candthesis, School of Engineering and Applied Sciences, The George Washington University, 2012.
- [171] Tuan Vo-Dinh, editor. *The Biomedical Photonics Handbook*, volume I. Fundamentals, Devices, and Techniques. CRC Press, Taylor & Francis Inc, 2 edition, 2014.
- [172] Anna N. Yaroslavsky and Ilya V. Yaroslavsky Ilya V. Yaroslavsky. 1: Handbook of Optical Biomedical Diagnostics: Light-tissue Interaction (Press Monograph), chapter 2. Optics of Blood, pages 161 – 219. SPIE Press, 2 edition, 2016.
- [173] W.F. Cheong, S.A. Prahl, and A.J. Welch. A review of the optical properties of biological tissues. *IEEE Journal of Quantum Electronics*, 26(12):2166– 2185, 1990.
- [174] Anabela Da Silva. *Photon-Based Medical Imagery*, chapter 7. Optical Imaging, pages 267 – 324. ISTE LTD/Oxford : Wiley-Blackwell, 2011.
- [175] Jeremy C Hebden, Simon R Arridge, and David T Delpy. Optical imaging in medicine: I. experimental techniques. *Physics in Medicine & Biology*, 42(5):825, 1997.
- [176] Simon R Arridge and Jeremy C Hebden. Optical imaging in medicine: II. modelling and reconstruction. *Physics in Medicine and Biology*, 42(5):841– 853, May 1997.
- [177] Martin Schweiger and Simon Arridge. The Toast++ software suite for forward and inverse modeling in optical tomography. *Journal of Biomedical Optics*, 19(4):040801, Apr 2014.

- [178] Martin Schweiger and et al. Toast++. image reconstruction in diffuse optical tomography. http://web4.cs.ucl.ac.uk/research/vis/toast/ intro.html, 2014.
- [179] Tara Yates, Jeremy C Hebden, Adam Gibson, Nick Everdell, Simon R Arridge, and Michael Douek. Optical tomography of the breast using a multi-channel time-resolved imager. *Physics in Medicine and Biology*, 50(11):2503–2517, May 2005.
- [180] Lorenzo Spinelli, Alessandro Torricelli, Antonio Pifferi, Paola Taroni, Gianmaria Danesini, and Rinaldo Cubeddu. Characterization of female breast lesions from multi-wavelength time-resolved optical mammography. *Physics in Medicine and Biology*, 50(11):2489–2502, May 2005.
- [181] Louise Enfield, Gabriel Cantanhede, Michael Douek, Vernie Ramalingam, Arnie Purushotham, Jem Hebden, and Adam Gibson. Monitoring the response to neoadjuvant hormone therapy for locally advanced breast cancer using three-dimensional time-resolved optical mammography. *Journal of Biomedical Optics*, 18(5):056012, May 2013.
- [182] Jeremy C Hebden, Adam Gibson, Topun Austin, Rozarina Md Yusof, Nick Everdell, David T Delpy, Simon R Arridge, Judith H Meek, and John S Wyatt. Imaging changes in blood volume and oxygenation in the newborn infant brain using three-dimensional optical tomography. *Physics in Medicine and Biology*, 49(7):1117–1130, Mar 2004.
- [183] A P Gibson, J C Hebden, and S R Arridge. Recent advances in diffuse optical imaging. *Physics in Medicine and Biology*, 50(4):R1–R43, Feb 2005.
- [184] Sarah Lloyd-Fox, M. Papademetriou, M. K. Darboe, N. L. Everdell, R. Wegmuller, A. M. Prentice, S. E. Moore, and C. E. Elwell. Functional near infrared spectroscopy (fNIRS) to assess cognitive function in infants in rural Africa. *Scientific Reports*, 4(1), Apr 2014.

- [185] Sabrina Brigadoi and Robert J. Cooper. How short is short? optimum source-detector distance for short-separation channels in functional nearinfrared spectroscopy. *Neurophotonics*, 2(2):025005, May 2015.
- [186] P. van der Zee, M. Cope, S. R. Arridge, M. Essenpreis, L. A. Potter, A. D. Edwards, J. S. Wyatt, D. C. McCormick, S. C. Roth, E. O. R. Reynolds, and D. T. Delpy. Experimentally measured optical pathlengths for the adult head, calf and forearm and the head of the newborn infant as a function of inter optode spacing. In *Advances in Experimental Medicine and Biology*, pages 143–153. Springer US, 1992.
- [187] Photoelectric effect. Encyclopædia Britannica, https://www.britannica. com/science/photoelectric-effect, June 2017. Accessed 30/11/2017.
- [188] Dieter Meschede. Optics, Light and Lasers. Wiley-VCH Verlag GmbH & Co. KGaA, Nov 2006.
- [189] Hamamatsu. Photomultiplier tubes. Basics and Applications. https://www. hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf, 2007. Accessed 30/11/2017.
- [190] Didier Decoster and Joseph Harari, editors. Optoelectronic Sensors. ISTE, Jan 2009.
- [191] Hamamatsu Opto-Semiconductor Handbook, Chapter 03 Si APD, MPPC. https://www.hamamatsu.com/resources/pdf/ssd/e03_handbook_ si_apd_mppc.pdf.
- [192] Frank J. Duarte. *Tunable Laser Optics*. Academic Press, 2003.
- [193] Mark Csele. Fundamentals of Light Sources and Lasers. John Wiley & Sons, Inc., Jul 2004.
- [194] James R. Taylor. Tutorial on fiber-based sources for biophotonic applications. *Journal of Biomedical Optics*, 21(6):061010, Jun 2016.

- [195] Stephan Henzler. Time-to-digital converter basics. In *Time-to-Digital Converters*, pages 5–18. Springer Netherlands, 2010.
- [196] D.M. Santos, S.F. Dow, J.M. Flasck, and M.E. Levi. A CMOS delay locked loop and sub-nanosecond time-to-digital converter chip. *IEEE Transactions* on Nuclear Science, 43(3):1717–1719, Jun 1996.
- [197] S. Cova and M. Bertolaccini. Differential linearity testing and precision calibration of multichannel time sorters. *Nuclear Instruments and Methods*, 77(2):269–276, Jan 1970.
- [198] INL/DNL measurements for high-speed Analog-to-Digital Converters (ADCs). https://www.maximintegrated.com/en/app-notes/index. mvp/id/283, November 2001. Accessed 9/1/2018.
- [199] Davide Bucci. Analog Electronics for Measuring Systems. John Wiley & Sons, 2017.
- [200] Haideh Khorramabadi. Eecs 247 lecture 12: Data converters testing. http://www-inst.eecs.berkeley.edu/~ee247/fa09/files07/ lectures/L12_2_f09.pdf, 2009. Accessed 9/1/2018.
- [201] Pong P. Chu. FPGA Prototyping by VHDL Examples. John Wiley & Sons, Inc., Jan 2008.
- [202] Pong P. Chu. Embedded SoPC Design with Nios II Processor and VHDL Examples. John Wiley & Sons, Inc., Sep 2011.
- [203] Nios II processors. https://www.intel.co.uk/content/www/uk/en/ products/programmable/processor/nios-ii.html, 2020. Accessed 1/4/2020.
- [204] PicoBlaze 8-bit microcontroller. https://www.xilinx.com/products/ intellectual-property/picoblaze.html, 2011. Accessed 1/4/2020.

- [205] MicroBlaze soft processor core. https://www.xilinx.com/products/ design-tools/microblaze.html, 2020. Accessed 1/4/2020.
- [206] Elizabeth M. C. Hillman. Experimental and theoretical investigations of near infrared tomographic imaging methods and clinical applications. PhD thesis, Department of Medical Physics and biomedical Engineering, University College London, 2002.
- [207] Jeremy C. Hebden, Felipe M. Gonzalez, Adam Gibson, Elizabeth M. C. Hillman, Rozarina Md. Yusof, Nick Everdell, David T. Delpy, Giovanni Zaccanti, and Fabrizio Martelli. Assessment of an in situ temporal calibration method for time-resolved optical tomography. *Journal of Biomedical Optics*, 8(1):87, 2003.
- [208] Hideo Eda, Ichiro Oda, Yasunobu Ito, Yukihisa Wada, Yukio Oikawa, Yoshio Tsunazawa, Michinosuke Takada, Yutaka Tsuchiya, Yutaka Yamashita, Motoki Oda, Angelo Sassaroli, Yukio Yamada, and Mamoru Tamura. Multichannel time-resolved optical tomographic imaging system. *Review of Scientific Instruments*, 70(9):3595–3602, September 1999.
- [209] Elizabeth M. C. Hillman, Jeremy C. Hebden, Florian E. W. Schmidt, Simon R. Arridge, Martin Schweiger, Hamid Dehghani, and David T. Delpy. Calibration techniques and datatype extraction for time-resolved optical tomography. *Review of Scientific Instruments*, 71(9):3415–3427, Sep 2000.
- [210] F.T. Ulaby and U. Ravaioli. Fundamentals of Applied Electromagnetics. Pearson, 2015.
- [211] David M. Pozar. *Microwave Engineering*. John Wiley & Sons Inc, 2011.
- [212] James Sutherland. As edge speeds increase, wires become transmission lines. https://www.edn.com/ as-edge-speeds-increase-wires-become-transmission-lines/. Accessed 1/4/2020.

- [213] Ian Poole. Electronics notes. https://www.electronics-notes.com/. Accessed 10/12/2020.
- [214] Microwaves 101. https://www.microwaves101.com/, 2020. Accessed 10/12/2020.
- [215] Standard laser diodes. http://www.roithner-laser.com/ld_ standard.html#infra700. Accessed 9/1/2018.
- [216] Single Photon Counting Modules (SPCM). http://www.excelitas. com/Pages/Product/Single-Photon-Counting-Modules-SPCM.aspx, 2018. Accessed 9/1/2018.
- [217] SensL MicroFC-100XX-SMT. http://sensl.com/estore/ microfc-100xx-smt/, 2018. Accessed 9/1/2018.
- [218] 9307 pico-Timing Discriminator. http://www.ortec-online.com/ products/electronics/fast-timing-discriminators/9307, 2015. Accessed 9/1/2018.
- [219] SPC-130 »Low cost TCSPC module«. http://www.becker-hickl.com/ spc130.htm, 2018. Accessed 9/1/2018.
- [220] Nexys 4 DDR Artix-7 FPGA. http://store.digilentinc.com/ nexys-4-ddr-artix-7-fpga-trainer-board-recommended-for-ece-curriculum/, 2018. Accessed 9/1/2018.
- [221] Vivado Design Suite. https://www.xilinx.com/products/ design-tools/vivado.html, 2018. Accessed 9/1/2018.
- [222] Hansang Lim. Constant fraction discriminator involving automatic gain control to reduce time walk. *IEEE Transactions on Nuclear Science*, 61(4):2351– 2356, Aug 2014.
- [223] LTC6752/LTC6752-1/ LTC6752-2/LTC6752-3/ LTC6752-4. https: //www.analog.com/media/en/technical-documentation/ data-sheets/6752fc.pdf, 2014. Accessed 1/4/2020.

- [224] Pablo Pérez Tirador. Development of a portable time-domain system for optical tomography of the newborn infant brain. mathesis, University College London, August 2016.
- [225] Sreevatsan Bhaskaran. Direct Detection Time of Flight Lidar Sensor System Design and a Vortex Tracking Algorithm for a Doppler Lidar. phdthesis, Arizona State University, 2018.
- [226] Pooja Chandran, Ramesh Bhakthavatchalu, and P. Pradeep Kumar. Time of flight measurement system for an ultrasonic anemometer. In 2016 International Conference on Control, Instrumentation, Communication and Computational Technologies (ICCICCT). IEEE, December 2016.
- [227] Matthieu Chevrier and Thomas Mauer. It's about time: How TDR enables predictive maintenance for industrial Ethernet. https://www.ti.com/lit/ an/slyt755/slyt755.pdf, 2018. Accessed 12/1/2021.
- [228] Simon R. Arridge and M. Schweiger. Photon-measurement density functions part 2: Finite-element-method calculations. *Applied Optics*, 34(34):8026, Dec 1995.