Imperial College London
Department of Computing

A Photoplethysmography System Optimised for Pervasive Cardiac Monitoring

James Alwyn Cameron Patterson

February 2013

Supervised by Professor Guang-Zhong Yang

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy in Department of Computing of Imperial College London
and the Diploma of Imperial College London

Abstract

Photoplethysmography is a non-invasive sensing technique which infers instantaneous cardiac function from an optical measurement of blood vessels. This thesis presents a photoplethysmography based sensor system that has been developed specifically for the requirements of a pervasive healthcare monitoring system. Continuous monitoring of patients requires both the size and power consumption of the chosen sensor solution to be minimised to ensure the patients will be willing to use the device. Pervasive sensing also requires that the device be scalable for manufacturing in high volume at a build cost that healthcare providers are willing to accept. System level choice of both electronic circuits and signal processing techniques are based on their sensitivity to cardiac biosignals, robustness against noise inducing artefacts and simplicity of implementation. Numerical analysis is used to justify the implementation of a technique in hardware. Circuit prototyping and experimental data collection is used to validate a technique's application. The entire signal chain operates in the discrete-time domain which allows all of the signal processing to be implemented in firmware on an embedded processor which minimised the number of discrete components while optimising the trade-off between power and bandwidth in the analogue front-end. Synchronisation of the optical illumination and detection modules enables high dynamic range rejection of both AC and DC independent light sources without compromising the biosignal. Signal delineation is used to reduce the required communication bandwidth as it preserves both amplitude and temporal resolution of the non-stationary photoplethysmography signals allowing more complicated analytical techniques to be performed at the other end of communication channel. The complete sensing system is implemented on a single PCB using only commercial-off-the-shelf components and consumes less than 7.5mW of power. The sensor platform is validated by the successful capture of physiological data in a harsh optical sensing environment.

Declaration of Originality

I herewith certify that all material in this dissertation which is not my own work has been properly acknowledged.

James A. C. Patterson

Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

For Paul Summersgill, Nigel Sherman, Chris Fryer, Matt Hayes and Guang-Zhong Yang: my mentors.

For my brother, Doug, my benchmark.

For my parents, Olwyn and Ewen, who have provided more than a lifetime's worth of opportunities and support.

For my wife, Sam, who sacrificed more than I can ever repay.

For Callum, maybe someday you'll look up to me. Maybe not for this thesis though...

List of Abbreviations

AAF Anti-Aliasing Filter

ABP Arterial Blood Pressure

ACG Audiocardiography

ACHI Asymmetric Cubic Hermite Interpolation

ADC Analogue to Digital Converter

ANC Adaptive Noise Cancellation

ASIC Application Specific Integrated Circuit

bps Bits per second

BSN Body Sensor Networks

CCII Second Generation Current Conveyor

CHD Coronary Heart Disease

CHI Cubic Hermite Interpolation

CMOS Complementary Metal Oxide Silicon

COPD Chronic Obstructive Pulmonary Disease

COTS Commercial Off-The-Shelf

CT Continuous Time

CTIA Conventional Transimpedance Amplifier

CVD Cardiovascular Disease

DAC Digital to Analogue Converter

DDS Direct Digital Synthesizer

DFT Discrete Fourier Transform

DSP Digital Signal Processing

DT Discrete Time

ECG Electrocardiography

EMC Electromagnetic Compatibility

EMI Electromagnetic Interference

FIR Finite Impulse Response

FVF Flipped Voltage Follower

GBW Gain-Bandwidth Product

HR Heart Rate

IC Integrated Circuit

ICA Independent Component Analysis

ICG Impedance Cardiography

IIR Infinite Impulse Response

ITIA Integrating Transimpedance Amplifier

LED Light Emitting Diode

LI Linear Interpolation

LUT Look-up Table

LWDF Lattice Wave Digital Filter

MCU Microcontroller Unit

OTA Operational Transconductance Amplifier

PCB Printed Circuit Board

PPG Photoplethysmography

PWM Pulse Width Modulation

RMS Root Mean Square

ROM Read Only Memory

S_PO_2 Oxygen saturation

SNR Signal to Noise Ratio

TDM Time Division Multiplexing

TIA Transimpedance Amplifier

VCCS Voltage Controlled Current Source

WHO World Health Organisation

Contents

1 Introduction **19**

1.1 Requirements for Pervasive Healthcare Monitoring 22
1.2 Thesis Outline 25
1.3 Original Contributions 27
1.4 Publications 28

2 Pervasive Cardiac Monitoring **29**

2.1 Wireless Sensor Networks for Pervasive Sensing 30
2.1.1 Body Sensor Networks: On-Node Intelligence 33
2.1.2 Cardiovascular Monitoring 35
2.2 Introduction to Photoplethysmography 36
2.2.1 Signal Chain Model 37
2.2.2 Pulse Oximetry 39
2.3 Transducer Construction and Location 41
2.3.1 Light Emitting Diode 42
2.3.2 Photodiode 43
2.3.3 Time Division Multiplexing 43
2.3.4 The Transimpedance Amplifier 44
2.4 Noise and Artefacts 50
2.4.1 Ambient Artefacts 50
2.4.2 Motion Artefacts 57
2.5 State-of-the-Art in Wearable Photoplethysmography Systems . 60
2.6 Challenges 61
2.7 Conclusions 64

3 Integrating Photocurrent Amplifier **65**

3.1 Circuit Analysis 65
3.1.1 Transfer Function 66
3.1.2 Amplifier Gain 70
3.1.3 LED Driver Implications 72
3.1.4 Noise 74

3.2 PCB Implementation 78
3.2.1 Design Specification 78
3.2.2 Op-Amp Choice 79
3.2.3 Gain Comparison 79
3.2.4 Noise Comparison 80
3.3 Discrete-Time Implications 82
3.3.1 Converting to Discrete-Time 82
3.3.2 Timing Resolution vs. Bandwidth 84
3.3.3 Aliasing 84
3.4 Prototypes 85
3.4.1 Prototype Version 1 86
3.4.2 Prototype Version 2 87
3.4.3 Transducer 88
3.5 Conclusions 90

4 Optical Homodyne Sensing **92**
4.1 Introduction to Signal Mixing 92
4.1.1 Product Detection 94
4.1.2 Envelope Detection 95
4.2 Mixed-Signal Homodyne Detector 98
4.2.1 Carrier Synthesis 99
4.2.2 LED Drive 101
4.2.3 Timing for Multi-Channel Homodyne 102
4.2.4 Photodetector 104
4.2.5 Product Detection 104
4.2.6 Lattice Wave Digital Filter 108
4.3 Mixed-Signal Homodyne Sensing Validation 116
4.3.1 Envelope vs. Product Detection 116
4.3.2 Homodyne vs. CDS 117
4.3.3 Reduced Filter Order 118
4.4 Conclusions 119

5 Multi-mode Artefact Rejection **121**
5.1 Additive Artefacts 121
5.1.1 DC Current Stabilisation 122
5.1.2 Continuous Time Analysis 123
5.1.3 Differentiated Solution 124
5.1.4 Validation 127
5.1.5 Implementation Details 128
5.1.6 Optimised Current Stabilisation 132

5.2 Multiplicative Artefacts . . . 134
5.2.1 Theory . . . 134
5.2.2 Validation . . . 137
5.3 Conclusions . . . 145
5.3.1 Additive Artefacts . . . 145
5.3.2 Multiplicative Artefacts . . . 146

6 On-Node Signal Processing 147
6.1 Quadrature Phase Segmentation . . . 149
6.2 Extrema Detection . . . 152
6.2.1 Finite State Machines . . . 152
6.2.2 Relative Thresholding . . . 153
6.2.3 Extension to Quadrature . . . 154
6.2.4 Choosing a Threshold . . . 155
6.3 Signal Reconstruction . . . 158
6.4 Extension of Delineation to Pulse Oximetry . . . 160
6.5 Conclusions . . . 165

7 A Complete Wearable Photoplethysmography System 167
7.1 Smart-Patch Compatible PCB . . . 167
7.1.1 Modular Platform . . . 167
7.1.2 PPG Module . . . 168
7.1.3 PCB Summary . . . 171
7.2 Towards an ASIC Implementation . . . 172
7.3 Photocurrent Amplification . . . 173
7.3.1 Current Conveyor . . . 173
7.3.2 CCII Implementation . . . 176
7.3.3 Simulation . . . 178
7.3.4 Alternative Transfer Functions . . . 182
7.3.5 Differential Amplification . . . 182
7.3.6 Homodyne Detector . . . 183
7.4 IC Summary . . . 186

8 Conclusions 188
8.1 Contributions . . . 188
8.2 Future Work . . . 190

List of Tables

1.1 Requirements of a pervasive healthcare monitoring system and possible solutions. 23

2.1 Physical characteristics of a WSN 31
2.2 Logical characteristics of a WSN 31
2.3 Photoplethysmography Signal Chain Glossary 38
2.4 Comparison of Transmission (TX) and Reflection (RX) Mode PPG 42
2.5 Transimpedance Amplifier Design Trade-offs. 47
2.6 Comparison of Wearable PPG Sensors 62

3.1 Wearable PPG System Specification 79
3.2 Commercially Available Low Power Operational Amplifiers . . 80
3.3 Conventional transimpedance amplifier component values . . . 80
3.4 Integrating transimpedance amplifier component values 80

4.1 Filter Specification for Homodyne Product Detector 107

5.1 System Design Parameters 126
5.2 Absorption Glossary 135

6.1 Quadrature Phases of a Non-Stationary Signal 151
6.2 Effect of varying threshold on data rate on a 17s data set. . . . 156
6.3 Accuracy of three simple interpolation algorithms for a range of thresholds. 161

List of Figures

1.1 Transistions in size of pulse oximeters. From top to bottom: Nellcor N-180, Masimo Rad-5 and Nonin Wrist-Ox 3150. 21

2.1 Multi-tiered Network Architecture. 32
2.2 Illustration of transmission mode photoplethysmography with a break-down of contributors to optical absorption. 37
2.3 Photoplethysmography Signal Chain. 37
2.4 Absorption spectra for haemoglobin (Hb) and oxyhaemoglobin (HbO_2). 40
2.5 Illustration of reflection mode photoplethysmography. 42
2.6 Responsivity of a Vishay TEMD5010X01 silicon photodiode. . 43
2.7 Illustration of optical time division multiplexing. 44
2.8 A standard transimpedance amplifier with the equivalent circuit model of a photodiode consisting of a current source, shunt resistor and shunt capacitor. 45
2.9 An integrating transimpedance amplifier with reset switch. . . 48
2.10 A logarithmic transimpedance amplifier implemented with an op-amp and a bipolar transistor. 49
2.11 Infra-red and ambient light captured by a reflection mode transducer worn on the forehead showing noise created by artificial light at 50Hz and its harmonics. The noise is equivalent in magnitude to the PPG signal in both DC offset and AC components. The data shown here was captured from a reflection mode PPG transducer worn on the author's forehead. 52
2.12 Correlated Double Sampling photocurrent waveform. The crosses indicate where samples are to be taken for subsequent subtraction to remove the influence of ambient light. 53
2.13 IR with ambient light (both natural and artificial) rejected using CDS. Note the residual high frequency noise. Ambient light data is included to illustrate relative intensity. 53

2.14 A DFT of a photodetector output with low frequency content removed. The data shown here was captured from a reflection mode PPG transducer worn at the tip of the author's index finger. 54
2.15 Basic DC photocurrent compensation control loop. 55
2.16 A sample and hold circuit is used to only feedback photocurrent reading when the LED is on. Adapted from figure 3 in [1]. . . . 56

3.1 Integrating transimpedance amplifier for photocurrent amplification. Reset switch and parasitic shunt capacitances are included. 67
3.2 Amplifier response for a range of C_f with C_{sh} is fixed at 10nF and GBW is 500kHz. 68
3.3 Amplifier response for a range of C_{sh} with C_f is fixed at 100pF and GBW is 500kHz. 69
3.4 Amplifier response for a range of GBW. C_f is fixed at 100pF and C_{sh} is fixed at 10nF. 69
3.5 ITIA with passive feedback. 72
3.6 Energy content in cumulative harmonics of pulse width. 73
3.7 Amplifier noise spectrum for a range of C_f. C_{sh} is fixed at 1nF and GBW is 1000kHz. 77
3.8 Amplifier noise spectrum for a range of C_{sh}. C_f is fixed at 100pF and GBW is 1000kHz. 77
3.9 Amplifier noise spectrum for a range of GBW. C_f is fixed at 100pF and C_{sh} is fixed at 10nF. 78
3.10 Comparison of CTIA response vs. differentiated ITIA response designed for 3MV/A gain with a Microchip MCP6001 op-amp. 81
3.11 Comparison of CTIA and ITIA noise spectra. 81
3.12 Timing diagram to illustrate how time-division-multiplexing is used to measure both red and IR absorption as well as measuring the ambient light level. Included are the timing of the synchronised analogue-to-digital conversions and the resetting of the integrator. 83
3.13 Hand built prototype of the ITIA. 86
3.14 Version 1 of the wearable PPG system prototype PCB. 87
3.15 Version 2 of the wearable PPG system prototype PCB. 88
3.16 A hand-built flexible PPG transducer detailing the type of sensor used throughout this thesis for testing. 89
3.17 A comparison of the PPG signal quality achieved at a range of sensing locations. 90

4.1 The consequence of mixing two signals represented in both the time and frequency domains. 93
4.2 Mixed-signal homodyne PPG sensor system architecture. . . . 98
4.3 Presence of noise from ambient light sources in the photodetector spectrum. The data shown here was captured from a reflection mode PPG transducer worn at the tip of the author's index finger. 100
4.4 A two channel LED driver supporting both pulse width and amplitude modulation. 102
4.5 Timing diagram for a 2 channel homodyne detection system with integrating photocurrent amplification, pulse width modulation of LED intensity and synchronised analogue to digital conversion. 103
4.6 Circuit diagram for PCB implementation of the integrating transimpedance amplifier. 104
4.7 The spectrum on an unfiltered homodyne signal shows at what frequency the PPG signal intensity is equivalent in magnitude to the noise floor of the sensor circuitry. The data shown here was captured from a reflection mode PPG transducer worn at the tip of the author's index finger. 106
4.8 Gain, phase and group delay of the filter specified for the homodyne product detector. 107
4.9 Block diagram of a 5th order low pass filter implemented with a LWDF. Included is an adaptor element to match equation 4.9. Single sample delay elements are labelled using z-domain nomenclature: z^{-1}. 109
4.10 Two port adaptor types for lattice wave digital filters for compressed multiplication coefficient range. 112
4.11 Comparison of IIR filter and fixed-point LWDF in the frequency domain. This data is generated by taking the discrete Fourier transform of each filter's impulse response. The noise in the group delay plot is due to quantisation error. 114

4.12 Comparison of floating-point IIR filter and fixed-point LWDF on a data set captured from a reflection mode PPG transducer on the author's finger tip. The top plot shows the raw photocurrent data (as output from the MCU ADC) corrupted by intense ambient noise as shown in figure 4.3. The middle plot is a PPG biosignal recovered from the baseband using both a LWDF and an IIR filter. The bottom plot is the difference between the two filter outputs. 115
4.13 When truncated to 16 bits the LWDF is still bound to within ±0.5LSB of the IIR filter. 115
4.14 Product detection compared against envelope detection for implementing a homodyne detector. The raw input signal and the conventional baseband signal output are included for comparison. 117
4.15 Comparison of homodyne sensing against correlated double sampling. As the signal amplitudes are so similar the difference between the two signals is plotted underneath. 118
4.16 Product detection implemented with a 3rd order Butterworth filter (solid blue) instead of the 5th order Cheybshev (dashed black). This shows that the higher order filter is not necessarily required for sensing applications where ambient light does not dominate at higher frequencies. 119

5.1 A single channel PPG homodyne detection system with DC stabilisation of the input current. 123
5.2 Integrating transimpedance amplifier with mixed-signal input current and voltage biasing.. 124
5.3 Frequency response of the DC stabilised photodetector. Passband start frequency tuned to 101Hz. DC gain is 52.1dB less than pass band gain. 127
5.4 Without any input current stabilisation the detector saturates and no PPG biosignal is detectable in either the baseband or passband. 129
5.5 Enabling the current stabilisation control loop allows the system to operate with high ambient light levels. However, both baseband and homodyne outputs are corrupted while the input current level stabilises. 129

5.6 The high bandwidth input current compensation loop removes most PPG biosignal content from the baseband signal, but the homodyne output is now robust against ambient light disturbances. 130
5.7 A DFT of the discrete-time photodetector output with offset compensation. 131
5.8 The optimised DC current rejection system 133
5.9 Frequency response of the optimised DC current rejection system 133
5.10 Simplified absorption model of transmission mode PPG. 135
5.11 The three signal processing techniques used to compare artefact rejection performance: (a) simple low-pass filtering, (b) ratiometric measurement calculated from baseband signals, and (c) ratiometric measurement calculated from amplitude modulated signals. 139
5.12 Sensor system output from when twitching forehead muscles: (a) is the X-axis output of accelerometer, (b) is the output from the IR ADC channel after the carrier has been removed with low-pass filtering, (c) is the output of the baseband ratiometric algorithm, and (d) is the output of amplitude modulated ratiometric algorithm. 140
5.13 Comparison of the fPMR of each algorithm for a given type of motion induced noise. 141
5.14 Spectrograms from the baseband and amplitude modulated filter outputs of the red PPG channel from the natural motion experiment. Removal of additive optical signal components has significantly reduced presence of motion induced noise during jogging although there is still some residual artefact. 143
5.15 Spectrograms comparing non-ratiometric with ratiometric noise rejection. There is a small reduction in the relative intensity of motion noise from jogging. Note that the constant signal at 102BPM is the coupling of power supply noise in to the IR PPG channel. 144
5.16 LED light bypassing arterial blood in reflection mode PPG. . . 146

6.1 The four phases of a cosine wave. 150
6.2 Quadrature phases of a typical PPG cycle. 151
6.3 Simple Binary State Extrema Detector 152
6.4 Relative Thresholding Extrema Detector 153

6.5 Without thresholding the dicrotic wave is classified as its own waveform rather than a subset of the cardiac cycle. 154
6.6 FSM for the $f(t)$ half of a quadrature extrema detector. Note the extra condition on the state of the $f'(t)$ half of the detector required to transition from a validation phase into a tracking phase. 155
6.7 Overshoot caused by the cubic Hermite interpolation polynomial with a 2nA threshold. 160
6.8 Asymmetric cubic Hermite interpolation polynomial with a 2nA threshold compared. The gradient is scaled down by 2.5 in phase ϕ_3. 161
6.9 S_PO_2 calculated at each of the four quadrature extrema. 165

7.1 A prototype a smart-patch with a cavity for electronics. 168
7.2 An example PCB stack for use in a wireless smart-patch sensor. 168
7.3 Architecture of the complete PPG sensor system. 169
7.4 Version 3 of the wearable PPG system PCB. 169
7.5 Data output by the sensor patch PCB to validate operation. Measured at the finger tip with total LED current consumption <2mA. 171
7.6 A first generation current conveyor circuit and its application as a conventional transimpedance amplifier. 174
7.7 A second generation current conveyor circuit implemented using an op-amp to servo the voltage on node X. Complementary output nodes are included. 175
7.8 The *super-OTA*: a class-AB balanced OTA built using Flipped Voltage Followers and local common-mode feedback. 177
7.9 Simulated gain response of the proposed CCII based transimpedance amplifier. 179
7.10 Simulated transient response of the proposed CCII based transimpedance amplifier. 180
7.11 Simulated noise voltage spectrum at the output of the proposed CCII configured as a resistive transimpedance amplifier. 181
7.12 A transimpedance amplifier with a choice of linear, integrating and logarithmic transfer functions. 182
7.13 Wire mesh electromagnetic shielding required to protect singled ended transimpedance amplifier. 183

7.14 A differential transimpedance amplifier using a pair of current conveyor circuits. A capacitor is connected to the output to perform integrating transimpedance functionality. 184
7.15 Reconfigurable current conveyor cell. Complimentary outputs enable the implementation of the coarse product detection required for homodyne sensing. 185

1 Introduction

Organismal senescence in humans reduces the body's ability to combat disease and maintain homeostasis as it gets older; generally leading to a gradual decline in health. In the past, lower life expectancy and higher birth rates meant the percentage of the population at risk of age related illnesses was low; the healthcare demands of this demographic were not a significant concern. However, the average age of the world's population is increasing at an "unprecedented" rate [2] due to the double effect of falling fertility rates and rising life expectancy which has created a trend towards 21% of the world population being over the age of 60 by the year 2050, equalling the the population aged below 15 years old. The trend is more severe in developed nations where the over 60s are expected to make up 34% of the population, up from 19% in 2000, and healthcare systems must evolve to meet the demands of this new demographic profile. This has encouraged research into new medical practices that can efficiently and effectively monitor and treat age related illnesses while reducing the cost of healthcare delivery. The concept of pervasive healthcare aims to remove the constraints on where and for how long a patient can be monitored by integrating the latest sensor, wireless communication and ubiquitous computing technology with Body Sensor Networks (BSN) [3].

Continuously monitoring the physiology of a patient provides two major benefits over traditional episodic measurements: the ability to detect temporally brief events (e.g. heart attacks and strokes) and the acquisition of a rich data set for long term analysis of either a patient's recovery from an operation or general well-being. However, implementing a pervasive monitoring system is not a trivial task. Such a system must satisfy the requirements of the patient as well as the healthcare professionals and financial stakeholders. Healthcare systems are very risk averse and must conform to stringent specifications to ensure that no device will ever compromise the health or safety of a patient or carer at any time. While less critical than safety requirements, usability requirements are still important and must be addressed to ensure that the patients actually want to use the monitoring devices and that the carers can gather useful information from the system. It is these base requirements that lead us towards small, low power, wireless sensor design to encourage patients

to embrace the technology.

Modern sensor technology allows us to measure a wide range of physiological symptoms, but cardiovascular function is the most indicative of patient health. Cardiovascular diseases (CVDs) are responsible for 31% of all deaths each year and in 2008, of the 17 million deaths globally due to CVDs, 15 million victims were aged over 60 [4]. It is therefore crucial for a pervasive healthcare monitoring system to have the ability to sense cardiovascular function continuously. Phonocardiography and electrocardiography are both non-invasive methods of monitoring the heart, but are restricted to monitoring the mechanical and electrical properties of the heart respectively. Photoplethysmography (PPG) is another non-invasive sensing method which measures how the optical absorption and scattering properties of blood in the capillaries varies during the systole phase of the cardiac cycle. The dynamic changes in volume of blood allow heart rate to be measured by timing the interval between pressure waves propagating through the vascular system. Measuring the transit time of a pressure wave between two sites also allows blood pressure to be inferred. The spectral properties of the blood allow the percentage of oxygen saturation in the blood to be determined (S_PO_2) - a technique called pulse-oximetry. These sensing modalities have made PPG a popular bed-side monitoring solution. The ability to sense oxygenation as well as blood volume dynamics is so useful that WHO have begun the "Patient Safety Pulse Oximetry" project to delivery pulse oximeters into operating theatres in developing nations.

This brings up the issue of the global divide between nations with the wealth to deliver effective healthcare systems and those without. In context, CVDs are both treatable and preventable although the costs of delivering the treatment means that not everywhere or everyone has access to the same level of healthcare. Only 4% of CVDs deaths were considered premature for high income countries, but this rises to 42% for low income countries [4] which means there is significant motivation to produce pervasive cardiac monitoring systems that are also low in cost such that wealth is not a barrier to the health benefits that these system can deliver.

Pulse oximeters are experiencing the same trend towards miniaturisation as other electronic consumer goods as can be observed in figure 1.1. However, transferring pulse oximetry to a truly wearable, or even implantable, solution is non-trivial due to requirements on performance, size, longevity of operation - all while operating in environments which can be hostile to optical sensing techniques.

A truly pervasive healthcare monitoring system will provide a framework for analysing and interpreting data as well as providing a continuous sens-

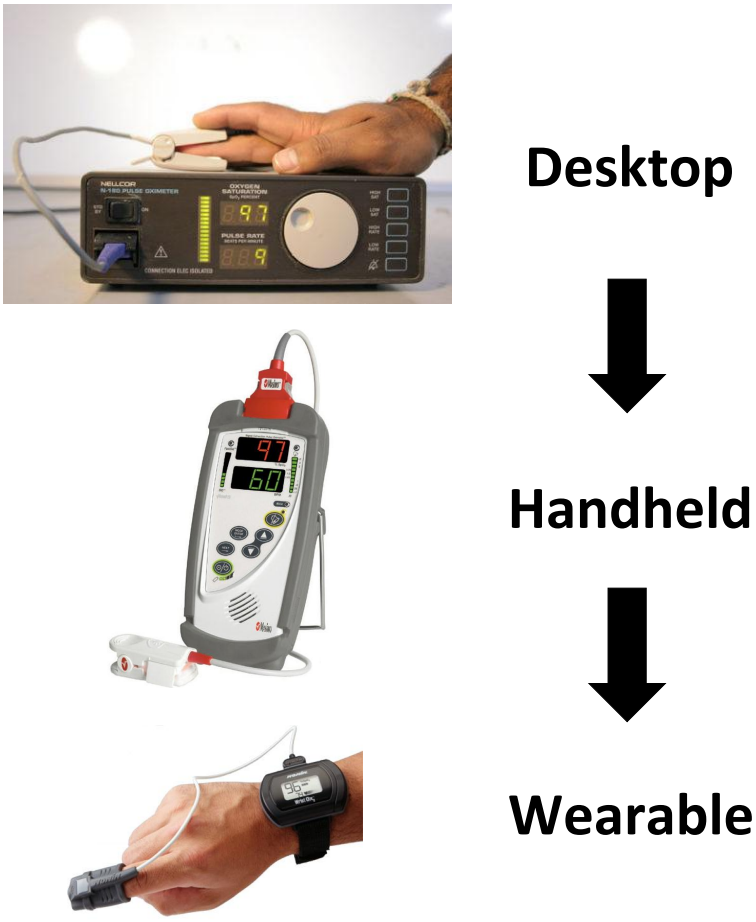


Figure 1.1: Transistions in size of pulse oximeters. From top to bottom: Nellcor N-180, Masimo Rad-5 and Nonin Wrist-Ox 3150.

ing solution. The traditional acquire-and-transmit method for gathering and processing data is far too reliant on communications. Transmitting data wirelessly is the prime consumer of energy and needs to be minimised. Moving the data processing tasks on to the sensor node will reduce the usage of power hungry communications and also allows the node to respond directly to its ambient environment if required to do so. This research aims to address the requirements of a pervasive sensing system through the application of novel techniques at both the circuit and the system level.

1.1 Requirements for Pervasive Healthcare Monitoring

The first and foremost requirement of any healthcare delivery system is that it must improve the well-being of the patient. A monitoring system may not have any direct effects on the health of a patient, but supplying useful information allows healthcare professionals to make more informed decisions and hence prescribe more effective treatments. In some cases the monitoring system may provide direct feedback to the patient allowing them to observe how their health is affected by their lifestyle.

Equally as important is the effect a device has on the safety of a patient. It is rarely feasible to guarantee any given device will not fail to perform its intended task, but the negative impact of a failure on the well being of the user must be minimised - and must never outweigh the benefits of using the treatment in the first place. To this end there are regulations by which medical devices must conform before they can be sold and used publically. In the UK the Medicines and Healthcare products Regulatory Agency (MHRA) enforce the European Community regulations regarding the conformity of medical devices to safety guidelines.

Following safety are the more aesthetic requirements of the end user. No patient will want to use a device that is uncomfortable to wear or that requires regular attention for maintenance issues such as battery charging or cleaning. Hence, it is directly from the user requirements that the primary physical requirements of size and operational lifetime are determined.

Finally the commercial aspects need to be considered. An intentional benefit of pervasive sensing is that it reduces the labour costs involved with healthcare monitoring by allowing the sensor networks to perform the role autonomously. This approach is feasible as the cost of manufacturing these sensor devices can be minimised by taking advantage of the economies of scale associated with producing integrated-circuit based electronic solutions.

There are many technical requirements which derive from these fundamental requirements. When specifying a healthcare delivery system it is important to accurately capture these requirements to ensure that all stakeholders and regulators are satisfied with the end product. The table below summarises the requirements which pose some of the key research challenges in the development of BSN for pervasive healthcare.

Table 1.1: Requirements of a pervasive healthcare monitoring system and possible solutions.

Requirement	**Approaches**
Intuitive data visualizations	Data abstraction, information rich graphical user interfaces, decision support systems
High quality information	Context-based sampling strategy, conformity with approved medical systems
Privacy	Data encryption, data anonymity, bio-inspired immune systems, privacy preserving data mining
Disease-specific sensing	Sensing modalities and algorithms targeted to specific patient conditions
Variable subject lifestyle	Sensing modalities and algorithms adaptable to different lifestyles, physically robust sensors, on-sensor data buffering for when out of network range
Scalability	Distributed processing, on-sensor processing, autonomic system design
Fault tolerance	Hardware and software failure recovery, feature redundancy in communications, autonomic system design
Communication reliability	On-sensor data buffering, autonomic system design
Ergonomics and wearability	Hardware miniaturization, tactile and unobtrusive devices, wireless communication

Adapted from the Wireless Sensor Network chapter written by the author and published in the Handbook of Healthcare Delivery Systems [5]

From that list of requirements the one that most urgently needs addressing for wearable PPG based cardiac monitoring system is *ergonomics and wearability* as this drives whether or not the patient is likely to adopt the technology. There is also the requirement for *high quality information* which any instrumentation system must abide by in order to ensure that the device delivers a benefit to the patient. With these two requirements in mind the following four key challenges have been outlined for research into a solution optimised for wearable PPG sensor systems.

1. **Low power opto-electronic instrumentation**
   The fact that PPG requires active illumination of the patient's skin means that it can never feasibly be made as low power as passive cardiac sensing modalities such as ECG, but it should be possible to reduce the power to the point where the benefits of PPG outweigh the costs. To minimise the amount of optical energy emitted by the illuminator the gain of the photodetector needs to be maximised. Normally this requires using a chain of multiple electronic amplifiers or relying on a full custom integrated circuit (IC) solution. Amplifier chains need too many discrete components to allow them to be implemented in a small form factor PCB and the cost of developing an IC can often be too prohibitive to consider using. By moving away from standard photodiode amplifier circuits it should be possible to find a solution that can be implemented with the minimum number of components.

2. **Robust physiological signal extraction**
   Ideally a PPG system would only measure the change in light levels due to light from the sensor's illumination module being absorbed by the wearer's arterial capillary blood, but there are countless independent light sources which the PPG will detect photons from and the optical path of the sensor contains a complex mix of different tissue and fluid components. The wearable PPG system will experience these complexities to a more severe degree than that of a PPG system used in the relatively controlled environment of a medical clinic. Investigation is required into techniques that can lock on to both the correct source of light and the correct mode of optical absorption.

3. **High resolution data at low data rates**
   A pervasive healthcare system requires that its constituent sensor nodes exhibit some degree of autonomy in choosing what and how much data to send to an observer without saturating the network with data that actually has little useful information content. General purpose techniques used for compressing data may not be suitable due to their requirement for memory and processing power that a low power sensor node might not be able to provide. As such it is necessary to select only the most relevant and accurate information from the raw physiological data based on the context of the measurement situation.

4. **Miniaturised physical form factor**
   It is not possible to accurately recreate the environment of an outwardly mobile patient in the confines of a research lab, nor is it possible to

expect a patient to wear an unwieldy, hand built, electronic instrumentation system on their body and still be able to carry on with their life unhindered. Whatever solution is chosen to solve the preceding challenges it must be synthesisable in to a form factor smaller than anything that has preceded it.

1.2 Thesis Outline

Here is an outline of the remaining chapters with a summary of each one's contribution to this thesis.

2. **Pervasive Cardiac Monitoring** introduces the reader to any background knowledge that is required as a precursor to the research presented in later chapters. Included is a summary of the most up to date technology and research that is used in the application of photoplethysmography to pervasive healthcare monitoring.

3. **Integrating Photocurrent Amplifier** presents an underused alternative to the form of transimpedance amplifier conventionally used for implementing the photodetection circuitry in a PPG system. The integrating form of the transimpedance amplifier is shown to be capable of achieving very high gain while providing superior bandwidth which permits measurement of physiological signals with high temporal resolution. Rigorous modelling of the gain and noise parameter, in comparison to the conventional transimpedance amplifier, simplifies the design procedure of what is an inherently unstable circuit.

4. **Optical Homodyne Sensing** investigates the separation of independent light sources from the PPG sensor's own illumination circuit using frequency domain techniques. Synchronous photocurent detection is a widely used technique in PPG, but homodyne sensing is a generalised solution which can be implemented almost entirely in the discrete-time domain. This means it can work in conjunction with time division multiplexing of multiple optical wavelengths and can be embedded in firmware on a microcontroller such that no additional analogue circuitry is required.

5. **Multi-Mode Artefact Rejection** builds on top of homodyne sensing to use the redundant region of the signal spectrum to cancel out any independent light sources that are so high in magnitude that they would saturate the analogue front end circuitry, rendering subsequent artefact rejection

useless. Having removed all linearly combined noise sources, a technique for rejection of non-linear noise sources, such as motion artefact, is investigated for its applicability to reflection mode PPG that is used predominantly in wearable solutions.

6. **On-Node Signal Delineation** aims to take the output signal from the artefact rejection system and select only the samples which best represent the underlying physiological waveform. Delineation is chosen as the non-stationary nature of biosignals may require high temporal resolution and very little frequency domain information. The cyclic nature of cardiac derived signals is embedded into the core function of the delineation machine which is also has programmable granularity to allow trade-offs between signal reconstruction accuracy and data rate.

7. **A Complete Wearable Photoplethysmography System** presents an embodiment of all methods previously described in a single printed circuit board that has been targeted for implementation in a wearable patch. The sensor provides best-in-class size and power consumption due to the possibility of implementing most of the system's functionality in a contemporary mixed-signal microcontroller. Also presented is an equivalent integrated circuit solution that uses general purpose analogue building blocks to implement not only the proposed PPG system, but a single chip that can be configured for many different sensing modalities.

8. **Conclusions** provides an executive summary of the main achievements of this thesis.

1.3 Original Contributions

The following items highlight how the research presented in this thesis has made a positive contribution to field of wearable PPG systems.

- Analysis of, and design methodology for, integrating transimpedance amplifier circuits with respect to large parasitic input capacitance and micro-power operational amplifiers.
- Homodyne light detection system with all signal processing performed in the discrete time domain.
- DC light rejection with high dynamic range and fast response using remainder bandwidth from homodyne detection.
- Efficient algorithm for reducing transmission data rate using delineation with complementary signal reconstruction method at receiver.
- Implementation of the PPG photometric instrumentation circuitry using general purpose integrated circuit current conveyor cell.

1.4 Publications

The following publications were produced from the research carried out in the process of creating this thesis.

1. J. A. Patterson, D. C. McIlwraith, and G.-Z. Yang, "A Flexible, Low Noise Reflective PPG Sensor Platform for Ear-Worn Heart Rate Monitoring," in *2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, pp. 286–291, IEEE, June 2009. [6]

2. J. A. C. Patterson, R. Ali, and G.-Z. Yang, "Wireless Sensor Network," in *Handbook of Healthcare Delivery Systems* (Y. Yih, ed.), ch. 46, pp. 46–1–46–18, CRC Press, 2010. [5]

3. J. A. C. Patterson and G.-Z. Yang, "Ratiometric Artefact Reduction in Low Power, Discrete-Time, Reflective Photoplethysmography," in *2010 International Conference on Body Sensor Networks*, pp. 174–179, IEEE, June 2010. [7]

4. J. A. C. Patterson and G.-Z. Yang, "Ratiometric Artifact Reduction in Low Power Reflective Photoplethysmography," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 5, pp. 330–338, Aug. 2011. [8]

5. J. A. Patterson and G.-Z. Yang, "Dual-Mode Additive Noise Rejection in Wearable Photoplethysmography," in *2012 Ninth International Conference on Wearable and Implantable Body Sensor Networks*, pp. 97–102, IEEE, May 2012. [9]

2 Pervasive Cardiac Monitoring

Parts of this chapter have been extracted, adapted or paraphrased from the *"Wireless Sensor Network"* chapter in the *"Handbook of Healthcare Delivery Systems"* [5]. Any figures that have been copied or modified from the original text will be identified.

Wireless Body Sensor Networks have been developed specifically to enable research into the technologies needed to address the requirements of a ubiquitous healthcare monitoring system [3]. Based on the wider field of Wireless Sensor Networks (WSN), a body sensor network (BSN) monitoring system is a distributed infrastructure of sensor nodes capable of autonomously gathering, processing and transmitting sensor data into a wireless network. The sensor data generated within the network can then be accessed externally via a gateway node. For consistency, this thesis will now refer to any person being monitored by sensor nodes as a *subject* and any entity accessing the sensor data via a gateway node as an *observer*.

These sensor nodes are a combination of sensor interfaces, embedded processing, wireless communications and power source, and are tailored specifically to address the low power and small size requirements of mobile healthcare monitoring devices. The viability of BSN platforms as a research tool has been made possible by the semiconductor industry's drive to reduce size and cost of integrated circuits. Advanced digital radio solutions, such as ZigBee, ANT+ and Bluetooth Low Energy devices as well as microcontrollers capable of running operating systems are all available as low cost Commercial-Off-The-Shelf (COTS) components with current consumptions low enough to be powered by coin cell batteries for hours, even days between charges. In the near future it should be possible for sensor nodes to operate indefinitely with energy harvested solely from the sensing environment using miniaturised power generators [10].

The main benefits of wireless communication over traditional wired solutions are the ease of deploying the network and the mobility it provides to the sensor node. What embedded processing provides is robustness; the sensor node is able to monitor its own environment and adapt its data processing, communication interface and power consumption to suit the context. Context-

awareness may be a relatively recent field of research, but it is clearly a more attractive solution for enabling pervasive sensing than having to model accurately the real-world environment of a BSN.

2.1 Wireless Sensor Networks for Pervasive Sensing

The development of Wireless Sensor Networks [11] began in the late 1990s as advances in silicon integrated circuit (IC), micro-electromechanical systems (MEMS) and packaging techniques made it feasible to implement autonomous monitoring systems that could be unobtrusively deployed throughout an environment. *Smart Dust* was one of the earliest systems which proposed a solution for monitoring movement over a large geographical area by integrating MEMS sensors, optical communications and solar power generation into a single device no larger than $1mm^3$ [12]. Since then many more applications for WSN have been proposed, but the majority of current implementations are based on COTS low power digital wireless communication solutions and microcontrollers.

The phrase "low power" is a pretty generic term, but in the context of this thesis it is intended to encompass any sensor system circuitry that can be powered from a battery weighing less than 5 grams and operate for a day between charging. While this definition is itself not very precise it does at least narrow the range in suitable power sources; an example battery is a 3.7V lithium polymer based cell which has a capacity of 110mAh which weighs 2.65g and has a space envelope of 4x15x28mm [13]. Lithium polymer is currently a popular choice of battery construction for rechargeable electronics devices, especially mobile phones, due to its high energy density and simplicity in construction. In reality only empirically acquired evidence will determine how long a battery can source energy to a given sensor system, but from these headline specifications a reasonable estimate would be capacity divided by demand, i.e. this cell could supply an average current of 4.5mA for a 24 hour period.

More generally, power consumption is a performance characteristic that is derived from more functional requirements like sensor node location (i.e. is it wearable, implantable, ambient?) and sensor node mobility (i.e. is it fixed in location or mobile with respect to the observer?). Characteristics of a WSN can be classified as either physical or logical. The tables below highlight some of these requirements and how they affect the detailed technical specification of a WSN.

Wireless communications give WSN many beneficial characteristics over tra-

Table 2.1: Physical characteristics of a WSN

Physical Characteristics	**Implications**
Node Location	Whether the node is ambient, wearable or implantable determines how large it can be and what environment it is likely to experience.
Communication Range	The larger the distance between devices the more power is required to establish a communications link.
Operating Lifetime	How long the device must operate before it requires servicing. In general this will determine maximum power consumption from a predetermined power source.
Information Density	The amount of data that needs to be transmitted in a given period of time; more data needs higher bandwidth which, in turn, requires more power.

Table 2.2: Logical characteristics of a WSN

Logical Characteristics	**Implications**
Node Dynamics	Mobility of sensor nodes may mean that the network infrastructure must be able to adapt in real time to maintain communication paths between all sensors.
Network Scope	The size of a network based on both the geographical area it covers and the number of active nodes associated with the network.
Network Capacity	How much information can propagate through the network without causing significant data loss or total communication failures.
Communication Reliability	The level of acceptable data loss determines the robustness of the communication protocol which can affect throughput and average power consumption due to resending data or embedding forward error correction.
Information Security	Privacy and incorruptibility of data is crucial in healthcare so encryption and authentication must be supported by the communication protocol.

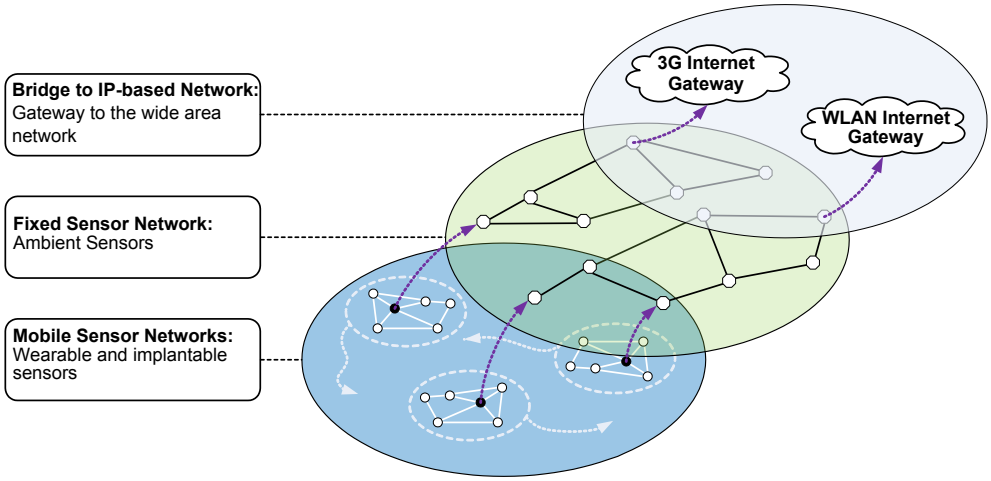


Figure 2.1: Multi-tiered Network Architecture.
Adapted from the WSN chapter in the Handbook of Healthcare Delivery Systems

ditional wired network solutions. One of the greatest benefits is that it makes physical cabling obsolete. The cost of wireless transceiver modules is continually dropping, but the price of copper conduit, per unit length, is relatively constant meaning that WSN is becoming the only financially justifiable sensor network solution for large installations. The small physical outline of sensor nodes in conjunction with the total lack of cabling makes the installation of the network minimally intrusive to the environment. Not being physically constrained by cables allows the network structure to be dynamic. While this makes communication between nodes more complicated it enables monitoring applications that were previously never practical due to the lack of node mobility in wired networks.

Even though wireless networks are not constrained in to a physical structure by cables they still require some form of logical structure, called a *topology*, to enable them to communicate with each other effectively. The choice of topology is highly dependent on the requirements of the monitoring system. Classical topologies, like ring and bus, are unlikely to be suitable due to their inability to support mobile nodes. For a pervasive healthcare monitoring system the chosen topology must support many different types of sensor nodes, a large geographical area and a high sensor node population. A multi-tiered architecture provides different levels of abstraction to represent the requirements of different sensor nodes. Figure 2.1 illustrates a typical healthcare monitoring scenario where multiple subjects, each wearing multiple sensors multiple sensors, are being monitored by multiple ambient sensors whose data is to be viewed by multiple observers.

In this example mobile sensor nodes form into small, ad hoc mesh networks

with the other nodes being worn by, or implanted in, the subject. These clusters of nodes pass data up the hierarchy, via a gateway node, to an ad hoc mesh network of sensor nodes in fixed locations. Fixed sensors can often afford larger power supplies and hence support higher data rates to support the increased volume of data that the sensors typically produce (e.g. video cameras). The increased power supply capability also enables the sensors to communicate over a larger range. A bridge is created between the fixed sensor network and the existing infrastructure of the internet to allow observers to access the sensor information from anywhere in the world (well, wherever internet access is available).

2.1.1 Body Sensor Networks: On-Node Intelligence

Traditional wireless sensor networks work on an acquire-and-transmit methodology where data captured continuously from a sensor is sent into the wireless network with minimal processing being performed on the data beyond basic spectral noise reduction (i.e. averaging). Having a large accumulation of raw sensor data is fine for research purposes where it is required to test and validate new analytical techniques, but it is too costly, in terms of power consumption and robustness, for a clinically viable healthcare delivery solution. Moving the role of data processing from a PC to the sensor node itself can minimize the usage of the wireless data connection allowing conservation of energy reserves and increased data integrity. BSNs aim to provide a healthcare specific solution by addressing the issues of deciding how to interpret sensor data, if and how to react to sensed events, and how to actually implement the data processing on the sensor node.

Context Aware Sensing

Context-awareness is the ability of a sensing system to observe how underlying trends affect an object. In the case of healthcare monitoring these trends can indicate how the lifestyle of a patient is affecting their health or even how a course of treatment is affecting the subject's lifestyle. The raw data provided by a sensor, or a combination of sensors, might not directly present contextual information although the data is influenced by the underlying context. As an example, an electrocardiogram (ECG or EKG) provides a measurement of heart rate but to determine the fitness of the person requires knowing how their cardiac function is affected by physical activity. In this case the sensor system is made context-aware by adding activity recognition [14].

For BSNs, detecting underlying health trends is not the only purpose of

context-aware sensing. Knowing how the ambient environment is affecting a measurement gives the sensor system the ability to compensate for the induced artifacts. One of the most studied environmental influences on sensors is temperature drift. Resistors and semiconductor devices can be particularly sensitive to temperature although compensating for the drift is a well understood process. However, the ubiquitous nature of BSNs means they will experience a much wider range of physical effects than just temperature drift.

Embedded Data Enhancement

In practice, rarely is the data acquired direct from a sensor a direct representation of the underlying event being monitored. For example, a PPG sensor outputs a signal representative of the cardiac cycle, but to measure the heart rate requires an algorithm to determine the periodicity of that cycle. A software based solution typically uses some form of peak detection algorithm in conjunction with a timer to measure the waveform's period. In this case the algorithm is simple enough for even the most basic microcontroller to implement, but more mathematically intensive feature extraction methods will require the sensor node transmitting raw data over the wireless connection to a more powerful processor. This approach is unsuitable as a fundamental requirement of BSNs is to minimize the usage of such resources in order to maximise operating lifetime.

Performing the feature extraction on the sensor node is not solely for reducing the power consumption, but is essential for autonomy. A system that does not have to transmit sensor data for analysis will be far more responsive to external stimuli. In addition to improved response time the node will be able to acquire and analyse useful data even if the node moves out of the communication range. Closing the control loop on-node means finding a solution to the opposing requirements of low power consumption and strong mathematical computational power. The problem lies with relying on the digital representations of data as is required by the memory storage and wireless communications components of a sensor node. To increase the computational power of a digital processor, it requires either more processing time to implement complex functions from combinations of more simple programs (e.g. calculating a logarithm using the Taylor's series), or using bespoke digital circuitry to perform the calculation (e.g. performing a multiplication with a binary multiplier circuit [15] as opposed to an accumulator loop in software). Both of these methods require additional transistors for the required extra program memory or custom circuitry respectively. Generally a balance is struck between the two methods with the software solution being cheaper

and the hardware method being faster. For autonomic sensing, neither solution is suitable as the inherent mathematical complexity in feature extraction would require too much memory and take too long to calculate if performed in software and the extra circuitry required to perform it digitally in hardware would blow the budget on device power consumption and size.

Compared to digital electronics, analogue circuitry is rich in mathematical properties. Addition and subtraction are implemented purely by Kirchhoff's Current Law, derived from the conservation of charge principle, where the sum of the currents flowing into a circuit node must equal the sum of the currents flowing out of the node. Multiplication and division can be performed by converting linear signals into logarithmic signals using either bipolar transistors, or CMOS transistors biased in the sub-threshold region. In the log domain addition and subtraction are equivalent to multiplication and division in the linear domain. This translinear process is the basis of the Gilbert multiplier cell which only requires 8 transistors to perform a full four-quadrant multiplication [16].

From the late 1980s, analogue and mixed-signal Artificial Neural Networks (ANN) have received a considerable amount of interest and led to the development of a number of analogue implementations such as the Cellular Neural Network, self-organizing neural networks and the SeeHear chip [17, 18, 19]. These forms of ANNs have not yet found widespread use in commercial products due to their complexity in the design phase and relative specificity in feature extraction ability. To implement a useful mixed-signal ANN requires 10's to 100's of neurons which may require thousands of transistors [20].

2.1.2 Cardiovascular Monitoring

CVDs have already been identified as posing a significant health risk, but also as preventable with both early detection of illness and lifestyle choices. The British Heart Foundation identify smoking, poor diet, low activity levels, alcohol and obesity as primary risk factors in the cause of coronary heart disease (CHD) [21].

This still leaves the problem of choosing how to actually take measurements of cardiovascular function. Auscultation, or listening directly to the sounds of the body, has historically been the only non-invasive cardiovascular monitoring method. These days, the two primary non-invasive methods are electrocardiography and photoplethysmography although there are others such as impedance cardiography and audiocardiography which is the electronic form of cardiovascular auscultation.

Electrocardiography: (ECG) Electrical measurement of the electrical activity of the subject's heart via electrodes attached to the subject's chest.

Photoplethysmography: (PPG) Optical measurement of the volume of blood in a localised region of the vascular system.

Impedance Cardiography: (ICG) Electrical measurement of the changes of impedance in the subject's thorax due to changes in blood volume.

Audiocardiography: (ACG) Acoustic measurement of sounds due to the mechanical function of the subject's heart.

ECG is particularly popular for translation to wearable sensing as it can be made very low power and has been adopted as the primary heart rate monitoring solution for athletes when implemented as a chest strap with wireless telemetry to an associated watch. However, the chest straps are noticeable by the subject and generally not worn beyond the duration of exercise which has led to the development of ultra-low power patch based solutions [22, 23, 24].

Where PPG has an advantage over the other sensing methods is that it can infer chemical content in the blood, arterial oxygen saturation in particular using pulse oximetry theory. This means that it also gives some indication of lung function which can aid in the monitoring of long-term illnesses such as chronic obstructive pulmonary disease (COPD) [25]. Furthermore, PPG is more suited to wearable systems as it only needs a single site to attach the sensor whereas ECG needs at least 2 connections to the body, separated by centimetres. PPG can also be placed at a number of locations about the body [26], but ECG is constrained to being located on the chest. For these reasons, this thesis focuses exclusively on photoplethysmography as a solution to pervasive cardiac monitoring.

2.2 Introduction to Photoplethysmography

The instantaneous volume of blood in a given location can be considered as a function of cardiac cycle; volume increases in the systole phase and decreases in the diastole phase. This relationship allows the direct measurement of heart rate as well as inferring many other characteristics of cardiovascular function based on dynamic changes in blood pressure.

A basic PPG sensor system consists of a light source to illuminate the blood vessels and a photodetector to measure the combined optical absorption or scattering properties of the blood, tissue and bone. As PPG is an optical measurement it is non-invasive and the transducer can be made at low cost

by using a light emitting diode (LED) for the illuminator and a photodiode for the detector. Figure 2.2 illustrates a typical PPG set up where an LED is used to shine light through a finger and the placement of a photodiode to measure how much light is received.

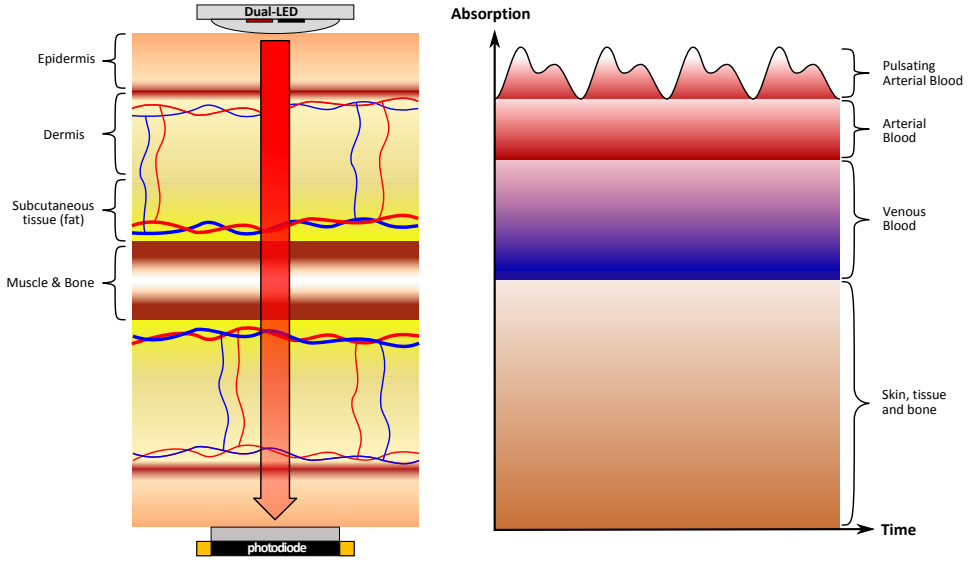


Figure 2.2: Illustration of transmission mode photoplethysmography with a break-down of contributors to optical absorption.
(Not to scale)

2.2.1 Signal Chain Model

Figure 2.3 illustrates the end-to-end signal chain of a PPG sensing system and table 2.3 describes each of the components in the chain. Note that geometrical parameters, such as the area of the photodiode and the beam angle of the LED, have been omitted to keep the system description as simple as possible.

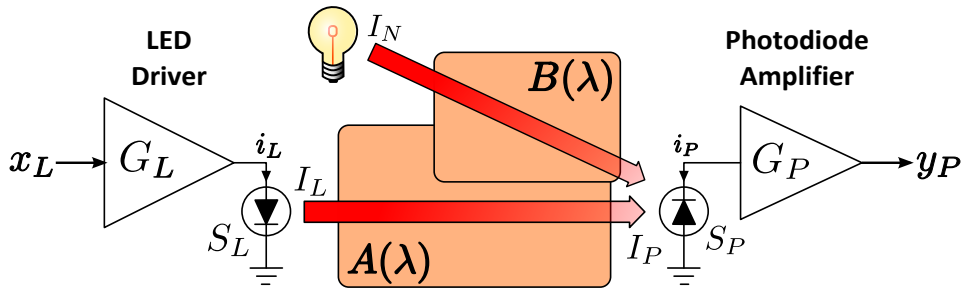


Figure 2.3: Photoplethysmography Signal Chain.

The Beer-Lambert law is often used as an approximate model for the absorption of LED light in a PPG sensing application [27]. The equation below

Table 2.3: Photoplethysmography Signal Chain Glossary

Symbol	Description	Unit
x_L	System Input	V
G_L	LED driver transconductance	Ω^{-1}
i_L	LED drive current	A
$S_L(\lambda)$	LED conversion efficiency	$(W \cdot m^{-1})/A$
$I_L(\lambda)$	Emitted spectral power from LED	$W \cdot m^{-1}$
$A(\lambda)$	Coefficient of absorption of blood, tissue, etc.	
$I_{N,j}(\lambda)$	Emitted spectral power from ambient light source j	$W \cdot m^{-1}$
$B_j(\lambda)$	Coefficient of absorption for each of $I_{N,j}(\lambda)$	
$I_P(\lambda)$	Spectral irradiance of the photodiode	$W \cdot m^{-1}$
$S_P(\lambda)$	Photodiode responsivity	$A/(W \cdot m^{-1})$
i_P	Photocurrent output of photodiode	A
G_P	Photodiode amplifier transimpedance	Ω
y_P	System Output	V

defines the application of this model to a PPG system in terms of time, t, and optical wavelength, λ:

$$I_P(\lambda, t) \equiv I_L(\lambda, t) \cdot e^{-A(\lambda,t)}, \tag{2.1}$$

$$A(\lambda, t) = \sum_{\forall i} \varepsilon_i(\lambda) c_i d_i(t) \tag{2.2}$$

where $\varepsilon_i(\lambda)$, the extinction coefficient of material i at wavelength λ; c_i, the concentration of material i; $d_i(t)$ is the optical path length; and $A(\lambda, t)$ is the collective term for the absorption coefficients of the tissue being sensed. The Beer-Lambert law is not the most accurate model for a PPG sensing system [28, 29], but it is a good enough approximation to illustrate the non-linear relationship between the amount of light detected in a PPG system and the variable of primary interest: $d_i(t)$, the time varying volume of the blood vessels which allows the detection of the cardiac cycle. For an approximation of the entire signal chain we can use the following equation:

$$y_P(\lambda, t) = G_P S_P(\lambda) \left(e^{-A(\lambda,t)} S_L(\lambda) G_L x_L(t) + \sum_{\forall j} \left[e^{-B_j(\lambda,t)} I_{N,j}(\lambda, t) \right] \right) \tag{2.3}$$

It is worth noting than terms A and B_j are likely to have significant correlation, i.e. some ambient light will be absorbed by the subject's tissue before being

detected by the photodiode.

A lot of coefficients have been included in this equation for completeness and each are defined by specific choices made during the design of the PPG transducer and by the choice of sensing location. However, a representative attenuation between the LED driver and the photodiode amplifier for the experiments performed in this thesis is in the order of 10,000 (i.e. 1mA supplied to the LED will result in 100nA generated by the photo diode). In terms of attenuation of optical power one needs to know the conversion efficiency of the LED and the photodiode at the chosen wavelength. For example, the IR LED from the OSI-Optoelectronics DLED-660/905-LLS-3 dual emitter and the Vishay TEMD5010X01 silicon photodiode with active area of 7.5mm^2 are components that were used repeatedly throughout the development of this thesis. The LED emits about 0.1mW per 1mA of drive current at 905nm and the photodiode creates about 0.7 Amps per Watt of received optical power. So the overall optical attenuation is in the order of 700 which is due to the combined effects of scattering, absorption and physical geometry of the transducer.

Typically an infra-red (IR) wavelength is chosen for the illuminator as it is absorbed less by blood than visible wavelengths which means less optical power is required to achieve a desired overall signal to noise ratio. In general, the dynamic signal component attributed to pulsating arterial blood is responsible for about 1% of the detected optical signal making dynamic range one of the more important requirements of a PPG photodetector.

For a more complete introduction to PPG technologies and their clinical application see the topical review written by John Allen [30].

2.2.2 Pulse Oximetry

In addition to measuring heart rate it is also possible to measure the oxygen saturation level of the arterial blood by measuring the absorption at 2 wavelengths. This is the concept of *pulse oximetry* , invented in the 1970s by Takao Aoyagi [27]. As a non-invasive method of measuring the percentage of blood oxygen saturation (commonly know as S_PO_2) pulse oximetry has proven valuable for many medical fields; anaesthesiology in particular. To understand how the measurement of S_PO_2 can be derived from a PPG measurement observe the different absorption characteristics of haemoglobin (Hb) and oxyhaemoglobin (HbO_2) as illustrated in figure 2.4.

Included in figure 2.4 are approximate emission spectra of a red and an IR LED. As the LED emission wavelengths are either side of the isobestic point (the point where the two absorption spectra cross) a ratiometric measurement can be made from which the S_PO_2 value can be derived. First of all, S_PO_2 is

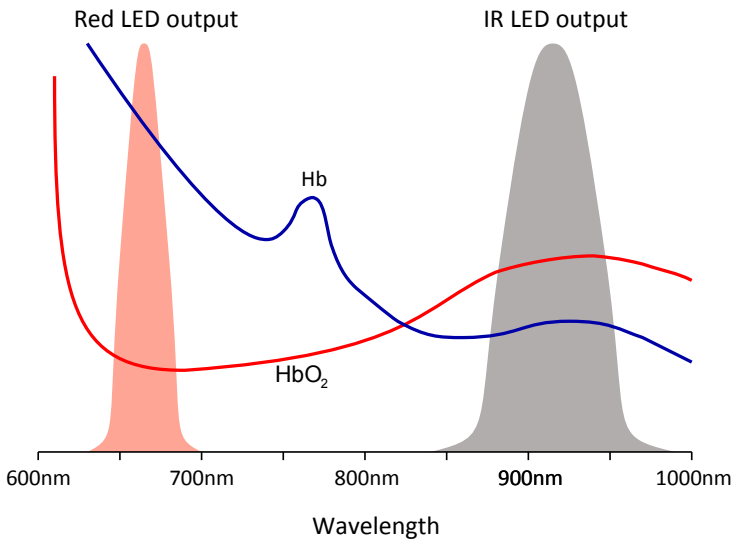


Figure 2.4: Absorption spectra for haemoglobin (Hb) and oxyhaemoglobin (HbO_2).
(Not to scale)

defined as the percentage of oxyhaemoglobin that makes up the total amount of haemoglobin in the blood.

$$S_pO_2 = 100\% \times \frac{HbO_2}{HbO_2 + Hb} \tag{2.4}$$

S_PO_2 can then be redefined in terms of the spectral absorption coefficients of haemoglobin and oxyhaemoglobin at the two chosen wavelength of LED and another term R; the "ratio of ratios" calculated directly from photocurrent measurements. Details of how the equations are derived can be found in and introduction to pulse oximetry by Aoyagi [27] and a summary of reflectance pulse oximetry by König et al [31].

$$S_pO_2 = \frac{\varepsilon_{Hb}(\lambda_{\text{Red}}) - \varepsilon_{Hb}(\lambda_{\text{IR}})R}{\varepsilon_{Hb}(\lambda_{\text{Red}}) - \varepsilon_{HbO_2}(\lambda_{\text{Red}}) + [\varepsilon_{HbO_2}(\lambda_{\text{IR}}) - \varepsilon_{Hb}(\lambda_{\text{IR}})]R} \tag{2.5}$$

R, which is derived from the Beer-Lambert model of the blood and tissue in the optical path, requires measuring the amplitude of the pulsatile (AC) component of the photocurrent at a given wavelength and dividing it by the static (DC) component of the photocurrent at the same wavelength. The ratio of the logarithms of these ratios defines R. Essentially this method is used to focus only on the spectral absorption profile of the blood (i.e. AC component)

and cancel out the contribution of any other substances (e.g. tissue and fat).

$$R = \frac{\log\left(\frac{\Delta I_P(\lambda_{\mathrm{IR}})}{I_P(\lambda_{\mathrm{IR}})}\right)}{\log\left(\frac{\Delta I_P(\lambda_{\mathrm{Red}})}{I_P(\lambda_{\mathrm{Red}})}\right)} = \frac{\varepsilon_{Hb}(\lambda_{\mathrm{IR}})c_{Hb} + \varepsilon_{HbO_2}(\lambda_{\mathrm{IR}})c_{HbO_2}}{\varepsilon_{Hb}(\lambda_{\mathrm{Red}})c_{Hb} + \varepsilon_{HbO_2}(\lambda_{\mathrm{Red}})c_{HbO_2}} \tag{2.6}$$

2.3 Transducer Construction and Location

The fundamental component of any sensor is the actual transducer that transforms energy from the physical or chemical event under observation into another form, typically electrical, so that it can be measured in a quantifiable manner. In the case of PPG the transducer consists of both an *actuator* and a *sensor*. The actuator is the LED light source that illuminates the capillary bed and the sensor is the photodetector which measures how variations in blood volume modulate the absorption and/or scattering of photons by the blood in the capillaries.

There are two main aspects to be considered in the implementation of a wearable sensor: where is the sensor to be located and how is it to be affixed to the subject? PPG provides a large amount of freedom in where the sensor can be placed as it will work, to varying degrees, wherever there is blood flow. The choice in placement of a PPG sensor is a balance between subject comfort, sensitivity to cardiac function and power consumption. Mendelson et al have studied the relationship between sensor size, location and power consumption in reflective PPG sensors [26].

Sensor location also brings up the issue of whether the sensor should be transmission mode or reflection mode:

Transmission mode: The LED shines light through an appendage (typically a finger or earlobe) and the photodetector measures the amount of light that is *not* absorbed as illustrated in figure 2.2.

Reflection mode: The LED is placed adjacent to the photodetector which measures the amount of light that is back-scattered by the capillaries immediately below the surface of the skin as illustrated in figure 2.5.

The majority of commercially available PPG system use transmission mode due to the fact that it is easier to constrain the volume of blood vessels being inspected. Reflection mode sensors are sensitive to transducer dimensions, application pressure and location.

However, a wearable PPG transducer has very different requirements to a transducer for use in a clinical setting. Foremost of the additional requirements are low power consumption and comfort when worn for extended period of

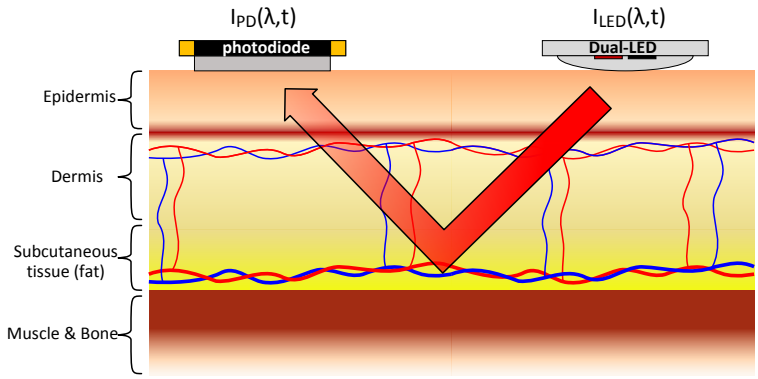


Figure 2.5: Illustration of reflection mode photoplethysmography.

Table 2.4: Comparison of Transmission (TX) and Reflection (RX) Mode PPG

PPG Mode	**TX**	**RX**	Comment
Power Consumption	High	**Low**	Shorter optical path for reflection mode.
Size/Convenience	Large	**Small**	Transmission needs to wrap around an appendage
Signal Reliability	**High**	Low	Constrained optical sensing volume.
Clinical Adoption	**High**	Low	Due to reliability.

time. Table 2.4 compares the two modes of operation with respect to a short-list of the more important characteristics. That table might suggest that reflection mode is the obvious choice for wearable PPG sensor systems, but both methods are being used in research as can be seen in table 2.6 at the end of this chapter.

2.3.1 Light Emitting Diode

The optical output power an LED is approximately proportional to the current flowing through the device at the low power end of the scale. This makes designing an LED driver the relatively trivial task of implementing a voltage controlled current source (VCCS). The main concern when designing an LED driver for a low power system is ensuring that it consumes as little power as possible and can operate with as little voltage headroom as possible to allow it to run directly from battery. For example, red LEDs, such as the 660nm ones in OSI-Optoelectronics DLED-660/905-LLS-3 device, can require as much as 2.4V voltage drop which would leave only 0.6V headroom from a fully charged 3V coin cell supply. Switching regulators can be used to boost the voltage available for powering the LED, but where possible added circuit complexity

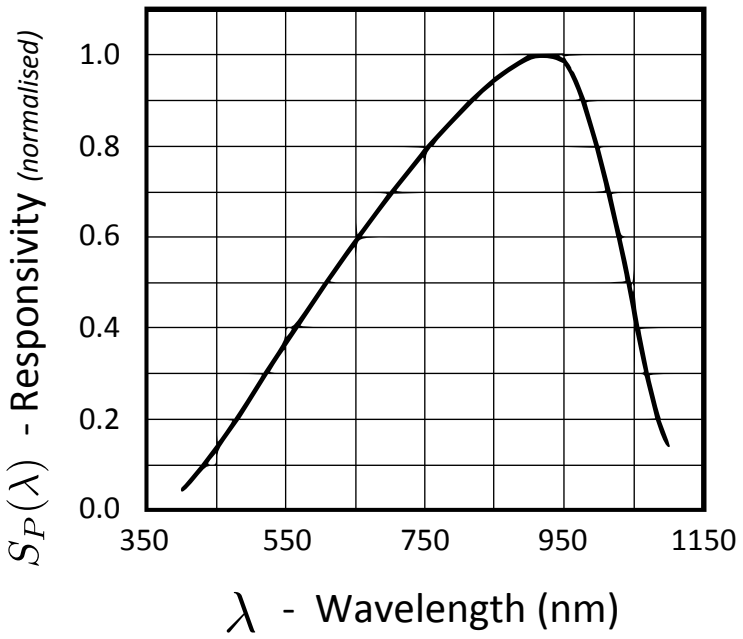


Figure 2.6: Responsivity of a Vishay TEMD5010X01 silicon photodiode.

should be avoided. Overall the LED driver is a low complexity circuit and will not be discussed in any further detail in this chapter.

2.3.2 Photodiode

Like the reverse of the LED, a photodiode outputs a current proportional the incident optical power. The amount of current that the photodiode outputs for a given incident power is dependent on the diode's responsivity (or quantum efficiency) for a given wavelength of light and the area of the diode. Figure 2.6 shows a typical responsivity curve of a silicon PIN photodiode - in this case a Vishay TEMD5010X01 device. They are a popular choice of detector as they have a wide enough responsivity curve to detect visible light as well as a peak responsivity in the near IR range which is particularly suited for pulse oximetry.

2.3.3 Time Division Multiplexing

In order to implement pulse oximetry a PPG system needs to be able to measure absorption at two wavelengths. Traditionally this would be been done by using wide spectrum illuminator and using two independent photodetectors, each optically tuned to have spectrally narrow responsivity. These broadband illuminators are not conducive to implementation in a wearable system due to their high power requirements. It is more common to use two LEDs, one for

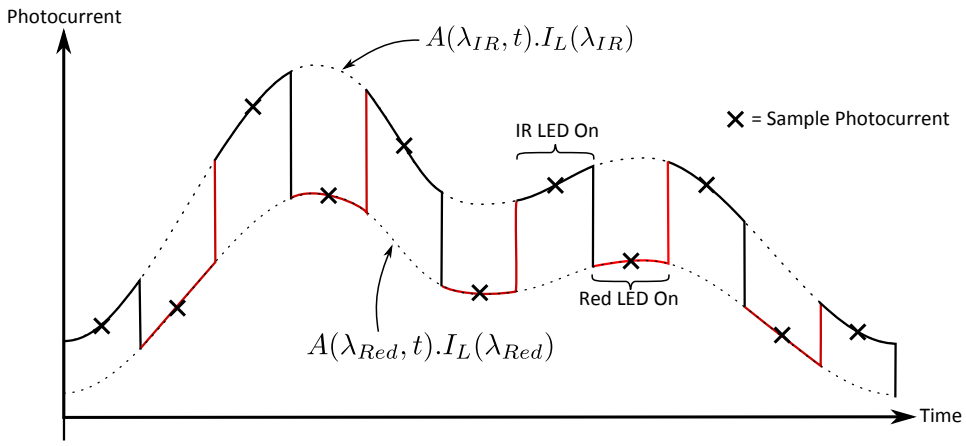


Figure 2.7: Illustration of optical time division multiplexing.

each of the desired wavelengths. Instead of having a photodetector matched to each LED wavelength it is more common to use *time division multiplexing* (TDM) to allow a single photodetector to detect both wavelengths. TDM requires that only one LED is turned on at any given time and then the output of the photodetector can be labelled with the wavelength of the LED as shown in figure 2.7. It is required that the rate at which the LED driver switches between LEDs is much higher that the highest frequency component of the absorption being detected such that no aliasing occurs. If more optical channels are required then more LEDs can be employed and the switching rate increased.

The trade-off is that both the LED drivers and photodetector have to support higher bandwidths than if two independent photodetectors were being used. In PPG, this is an acceptable sacrifice as the bandwidth of biosignals, such as the cardiac cycle, are very low compared to the achievable bandwidths of the most commonly available low power semiconductor amplifier components.

2.3.4 The Transimpedance Amplifier

Typically, the photocurrent output from the photodiode is converted into a voltage with a transimpedance amplifier circuit [32]. Unlike an LED driver, the design of a transimpedance amplifier is a much more involved process and is also the key component of any PPG instrumentation system. A conventional transimpedance amplifier (TIA) is constructed from a single operational amplifier (op-amp) with a resistor in the negative feedback path as shown in figure 2.8. This circuit performs two functions at once: stabilisation of the

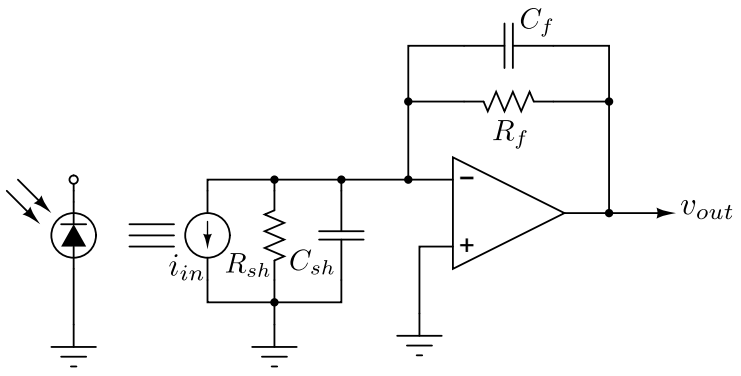


Figure 2.8: A standard transimpedance amplifier with the equivalent circuit model of a photodiode consisting of a current source, shunt resistor and shunt capacitor.

voltage across the photodiode and conversion of the input current to voltage.

The shunt resistance, R_{sh}, is generally large enough such that it can be approximated as open circuit which simplifies the transfer function. In the absence of R_{sh} the Laplace domain transfer function of the transimpedance amplifier is:

$$\frac{V_{out}(s)}{I_{in}(s)} = \left[s^2 \frac{C_f + C_{sh}}{2\pi f_0 A_0} + s\left(C_f + \frac{C_f + C_{sh}}{A_0} + \frac{1}{2\pi f_0 A_0 R_f} \right) + \frac{1}{R_f}\left(1 + \frac{1}{A_0} \right) \right]^{-1} \tag{2.7}$$

Note that A_0 is the open loop gain of the op-amp and f_0 is its open-loop cut-off frequency which together set the gain-bandwidth product. When the op-amp gain is very large the transimpedance amplifier gain is approximately proportional to the feedback resistance R_f.

$$A_P \approx \frac{y_P(t)}{i_P(t)} = \frac{v_{out}(t)}{i_{in}(t)} = R_f \tag{2.8}$$

Increasing the gain should be a case of increasing R_f which has the added benefit of reducing the input referred thermal noise current of the resistor ($\sqrt{4kT/R_f}$). However, what this equation omits is the dependence of the circuit on input shunt capacitance and resistance (C_{sh} and R_{sh} respectively), and the gain and bandwidth limitations of a non-ideal op-amp (defined by the gain-bandwidth product, GBW). Increasing the area of a photodiode increases its sensitivity, but also increases its capacitance which makes it harder for the TIA to stabilise the voltage across the photodiode. Another way of thinking about it is that the shunt capacitance forms an RC low pass filter with the feedback resistor R_f which limits the speed at which the op-amp output can

stabilise the input voltage.

To boost the bandwidth of the feedback loop a capacitor, C_f, is put in parallel with R_f. The choice of C_f depends on all the other parameters of the circuit and details of the relationships can be found in both [33] and [34]. Using the equation from [34] to provide critical damping C_f is given by:

$$C_f = \sqrt{\frac{2C_{sh}}{\pi f_0 A_0 R_f}} \tag{2.9}$$

which places a double pole at frequency $1/(\pi C_f R_f)$. Substituting in the choice of C_f locates the double pole at:

$$f_{Cf} = 1.3\sqrt{\frac{f_0 A_0}{2\pi C_{sh} R_f}} \tag{2.10}$$

If C_f is always chosen dependent on the value of R_f then there is a clear relationship between increasing the gain of the TIA at the expense of bandwidth.

Performance Metrics

Outlined below are some performance metrics which can be used to aid the analysis of a given photodiode amplifier solution and table 2.5 shows how altering one design variable can alter the performance metric of the circuit.

Gain: How much the photodetector amplifies the input current when converting it to a voltage. The gain defines the sensitivity of the circuit to both light and noise.

Noise: The main noise metric is the input referred noise current that the electronics contribute to the sensing system. Referring to the input allows a direct comparison with the input photocurrent level. Another source of noise is instability in the photodetector circuit.

Bandwidth: A measure of the frequency range at which the photodetector can sustain the desired signal to noise ratio.

Dynamic Range: The ratio of the maximum signal amplitude that the system can measure to the input referred noise current.

Quiescent Current: Power consumption is a function of the quiescent current of the photodetector and will generally be determined by the desired GBW and slew rate of the op-amp(s) used to create the circuit.

Table 2.5: Transimpedance Amplifier Design Trade-offs.

	Gain	Noise	Dynamic Range	Bandwidth	Quiescent Current
$R_f \Uparrow$	$\Uparrow$	$\Downarrow$	$\Downarrow$	$\Downarrow$	
Photodiode area $\Uparrow$	$\Uparrow$	$\Uparrow$	$\Downarrow$	$\Downarrow$	
$C_f \Uparrow$		$\Downarrow$		$\Downarrow$	
Op-Amp GBW $\Uparrow$				$\Uparrow$	$\Uparrow$

Integrating Amplifiers

One approach to circumventing the issue of input capacitance limiting TIA peformance is to use purely capacitive feedback which turns the CTIA into an input current integrator. This has been done before for optical detection systems [35, 36, 37, 38], but has had only limited application in PPG systems. There are even commercially available devices that rely on this method - the Texas Instruments IVC102 and DDC112 [39, 40]. The main complexity of this design is that as the current from the photodiode is always greater than or equal to zero the output of the integrator will be monotonically increasing. A reset switch allows the capacitor voltage to be reset to zero after a predetermined interval making the effective transimpedance gain dependent on both feedback capactior C_f and integration time T_s.

$$v_{out}(t) = \frac{1}{C_f}\int_0^{T_s} -i_{in}(\tau)\,d\tau \approx -\frac{T_s}{C_f} i_{in}(t) \tag{2.11}$$

This circuit arrangement means that the op-amp output only needs to track the integrated photocurrent which is much lower in bandwidth hence an op-amp with a much lower GBW and slew rate, and hence lower quiescent current, can be chosen to achieve a given photodetector bandwidth. It also allows for non-linear control of the LED driver current as the integrator effectively averages the received photocurrent over the integration period which even further simplifies the design of the LED driver. Figure 2.9 illustrates the circuit arrangement. The difficulty in designing an integrating TIA is the reset circuitry. Parasitic capacitance in the reset switch allows charge to be injected into C_f when the switch is released which adds an offset to the output voltage. De-

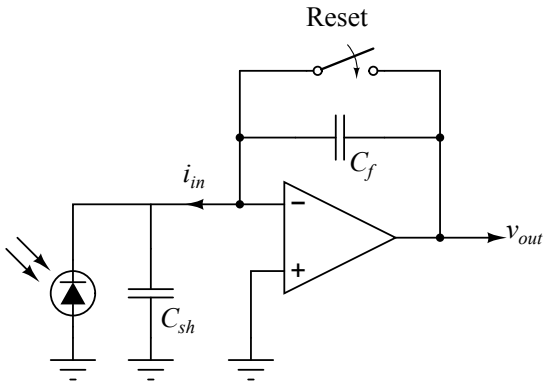


Figure 2.9: An integrating transimpedance amplifier with reset switch.

vices like the IVC102 will use low capacitance switches as part of the integrated circuit in order to keep charge injection to a minimum, but building an integrating TIA from discrete components does not have the luxury of nanoscale device geometries that integrated circuit implementations have. Further to the issue of charge injection, the op-amp needs to be able to discharge C_f to a stable state within a specified amount of time. The longer that reset time is, due to low op-amp slew rates and low stabilising bandwidth, the less time the circuit can spend integrating in a given sample period which reduces the sensitivity of the photodetector. However, if these design complexities can be overcome then an integrating TIA is an attractive photodetection solution for wearable PPG.

Logarithmic Amplifiers

A logarithmic amplifier uses the non-linear properties of semiconductor transistors to simultaneously amplify a signal and perform a logarithmic conversion on it. Figure 2.10 shows a classical implementation of a log amplifier. It is also possible to implement a log amplifier with purely CMOS transistors by exploiting the sub-threshold operating region. Equation 2.12 shows the logarithmic relationship between the input current and the output voltage where k is the Boltzmann constant, T is the absolute temperature of the bipolar transistor, q is the charge of an electron and i_{sat} is the reverse saturation current of the transistor.

$$v_{out}(t) = \frac{kT}{q} \ln\left(\frac{i_{in}(t)}{i_{sat}}\right) \tag{2.12}$$

There are two good reasons for performing a logarithmic conversion on the photocurrent produced by the photodiode. The first is that the relationship between the detected light and the variables of interest in a PPG system

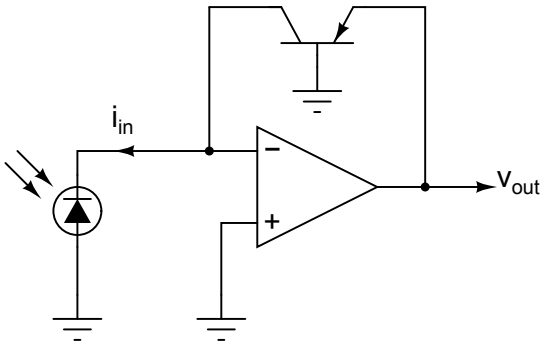


Figure 2.10: A logarithmic transimpedance amplifier implemented with an op-amp and a bipolar transistor.

can be modelled as exponential as proposed in equation 2.1. Calculating the logarithm allows a more direct calculation of R to be used in the measurement of S_PO_2 (equation 2.5). Applying the PPG system model to a log amplifier creates the following term:

$$v_{out}(t) = \frac{kT}{q} \ln \left(\frac{S_L(\lambda) I_L(\lambda, t) \cdot e^{-A(\lambda,t)}}{i_{sat}} \right)$$

which can be separated out into its constituent components as follows:

$$\begin{aligned} v_{out}(\lambda, t) = & \frac{kT}{q} \ln \left(S_L(\lambda) I_L(\lambda, t) \right) \\ & + \frac{kT}{q} \ln \left(e^{-A(\lambda,t)} \right) \\ & - \frac{kT}{q} \ln \left(i_{sat} \right) \end{aligned} \tag{2.13}$$

If all static (or DC) components are removed from the output voltage then we are left with a signal that is dependent solely on the absorption properties of the blood and independent of photodiode responsivity and LED radiated power:

$$\begin{aligned} AC \left(v_{out}(t) \right) &= \frac{kT}{q} \ln \left(e^{-A(\lambda,t)} \right) \\ &= -\frac{kT}{q} \sum_{\forall k} \varepsilon_k(\lambda) c_k d_k(t) \end{aligned} \tag{2.14}$$

where k is an index of each dynamic component of the blood. This logarithmic separation terms in the photocurrent signal is applied in more detail in the future section on motion artefact based on the research of Hayes and Smith [41]. This technique has recently been used to create an ultra-low-power pulse

oximeter using purely analogue integrated circuitry [42].

A second purpose for a logarithmic transimpedance amplifier is that it compresses the photocurrent signal which allows for an extended dynamic range [43] by reducing the sensitivity of the detection system as the incident photometric intensity increases. The consequence of having high linear gain is that high light levels are likely to saturate the front-end amplification circuitry and conversely low linear gain is unlikely to provide enough signal to noise ratio for biosignal detection at low light levels. A logarithmic photodetection amplifier has high gain at low light levels where the biosignal is weakest in intensity and reduced gain at high light levels. The net effect is that the intrinsic noise level stays roughly constant across a wide range of signal intensities providing a relatively uniform signal to noise ratio.

2.4 Noise and Artefacts

Any electronic instrumentation system is subject to electrical noise and it is generally the responsibility of the circuit designer to ensure that the electrical noise of their circuitry is small enough to enable detection of the signal of interest. These electrical noise sources tend to be linearly additive, Gaussian in distribution, time invariant and statistically independent of each other. However, the designer of a photoplethysmography sensor system has some additional noise sources which are non-stationary in nature - just like the signals that the system is trying to measure. These noise sources cause opto-electronic events, or *artefacts*, which distort the measured waveform. In photoplethymography these artefacts can be divided into one of two classes: *ambient* or *motion* based.

2.4.1 Ambient Artefacts

Ambient light artefacts are induced by light sources other than the LEDs of the PPG sensor. In a clinical or laboratory environment this is likely to be due to indoors fluorescent/incandescent lighting. The other main source for ambient light is the Sun; either coming through a window or from the sensor being worn while the subject is out doors.

For the case of artifical light the sources are likely to be generated from the electrical mains supply which is nominally a 50Hz sine wave in the UK. Lightbulbs often have a non-linear relationship between input current and output light intensity so the observed artificial light will consist of harmonics of 50Hz - primarily 100Hz. While this specific light source can be considered as stationary in frequency it is far from the only artificial light. A common

activity of the subject will be watching television which not only has the non-stationary light intensity output of the video image, but modern LCD/LED television use pulse width modulation (PWM) of the backlight to control intensity which creates spectral peaks at a number of 100Hz harmonics.

While artificial light is generally created from periodic electrical signals natural light is much more static in nature so primarily exists at DC (well, one primary oscillation per day). Sun light also tends to have a much higher intensity than LED light sources so is has to be a significant consideration in the design of a pervasive PPG system.

The first line of defence against ambient artefacts is good optical screening which is part of the transducer design. With the push for thinner and more unobtrusive sensors there is less and less material available for absorbing ambient light. Additionally, optical screening is difficult to maintain on a moving subject without resorting to strapping on large amounts of dark material which is against the ethos of unobtrusive sensor design. A truly unobtrusive device cannot rely on optical shielding to stop ambient lighting from corrupting the signal. To show this effect a few seconds of data were captured from a forehead worn pulse oximeter transducer where infra-red and ambient light was collected using time division multiplexing as described before. In figure 2.11 a typical PPG waveform shape can be discerned from the IR data, but the ambient light is equivalent in intensity to that of the IR signal in both DC offset (due to natural light levels) and AC components (due to artificial lights).

Correlated Double Sampling

A popular approach to cancelling out ambient artefacts relies on the fact that the frequency content of these artefacts be made to be significantly lower than the modulation rate of the LED(s). When measuring S_PO_2 the two LEDs of different wavelengths must be time-division-multiplexed in order to allow a single photodetector to detect both emission wavelengths. This technique can be extended by adding an extra TDM phase when both LEDs are extinguished. If the period between TDM phases is short enough then the low frequency ambient signal will not have changed much from sample to sample. The difference between the LED on phase and the LEDs off phase should indicate the true PPG signal independently of any detected ambient light. This technique is known as *Correlated-Double-Sampling* (CDS) and is commonly used in instrumentation circuits to remove fixed pattern noise. Figure 2.12 illustrates how CDS would be implemented with a single wavelength PPG system. Designing a PPG system with CDS requires careful analysis of the expected frequency

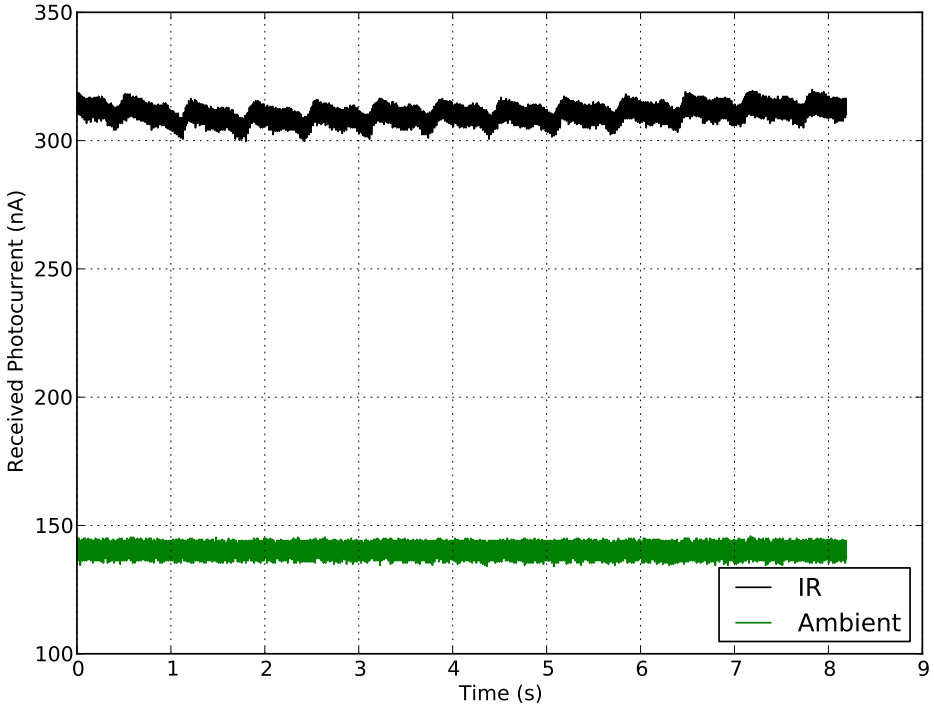


Figure 2.11: Infra-red and ambient light captured by a reflection mode transducer worn on the forehead showing noise created by artificial light at 50Hz and its harmonics. The noise is equivalent in magnitude to the PPG signal in both DC offset and AC components. The data shown here was captured from a reflection mode PPG transducer worn on the author's forehead.

content of the ambient light. The sampling rate must be set high enough such that the ambient signal does not change significantly between samples. The trade-off is that higher sampling rates require higher bandwidth LED drivers and photodetectors which inherently are more complicated and have higher current consumptions.

Using the same set of data as shown in figure 2.11 the efficacy of CDS is illustrated in figure 2.13. Note that while CDS has, as expected, removed the DC offset there is still some higher frequency noise present in the data although this can be removed easily using linear filtering.

Homodyne Detection

Any system that oscillates its emitter to create an amplitude modulated waveform and then synchronises the detection system to the same oscillator can be considered to be a *homodyne* detection system. By modulating the LED drive current the PPG signal is being up-mixed to a higher frequency band - above

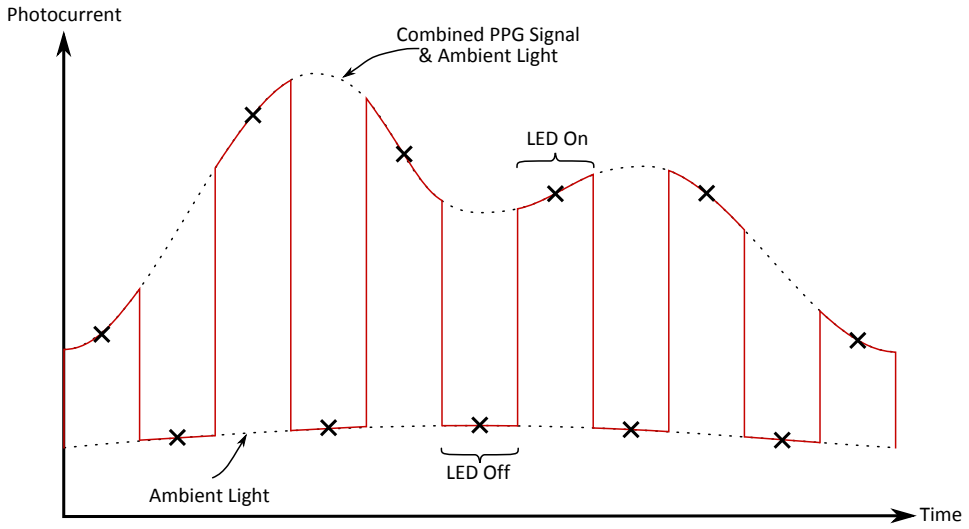


Figure 2.12: Correlated Double Sampling photocurrent waveform. The crosses indicate where samples are to be taken for subsequent subtraction to remove the influence of ambient light.

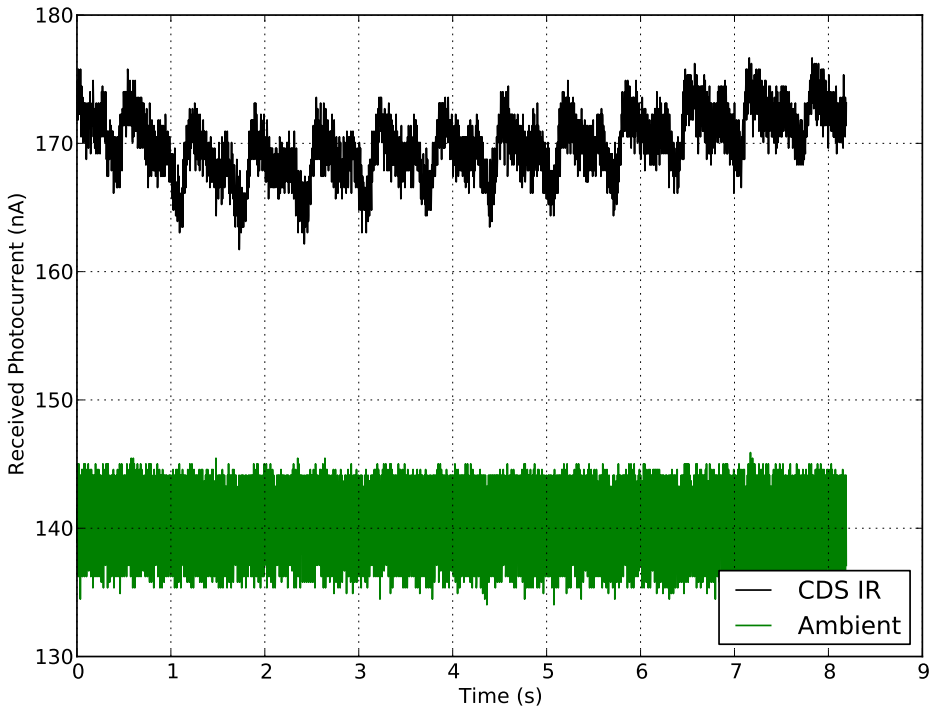


Figure 2.13: IR with ambient light (both natural and artificial) rejected using CDS. Note the residual high frequency noise. Ambient light data is included to illustrate relative intensity.

the influence on flicker noise that is prevalent in low frequency semiconductor circuits. Theoretically, by using a high enough LED modulation rate, low frequency ambient components will be discarded when down-mixing the received photocurrent to extract the PPG signal. Figure 2.14 shows how modulating the LED drive current creates a high frequency copy of the PPG waveform. Other terms that are used for this approach are *synchronous detection* and

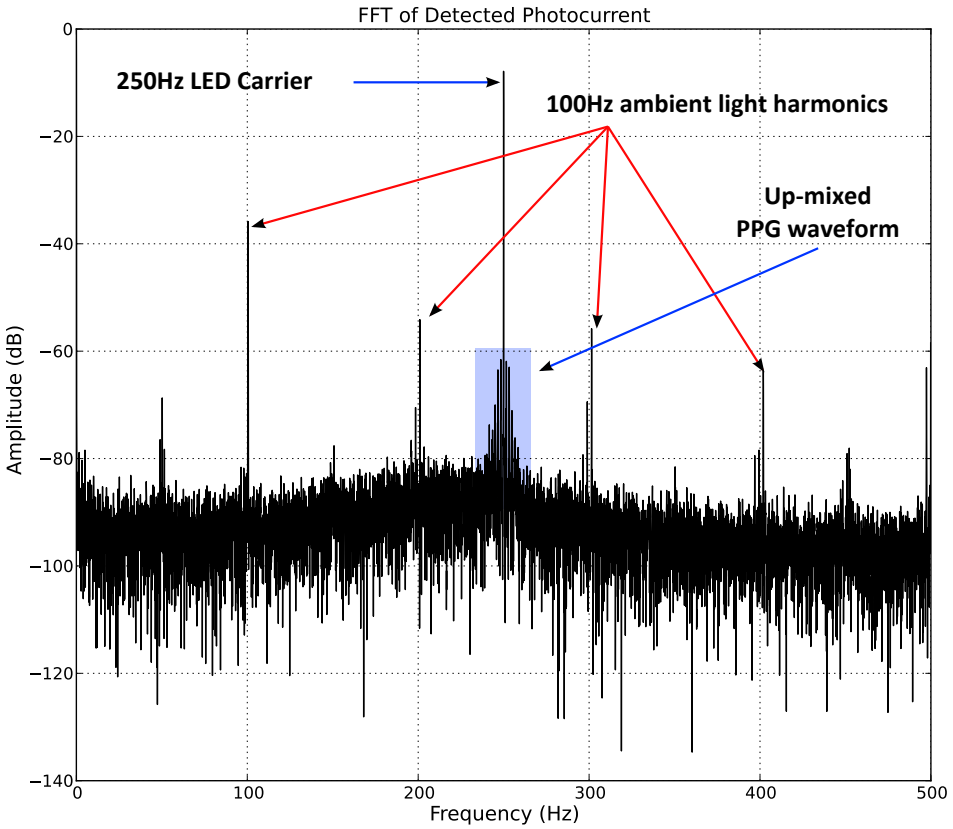


Figure 2.14: A DFT of a photodetector output with low frequency content removed. The data shown here was captured from a reflection mode PPG transducer worn at the tip of the author's index finger. *Adapted from BSN2012 submission.*

lock-in amplification and they have also been applied to PPG [44, 45, 46]. Even CDS can be considered as a specific form of homodyne detection. However, most PPG systems implement square-wave modulation of the LEDs and square-wave demodulation of the photocurrent which complicate the spectral response so it cannot be automatically assumed that homodyne detection will reject the ambient signals.

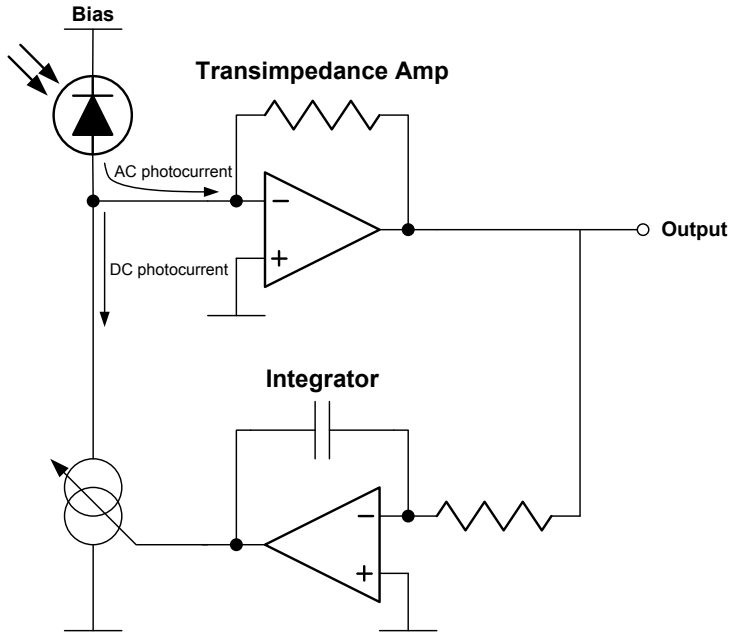


Figure 2.15: Basic DC photocurrent compensation control loop.

DC Current Compensation

A major shortcoming of CDS is that it can only reject ambient signals after the photodetector has converted the photocurrent into a voltage. While the ambient light is similar in magnitude to the PPG signal this is fine, but if the ambient light component is much larger then there is risk of the photodetector saturating. One approach is to reduce the gain of the transimpedance amplifier used in the photodetector circuit. This solution is not optimal as this compromises the signal-to-noise ratio (SNR) [32]. A better approach is to rely on the fact that the highest intensity light sources usually consist of low frequency components (e.g. sun light) and use a low frequency current feedback circuit to subtract the DC current directly from the photodiode so the transimpedance amplifier only ever detects AC signals. Figure 2.15 illustrates the basic structure of a DC current compensation circuit. A negative feedback loop is created by using an integrator to provide high gain at very low frequencies, to cancel out the high gain of the transimpedance amp, but low gain at higher frequencies allowing the forward path to function as normal. This current-mode negative feedback configuration extends the dynamic range of the photodetector by removing unwanted photocurrent directly from the photodiode allowing the gain of the transimpedance amplifier to be optimised for noise and bandwidth performance. This technique is not applicable to PPG systems where it is desirable to measure DC or very low bandwidth

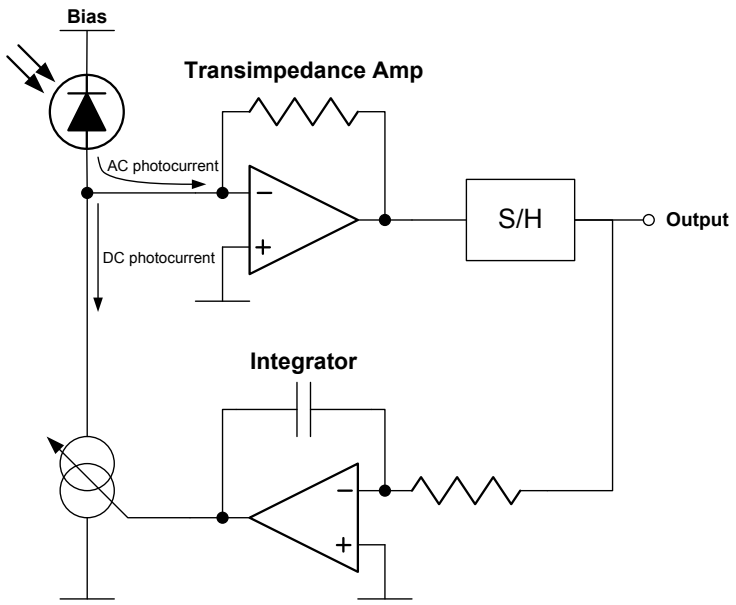


Figure 2.16: A sample and hold circuit is used to only feedback photocurrent reading when the LED is on. Adapted from figure 3 in [1].

signals such as S_PO_2, respiration and vasomodulation. Also, such high time constants (>1s) are not trivial to implement in analogue circuitry or require a high memory to sampling rate ratio in digital signal processing (DSP) and intrinsically take a long time to settle from ambient light level disturbances.

There has been some effort put into advancing DC photocurrent subtraction. Wong et al. [1] proposed an integrated circuit design that uses a sample and hold circuit to capture the photocurrent when the LED is on. By feeding back this value the dynamic range can be extended to include what would be the continuous value of the PPG photocurrent were the LED turned on continuously. This circuit performs best when the duration of the LED *ON* phase is short when compared to the period of the modulation rate - a common technique for reducing the average power consumption of the LED. However, as mentioned in the previous paragraph, this solution is not suitable for S_PO_2 measurements unless the DC value is recovered from the output of the feedback integrator which requires an additional analogue data channel.

Interference

Aside from ambient light sources a PPG system is prone to the same electrostatic and electromagnetic interference sources as any other sensor. Electrostatic interference is a particularly significant issue as the very high gain of the transimpedance amplifier makes the photodetector very sensitive to charge

injection at the transducer interface. Metal screening of the transimpedance circuit and wires to the transducer's photodiode is generally required to stop interference from dominating the output signal (i.e. creating a Faraday cage around the sensitive part of the circuit). Wireless device manufacturers have successfully miniaturised surface mounted "cans" for meeting electromagnetic compliance (EMC) regulations (less than 4 mm high), but wearable devices often need to be an order of magnitude smaller than typical mobile devices so screening of the circuit may not be possible. One solution to interference is to implement a pseudo-differential transimpedance amplifier to cancel out common-mode electromagnetic intereferece (EMI) [32], but this requires extra operational-amplifiers, good matching of components, does not allow for reverse biasing of the photodiode and over-complicates the possible introduction of DC photocurrent subtraction. A simpler solution to interference needs to be found.

2.4.2 Motion Artefacts

Motion artefact is any corruption of the photoplethysmogram due to motion of the subject. There are two main mechanisms for motion artefact: mechanical distortion of the optical path between the LED and photodiode of the PPG transducer, and distortion of the coupling between ambient light sources and the photodiode. The latter of these two mechanisms can also be classified as ambient artefact and in the scope of this thesis will be treated solely as an ambient noise problem. A dynamic optical path, which consists of tissue, bone and blood of varying thickness and density, is a vastly complicated system and makes motion artefact an incredibly difficult problem to solve. Many commercial solutions have been developed and analysed [47] with respect to their ability to detect and compensate for motion artefacts, but these are constrained to large devices that are only ever used in clinical settings with a limited range of transducers and sensor locations. By limiting the operational range of the sensor assumptions can be made about how patient motion will affect the plethysmogram.

Many techniques for solving the problem motion artefacts have been proposed including *independent component analysis* (ICA) [48] and *adaptive noise cancellation* (ANC) [49, 50, 51]. These techniques are general noise rejection methods in the sense that they can be applied to many sensing modalities; not just PPG. Specific to PPG and pulse oximetry are methods that are based on the spectral properties of blood. Introduced below are three different approaches to using optical spectrum information to cancel out motion artefacts.

Masimo is one of the leading commercial pulse oximeter manufacturers and their SET® technology uses a *Discrete Saturation Transform* (DST®) to isolate activity of blood in the veins from activity in the arteries [52]. The assumption is that venous blood is at a much lower pressure than arterial blood so its localised volume (and hence absorption properties) will be dominated by physical motion. In order to separate the motion based signal component from the cardiac based signal component the DST® creates a number of reference waveforms derived from the measured red and IR photocurrents ($I_P(\lambda_{\text{Red}})$ and $I_P(\lambda_{\text{IR}})$ respectively), and an assumed value for S_PO_2 (r).

$$N^{'}(r) = I_P(\lambda_{\text{Red}}) - I_P(\lambda_{\text{IR}}) \times r \tag{2.15}$$

For each value of r from 1 to 100 $N^{'}(r)$ is generated and used as a noise reference input to an ANC filter. The main signal input to the ANC is $I_P(\lambda_{\text{Red}})$. At the value of r which corresponds to the S_PO_2 of the venous blood, $N^{'}(r)$ will correspond to the cardiac component of the optical signal so the output of the ANC filter will be the motion component. Conversely, at the value of r which corresponds to the S_PO_2 of the arterial blood, $N^{'}(r)$ will correspond to the motion component of the optical signal so the output of the ANC filter will be the cardiac signal. In sweeping r it is at these two points where the output of the ANC filter is strongest. At all other values of r $N^{'}(r)$ will be quite similar to $I_P(\lambda_{\text{Red}})$ so the output of the ANC filter will be close to zero. By this process the motion component, cardiac component, arterial S_PO_2 and venous S_PO_2 have been determined. While Masimo's SET® technology has proven to be very effective at rejecting motion artefacts and detecting HR and S_PO_2 at low perfusion it is a very computationally intensive solution. Each time window of photocurrent signals requires 100 noise reference signals to be generated and 100 passes of the ANC filter. This is not conducive for implementation in a wearable sensor.

Hayes and Smith proposed an alternative method of reducing the effect of artefacts which used the photoplethysmogram from one wavelength of light to logarithmically equalise the signal at another wavelength [41]. First of all the wavelength dependent detected photon density $I_P(t, \lambda)$ is defined with respect to the transmitted photon density, the multiplicative component of a motion induced artefact $m(t)$, the dynamic physiological signal of interest $p(t)$ and dynamic changes in blood volume not due to pulse $b(t)$. The rest of the terms are optical coupling coefficients.

$$I_P(t, \lambda) = \sum_{j=1}^{n} I_j(t) \left(\alpha_j + \beta_j[1 + m(t)] \times \left[1 + b_j(t) + \frac{\gamma_j(\lambda)}{\beta_j(\lambda)} p_j(t) \right] \right) \tag{2.16}$$

The detected photon density is logarithmically converted to a voltage to give the following relationship:

$$\begin{aligned} v_j(t,\lambda) &= v_0 \ln\left(1 + I_P(t,\lambda)R(\lambda)\right) \\ &= v_0 \ln\left(I_j(t,\lambda)R_j(\lambda)\beta_j\right) + v_0 \ln\left(1 + m(t)\right) + v_0 \ln\left(1 + b_j(t) + \frac{\gamma_j(\lambda)}{\beta_j(\lambda)} p_j(t)\right) \end{aligned}$$

This shows how the logarithmic conversion separates out the multiplicative motion artefact term independently from wavelength and optical path. The claim that motion artefacts are multiplicative in nature is justified empirically from experimental evidence (see [41] for more details). If signals are collected from two different wavelengths, converted and subtracted then the following voltage is produced:

$$v_{jk}(t,\lambda) = v_0 \ln\left[\frac{I_j R_j \beta_j}{I_k R_k \beta_k}\right] + v_0 \ln\left[\frac{1 + b_j(t) + \frac{\gamma_j(\lambda)}{\beta_j(\lambda)} p_j(t)}{1 + b_k(t) + \frac{\gamma_k(\lambda)}{\beta_k(\lambda)} p_k(t)}\right] \tag{2.17}$$

If it is then assumed that the only changes in blood volume are physiological and that the contrast between the two wavelengths are large then the equation approximates at first order to:

$$v_{jk}(t,\lambda, b=0) \approx v_0 \ln\left[\frac{I_j R_j \beta_j}{I_k R_k \beta_k}\right] + v_0 \frac{\gamma_j(\lambda) p_j(t)}{\beta_k(\lambda)} \tag{2.18}$$

These equations all suggest that if the absorption of the optical path is significantly different at the two wavelengths then the non-linear equalisation technique will provide strong pulsatile signals in the presence of motion artefacts. Experimental data presented in the paper appears to support the theory. In a later paper the concept of non-linear equalisation is extended by using a third wavelength allowing two output signals representing two different wavelengths such that S_PO_2 can be measured while maintaining the motion artefact rejection performance of the two wavelength approach [53]. Compared to Masimo SET® the algorithm is much simpler to implement, but always requires one extra optical illumination wavelength in order to perform the equalisation and cancel out the motion artefact. Dual LED emitters are commonly produced for pulse oximetry purposes, but the 3 wavelength solution would require the construction of a custom transducer which will be larger and more obtrusive than the dual wavelength emitters.

More recently the concept of normalising the signal from one wavelength with another has been extended by performing it in a mixed time-frequency domain. Addsion and Watson performed a wavelet transform on each of a red

and infrared PPG signal and then divised the time-frequency surface map of the red signal by the infrared [54]. The result of the normalisation showed a region in the frequency domain where there was a strong correlation between the red and infrared channels which corresponded to the heart rate. This information could then be used to extract S_PO_2 from the motion corrupted signals. The wavelet transform is quite computationally intensive when performed digitally so not initially ideal for embedded implementation, but time-frequency analysis of the PPG signals could prove to be a powerful technique if a method of implementing it continuously in the analogue domain.

The choice of a motion artefact solution for a wearable PPG system involves finding a suitably accurate model for how motion couples in to the PPG signal for any given transducer construction and location, and then choosing an artefact detection and rejection technique that is suitably low in computational complexity such that it can be implemented in ultra-low power to fit with the requirements of a pervasive sensing system.

2.5 State-of-the-Art in Wearable Photoplethysmography Systems

A complete wearable PPG system will include an opto-electronic transducer for illuminating and measuring blood volume and one or more optical wavelengths, electronic instrumentation to amplify the opto-electronic signals from the transducer, signal processing to extract the desired biosignal from the raw photometric signal, and a communications interface to deliver the desired information to the observer.

Allen et al. studied the PPG signal quality at the ears, thumbs and toes [55] and Mendelson et al. have investigating measuring PPG signals at the forehead and on the wrist [26]. The choice of where to place the sensor has to be a trade-off between the quality of signal that can be obtained at the site and the likelihood that a subject is willing to adopt the technology in that manifestation. The latter requirement has pushed most developers of wearable PPG system towards reflection mode PPG although this does not in any way restrict the choice of where to place the sensor.

The sensor needs also to be packaged in a manifestation that enables the user to wear it and this is where there is significant diversity in proposed solutions. One of the first wearable devices developed by Sokwoo et al. [56] was integrated into a ring which takes advantage of the high quality PPG signals that can be detected in the finger. Branche et al. [57] built a reflection mode into a number of locations in a military helmet for healthcare monitoring in

the battlefield and found that the best oxygen saturation measurements were achieved at the forehead. Anliker et al. [58] integrate many different sensor technologies into their wrist-worn *AMON* device that resembles a watch. McCombie et al. [59] also use the wrist as a sensing location, but combine it with an extra finger PPG sensor in order to infer blood pressure by measuring the velocity of the blood pressure wave between the two locations. Vogel et al. [60] have developed in-ear PPG sensor which, like the ring design, used the mechanical transducer design to minimised the effect of motion on the PPG signal. Also in the ear location and lightweight transducer design for motion tolerance Poh et al. [61] have opted to place the sensor on the ear lobe using a magnet. Both Mendelson et al. [62] and Haahr et al. [63] propose patch based sensors which are flexible in their location, but are particularly suited to placement on the forehead where there the PPG signals are strong and most responsive to changes in oxygen saturation.

Table 2.6 lists out a subset of these wearable PPG systems and highlights some of their other features such as power consumption and communication method where available.

One of the most recently published system by Haahr et. al. uses reflection mode PPG in a patch platform which is one of the least intrusive sensor embodiments. Most of the novelty of that system is in the design of the opto-electronic transducer while the design's electronic instrumentation, signal processing and wireless communications receive less attention. To truly optimise a PPG system for wearable applications it is necessary to address all aspects of system. In particular this includes the specification of the radio system which only Asada et. al. considered in detail back in time when contemporary radio solutions such as ANT+ and Bluetooth Low Energy were not available - neither of which currently have widely publicised pulse-oximeter products associated with them.

2.6 Challenges

Although significant progress has been made in translating PPG from a clinical sensor technology into a wearable solution there are still some issues to be addressed before the technology is mature enough for wide scale adoption. Dhukaram [65] recently carried out a survey to gauge the interest levels of both patients and caregivers towards healthcare services based of pervasive cardiac monitoring. While the future of the technology was considered very positively there was little desire to adopt it in its current state. The two main factors that are barriers to adoption are trust in the system to treat

Table 2.6: Comparison of Wearable PPG Sensors

Solution	Type	PPG Mode	Sample rate	Wavelength	Radio	LED Current	Total Power	Motion Tolerance
MIT [64]	Ring	Transmission	100Hz	IR&Red	915MHz OOK	0.114mA	0.491mA	Mechanical isolation
ETHZ [58]	Watch	Reflection	n/a	IR&Red	GSM	n/a	25.4mA	n/a
RWTH [60]	Ear canal	Reflection	200Hz	IR&Red	Bluetooth	n/a	48mW	Mechanical isolation
MIT [61]	Ear lobe	Reflection	400Hz	IR	2.4GHz GFSK	n/a	n/a	Accelerometer based ANC
DTU [63]	Patch	Reflection	125Hz	IR&Red	Zigbee	2.8mA	18.2mA	n/a

the patient's information with due care and respect, and how difficult the system is to use. A good way of making the sensor device easy to use is to minimise the required interaction between the user and the device. This leads us toward *miniaturisation* and *long operating life*. But before a sensor can be miniaturised it still must be capable of providing valid data while operating out the in real world.

Any wearable sensor requires large scale user trials in order to validate that the system can reliably make a beneficial contribution to healthcare. Often solutions to individual problems within the field of PPG are tested on small cohorts (<100 subjects), but the results are of limited worth if they are not representative of the target solution that is to be delivered to the end-user. As such it is necessary to first address the following issues.

Use Cases: Making the sensor wearable makes it applicable to a wide range of monitoring situation (e.g. general lifestyle monitoring vs. athletic performance analytics) which often have divergent requirements (long operating life vs. high precision and accuracy data). This thesis will tend towards analysing the more long term lifestyle parameters that could aid in preventing cardiovascular disease, as was highlighted in first chapter.

Harsh Environments: While ambient light artefacts are well understood they are rarely explicitly addressed in wearable sensor solutions. A much less well defined problem is motion induced artefacts as there is no universally accepted model of how they couple in to the PPG signal - mainly due to the complexity of the coupling between the PPG transducer and the tissue that it is measuring. In fact there is the case where some motion artefacts are actually derived from ambient light sources being modulated by the subject motion.

Power Consumption: For BSNs in general energy can be conserved by reducing the amount of data that is wirelessly transmitted which means more data analysis is required on the sensor node in order to extract the information of interest. Of specific concern to PPG is the minimisation of the energy consumed by the LEDs. There is also the efficiency of the signal processing techniques used to remove artefacts.

Biosignal Extraction: The sensor needs to provide more than just heart rate and oxygen saturation. The list of desirable physiological metrics could be exceptionally long, but heart-rate variability and respiration rate are top of the priority list. Even if these measurements are not extracted

on the sensor node itself it should at least preserve enough signal information for them to be extracted further down the communications chain.

Miniaturisation: Reflection mode PPG requires only a small area [66] so is already well suited to wearable sensors, but there is also the associated circuitry and energy supply. The majority of wearable PPG system researchers are adopting the approach of embedding the entire system into a monolithic sensor device, but there is still a drive for further miniaturisation in all dimensions.

2.7 Conclusions

This chapter has discussed the contribution that BSN can make to improving cardiovascular well-being on a global scale and has championed PPG and pulse-oximetry as one possible sensing solution that can help deliver this goal (although not to the exclusion of other sensing modalities). The fundamental operation of each constituent part of a PPG system has been discussed and their integration into a wearable PPG system has been illustrated with a number of examples that have been produced by the academic community. This thesis aims to advance beyond the limitations of these systems by addressing the challenges mentioned. The first problem addressed is reducing both the power consumption and physical size of the analogue front-end circuitry required to capture PPG signals without having to rely on custom integrated circuit design.

3 Integrating Photocurrent Amplifier

The core component of any PPG system is its ability to detect light in order to infer absorption. As previously stated the best power efficiency is achieved by minimising the power output from the LED-based illumination. To maintain a consistent signal to noise ratio (SNR) requires maximising the sensitivity of the photodetector.

The original motivation behind using an integrating architecture for photocurrent amplification was to implement a circuit where the output at a discrete sampling event is proportional to the total number of photons incident to the photodiode since the previous discrete sampling event - regardless of when the photons arrived during that sampling interval. Subsequent research, as published in BSN 2009, 2010 and 2012 [6, 7, 9], as well as TBioCAS2010 [8], have used this circuit due to its benefits in power consumption, noise and bandwidth performance over the traditional passive transimpedance amplifier. This chapter will analyse the integrating transimpedance amplifier in more detail than in chapter 2 and will compare its sensitivity and noise performance to that of a conventional transimpedance amplifier. A set of performance requirements for a wearable PPG system photocurrent amplifier is presented. Three physical PCB iterations of an integrating transimpedance amplifier based PPG system are also presented.

3.1 Circuit Analysis

The desired output from a photodetector circuit may be a sampled signal, and hence naturally a discrete time (DT) domain signal, but the photocurrent amplifier is fundamentally a continuous-time (CT) circuit so the circuit analysis in this chapter is based on CT models. CT is also more commonly used for photodiode amplifier analysis [32] and so allows for more direct comparison between the integrator approach and the conventional method. As an ideal integrator conserves photoelectric energy collected between sampling events it is not necessary to resort to using zero order hold approximation in analysis.

Note that throughout this chapter all figures that present numerical relationships are generated from simulations performed using the NumPy library for Python with matplotlib for plotting the graphs.

3.1.1 Transfer Function

Figure 3.1 revisits the circuit of the integrating transimpedance amplifier. The analysis of the integrating transimpedance amplifier focuses on the effects of parameters such as sensitivity, bandwidth and current consumption. The practical implications of these design parameters are described here:

Photodiode Area: Larger photodiodes can collect more photons from a fixed luminous intensity as long as they are overfilled which is generally the case in PPG systems. This increased optical sensitivity comes at the cost of increases parasitic shunt capacitance, C_{sh}, which can affect circuit dynamics. An example device, which has been used throughout this thesis, is the Vishay TEMD5010X01 which has a typical junction capacitance of 70pF (at 0V bias).

Transimpedance Gain: In conjunction with the photodiode responsivity the amplifier gain sets how sensitive the circuit is to input current. For every amplifier configuration there is a trade-off decision to be made between high sensitivity, high bandwidth, and low noise.

Amplifier Quiescent Current Consumption: The op-amp is the only active component in the photocurrent amplifier front end circuit so defines the quiescent current consumptions of the sensor interface. In CMOS op-amp design improved offset, noise, bandwidth and slew rate can all be achieved at the expense of extra bias current [67]. For the purposes of simplicity this chapter focusses on the gain-bandwidth product.

The Laplace domain will be used throughout this chapter for CT analysis of dynamic circuits. The op-amp is modelled as an ideal differential amplifier with finite DC gain of A_0 and a single pole low-pass frequency response at f_0.

$$V_{out}(s) = (V_+(s) - V_-(s))\, A_0 \left[1 + \frac{s}{2\pi f_0}\right]^{-1} \tag{3.1}$$

The output voltage limitation of a real op-amp due to finite supply voltage is ignored for the purposes of frequency analysis. The transfer function of the

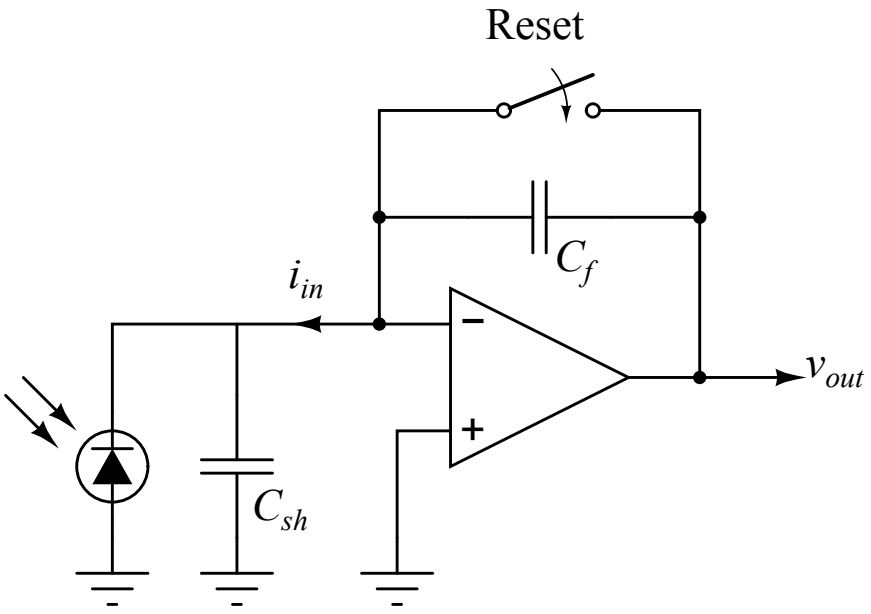


Figure 3.1: Integrating transimpedance amplifier for photocurrent amplification. Reset switch and parasitic shunt capacitances are included.

integrator is presented in equation 3.2.

$$\frac{V_{out}(s)}{I_{in}(s)} = \left[\left(\frac{C_f + C_{sh}}{2\pi f_0 A_0}\right)s^2 + \left(C_f + \frac{C_f + C_{sh}}{A_0} + \frac{1}{2\pi f_0 A_0 R_{sh}}\right)s + \frac{1}{A_0 R_{sh}}\right]^{-1} \tag{3.2}$$

The first approximation we make is that the shunt resistance, R_{sh}, is very large and can be approximated to open-circuit which reduces the system response to:

$$\frac{V_{out}(s)}{I_{in}(s)} = \left[s\left(\frac{C_f + C_{sh}}{2\pi f_0 A_0}s + \left(C_f + \frac{C_f + C_{sh}}{A_0}\right)\right)\right]^{-1} \tag{3.3}$$

To get the the ideal integrator transfer function of s^{-1} the input signal frequency range needs to be below the second system pole which gives the following limit on integration bandwidth.

$$f_P < \frac{f_0 A_0}{1 + C_{sh}/C_f} \tag{3.4}$$

A further simplification can be made in the case that $C_f \gg C_{sh}/A_0$ where the amplifier gain and photodiode shunt capacitance can be ignored giving the desired integrator response:

$$V_{out}(s) \approx \frac{I_{in}(s)}{sC_f} \;\Rightarrow\; v_{out}(t) \approx \frac{1}{C_f}\int_0^t i_{in}(\tau)\,d\tau \tag{3.5}$$

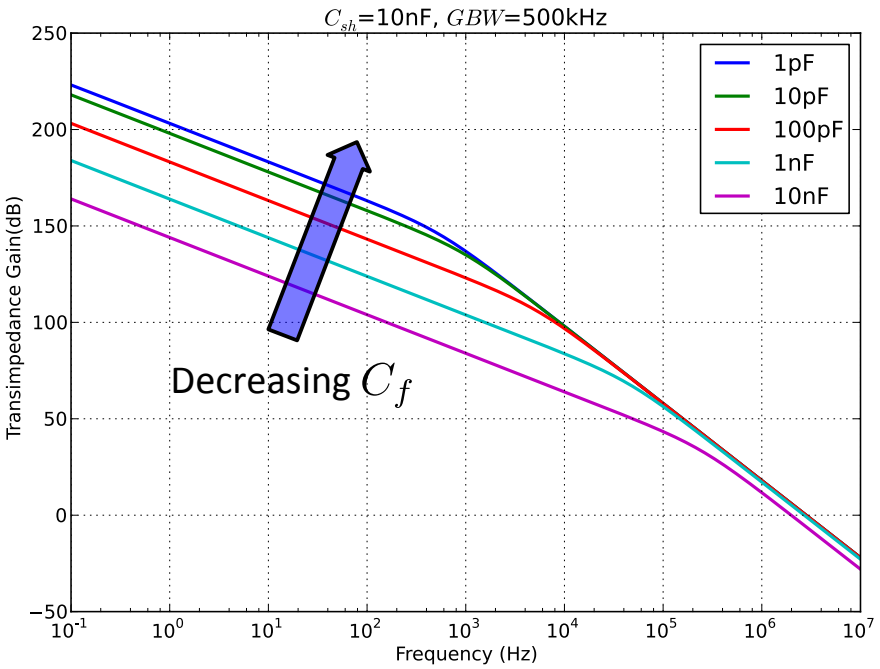


Figure 3.2: Amplifier response for a range of C_f with C_{sh} is fixed at 10nF and GBW is 500kHz.

Given the response in equation 3.3 it is possible to analyse the photocurrent amplifier performance for ranges of C_f, C_{sh} and GBW (which is equal to $f_0 A_0$). Figures 3.2, 3.3 and 3.4 illustrate the trade-offs for varying C_{sh} and GBW as these two parameters are fixed by available components, but C_f can be varied to the desire of the designer. In all cases the load capacitance on the output of the amplifier has been assumed to be negligible which can be assured by using a voltage follower to ensure any capacitive loads from subsequent circuitry do not affect the integrator.

As per equation 3.5 figure 3.2 shows that decreasing the value of C_f increases the proportional gain of the amplifier. Also shown is the reduction in bandwidth of the pure integration region as C_f decreases its the combination with the parasitic shunt capacitance C_{sh} and finite op-amp bandwidth begins to dominate the circuit gain. Figure 3.3 shows how large parasitic capacitance can limit the bandwidth of the integrator. There is also visible a minor effect of the overall proportional gain of the circuit which only occurs when C_{sh} is large enough that R_{sh} begins to influence the system's performance. Unsurprisingly, figure 3.4 shows that the gain bandwidth product of the op-amp used to implement the integrator circuit has a direct impact on its integration bandwidth.

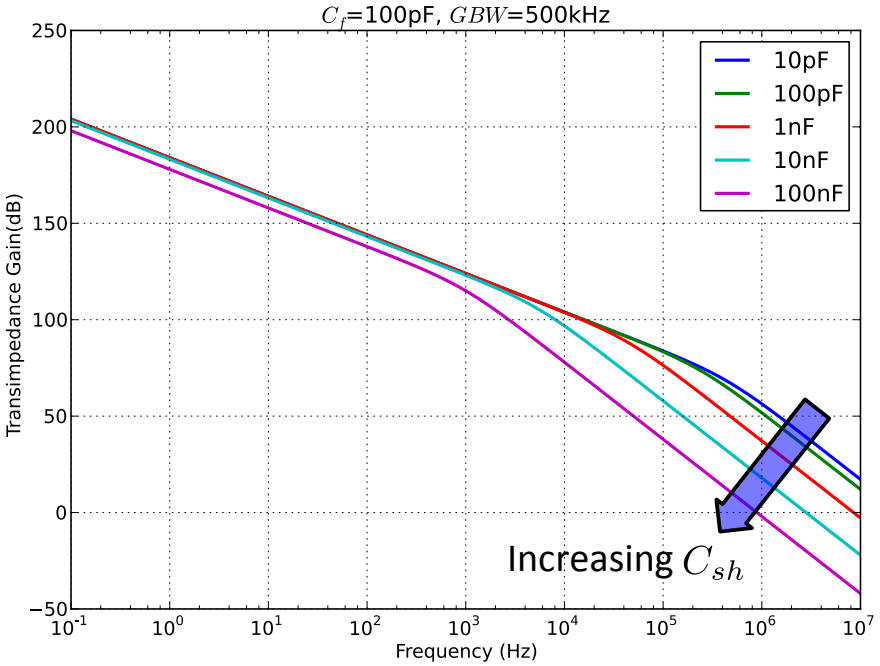


Figure 3.3: Amplifier response for a range of C_{sh} with C_f is fixed at 100pF and GBW is 500kHz.

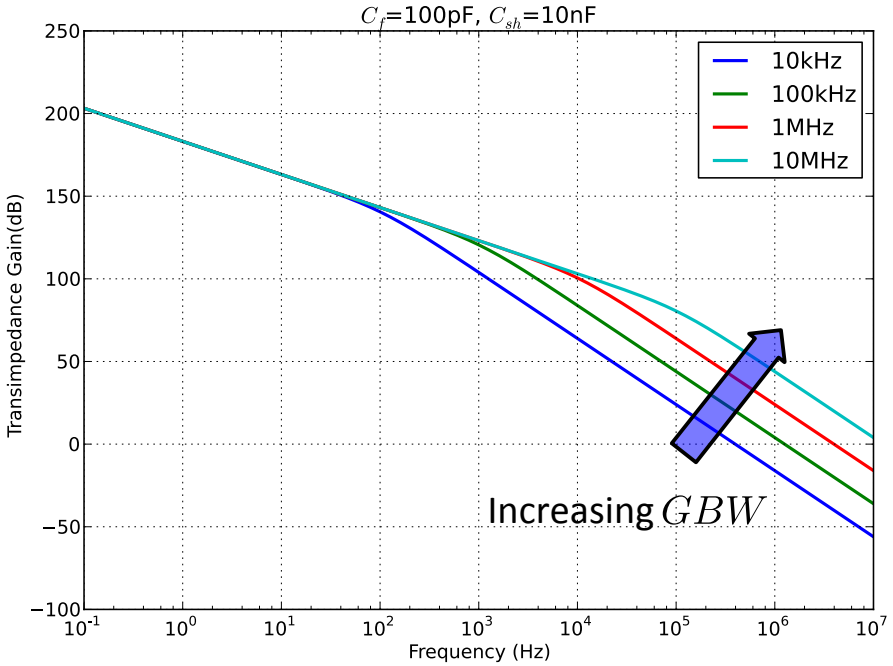


Figure 3.4: Amplifier response for a range of GBW. C_f is fixed at 100pF and C_{sh} is fixed at 10nF.

3.1.2 Amplifier Gain

To measure the amount of light received in a given period requires the calculation of the definite integral. For a digital sampling system the interval can be specified as the sampling period T_s. Each sample is given an index n and occurs at time $t = nT_s$. The output of the amplifier is redefined as:

$$\begin{aligned} v_{out}[n] &\equiv \int_{(n-1)T_s}^{nT_s} \frac{i_{in}(t)}{C_f}\, dt \\ &= \int_{0}^{nT_s} \frac{i_{in}(t)}{C_f}\, dt - \int_{0}^{(n-1)T_s} \frac{i_{in}(t)}{C_f}\, dt \\ &= v_{out}(nT_s) - v_{out}((n-1)T_s) \end{aligned} \tag{3.6}$$

An intuitive way of interpreting this is that calculating the difference between two consecutive samples is equivalent to differentiating the integrated signal in the discrete-time (DT) domain which recovers the original photocurrent signal. The DC gain of the system is now proportional to the ratio of sampling period to C_f which confirms the response stated in chapter 2:

$$A_P = \frac{v_{out}[n]}{i_{DC}} = \frac{1}{i_{DC}} \int_{(n-1)T_s}^{nT_s} \frac{i_{DC}}{C_f}\, dt = \frac{nT_s - (n-1)T_s}{C_f} = \frac{T_s}{C_f} \tag{3.7}$$

Together with equation 3.4 it is possible to determine the required gain-bandwidth for the op-amp from the desired gain A_P, bandwidth f_P, sampling period T_s, and the photodiode's parasitic capacitance C_{sh}:

$$GBW = f_0 A_0 = f_P \left(1 + \frac{A_P C_{sh}}{T_s}\right) \Rightarrow f_P \left(1 + \frac{C_{sh}}{C_f}\right) \tag{3.8}$$

The detected photocurrent is always greater than or equal to zero so the output of a true integrator will monotonically increase. A consequence of this is that it is necessary to take steps to ensure that the amplifier output does not saturate due to the finite supply voltage.

Resettable Integrator

The most common solution is to periodically reset the integrator voltage with a switch to dissipate the charge on C_f at the start of each sampling interval as shown in figure 2.9 in chapter 2. To minimise charge injection a purpose made analogue switch is used. An example device is the Texas Instruments TS5A3159 which has only 36pC of charge injection with an ON resistance of 1Ω. The amount of time that needs to be allocated to resetting C_f depends on its size, the load capacitance, the switch on resistance, op-amp slew rate and

op-amp phase margin at unity gain. This relationship is complex and won't be investigated in detail in this thesis. The most accurate method of determining the minimum requirement for the reset time, T_{rst}, is by experimental investigation. For a given sampling period this has the effect of reducing the gain:

$$A_P^* = A_P \left(1 - \frac{T_{rst}}{T_s}\right) \tag{3.9}$$

This now forces a trade-off between achieving a high gain and ensuring enough time has been allowed for the voltage across C_f to settle such that it does not induce significant noise in the current measurement.

Biased Integrator

An alternative to resetting the charge on C_f is to subtract a current from the input node which causes the current flowing through C_f to have zero mean. The basic DC photocurrent compensation control loop does just this as illustrated in figure 2.15 in chapter 2.

$$v_{out}(t) \equiv \frac{1}{C_f} \int_0^{nT_s} \left(i_{in}(t) - \overline{i_{in}(t)}\right) dt \tag{3.10}$$

In the analogue domain the mean calculation is performed by another integrator with a Laplace transfer function of A_i/s. This essentially adds a high-pass filter to the photocurrent amplifier which removes DC signal components in order to achieve an overall zero mean current through C_f. The Laplace transfer function is given by:

$$\frac{V_{out}(s)}{I_{out}(s)} = \frac{s}{C_f \left(s + \sqrt{\frac{A_i}{C_f}}\right)^2} \tag{3.11}$$

In the case that $s \gg \sqrt{A_i/C_f}$ then the response approximates to $1/sC_f$, but the zero at the origin means that the output of the integrator has zero mean. A_i/C_f needs to be set low enough to preserve as much low frequency information in the photocurrent signal as required as previously discussed in chapter 2. Higher order filters could be used in the feedback path to increase the width of the integration band instead of an integrator, but that complicates the design of the circuit to ensure stability.

An alternative to using an integrator in the feedback path is to use non-reactive feedback, i.e. a resistor, as shown in figure 3.5 which is actually a conventional transimpedance amplifier except that the feedback capacitance is the dominant element in determining dynamic system performance. The

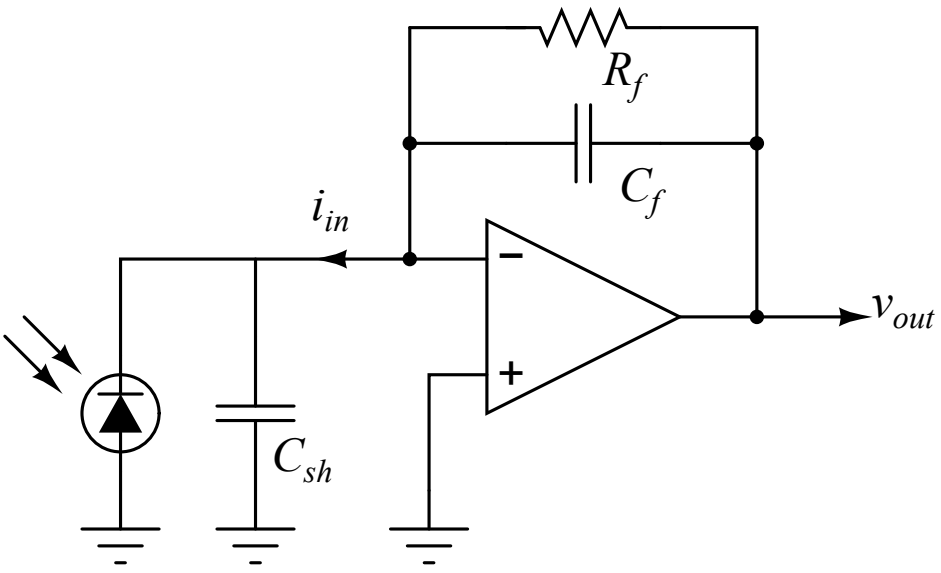


Figure 3.5: ITIA with passive feedback.

transfer function of the circuit is now no longer a true integrator, but first order low pass filter:

$$\frac{V_{out}(s)}{I_{out}(s)} = \frac{1}{sC_f + 1/R_f} \tag{3.12}$$

Below the pole frequency of $(2\pi C_f R_f)^{-1}$ the amplifier has a gain of R_f which is no longer proportional to the sampling period. For a desired integrating photocurrent gain of $1/sC_f$ that pole frequency needs to be below the lowest frequency of interest in the input photocurrent which means making the feedback resistor R_f as large as possible which increases the DC gain and actually increases the risk of saturating the amplifier.

3.1.3 LED Driver Implications

For the integrating photocurrent amplifier it is not important to preserve the shape of the incident optical waveform between sampling events - it is only necessary to preserve the energy contained in the waveform. In low power PPG systems it is common to pulse-width modulate the LED drive current so the photocurrent amplifier needs enough bandwidth to preserve enough of the signal energy contained within the current pulse to achieve the required signal to noise ratio. The energy in a rectangular pulse is unbounded in the frequency domain and for a unity amplitude and width of T the contained

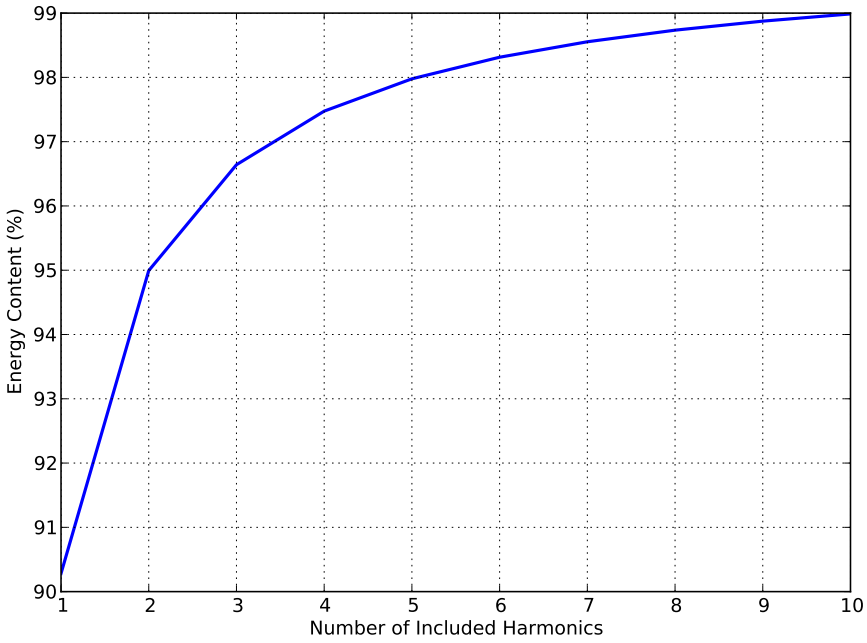


Figure 3.6: Energy content in cumulative harmonics of pulse width.

energy is described by:

$$E = \int_{-\infty}^{+\infty} T^2 \text{sinc}^2(\pi f T)\, df \tag{3.13}$$

The integrating transimpedance amplifier is constrained in frequency (equation 3.4) so it is not possible to capture all of the pulse's energy. Eduardo Aguilar Pelaez has analysed how many harmonics of the fundamental frequency of a pulse train are required to retain 95% of the signal power [68], but that analysis focused on the optimum choice of a fixed pulse width to maximise signal to noise ratio in a pulse oximeter system. In the following chapter the LED drive strength will be modulated in order to improve signal to noise ratio so we need to find the minimum photocurrent amplifier bandwidth required to detect a given percentage of energy in a single LED pulse. An alternative perspective is to consider the percentage of the total energy contained in a given number of harmonics of the pulse width as shown in figure 3.6.

For example, to achieve a 1kHz sampling rate for a two channel pulse oximeter system requires a 500μs sample period. If the duty cycle of the LED were to be for 50% of the sampling period then to preserve at least 95% of the energy requires the first two harmonics of the pulse width which requires a bandwidth of 4kHz. Dropping the duty cycle to 5μs pushes the required integrator bandwidth up to 400kHz. The practical limitation of how small the

pulse width can be set is by the slew rate and settling time of the LED driver design. However, the maximum bandwidth of the integrator for a given current consumption indirectly sets a minimum limit on the LED pulse width. Alternatively, the LED drive current can be modulated in a more linear fashion which would significantly reduce the bandwidth requirements, but the required circuitry is more complicated and less energy efficient.

3.1.4 Noise

Conventional Transimpedance Amplifier

There are four main noise sources in a conventional transimpedance amplifier (CTIA). All noise sources are referred to in terms of noise power spectral density (W/Hz). In all subsequent equations k is Boltzmann's constant and T is the temperature of the circuit in Kelvin.

Feedback Thermal Noise: Johnson-Nyquist noise created by the feedback resistor. When input-referred the Thevenin equivalent circuit is a current source and is inversely proportional to the feedback resistance R_f.

$$i_R^2 = 4kT/R_f$$

Op-Amp Voltage Noise: A combination of input-referred thermal noise and flicker noise from the op-amp which gives it a $1/f$ profile at low frequency.

$$e_n^2 = \text{Op-amp dependent}$$

Op-Amp Current Noise: Input-referred noise current of the op-amp. Generally induced by shot noise.

$$i_n^2 = \text{Op-amp dependent}$$

There are other noise sources in a photodiode amplifier, but these are due to shot noise from either the photodiode dark current or the generated photocurrent. As these are predominantly independent from the amplifier design they will be ignored in this section. The following equation combines the aforementioned noise sources and translates them to an output referred noise voltage using the amplifier gain, A_P. Note that flicker noise will not be included in subsequent noise models. Flicker noise is a significant noise contributor at low frequency, but is not trivial to model accurately and often needs to be measured empirically. Furthermore, future developments in this thesis will focus

on operating outside of the low frequency region in order to avoid the numerous artefacts that can occur there. For more details regarding noise analysis refer to Jerald Graeme's *Photodiode Amplifiers: Op Amp Solutions*[32]. Equation 3.14 gives the output referred noise voltage in a CTIA.

$$e_{ctia}(s) = \sqrt{\left(i_R^2 + i_n^2\right) A_P^2(s) + \left(\frac{e_n}{\frac{1}{A_0} + \frac{1+sC_fR_f}{1+s(C_{sh}+C_f)R_f}} \frac{1}{1+s\frac{A_0}{GBW}} \right)^2} \tag{3.14}$$

Substituting equation 2.8 in for A_P, along with the available noise definitions, gives the following simplified noise model for low frequencies:

$$e_{ctia}(s) = \sqrt{4kTR_f + (i_nR_f)^2 + e_n^2 \left(\frac{1+s(C_{sh}+C_f)R_f}{1+sC_fR_f} \right)} \tag{3.15}$$

The third term shows the relationship between parasitic shunt capacitance, transimpedance gain and the stabilisation capacitance. The purpose of including this detail is to show how the choice of C_f for stabilisation also affects the noise spectral profile of the amplifier.

Integrating Transimpedance Amplifier

The integrating transimpedance amplifier (ITIA) can be considered a reduced version of the switch-capacitor integrator - of which there have been numerous noise analysis studies [69, 70, 71, 72, 73]. These studies tend to be quite specific to the SC-integrator which is a voltage input circuit as opposed to the current-mode input of a transimpedance amplifier. This means the studies focus mainly on the discrete time nature of switched capacitor circuits whereas this chapter is focussing on a continuous time model to allow direct comparison with the conventional transimpedance amplifier. This does not mean that the cited studies do not apply to the proposed circuit - only that the level of detail required for the discrete-time analysis is unnecessary for analysis of the system model in continuous-time.

The noise model for the integrating transimpedance amplifier differs from the conventional transimpedance amplifier in that there is no thermal noise induced in the feedback path as there is no R_f and the spectral profile is translated by the integrating nature of the amplifier's transfer function. In the case that a switch is used to reset the feedback capacitor there is some noise induced by charge transfer during the action of switching defined as:

$$i_C^2 = kTC_f$$

If a reset switch is not employed then closed loop biasing is required as discussed previously. The current feedback will have its own noise contribution labelled as i_f when referred to the input of the amplifier. The value of i_f is determined by the implementation of the feedback loop. Both the switching noise and the bias control noise are included in the general noise model of the integrating transimpedance amplifier.

$$e_{itia}(s) = \sqrt{\left(i_C^2 + i_f^2 + i_n^2\right) A_P^2(s) + e_n^2 \left(\frac{C_{sh} + C_f}{C_f}\right)^2} \tag{3.16}$$

Substituting equation 3.5 in for A_P, along with the available noise definitions, gives the following noise model:

$$e_{itia}(s) = \sqrt{\frac{kTC_f + i_f^2 + i_n^2}{s^2 C_f^2} + e_n^2 \left(\frac{C_{sh} + C_f}{C_f}\right)^2} \tag{3.17}$$

If the output is differentiated then equations 3.7 and 3.9 can be substituted in for A_P which gives the following simplified noise model:

$$e_{itia}(s) = (T_s - T_{rst}) \sqrt{\frac{kTC_f + i_f^2 + i_n^2}{C_f^2} + e_n^2 \left(\frac{C_{sh} + C_f}{C_f}\right)^2 s^2} \tag{3.18}$$

To analyse the trade-offs between design variables and the resulting noise spectrum the amplifier is modelled with a differentiated output and arbitrarily chosen op-amp noise parameters of $28\text{nV}/\sqrt{Hz}$ and $0.6\text{fA}/\sqrt{Hz}$. Figure 3.7 shows the effect of varying the value of feedback capacitor C_f. As might be expected, increasing C_f to reduce the overall gain has the effect of decreasing the overall noise. At the highest value of C_f (10nF) there is a dominant low frequency noise component that is not affected by the change in gain which is due to the noise voltage of the op-amp.

Figure 3.8 shows how the size of the parasitic shunt capacitance mainly affects the high frequency range of the noise spectrum where increased C_{sh} reduces the overall bandwidth and hence reduces the output noise voltage. The effect of C_{sh} at low frequencies is the opposite due to the combination of C_{sh} and C_f acting as a non-inverting amplifier for the op-amp's input-referred noise voltage. Note that there is a frequency band in the middle of the spectrum where C_{sh} has no effect on the output referred noise voltage - here the op-amp's noise current, i_n, is the dominant noise source. This region will be taken advantage of later in subsequent chapters of the thesis.

Varying the op-amp gain bandwidth product has a much more significant

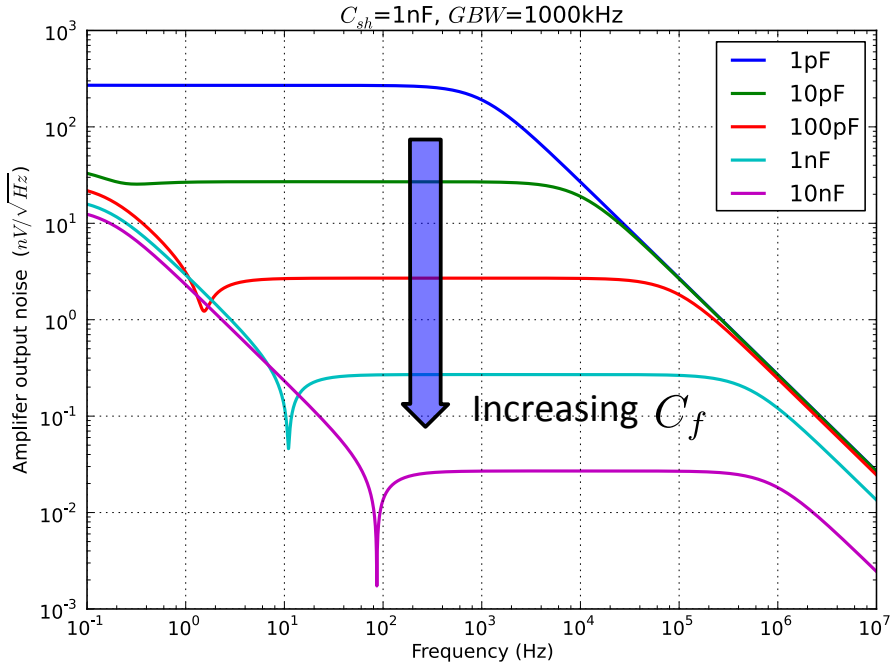


Figure 3.7: Amplifier noise spectrum for a range of C_f. C_{sh} is fixed at 1nF and GBW is 1000kHz.

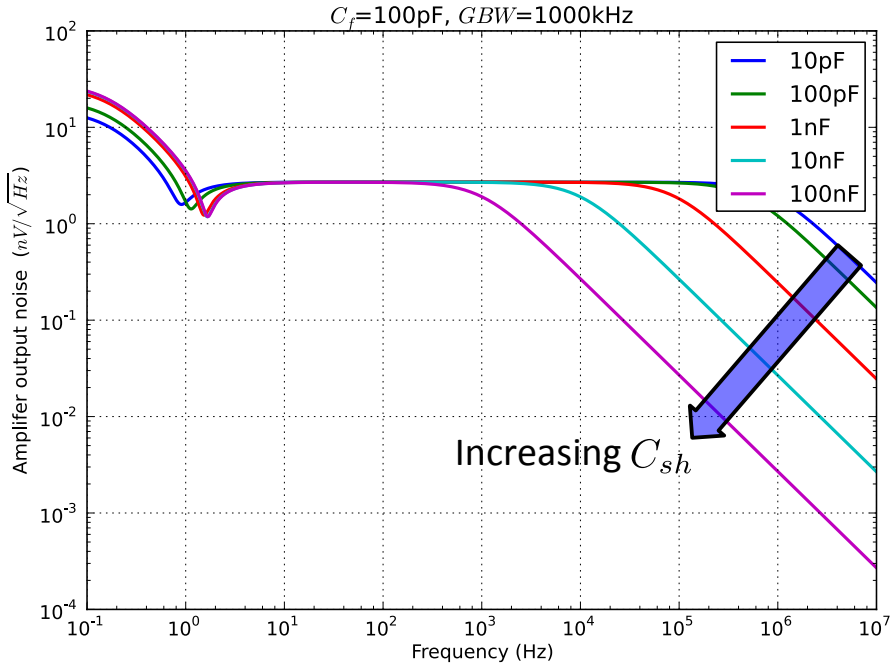


Figure 3.8: Amplifier noise spectrum for a range of C_{sh}. C_f is fixed at 100pF and GBW is 1000kHz.

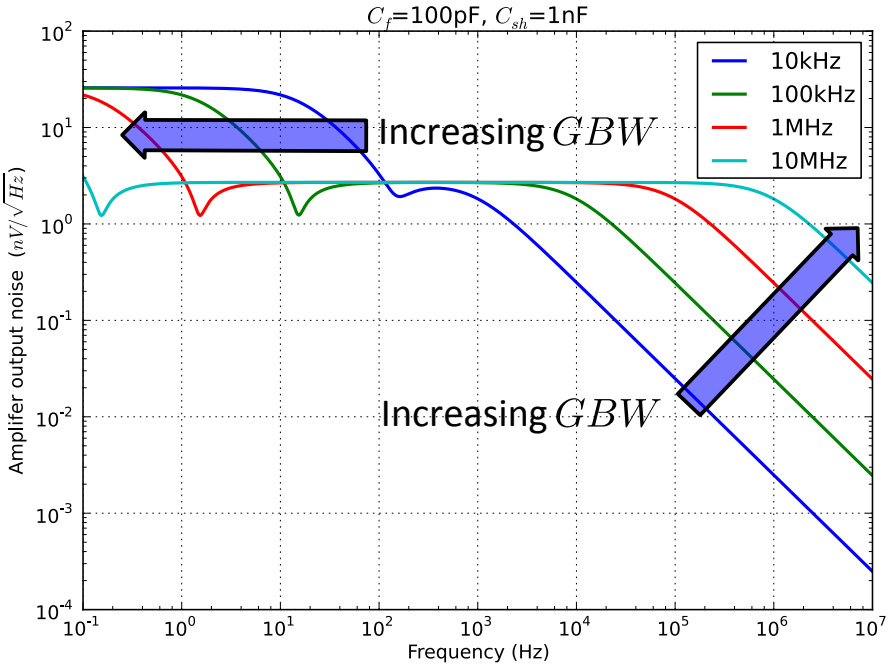


Figure 3.9: Amplifier noise spectrum for a range of GBW. C_f is fixed at 100pF and C_{sh} is fixed at 10nF.

effect on the noise spectrum as can be seen in figure 3.9. As with varying C_{sh}, there is a mid-range frequency band where the noise spectrum is flat and increasing the gain bandwidth increases the width of this frequency band, i.e. reduces the low frequency noise and increases the cut-off frequency at which high frequency noise is attenuated.

3.2 PCB Implementation

This section compares the two types of transimpedance amplifier architectures based on desired photocurrent gain specification and available PCB components using the gain and noise models described previously.

3.2.1 Design Specification

Specification of a PPG system is not a trivial task as it is not feasible to know *a priori* the expected range of absorption coefficients for all possible sensor sites, but a reasonable estimate can be made based on measurements from previous PPG sensors. Table 3.1 lists the specifications to be used for evaluation of the transimpedance amplifiers. The bottom part of the table lists requirements derived from the target performance specifications in the top part of the table.

Table 3.1: Wearable PPG System Specification

Specification	**Value**	**Comments**
Supply Voltage, V_D	3V	Common in low power MCU based systems
Photocurrent Range	1nA-1μA	Based on LED drive current <5mA
Sample Frequency, f_s	2kHz	Support ±1ms timing accuracy with 2 channels
Shunt Capacitance, C_{sh}	100pF	Depends on photodiode choice and PCB implementation but is typically 10pF-1nF
Photocurrent Gain, A_P	3MV/A	Maximum current converted to supply voltage (=130dB).
Input Current Noise, i_N	<1nA$_{\text{RMS}}$	Less than minimum expected current
Bandwidth, f_P	40kHz	To detect 2 harmonics of 10% LED dutycyle

3.2.2 Op-Amp Choice

Table 3.2 lists some possible choices of op-amp for use in the transimpedance amplifier. These devices have been chosen for their low quiescent current consumption and availability in small packages (i.e. SOT23-5 or smaller). Included in the table are the most relevant specifications of the devices to the requirements of a low power transimpedance amplifier. Current and voltage offset specifications have been ignored because part of the complete PPG system design will be to compensate for ambient light offsets which are considerably larger than systematic circuit offsets. Devices which start "OPA" are manufactured by Texas Instruments Inc. and those which start "MCP" are manufactured by Microchip Technology Inc. This is by no means an exhaustive list of suitable op-amps; it is a subset which is intended to be representative of available devices sorted by gain-bandwidth and current consumption.

3.2.3 Gain Comparison

Equation 2.7 is used to model the conventional amplifier and equation 3.3 is used to model the integrating amplifier. Tables 3.3 and 3.4 lists the chosen component values for those equations respectively. In both circuits the Microchip MCP6001 is chosen as a middle of the range representative op-amp to allow like for like comparison.

Table 3.2: Commercially Available Low Power Operational Amplifiers

Device	GBW (MHz)	I_Q (μA)	e_n $\left(\frac{nV}{\sqrt{\text{Hz}}}\right)$	i_n $\left(\frac{fA}{\sqrt{\text{Hz}}}\right)$	SR $\left(\frac{\text{V}}{\mu\text{s}}\right)$	I_B (pA)
OPA333	0.35	17	55	100	0.16	70
MCP6231	0.3	20	52	0.6	0.15	1.0
MCP6001	1.0	100	28	0.6	0.6	1.0
OPA237	1.4	170	28	60	0.5	10nA
OPA335	2.0	285	65	20	1.6	70
OPA340	5.5	750	25	3.0	6.0	0.2

Table 3.3: Conventional transimpedance amplifier component values

R_f	3MΩ	From A_P in table 3.1 and equation 2.8
C_f	4.6pF	From equation 2.9 and specified op-amp

It serves no purpose to compare the undifferentiated outputs from the specified ITIA and the CTIA output as they have no parity in the usability of the output signals. A more constructive comparison is that of the differentiated ITIA output to the CTIA output which is shown in figure 3.10. What is clear from figure 3.10 is the distinct advantage the ITIA has over the CTIA in terms of bandwidth for a common performance specification. This is mainly due to the dominant system pole for the ITIA being set by the reactive components as opposed to the CTIA where the sensitivity is dependent on a resistor whose lossy nature impedes the responsiveness of the circuit.

3.2.4 Noise Comparison

Equation 3.14 is used to model the CTIA noise with A_P specified by 2.8 and equation 3.16 is used to model the integrating amplifier with A_P set by 3.7. Figure 3.11 shows how the differentiated ITIA compares to the CTIA with regard to output referred noise voltage. It is clear from this graph that the improved bandwidth of the ITIA over the CTIA translates directly to having a wider noise bandwidth as well as an increased noise density in the low frequency part of the spectrum.

Table 3.4: Integrating transimpedance amplifier component values

T_{rst}	50μs	Chosen by empirical measurement
C_f	150pF	From A_P and equations 3.7 and 3.9

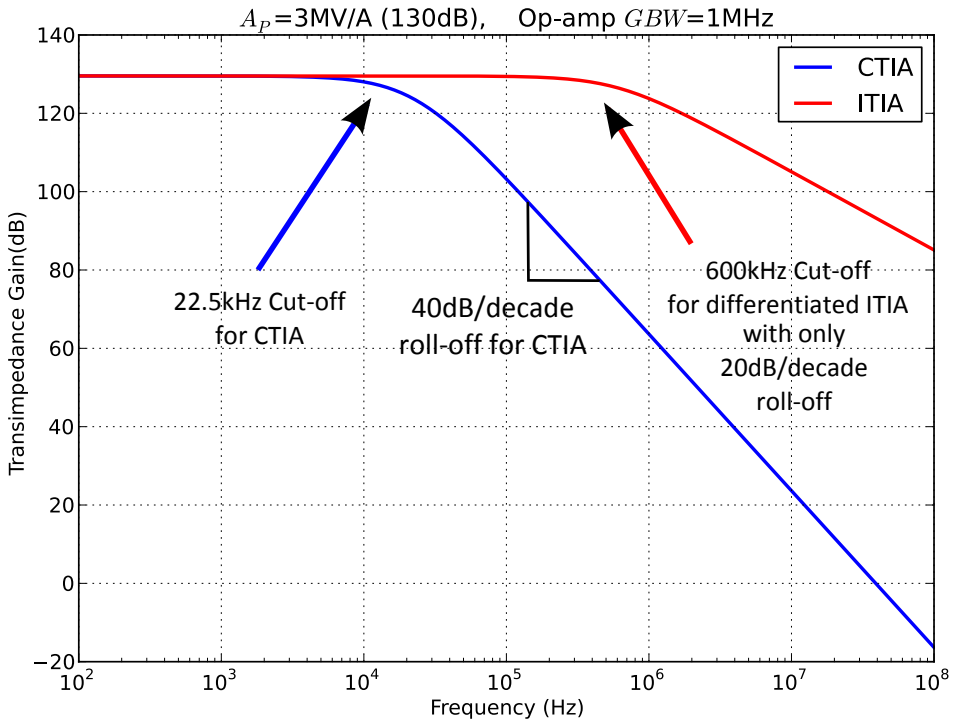


Figure 3.10: Comparison of CTIA response vs. differentiated ITIA response designed for 3MV/A gain with a Microchip MCP6001 op-amp.

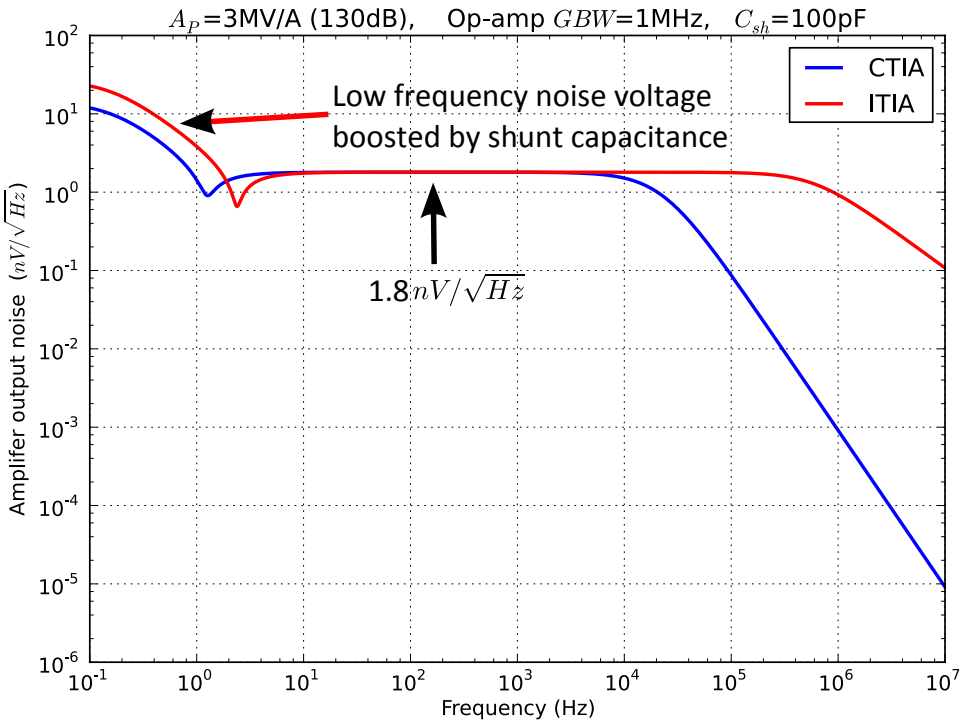


Figure 3.11: Comparison of CTIA and ITIA noise spectra.

To convert the noise spectrum to a root-mean-square (RMS) input current integrate the entire noise spectrum and divide by the amplifier gain.

$$i_n = \frac{1}{A_P}\sqrt{\int_0^\infty |e_{out}(j\omega)|^2 \, d\omega} \tag{3.19}$$

The model-based data shown in figure 3.11 is used to numerically calculate the RMS input noise current for each architecture. To simplify the calculation the integration is performed between 0.1Hz and 1GHz instead of between 0Hz and infinity. Using this method the CTIA has 80fA$_{\text{RMS}}$ for and the ITIA has 583fA$_{\text{RMS}}$, both of which meet the noise requirement from table 3.1.

3.3 Discrete-Time Implications

It is possible to create a pulse oximeter system that operates completely in the analogue and continuous time domain as Tavakoli did [42], but because the sensor is integrated as part of a wireless network the data has to be sampled and converted into digital format at some point in the signal chain. Sampling and converting at a constant time interval is the most common, and easiest to analyse, analogue to digital conversion procedure. In order to implement a pulse oximiter with a single photodiode receiver time division modulation (TDM) is used so there is already a discrete-time aspect to the system architecture. The resettable integrator adds another discrete time aspect to the system. It is logical to synchronise the analogue to digital conversion sampling event with both the LED TDM scheme and the integrator reset schedule.

3.3.1 Converting to Discrete-Time

Figure 3.12 provides a timing diagram to illustrate how the discrete-time nature of both the TDM LED drive and the analogue-to-digital conversion are synchronised with the integrating transimpedance amplifier. Observe during the integration phase the slope of the integrator voltage is proportional to the received photocurrent. In this diagram the IR photocurrent is larger than the red photocurrent which is in turn larger than the photocurrent due to ambient light. The bottom three rows of the diagram show when the discrete time samples are available at the output of the ADC. Note that timings in this diagram are meant to be illustrative and are not necessarily to scale.

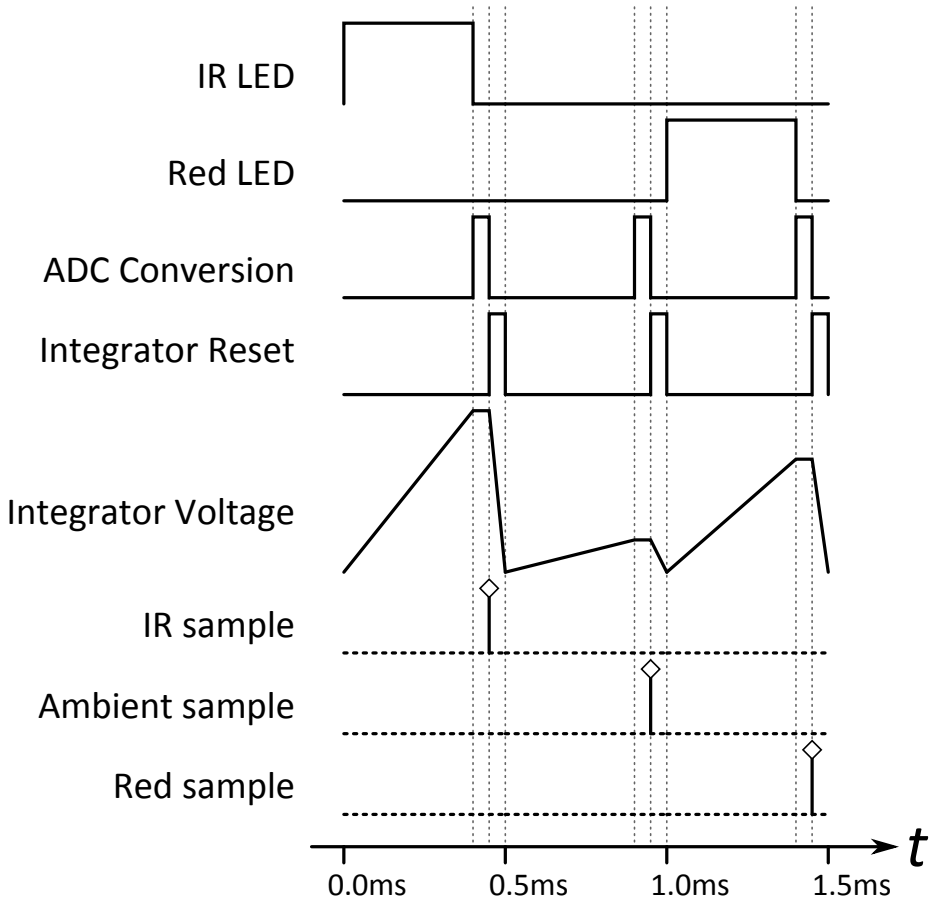


Figure 3.12: Timing diagram to illustrate how time-division-multiplexing is used to measure both red and IR absorption as well as measuring the ambient light level. Included are the timing of the synchronised analogue-to-digital conversions and the resetting of the integrator.

3.3.2 Timing Resolution vs. Bandwidth

It is important to re-iterate that while the proposed PPG system specification calls for timing resolution of ±1ms this does not necessarily mean that the system requires the full bandwidth provided by a 1kHz sampling frequency. This is a fundamental difference between systems for measuring stationary signals versus systems for measuring non-stationary signals. In the context of PPG the bandwidth of a plethysmogram is generally considered to be less than 15Hz [74] which, by Nyquist-Shannon sampling theory, only requires sampling at 30Hz for perfect reconstruction with the use of a sinc interpolator. However, such reconstruction requires both a powerful maths processor to perform the interpolation, which is not an unreasonable assumption if the recipient of the raw data is a PC, and perfectly bandwidth limited data which is a more risky assumption due to the constraints on the quality of filtering that can be implemented on an embedded system.

The direct approach to solving this problem is to increase the sampling rate to match the timing resolution requirement. Essentially the input signal is oversampled with respect to its required frequency resolution. The spare bandwidth provided by oversampling can be used for noise reduction by using digital low pass filters to constrain the bandwidth and hence reduce the overall noise as calculated by equation 3.19. The embedded filters can be implemented with standard digital signal processing (DSP) techniques. This approach is similar to using oversampling and decimation to increase the effective number of bits available from an ADC, but the downsampling step is skipped in order to preserve the timing resolution.

In summary, oversampling the photocurrent signal with respect to its bandwidth allows for increased timing resolution without the need for interpolation and increased signal to noise ratio at the cost on increasing the average current consumption due to increased sampling rate and extra processing time required to implement the DSP filters.

3.3.3 Aliasing

Equation 3.7 has shown how the sampling period affects the gain of an ITIA, but it also has an affect on the system noise due to the possibility of aliasing. The proposed system has a 1kHz sampling rate for a single channel, but as can be seen from figure 3.10, the bandwidth of the amplifier extends beyond that frequency so aliasing is a significant possibility. We already know from figure 3.11 that there is some noise content well above the sampling frequency, but as long as noise floor is low enough aliasing won't add significantly to the

overall noise power. However, as was highlighted in chapter 2, it is possible that there are ambient light signals at frequencies above the Nyquist sampling limit due to modern artificial illumination sources.

The conventional approach to limiting the effects of aliasing is to use an anti-aliasing filter which is simply an analogue low-pass filter designed to ensure that there is little signal content at or above the Nyquist frequency limit. Most often active filters are used, such as the Sallen-Key and Multiple-Feedback topologies, but these require op-amps and hence increase the current consumption of the circuit. Furthermore, this additional filtering in the analogue signal chain can induce phase delay, especially for filters with orders >1, which can be significant enough to affect the implementation of a homodyne detection system which is discussed in chapter 4.

Another approach would be to increase the sampling rate to the point that the Nyquist limit is above the highest frequency input signal and then decimate the signal in the digital domain. Oversampling comes at the cost of higher average current consumption. There is also an adverse affect on the system gain due if it is assumed that a fixed integrator reset duration is required. If reset is not implemented for the integrator then any gain lost by reducing the sampling period can be recovered by integrating within the decimation process, but if reset is implemented then reducing the sampling period causes the reset duration to become a larger percentage of the sampling period.

If the frequency of an ambient signal component can be predicted then it is possible to choose the sampling frequency such that any aliases are forced to occur at frequencies which do not significantly affect the detection of the PPG signal. For example, if ambient artefacts are known to occur at multiples of 100Hz then sampling at 1,000Hz will cause all aliases to also occur at 100Hz intervals - including at DC. This means that the DC signal is corrupted by aliases, but with a homodyne system it is possible to place the desired signal in the gap between aliases as shown back in figure 2.14. Again, more on homodyne detection in chapter 4.

3.4 Prototypes

The functionality of the ITIA has been validated with one hand built prototype and two PCB prototypes.

Prototype Version 0

Version 0 of the sensor system is a hand built electronic circuit created to validate that a resettable ITIA could be implemented using a low gain bandwidth

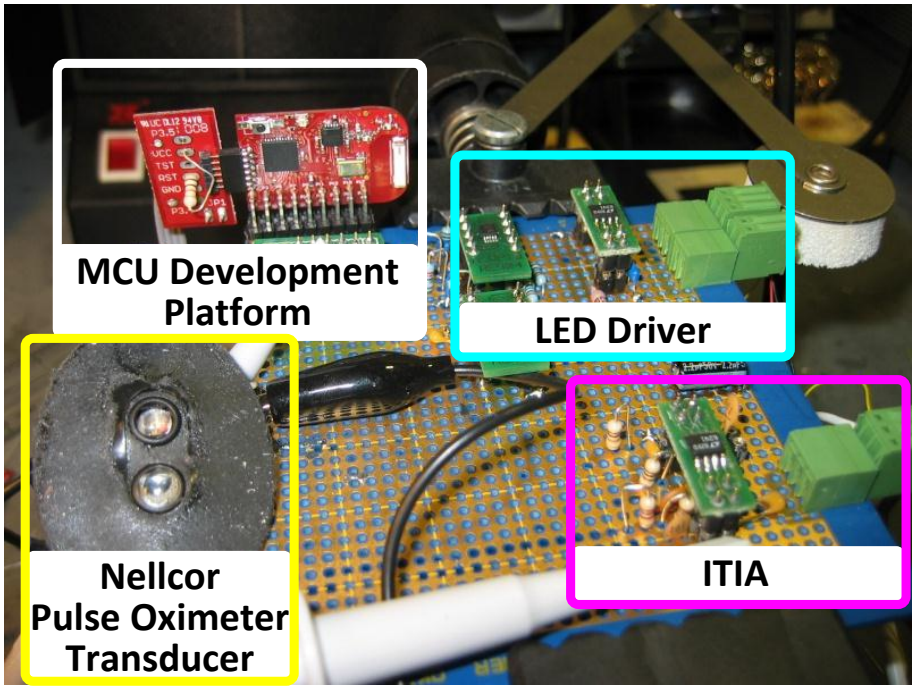


Figure 3.13: Hand built prototype of the ITIA.

op-amp and can be seen in figure 3.13. The op-amp chosen was the OPA333 as it has the lowest power consumption and the worst input noise current making it the most rigorous test of the ITIA architecture. The system control and data interface was provided by the Texas Instrument ez430-RF2500 which has an MSP430F2274 MCU. Resetting of the integrator was performed by a discrete MOSFET which exhibited significant charge injection, inducing a 500mV step across the 100pF integrating capacitor, so it was necessary to ensure that the ambient light level was very low to avoid saturation. The ADC was the National Semiconductor ADC121S085 which has 12 bits of resolution and an input track-and-hold amplifier directly controllable by the SPI chip select line which allowed accurate synchronisation of the converter with the integrate and reset phases of the ITIA. Even with the high level of charge injection and intrinsic weaknesses of a hand built sensor instrumentation it was possible to capture a PPG biosignal which justified the cost of developing of a proper PCB implementation of a wearable PPG sensor system.

3.4.1 Prototype Version 1

Version 1 of the prototype (figure 3.14) was built on a 4 layer, 1.6mm PCB and designed to be compatible with version 3 of the BSN development kit [75]. Control and data interface was provided by the MSP430F2122 which

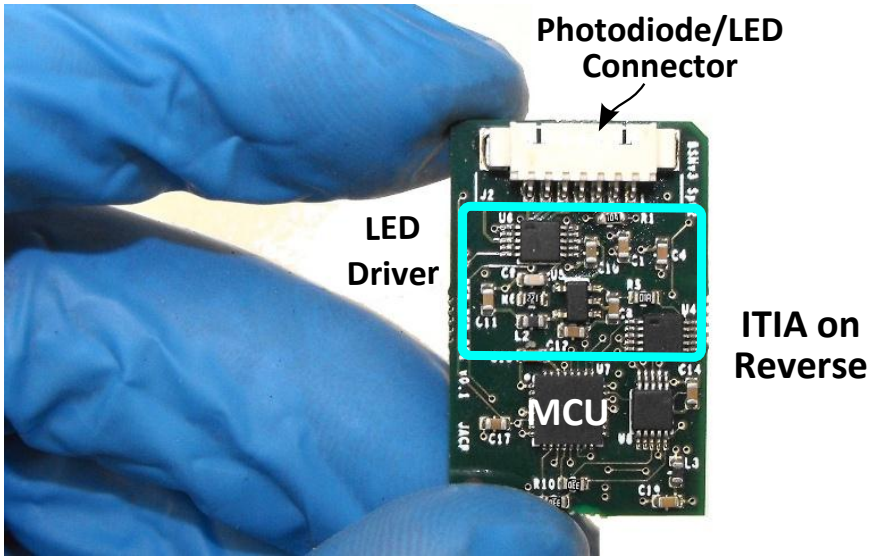


Figure 3.14: Version 1 of the wearable PPG system prototype PCB.

was chosen as it provided the bare minimum peripherals and memory required in order to support the testing of the ITIA. A Texas Instrument TS5A3519 analogue switch which has an ON resistance of 36pC which translates to only a 0.36mV step in potential across the integrating capacitor.

This version was used in the BSN2010 publication [7] and its equivalent journal version [8]. These investigated noise and artefact reduction as will be discussed in later chapters. During the process of gathering results for these publications it was discovered that the lack of an output buffer on the ITIA meant that the large capacitance of an included passive anti-aliasing filter induced oscillations which reduced the signal to noise ratio and induced inter channel interference between the red and IR optical channels. Regardless of this shortcoming version 1 prototype proved that the ITIA is capable of extracting biosignals.

3.4.2 Prototype Version 2

Functionally, version 2, pictured in figure 3.15, is identical to version 1 except with an output buffer for the ITIA to separate it from the anti-alias filter's capacitance. The output of the buffer was also connected to the MSP430F2122's internal 10 bit ADC to investigate the possibility of discarding the external ADC and accepting a lower resolution.

Physically, the form factor of the PCB was very different to version 1. The data interface was changed to a general purpose connector instead of the BSN specific interface. The transducer was also implemented as a flexible PCB

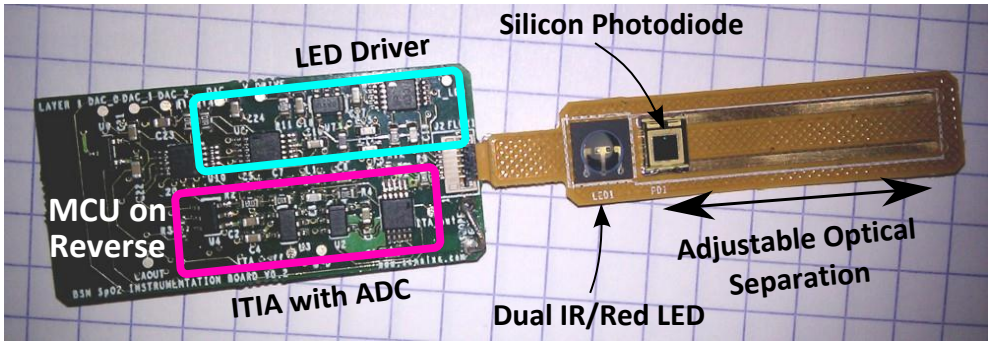


Figure 3.15: Version 2 of the wearable PPG system prototype PCB.

instead of a hand built one as was used in the BSN2009 publication[6].

The full circuit diagram of prototype version 2 can be found in Appendix A on page 206.

This version of the circuit was used as the platform for the research presented in the BSN2012 publication of high-dynamic range noise rejection [9]. This publication did not focus on the integrator circuit itself, but the low power and high bandwidth performance of the ITIA provides was essential for the implementation of the signal processing which is covered in chapters 4 and 5.

3.4.3 Transducer

In order to test any photodetection circuitry for PPG applications it is necessary to have a transducer to house the photodiode and LEDs. As can be seen in figure 3.13 the very first prototype used a Nellcor^{TM}forehead sensor with the connector adapted to connect to the prototype circuitry (the green connectors in the photo). The Nellcor transducer is a commercial product and hence has passed all required compliance testing so seems a suitable choice, but its large size and long (i.e. >10cm) cable to the instrumentation electronics mean it is not suitable for wearing it anywhere other than the forehead.

The key challenges in designing a transducer for wearable applications are keeping the height profile of the patch minimised and allowing the transducer to flex enough to meet the contours of the surface that it is being placed on. The requirement for a flexible PPG transducer was investigated for the 2009 BSN conference [6] and the hand-built transducer is illustrated in figure 3.16. A substrate of polyester tape makes it simple to place the photodiode and LED with the desired separation while retaining high flexibility (i.e. a bend radius of <10cm). Black foam is cut to shape to surround the opto-electronic components to provide optical screening between them while also making the sensor comfortable to wear before another layer of polyester tape is used to

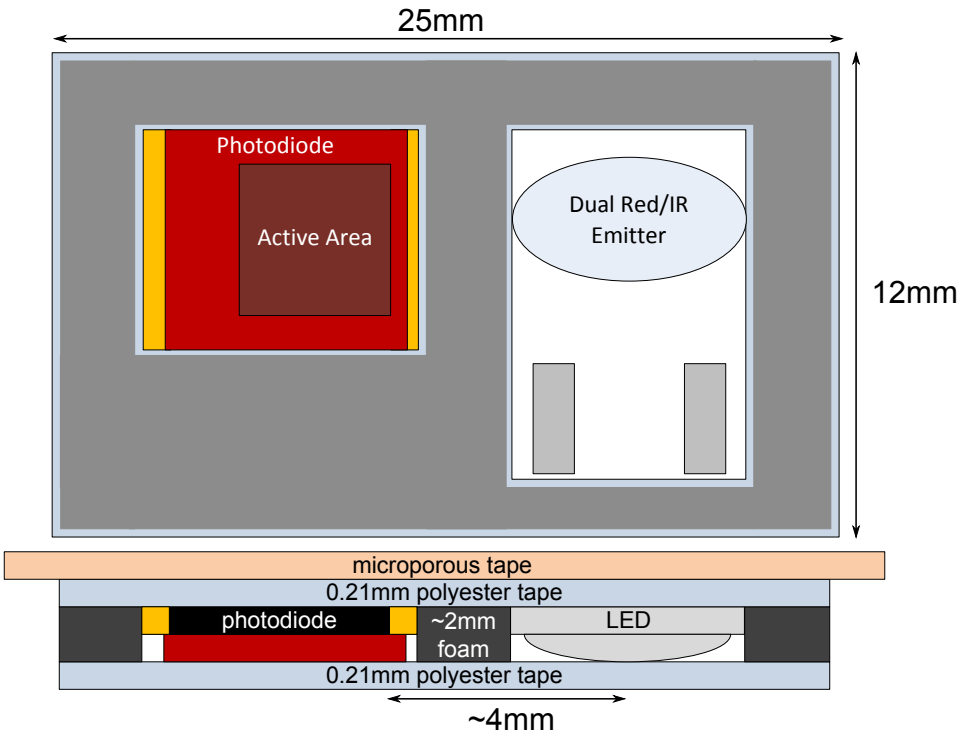


Figure 3.16: A hand-built flexible PPG transducer detailing the type of sensor used throughout this thesis for testing.

seal the components within the transducer.

Using the rudimentary signal strength metric of AC to DC ratio of photocurrent it was shown that it is possible to measure cardiac signals in the very shallow tissue around the ear where a rigid transducer would have blocked all capillary blood flow. Figure 3.17 shows the locations tested and how well each location performed, with and without supporting pressure. The ability to measure without supporting pressure (i.e. taped or glued on) is important as adhesive is the primary method of attaching patches to tissue.

The transducer was subsequently constructed using flexible PCB technology (as shown in figure 3.15), but the vias in the 2-layer design proved to be too delicate and kept breaking during experimentation so it was decided to avoid flexible circuits and construct all transducers by hand in the same form factor as shown in figure 3.17.

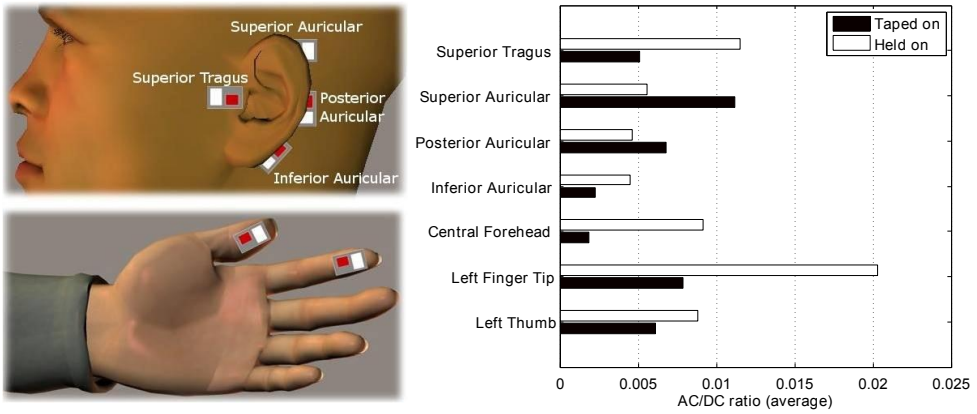


Figure 3.17: A comparison of the PPG signal quality achieved at a range of sensing locations.

3.5 Conclusions

This chapter has presented the integrating transimpedance amplifier as a more suitable alternative to the conventional transimpedance amplifier due to the ability to implement high gain and bandwidth in a single amplifier stage. A reduction in power is enabled by reducing the gain bandwidth product required of the op-amp around which the transimpedance function is implemented which is directly related to the current consumption of the op-amp.

A further benefit of the ITIA over the CTIA is that its bandwidth is comparatively immune to large parasitic shunt capacitances and nor does it require a feedback compensation capacitor whose value is dependent on the specific value of the shunt capacitance. This means a single ITIA, designed to a gain and bandwidth specification, can be used for a wide range of photodiode sizes as well as being tolerant to manufacturing variability in both the opto-electronics and the PCB assembly.

Furthermore, by converting from continuous time to discrete time with an integrator means that the constituent photometric content between samples is preserved allowing for non-linear drive of LEDs while preserving a linear measurement relationship which simplifies the design of the PPG system's illumination circuitry. It also means signal processing can be performed in the digital domain without having to take into consideration complexities such as zero order holds

Design of an ITIA is not necessarily as straight-forward as designing a CTIA, but the simplified system models derived in this chapter reduce the complexity of translating a set of instrumentation requirements into circuit parameters.

All the benefits of the ITIA come at a cost of increased circuit noise, both

at low frequencies and from the extended frequency response, but the noise level is still well below that required of a wearable PPG system. The ITIA can also be designed to have a frequency band which has flat spectral noise density identical to that of the equivalent CTIA which means any subsequent bandpass separated signal will achieve identical noise performance to that of the CTIA without the cost of increased op-amp current consumption. In order to take advantage of this low noise region of the ITIA spectrum means a way must be found to operate away from DC.

4 Optical Homodyne Sensing

Having chosen a high transimpedance gain in order to minimise the LED current it is necessary to make sure LED light is detected in preference to light from independent sources. Chapter 2 introduced the concept of synchronising the photodetector to the LED modulator in order to separate the PPG system from the ambient environment in the phase domain (e.g. correlated double sampling). All previous solutions have required additional circuitry in the detection chain in order to perform the required phase locking and signal conditioning. They also suffer from a mismatch between the discrete time, spread spectrum output of a pulse width modulated LED and the continuous time photodetection circuitry [68, 76]. This chapter will focus on the details of the modulation and demodulation of an optical absorption signal from a frequency domain perspective, and propose one synchronous and one asynchronous detection method, both of which can be implemented in the digital domain with a microcontroller with no extra analogue domain circuitry.

4.1 Introduction to Signal Mixing

Homodyne sensing is based on the the *product-to-sum* trigonometric identities for the multiplication of two time varying signals. The most simple case is the multiplication of two cosine waves which produces a new signal which is the sum of two cosine waves with frequencies at the sum and the difference of the two original signal frequencies.

$$\cos(\omega_A t)\cos(\omega_B t) = \frac{\cos([\omega_A - \omega_B]t) + \cos([\omega_A + \omega_B]t)}{2} \tag{4.1}$$

The product-to-sum identity can still be applied to signals which are not pure sine waves by considering the signals in terms of their Fourier transforms. In the frequency domain the resultant signal is still the sum and difference of all constituent frequency components for each of the original signals. This concept is illustrated in figure 4.1 showing the effect of multiplying, or *mixing*, a multiple frequency signal with a pure tone sinusoidal signal - or *carrier*.

To apply this to optical sensing one of the sensors has a local oscillator to generate a sinusoidal signal which drives an LED with frequency ω_c, amplitude

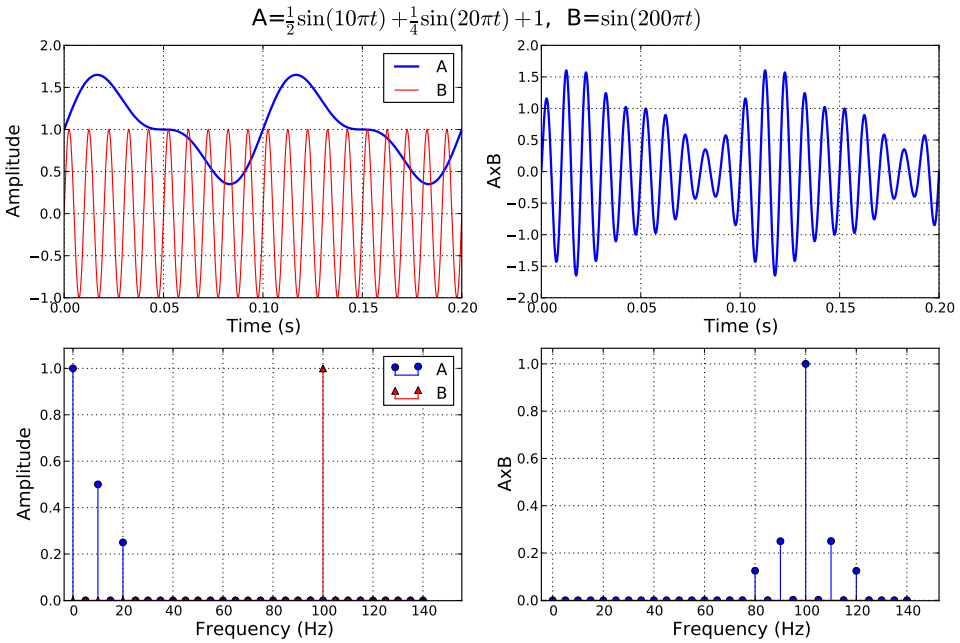


Figure 4.1: The consequence of mixing two signals represented in both the time and frequency domains.

A_L and offset C_L to ensure that the LED drive current is always positive. The LED optical output is absorbed by an arbitrary substance which, for simplicity, is ascribed the frequency ω_A, amplitude A_A and offset C_A. The irradiance of the photodetector is now the product of two non-zero mean oscillating signals plus an additive noise term which represents light from sources independent from the sensor:

$$I_P(t) = (\underbrace{A_L \cos(\omega_c t) + C_L}_{\text{LED output}})(\underbrace{A_A \cos(\omega_A t) + C_A}_{\text{absorption signal}}) + I_N(t) \tag{4.2}$$

As both signals have DC components the product signal has 4 frequency components, which includes the original signal frequencies, and a DC component:

$$\begin{aligned} I_P(t) = &\frac{A_L A_A}{2} \cos([\omega_c - \omega_A]t) \\ &+ \frac{A_L A_A}{2} \cos([\omega_c + \omega_A]t) \\ &+ A_L C_A \cos(\omega_c t) \\ &+ A_A C_L \cos(\omega_A t) \\ &+ C_L C_A + I_N(t) \end{aligned} \tag{4.3}$$

Now most of the signal content is centred around the ω_c sinusoid which from now on will be referred to as the *carrier*. The desired outcome of the detection process is to recover the absorption coefficient of the substance under investigation which in the current model has coefficients A_A, C_A and ω_A. There are two main approaches to recovering a modulated signal: *product detection* and *envelope detection*.

4.1.1 Product Detection

Product detectors recover the absorption signal by multiplying the received signal by a signal which has the same frequency as the signal which was used to modulate the substance under investigation. In a homodyne detector the original modulation signal is available to the detector to be used for product detection. In a system where the modulating signal is not available locally it has to be recovered somehow from the incoming signal - this is *heterodyne* detection. Multiplying the received signal by the LED modulation signal (without the DC offset) produces the following signal:

$$\begin{aligned} I_P(t)\cos(\omega_c t) = &\frac{A_L A_A}{4}\left(\cos([2\omega_c - \omega_A]t) + \cos([2\omega_c + \omega_A]t)\right) \\ &+ \frac{A_A C_L}{2}\left(\cos([\omega_c - \omega_A]t) + \cos([\omega_c + \omega_A]t)\right) \\ &+ \frac{A_L C_A}{2}\cos(2\omega_c t) \\ &+ (C_L C_A + I_N(t))\cos(\omega_c t) \\ &+ \frac{A_L}{2}(\underbrace{A_A\cos(\omega_A t) + C_A}_{\text{absorption signal}}) \end{aligned} \tag{4.4}$$

The resulting signal from the receiver product calculation has many frequency components, but the one of primary interest is the final term in equation 4.4 which contains the absorption signal. By choosing the value of ω_c carefully the absorption signal can be separated in the frequency domain from all other signal components. For example, choosing $\omega_c \gg \omega_A$ allows low-pass filtering to be used to reject all undesired signal components. Another important constraint on the choice of ω_c is that is should not overlap with any frequencies present in the noise signal $I_N(t)$ otherwise the noise will also be mixed down into the low frequency band and corrupt the absorption signal.

It is important to note that this model assumes that there is negligible phase difference between the carrier signal received at the output of the photodetector and the ω_c based signal used to demodulate it. When the phase difference cannot be ignored then *quadrature demodulation* is required - also known as

direct-conversion in radio terminology. This entails implementing two parallel detector circuits: one mixes the incoming signal with $\cos(\omega_c t)$ to create a "real" component and the other mixes it with the local oscillator phase shifted back by $\pi/2$ (i.e. $-\sin(\omega_c t)$) to create an "imaginary" component. From this complex representation both the absorption signal magnitude and the phase delay in the carrier can be calculated. This is the approach used by Grubb et al. for their lock-in detector technique [46]. However, it is possible to keep the phase difference low enough such that quadrature demodulation is not required by using the discrete time nature of the integrating transimpedance amplifier.

4.1.2 Envelope Detection

Envelope detection is an alternative to product detection which is desirable due to the simplicity of its implementation. It is possible to recover the absorption signal by inspecting how far the output from the photodetector deviates from the mean throughout the cycle of the carrier. Classical radio receivers do this by rectifying the received radio signal and low-pass filtering it to recover the audio "envelope". Unfortunately, due to the non-zero mean of optical signals rectification will not work and hence this technique cannot be directly applied to PPG.

Instead of relying on rectification which is a predominantly analogue design approach a simple alternative is to measure the signal variance which was used successfully in the paper "Ratiometric Artefact Reduction in Low Power, Reflective Photoplethysmography" which was published in a special edition of the IEEE's Transaction on Biomedical Circuits and Systems in 2011 [8]. Calculating the variance has the effect of mixing the received signal with itself, removing the mean and low-pass filtering the signal all in one process. As the variance can be calculated from the "mean of the square minus the square of

the mean" the analysis starts by considering the square of the input signal:

$$\begin{aligned}
I_P(t)^2 &= \cos(\omega_c t) \times (2[C_L C_A + I_N(t)]A_L C_A + A_L A_A A_A C_L) \\
&+ \cos(\omega_A t) \times (2[C_L C_A + I_N(t)]A_A C_L + A_L A_A A_L C_A) \\
&+ \cos(2\omega_c t) \times \left[\left(\frac{A_L A_A}{2}\right)^2 + \frac{(A_L C_A)^2}{2}\right] \\
&+ \cos(2\omega_A t) \times \left[\left(\frac{A_L A_A}{2}\right)^2 + \frac{(A_A C_L)^2}{2}\right] \\
&+ \cos([\omega_c - \omega_A]t) \times (A_L A_A C_L C_A[1 + I_N(t)]) \\
&+ \cos([\omega_c + \omega_A]t) \times (A_L A_A C_L C_A[1 + I_N(t)]) \\
&+ \cos([2\omega_c - \omega_A]t) \times \left(A_L^2 A_A C_A\right)/2 \\
&+ \cos([2\omega_c + \omega_A]t) \times \left(A_L^2 A_A C_A\right)/2 \\
&+ \cos([\omega_c - 2\omega_A]t) \times \left(A_L A_A^2 C_L\right)/2 \\
&+ \cos([\omega_c + 2\omega_A]t) \times \left(A_L A_A^2 C_L\right)/2 \\
&+ \cos(2[\omega_c - \omega_A]t) \times \frac{1}{2}\left(\frac{A_L A_A}{2}\right)^2 \\
&+ \cos(2[\omega_c + \omega_A]t) \times \frac{1}{2}\left(\frac{A_L A_A}{2}\right)^2 \\
&+ \left(\frac{A_L A_A}{2}\right)^2 + \frac{(A_L C_A)^2}{2} + \frac{(A_A C_L)^2}{2} + [C_L C_A + I_N]^2
\end{aligned} \tag{4.5}$$

The next step is to estimate the mean of the squared signal. In a real-time system the mean is calculated over a moving time frame, or window, and if the width of that window is chosen carefully enough then unwanted high frequency components in equation 4.5 are removed to give the following expression:

$$\begin{aligned}
\overline{I_P(t)^2} &= \cos(2\omega_A t) \times \left[\left(\frac{A_L A_A}{2}\right)^2 + \frac{(A_A C_L)^2}{2}\right] \\
&+ \cos(\omega_A t) \times (2[C_L C_A + I_N(t)]A_A C_L + A_L A_A A_L C_A) \\
&+ \left(\frac{A_L A_A}{2}\right)^2 + \frac{(A_L C_A)^2}{2} + \frac{(A_A C_L)^2}{2} + [C_L C_A + I_N]^2
\end{aligned} \tag{4.6}$$

The square of the mean can be directly inferred from the low frequency com-

ponents for equation 4.3.

$$\begin{aligned}\overline{I_P(t)}^2 &= (A_A C_L \cos(\omega_A t) + [C_L C_A + I_N(t)])^2 \\ &= \cos(2\omega_A t) \times \frac{(A_A C_L)^2}{2} \\ &+ \cos(\omega_A t) \times 2[C_L C_A + I_N(t)] A_A C_L \\ &+ \frac{(A_A C_L)^2}{2} + [C_L C_A + I_N]^2\end{aligned} \tag{4.7}$$

The output of the variance calculation can now be expressed in terms that are independent of the carrier.

$$\begin{aligned}\sigma^2 &= \overline{I_P(t)^2} - \overline{I_P(t)}^2 \\ &= \cos(2\omega_A t) \times \left(\frac{A_L A_A}{2}\right)^2 \\ &+ \cos(\omega_A t) \times (A_L A_A)(A_L C_A) \\ &+ \left(\frac{A_L A_A}{2}\right)^2 + \frac{(A_L C_A)^2}{2} \\ &= \left(\frac{A_L}{\sqrt{2}}\right)^2 \cdot (\underbrace{A_A \cos(\omega_A t) + C_A}_{\text{absorption signal}})^2\end{aligned} \tag{4.8}$$

This has shown that in the case that $\omega_c \gg 2\omega_A$ the absorption signal can be be recovered from the modulated photocurrent signal by calculating the square root of the variance, i.e. the standard deviation of the signal, over a time period where the absorption signal is relatively constant but the carrier has completed at least one full oscillation. A significant advantage of self-mixing over the product detector method is that it is immune to phase lag in the detector chain.

There are two weaknesses to this method. The first is that self-mixing allows the additive noise signal $I_N(t)$ to mix with *both* the absorption signal *and* the carrier. Depending on the spectral content of $I_N(t)$, it is possible that noise will mix into the same frequency band as the absorption signal. With the product detector it is only necessary to avoid sum-and-difference frequencies between the ω_c and I_N whereas with the envelope detector all combinations of ω_A, ω_c and I_N, as well as their double frequencies, need to be avoided which complicates the design. The second weakness is that the carrier needs to be significantly larger than the *double* of the absorption signal which either means using a high carrier frequency or using higher order filters to calculate the mean to ensure satisfactory spectral separation. A higher carrier frequency requires a higher bandwidth photodetector and higher order filter requires

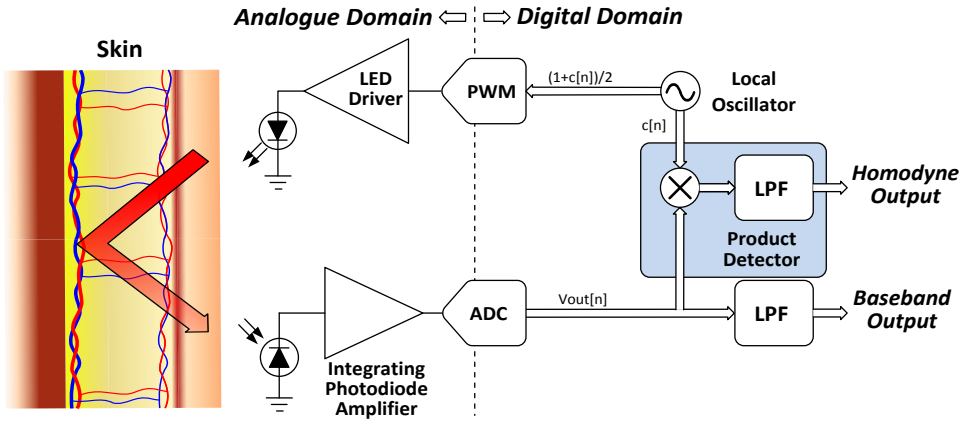


Figure 4.2: Mixed-signal homodyne PPG sensor system architecture.

more processing elements - both of which increase current consumption.

In [8] the rolling mean was calculated as the average of a contiguous set of N samples where the carrier frequency, f_c, was chosen to be $2f_s/N$ where f_s was the sampling frequency of the detector. Choosing the width of the mean window to be an integer multiple of the number of samples in the carrier frequency meant that the carrier is completely cancelled out. The square root of the variance was never calculated due to a subsequent ratiometric calculation rendering it unnecessary. Research following the publication of that paper indicated that an FIR based rolling average calculation does not provide enough spectral rejection of high frequency mixture signals at reasonable choices of carrier frequency for a low-power opto-electronic sensor system (e.g. <500Hz).

It was decided to focus on using a product detector for implementing a wearable PPG system due to its simplicity in implementation, superior noise rejection and the availability of the carrier locally with a low phase lag detector chain.

4.2 Mixed-Signal Homodyne Detector

The PCB implementation of a PPG system based on homodyne sensing places most of the signal processing in the digital domain to keep PCB size to a minimum. The main signal processing functions are multiplication, low-pass filtering and sine wave generation. All of which can be implemented on a low-power MCU which means the only analogue components are the LED driver and the photodiode amplifier. Figure 4.2 shows the distribution of functionality between the analogue and digital domains. The only required mixed signal components are one analogue-to-digital converter (ADC) photodetec-

tor and two digital-to-analogue converters (DAC). Most contemporary MCUs have ADCs and some include DACs as well, but for flexibility in choice of MCU the DAC will be implemented with a discrete, 4-channel DAC which the MCU controls via a serial interface. The extra DAC channels will be used to implement features discussed in the next chapter.

4.2.1 Carrier Synthesis

The mixed-signal homodyne sensor architecture relies on the MCU to internally generate the carrier signal as it already has an intrinsic oscillator which can be used as a reference for sinusoid synthesis. Using a discrete oscillator would not only require more PCB area, but would need to be interfaced with both the LED driver and product detector, complicating the design unnecessarily. *Direct Digital Synthesizers* (DDS) is a term often used to describe systems which rely on digital signal processing (DSP) and a high frequency reference oscillator to generate a waveform in the discrete time domain and then use a DAC to convert to analogue. An example using of DDS in sensing is the Analog Devices AD5934 impedance converter network analyser where a DDS drives a voltage across an unknown impedance and a transimpedance amplifier measures the resulting current before demodulating it in the digital domain to measure the amplitude and phase characteristics of the device under test. This is actually very similar to what this thesis is presenting as a PPG sensing solution, but instead of measuring electrical impedance we measure "optical resistance" i.e. absorption.

There are a few approaches to generating a sinusoidal oscillator in discrete-time systems; here are three of the most popular solutions: recursive oscillators [77], numerically controlled oscillators and the CORDIC vector rotation algorithm [78, 79]. Recursive oscillators are state space systems designed deliberately such that the internal state variables oscillate at a stable frequency and fixed amplitude. However, the state transition matrix requires multiple multiplications per sample and truncation errors in fixed-point approximations of the transition matrix both make recursive oscillators non-trivial to implement on an MCU. The same set of issues also make CORDIC difficult to implement on an MCU as well as it being a successive approximation algorithm which further increases the required processing time and hence power. The simplest solution for MCU implementation is a numerically controlled oscillator which uses an internal timer to generate a phase accumulation oscillator and then a look-up table (LUT) to calculate the sinusoidal output from the phase. The relatively low frequencies (i.e. <1kHz) required to implement homodyne sensing for PPG mean that it is acceptable that the period of the carrier has to

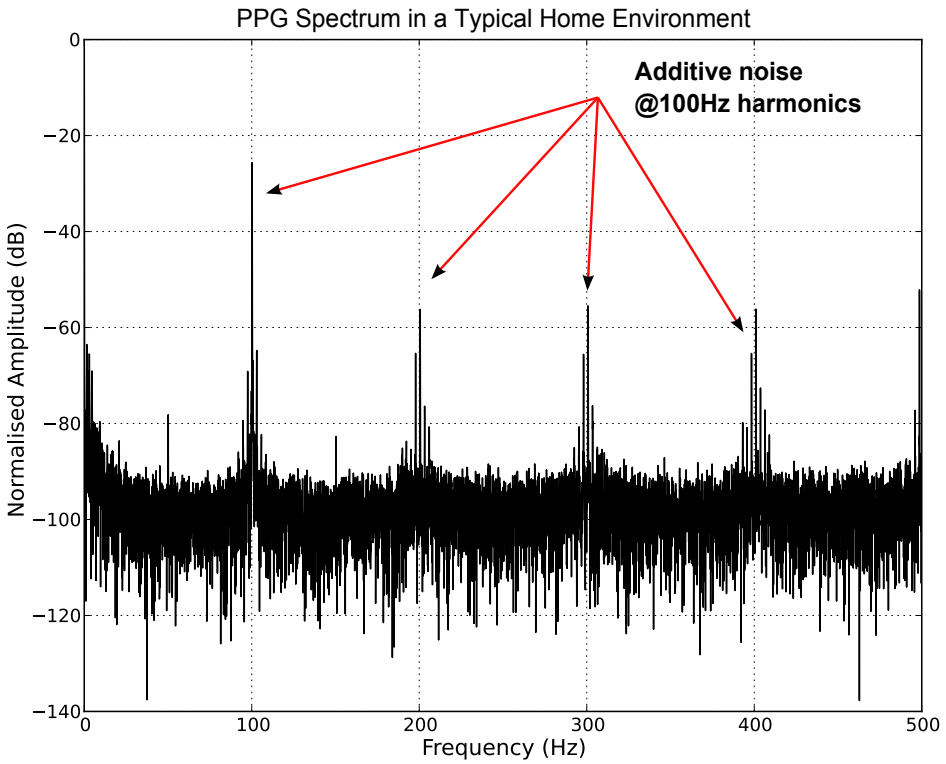


Figure 4.3: Presence of noise from ambient light sources in the photodetector spectrum. The data shown here was captured from a reflection mode PPG transducer worn at the tip of the author's index finger.

be an integer multiple of the sampling period. The upper limit on the carrier frequency is half of the sampling frequency which for a single channel system is equivalent to correlated-double-sampling (CDS) as the only two modulation states of the LED are fully on and fully off. Creating the carrier at the Nyquist limit of the sampling frequency is possible due to both the photodetector and product detector both operating with the same discrete time phase. As such, aliasing with respect to the carrier cannot occur as long as the absorption waveform frequency range is well below half the carrier frequency. An additional feature is that if the length of the LUT is a multiple of four then it is simple to generate full quadrature oscillator output which can be used if it is deemed that the analogue signal path induces more than a single sample's worth of phase delay and quadrature amplitude demodulation is required.

Choice of what frequency to set the carrier at depends on where the least interference from ambient light sources occurs. The presence of 100Hz harmonic from ambient light sources (empirically determined to be from a combination of low energy light bulbs and LCD TV backlights) are shown in figure 4.3. To maximise the spectral separation of the carrier from the noise peaks it is

necessary to bisect the frequency gap between them which occur at 150Hz, 250Hz, 350Hz and 450Hz etc. From a sampling frequency of 1kHz the easiest frequency to synthesise is 250Hz by using a LUT of length 4.

4.2.2 LED Drive

Chapter 2 stated that the LED driver circuitry is trivial, but for completeness it is presented here. The primary requirements of the LED driver for the proposed homodyne sensing system are:

Minimal Components: Time division multiplexing of multiple optical channels for detection by a single photodetector means that only one LED is turned on at any given time. In turn, this means that only one current source is required to drive multiple channels by using a switch matrix to independently connect the current source to individual LEDs. All LEDs which are not turned on are reverse biased to ensure that no leakage current allows the emission of any undesired photons causing cross-channel interference.

Channel Programmable Current: Different types of LEDs have different conversion efficiencies and voltage requirements so for a multi-channel sensor system it is necessary to support independent drive currents for each LED. For flexibility a DAC is used to drive a voltage controlled current source (VCCS) to allow the MCU to specify the peak drive current for each LED. For a mass produced design it would be sufficient to build independent, fixed voltage references for each optical channel and do away with the DAC.

PWM LED Intensity Modulation: The integrating photodetector means that it is only necessary to modulate the average LED current during the sampling period. The instantaneous current through the LED is set by the VCCS and DAC, but instead of synthesising the carrier by modulating the DAC output it is simpler to modulate the percentage of time that the LED is turned on during the sampling period using pulse width modulation. The simplicity comes from only needing a logic level input to determine whether an LED is on or off and only a single synthesizer output is needed to generate the carrier for all optical TDM channels while the DAC can be used to implement channel independent LED intensity.

The circuit in figure 4.4 shows a two channel LED driver which meets all of these requirements with a single op-amp, single DAC channel and two discrete analogue switches. A DAC resolution of 10 bits or more provides suitable

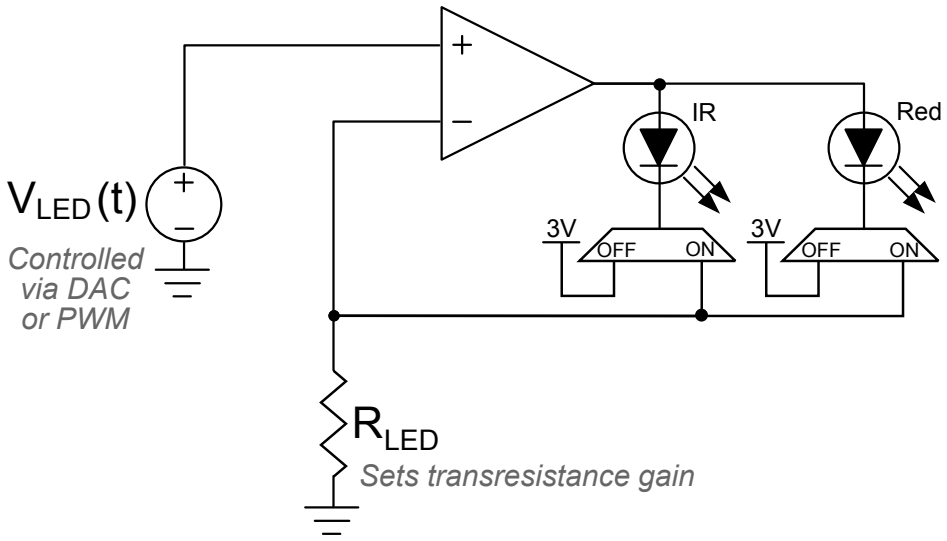


Figure 4.4: A two channel LED driver supporting both pulse width and amplitude modulation.

tunability in LED drive current. The MSP430F2122 used in prototype PCB versions 1 and 2 can run at 8MHz from the internal oscillator which for a 2kHz sampling rate allows for 4000 duty cycle settings which is just short of a full 12 bit resolution. To achieve 14 bits would require quadrupling the MCU core clock frequency to 32MHz which is beyond the capabilities of the chip. However, if the carrier synthesiser LUT length is only 4 then only 2 bits of resolution are required and the PWM timer only needs to be 8kHz which mean both the synthesiser and LED PWM modules can be referenced from a low frequency (e.g.32kHz) crystal which can lead to further MCU power reductions. Unfortunately, a low frequency crystal was not designed into either version 1 or 2 of the prototype sensor PCBs so it is not worth critically analysing the power consumption of homodyne sensing on these boards.

4.2.3 Timing for Multi-Channel Homodyne

Figure 4.5 illustrates how the LED drive, transimpedance amplifier and the ADC are synchronised in order to implement a completely discrete-time signal chain. The conversion and reset phases of the integrating amplifier have been omitted for simplicity, but what is shown is the ever present effect of ambient light during the period when neither LED is turned on. What is also shown in this diagram is how the LED intensity per sample period is modulated between zero and maximum drive current, mapped from the -1 to +1 range of the local oscillator.

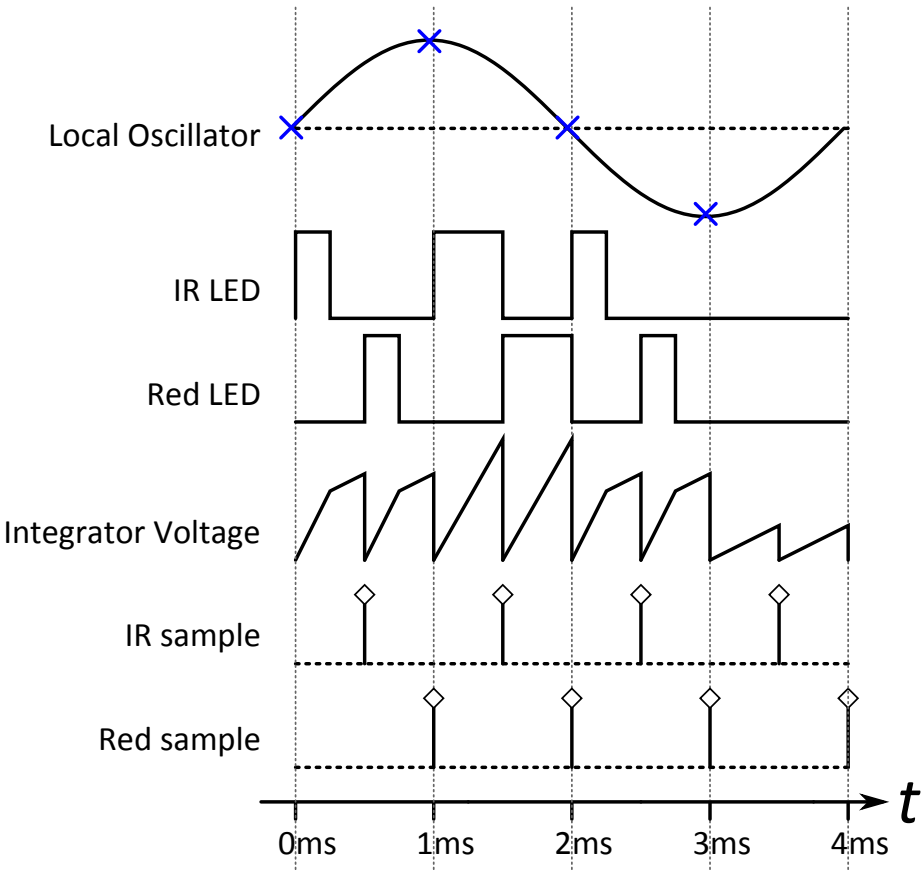


Figure 4.5: Timing diagram for a 2 channel homodyne detection system with integrating photocurrent amplification, pulse width modulation of LED intensity and synchronised analogue to digital conversion.

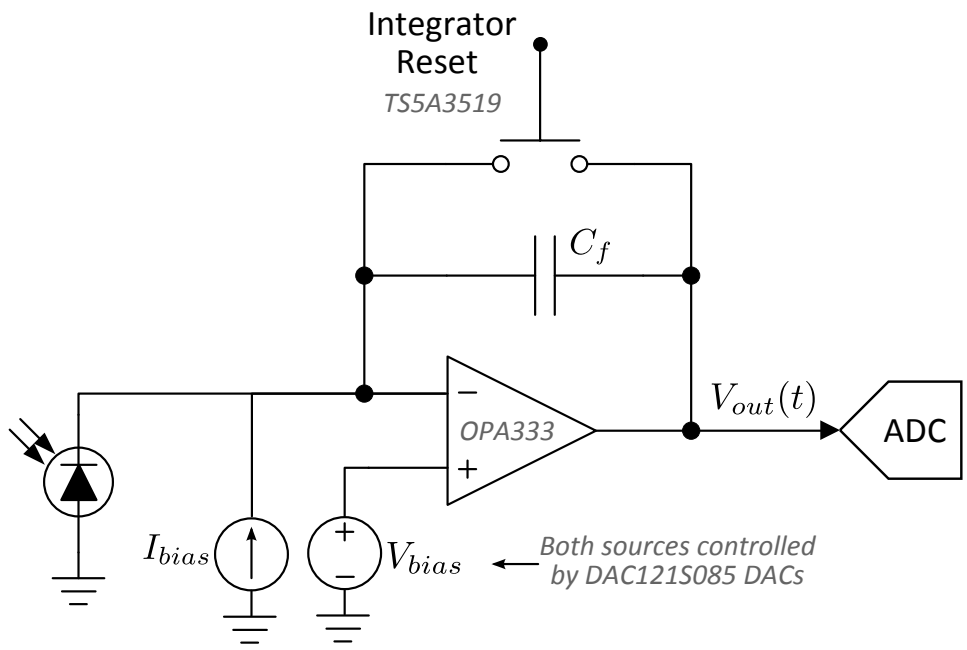


Figure 4.6: Circuit diagram for PCB implementation of the integrating transimpedance amplifier.

4.2.4 Photodetector

In keeping with the discrete time nature of the carrier synthesis, the photodetector used in the system is the resettable integrating transimpedance amplifier as presented in section 3.1.2 and specified in section 3.2. The circuit diagram for the photodetector used for homodyne sensing is shown in figure 4.6. The ADC is the one included in the MCU which for PCB version 2 is the Texas Instruments MSP430F2122 which has a resolution of 10 bits. Bias current and voltage is provided by two channels of the National Semiconductor DAC125S085 4 channel DAC.

4.2.5 Product Detection

The demodulator is performed by multiplying the output of the ADC with the carrier sample which was used to generate the LED duty cycle for the integration period which just completed and then low-pass filtering the resultant signal. A 16x16 bit multiplier is sufficient to calculate the product of a 4 bit carrier and a 10 bit ADC output. The MSP430F2122 does not have a hardware multiplier so the multiplication is implemented in software using Horner's method. A software multiply takes much longer to execute than a

hardware multiply so neither prototype version 1 or 2 can be used to fairly determine the execution speed and hence the average power consumption of an embedded homodyne PPG sensor. Versions 1 and 2 of the system were only ever intended to aid research into the ITIA so a requirement for hardware multiplication was not evident when the PCBs were designed.

Desired Signal Bandwidth

The low-pass filter is required to remove the undesired high frequency components from equation 4.4 to recover the absorption signal. The shape of the waveform is important so it is crucial to use a filter which is flat in the pass-band such as the Butterworth, Bessel or Chebyshev type II filter structures. The choice of filter cut-off frequency is a trade-off between maximising the amount of PPG signal spectrum that is preserved by the detector and minimising the amount of Gaussian additive noise that the filter allows to remain in the signal. The bandwidth of a PPG signal is a physiological parameter and independent of the sensor system, but the noise is dependent on the design and hence will determine what the filter cot-off frequency is. To determine this it is necessary to inspect the DFT of a homodyned signal before it has been filtered as shown in figure 4.7. This shows a signal to noise ratio of 1 (or 0dB) at about 10Hz which indicates that it there would be no benefit in setting the filter cut-off frequency above 10Hz - even for a very high order filter. A low order filter will be restricted to having an even lower cut-off although even 5Hz would preserve the first two harmonics of the PPG signal shown in figure 4.7. Due to the non-stationary nature of the signals being captured restricting the signal bandwidth to 5Hz does not necessarily mean that there is no benefit to shorter sampling periods that 100ms. Sampling faster than the Nyquist frequency means that it is possible to extract higher resolution phase information of the measured signal without having to resort to powerful interpolation algorithms to recover the information. Described another way, the simplest way to extract heart rate from a PPG signal is to measure the peak-to-peak interval which can vary by a very small amount from pulse to pulse. This needs a high sampling rate to extract the phase variation precisely enough to differentiate between similar frequency - a matter of using high resolution in the time domain over a small range in the frequency domain.

Filter Order

To specify the stop-band it is necessary to consider the choice of carrier frequency with respect to frequencies of the ambient noise sources. The choice

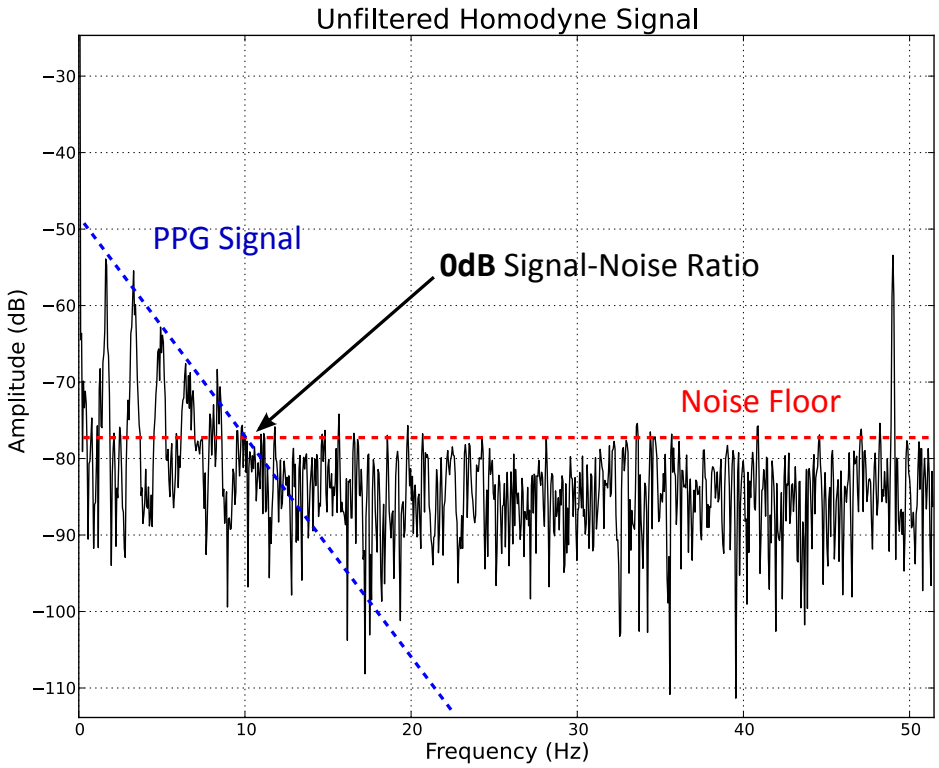


Figure 4.7: The spectrum on an unfiltered homodyne signal shows at what frequency the PPG signal intensity is equivalent in magnitude to the noise floor of the sensor circuitry. The data shown here was captured from a reflection mode PPG transducer worn at the tip of the author's index finger.

Table 4.1: Filter Specification for Homodyne Product Detector

Specification	**Value**
Pass-band	5Hz
Stop-band	45Hz
Maximum pass-band ripple	0.1dB
Minimum stop-band attenuation	96dB
Type	Chebyshev II
Required filter order	5

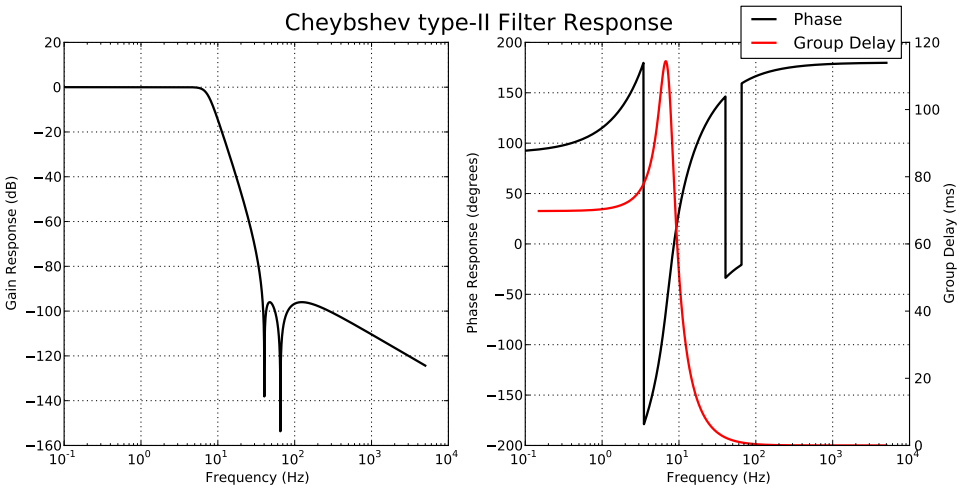


Figure 4.8: Gain, phase and group delay of the filter specified for the homodyne product detector.

of 250Hz for the carrier frequency places the nearest noise peaks at ±50Hz (200Hz and 300Hz). The ambient light also mixes with the absorption signal which pulls the edge of the noise band in to ±45Hz. It is not possible to accurately specify the required stop-band attenuation without knowing beforehand what the amplitude of the noise sources are with respect to amplitude of the received carrier. A cautious choice of 96dB of attenuation means that it is possible to have 16 bits of dynamic range between the carrier and any out-of-band noise that was received at the same amplitude as the carrier. In keeping with the minimalistic system design methodology the Chebyshev type II filter is chosen as it provides the steepest roll-off for a given order of filter.

A summary of the filter specification can be found in table 4.1 and the gain response is illustrated in figure 4.8.

4.2.6 Lattice Wave Digital Filter

The filter for the homodyne detector is implemented with a lattice wave digital filter (LWDF) structure instead of using either finite or infinite impulse response filters (FIR & IIR). LWDFs are attractive for embedded filtering as they are computationally efficient, are robust against rounding errors and recover quickly from saturation and wrapping errors [80, 81]. The underlying architecture of a low-pass LWDF is the calculation of the difference between two versions of the input signal, each delayed to a different degree with two-port, loss-less, all-pass filter elements. By working exclusively in the time domain it is possible to avoid the limitations of the bilinear transform often used when implementing an IIR filter from an analogue filter specification.

Each two port element in a wave digital filter, called an *adaptor*, is formally defined by the equations in 4.9 which specify the relationship between the incident waves (x_0 and x_1) and the output waves (y_0 and y_1). Each adaptor, indexed i, has only one coefficient, γ_i, which is bound between -1 and 1.

$$\begin{aligned} y_0 &= x_0 + \gamma_i(x_0 - x_1) \\ y_1 &= x_1 - \gamma_i(x_0 - x_1) \end{aligned} \tag{4.9}$$

The LWDF is constructed from two parallel chains of these adaptors combined with single sample delay elements as shown in figure 4.9.

Filter Design

To implement a low-pass filter the average of the output of the two chains is calculated; the difference of the outputs is the high-pass equivalent. Each adaptor has its own coefficient, γ_i, and the choice of each coefficient affects the filter's properties. The specific techniques required to design wave digital filters rely on network synthesis theory and is beyond the scope of this thesis but is covered in detail by Fettweis [81]. Instead set of equations provided by Gazsi is used in order to synthesise the Chebyshev filter. The same technique can be applied for the implementation of Butterworth and elliptic filters. For some reason, as of yet undetermined, Bessel filters cannot be implemented using the same procedure, but this limitation does not hinder the use of LWDFs for PPG homodyne sensing. To determine the γ_i values for Chebyshev II filter first determine the continuous-time transfer function of the filter in the Laplace domain. Details of how to calculate filter coefficients can be found in text books such as *Electronic Filter Design Handbook* by Williams and Taylor [82] or *Continuous and Discrete Time Signals and Systems* by Mandal and Asif [83]. The filter transfer function is generated from a characteristic polynomial

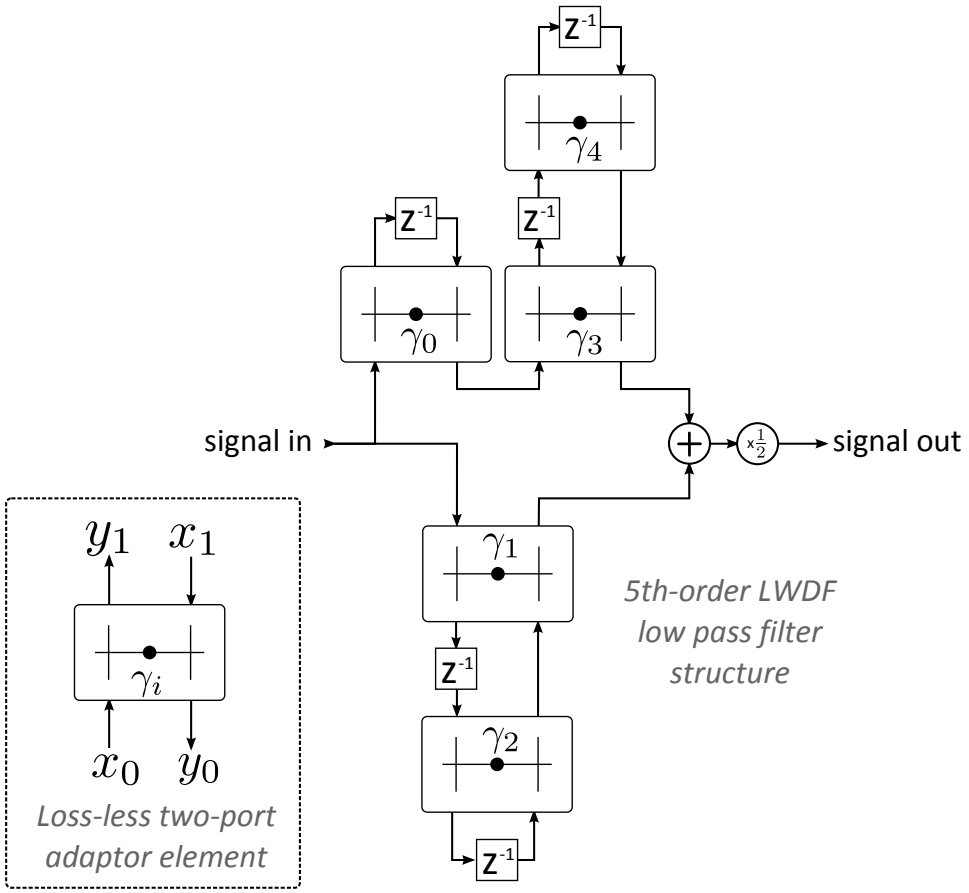


Figure 4.9: Block diagram of a 5th order low pass filter implemented with a LWDF. Included is an adaptor element to match equation 4.9. Single sample delay elements are labelled using z-domain nomenclature: z^{-1}.

in its denominator which is generated from the cut-off, or "-3dB frequency", of the filter which in radians is labelled ω_c. For a Butterworth filter this is the frequency at which the amplitude of the output is $\sqrt{2}$ of the input which is approximately 3dB. The resulting polynomial has roots, or poles, and a LWDF must always have one real pole and the remaining poles must be pairs of complex poles to retain symmetry about the real axis. This means that LWDFs are constrained to only ever have an odd number of poles. To use Gazsi's explicit formulas for LWDF synthesis it must be possible to represent the characteristic polynomial of an N order filters by the following equation:

$$C(s) = (s + B_0) \prod_{i=1}^{(N-1)/2} (s^2 + sA_i + B_i) \tag{4.10}$$

For a LWDF the polynomial needs first to be translated into the ψ domain. The conversion between the Laplace domain and ψ uses the following relationship (where T_s is the sampling period of the discrete-time system.):

$$\psi = \tanh\left(\frac{sT_s}{2}\right) \tag{4.11}$$

The Laplace parameter s can be set to $j\omega$ for frequency domain analysis so similarly ψ can be represented by the imaginary frequency $j\varphi$. Putting these into equation 4.11 give the following relationship between the cut-off φ_c in the ψ domain and the ω_c cut-off frequency in the Laplace domain:

$$\varphi_c = \tan\left(\frac{\omega_c T_s}{2}\right) \tag{4.12}$$

The characteristic polynomial of the odd order filter in the linear domain can now be converted directly to the ψ domain.

$$C(\psi) = (\psi + B_0) \prod_{i=1}^{(N-1)/2} (\psi^2 + \psi A_i + B_i) \tag{4.13}$$

Determining the values of B_i and A_i from the specified φ_c can be done using tools such as the MATLAB Signal Processing Toolbox or the SciPy signal processing library for use with Python to provide values for the system poles from which the desired constants can be calculated. Finally the γ coefficients for each adaptor can be calculated from the values of B_i and A_i. For the real pole the γ is:

$$\gamma_0 = \frac{1 - B_0}{1 + B_0} \tag{4.14}$$

As the rest of the poles are paired into complex conjugates their γ values can be calculated from the following equations:

$$\gamma_{2i-1} = \frac{A_i - B_i - 1}{A_i + B_i + 1} \tag{4.15}$$

$$\gamma_{2i} = \frac{1 - B_i}{1 + B_i} \tag{4.16}$$

The LWDF essentially assigns each system pole its own processing element whereas an IIR implementation is a purely mathematical implementation of the polynomial where each coefficient is dependent on multiple pole values as well as the accuracy of the transform used to translate between the continuous-time Laplace domain and the discrete-time z-domain during filter design. This even distribution of poles throughout the structure of the LWDF is what lowers its sensitivity to truncation errors in coefficient values and makes it resilient against non-linear events such as internal saturation and limit cycling which are always concerns in embedded DSP when floating point types are not available.

Filter Implementation

Referring back to the equations in 4.9 shows that each adaptor consists of three additions and a single multiplication by a constant in the range of -1 to +1. This means that only one multiplication is required per filter order whereas the direct-form IIR filter structures need $2N+1$ multiplies for an N^{th} order filter. Each LWDF adaptor element requires a single memory register to implement the sample delay which matches the direct-form II IIR filter structure. The cost of the LWDF adaptor is the requirement for 3 additions per filter order and a final addition for the combinations of the two signal paths as shown in figure 4.9 giving $3N + 1$ additions for an N^{th} order filter compared to $2N$ for the direct-form II. However, a multiplication procedure is much more complicated than an addition and either needs more instruction cycles or more transistors, and costs more energy to execute than an addition so the LWDF is still the more attractive solution for implementing embedded filters.

To further simplify the multiplication process it is common to split the -1 to +1 range into four quadrants with each quadrant having its own individual structure. Requiring four different adaptor structures may require more programme memory in software or more transistors in a physical implementation, but the advantage is that the multiplication coefficient is always in the range of 0 to 0.5 which reduces the number of bits required to represent γ for a

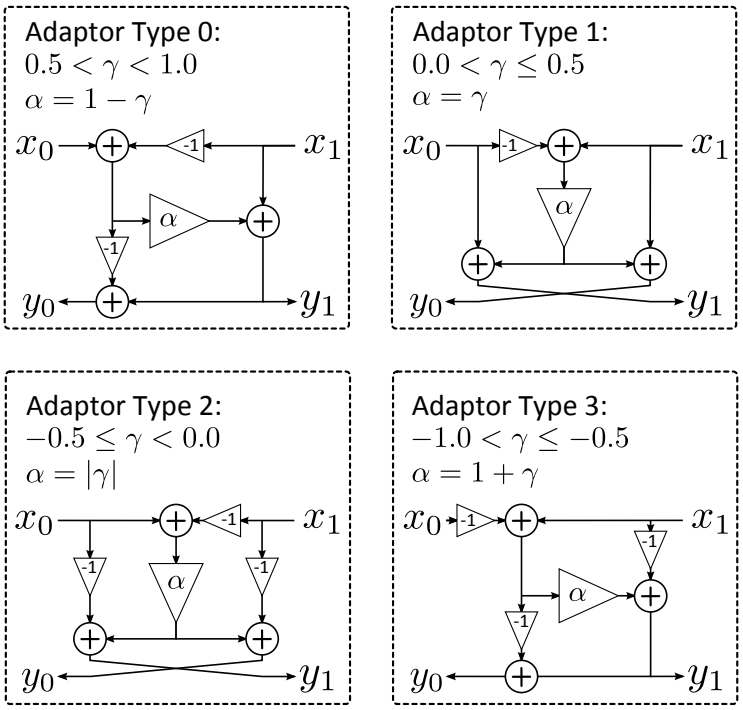


Figure 4.10: Two port adaptor types for lattice wave digital filters for compressed multiplication coefficient range.

given quantisation error by 2 which increases the possible dynamic range of the process. Also, the multiplication only needs to operate over two quadrants which reduces the number of transistors required to implement the processing element. These four adaptor configurations and their respective ranges of γ are shown in figure 4.10.

Implementation of any signal processing algorithm in fixed-point arithmetic needs carefully defined variable resolutions to ensure stability is maintained, limit-cycling is avoided, dynamic range is maximised and the underlying arithmetic processes are matched to the CPU architecture. The MSP430 family of MCUs used in the first three versions of the PPG system have native 16 bit data widths so all additions and multiplications are targeted at integer multiples of 16. Therefore the coefficients are all rounded to 16 bit unsigned integers and the filter input is rounded to 16 bit signed integers. The nature of the homodyne sensing system is to oversample with respect to the highest PPG biosignal frequency component so it is possible to achieve higher than 16 bit signal to noise ratios at the output of the filter therefore the output is 32 bits wide and all internal registers are also 32 bits wide. Restrictively specifying data types for all constants and variables aids in implementation of the filter in software. The C programming language is the primary language

for implementing the LWDF as the same code base can be compiled on a PC for prototyping and debugging the filters as well as being compiled for downloading onto an MCU for final implementation. Using Texas Instruments' Code Composer Studio version 4.1.0 integrated development environment the 5th order Chebyshev type II filter executes in a maximum of 3804 instruction cycles on the MSP430F2122. If the core clock frequency of the MCU is 8MHz then this converts to a processing time of 475.5μs which takes up 95.1% of the sample period which means it can operate in real time at the required 2kHz sampling frequency. As the MSP430F2122 does not have a hardware multiplier the 95.1% is much higher than is suitable for a low power system where the processor spends the majority of its time in a sleep mode. By moving to a processor which has a hardware multiplier, e.g. the MSP430F1611 used in the BSN modules [75], reduces the execution time to less than 2000 instruction cycles. Moving to a native 32 bit processor, e.g. the Energy Tiny Gecko MCU which has an ARM Cortex-M3 core, further reduces the execution time to 405 cycles. At 3.3V and 8MHz the MSP430F2122 core is drawing about 2.8mA which means that the 5th order LWDF consumes 4.39μJ per execution. The Tiny Gecko running at 14MHz would take 28.9μs to execute the filter which at a current draw of 2.1mA consumes only 0.2μJ - over a factor of 20 reduction from the MSP430F2122.

LWDF Validation

To validate the LWDF filter structure it has been implemented in the C programming language library and is compared against the IIR filter library available in the Python SciPy signal library. A Python version of the LWDF is included in Appendix B on page 208 to demonstrate how the filter is constructed and used - the C language version is less instructive. The SciPy based IIR filter uses floating-point data so can be used as a reference to observe the effects of truncation error on the integer based LWDF. Figure 4.11 shows the frequency domain response of the LWDF specified in table 4.1 as well as the equivalent IIR response. As both filters are implemented in the time domain it is necessary to use the impulse response of each filter to allow frequency domain analysis. Note that the as LWDF has an effective gain of 2^{16} to allow oversampling to provide extra bits of information. For comparison the LWDF output has been divided by 2^{16} to allow direct comparison with the IIR impulse response. The top plot in figure 4.11 shows that the LWDF filter is a good match to the floating point IIR filter until the the attenuation exceeds 60dB. This mismatch is of no concern as the important metric is that the minimum attenuation beyond 45Hz is greater than 96dB and the LWDF

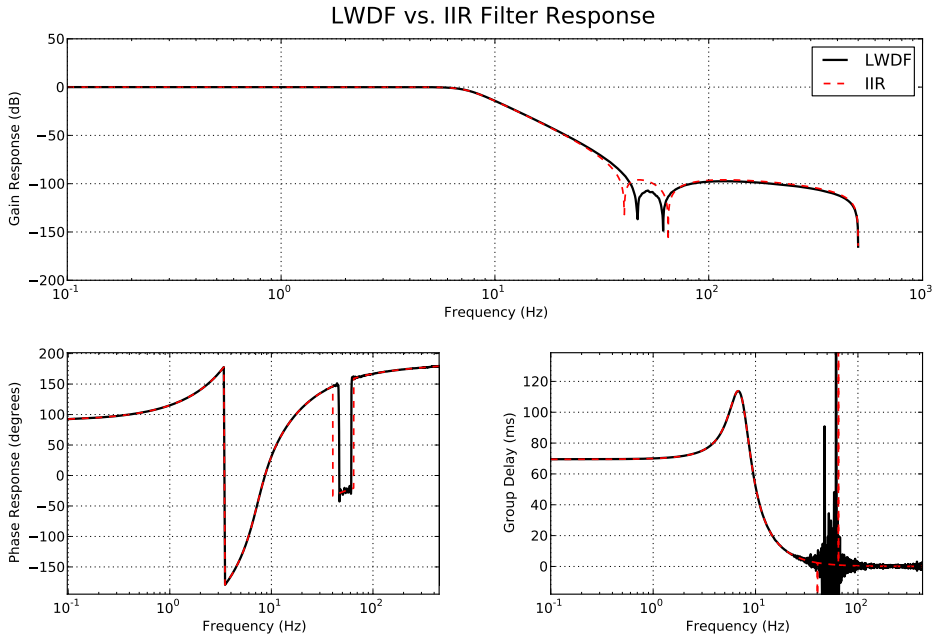


Figure 4.11: Comparison of IIR filter and fixed-point LWDF in the frequency domain. This data is generated by taking the discrete Fourier transform of each filter's impulse response. The noise in the group delay plot is due to quantisation error.

in fact achieves this all the way up to the Nyquist frequency. The phase and group delay plots are included to show that the LWDF is an excellent match to the floating point IIR filter up to the first phase nulling frequency by which point the signal attenuation is exceeding 80dB. There is some noise in the group delay plot for the LWDF which is due to quantisation error in the denominator of the phase differentiation.

Figure 4.12 compares the LWDF against the SciPy IIR filter when applied to a real PPG signal captured from a reflection mode PPG transducer worn on the author's fingertip. This 30 second data set is used throughout the remainder of this chapter. The data used for this analysis was captured by prototype version 2 at 1kHz where no detection algorithms were implemented on node in order to allow different detection algorithms to be compared off-line. Figure 4.13 shows the effect of truncating the LWDF output to 16 bits. As can be seen in the error plot the truncation is always within ± 0.5 least significant bits of the floating point based IIR filter. Retaining this fixed-point arithmetic accuracy throughout the truncation process is an important feature if closed-loop control were to be implemented across the analogue-digital divide as any systematic offsets can lead to instability if integrators are used anywhere in the feedback path. The need for this requirement will become clear in the

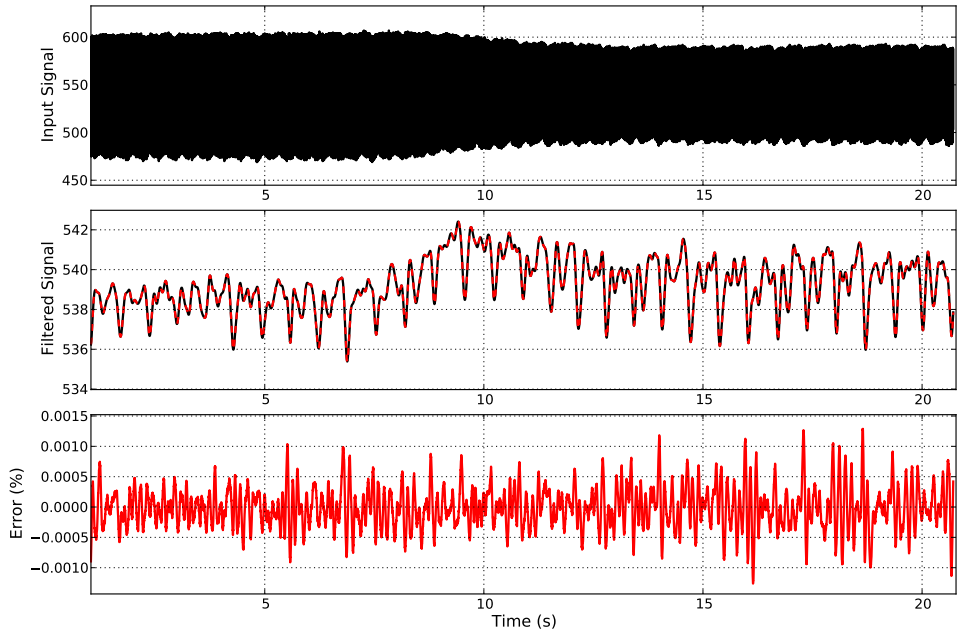


Figure 4.12: Comparison of floating-point IIR filter and fixed-point LWDF on a data set captured from a reflection mode PPG transducer on the author's finger tip. The top plot shows the raw photocurrent data (as output from the MCU ADC) corrupted by intense ambient noise as shown in figure 4.3. The middle plot is a PPG biosignal recovered from the baseband using both a LWDF and an IIR filter. The bottom plot is the difference between the two filter outputs.

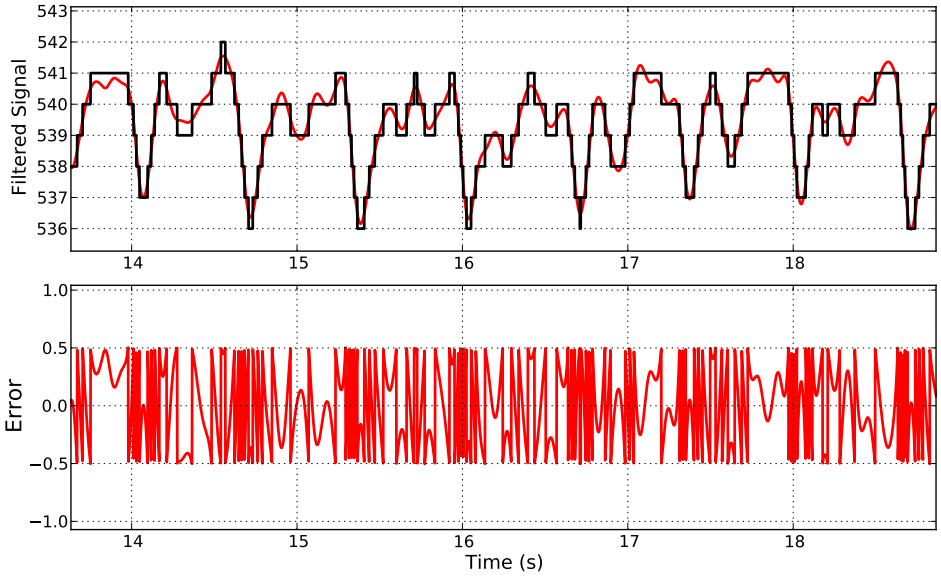


Figure 4.13: When truncated to 16 bits the LWDF is still bound to within ±0.5LSB of the IIR filter.

next chapter.

4.3 Mixed-Signal Homodyne Sensing Validation

4.3.1 Envelope vs. Product Detection

Prototype version 2 was used to capture a short data set to prove the validity of applying homodyne sensing to wearable PPG. Three signal conditioning processes are used on the same set of data:

Baseband: The low-pass filter specified in table 4.1 is used to extract the original low frequency PPG signal and reject higher frequency noise and carrier optical components. This is the conventional method used for PPG.

Envelope Detection: The LED is modulated with a 250Hz carrier and envelope detection is used to extract the PPG signal image centred around the carrier. The envelope detector is implemented with a moving standard deviation calculation where the window width of the mean calculations is the same length as the carrier in samples in order to block the carrier from appearing in the output. This simple FIR is not enough to block ambient noise harmonics so a 5th order low pass filter, identical to that used in product detection, is required.

Product Detection: Product detection is used to implement the homodyne architecture with LWDFs used to provide improved noise rejection over the mean used in the envelope detection process.

Figure 4.14 shows how these three methods compare when used on the same data set. Note the the envelope and product detection signals have been scaled according to equations 4.8 and 4.4 such their outputs should ideally have the same amplitude.

The comparison of the baseband output to the homodyne outputs shows that there is a significant visible difference in both the DC level of the waveform and the amount of noise present in the output signal. A typical PPG waveform is visible in both homodyne detector outputs. The only visible difference between the envelope detection and product detection methods is a small increase in offset in the envelope output. This is likely to be due to the additional simplification made in order to achieve equation 4.8 in its compact form - i.e. not all signal terms created by the signal squaring will have been successfully rejected or cancelled out. The baseband signal is higher than the homodyned signals due to the presence of systematic current and voltage

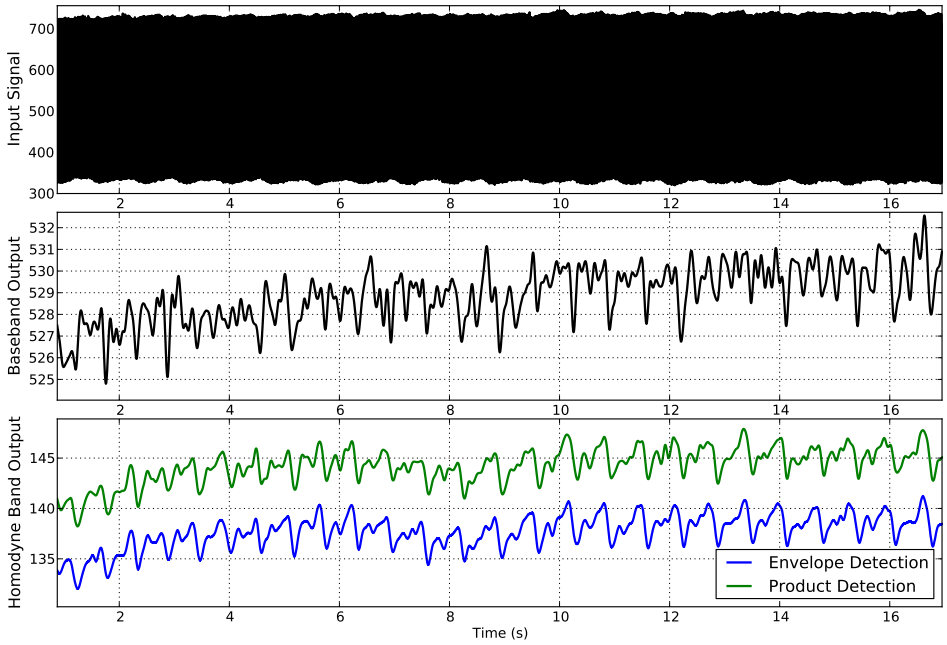


Figure 4.14: Product detection compared against envelope detection for implementing a homodyne detector. The raw input signal and the conventional baseband signal output are included for comparison.

offsets - including the bias voltage placed across the photodiode to put it in photoconductive mode. Correlated double sampling is the more conventional method of removing this offset. It is possible to approximate CDS removal of offset using the same data set used for validating homodyne detection. This is because the carrier has at least one sample where the LED is turned off which can be used to subtract offset from the other 3 samples in the carrier sequence.

4.3.2 Homodyne vs. CDS

Figure 4.15 compares the output of this pseudo-CDS process and the homodyne PD process as shown before. The signals are so similar in amplitude that it is necessary to separately plot the difference between them. These results show that homodyne detection is a match to CDS to within 0.15%. As CDS is simpler to implement than homodyne detection (one add versus one multiply respectively) it would appear that CDS is the more suitable method to use. However, CDS does not scale as well as homodyne detection does when extra optical channels are required because each channel also requires its own CDS channel doubling the number of TDM channels.

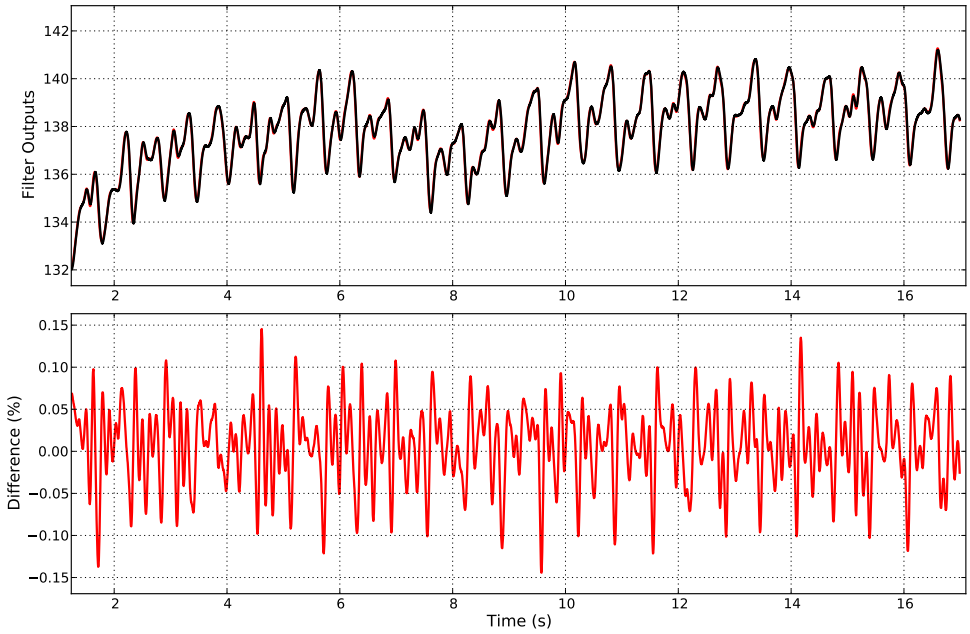


Figure 4.15: Comparison of homodyne sensing against correlated double sampling. As the signal amplitudes are so similar the difference between the two signals is plotted underneath.

4.3.3 Reduced Filter Order

The filter specified earlier in this chapter is specified for the absolute worst case ambient noise levels. No other synchronous detection system use a filter as high order as the one proposed here which raises the question: is a 5th order filter really needed?

If it can be assumed that the ambient noise within 50Hz of the carrier frequency is lower in magnitude than the detected PPG biosignal intensity then it is possible to specify a filter that only needs to reject the baseband signal which is 250Hz away from the carrier. As such a Butterworth filter, where the attenuation approaches infinity as frequency increases, would be a better choice than the Chebyshev type II used earlier which has finite stopband attenuation.

This is easily investigated by trying to implement the product detector with less highly specified filters. All the data presented earlier in this chapter was captured directly from the output of the MCU's ADC so it is possible to prototype filters in post processing. Figure 4.16 shows a zoomed in section of the dataset used previously in this chapter where the 5th order filter is swapped for a 3rd order Butterworth filter with a 5Hz cut-off frequency - the dotted line shows the output from the filter specified earlier in this chapter.

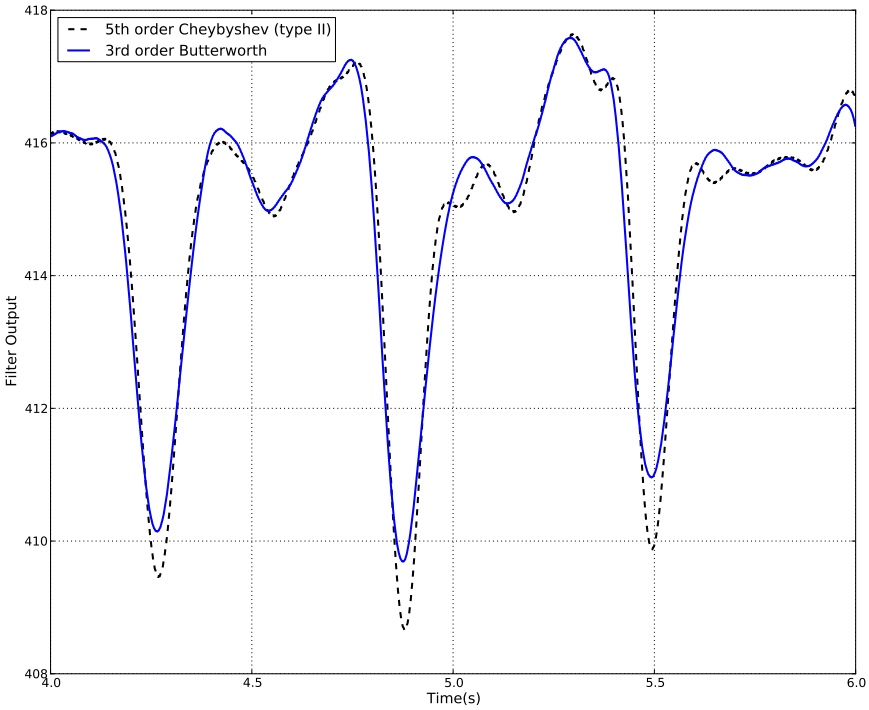


Figure 4.16: Product detection implemented with a 3rd order Butterworth filter (solid blue) instead of the 5th order Cheybshev (dashed black). This shows that the higher order filter is not necessarily required for sensing applications where ambient light does not dominate at higher frequencies.

Being able to reduce the order of the filter by two means two less adaptors are required to implement the LWDF which could, in theory reduce the required processing time, and hence power consumption, of the product detector by 40%.

4.4 Conclusions

This chapter applied the theory of signal mixing to the application of detecting PPG waveforms in a background of both ambient noise and systematic offsets. Synchronising the photodetection sampling to the drive of the illumination circuitry allows it to lock on to the desired PPG optical absorption waveform while significantly attenuating any signals from uncorrelated light sources. Synchronous detection is a common solution to the problem of additive light, but the proposed homodyne method is particularly efficient in its implementation and hence makes it an ideal solution for wearable sensing applications. Again, by considering the PPG system as a whole it has been

possible to choose a noise rejection technique that makes the most of how the LEDs are already being driven and how the photocurrent is detected.

Homodyne sensing is capable of matching the performance of correlated double sampling in rejecting systematic offsets and ambient light, but is more flexible in its implementation and can be scaled more easily to higher numbers of optical channels as there is no need for an extra ambient light measurement sample in the TDM scheme. Nor is it required to have an ambient measurement sample adjacent to each LED channel sample in order to reduce phase difference errors. It has also been shown that it is possible to implement homodyne detection without any additional circuitry in the detection chain if an integrating transimpedance amplifier is used. Implementing all of the signal processing in the digital domain means that the algorithm can be added to the firmware of an embedded MCU that all wearable PPG systems rely on already so no extra electronic components are required.

Of the two detection methods product detection, the computationally simplest of the two, has been shown to be the most accurate and comparable to the performance of CDS within 0.15%. Careful choice of the carrier frequency in the context of the likely operating environment of a wearable PPG system has made it possible to synthesise the carrier signal with a 4 sample LUT which has 0, +1 and -1 as the only multiplication coefficients of the detection mixed and makes it possible to implement complex demodulation if it were desirable to do so.

This chapter has also shown the necessity of a having a high performance low pass filter to remove ambient noise signal components which are present throughout the entire optical signal spectrum - whether using homodyne sensing or CDS. Lattice wave digital filters have been chosen as a highly efficient and robust architecture for implementing linear signal filters using fixed point arithmetic on resource limited MCUs for embedded systems. The performance of the LWDF has been shown to be a match to that of a conventional IIR filter implemented using floating point arithmetic to within 0.002% on a sample by sample basis in the time domain. However, when it is known that the ambient light at higher frequencies is low in intensity it is possible to reduce filter performance in order to reduce the processing requirements of the sensor's MCU and hence the overall power consumption.

Synchronous detection techniques are still vulnerable to failure when the system is exposed to input signals large enough in magnitude to saturate the front end amplifier. Another weakness is that only noise that is additive in nature can be rejected by homodyne detection. The next chapter investigates solutions to both of these shortcomings.

5 Multi-mode Artefact Rejection

Up until this point the thesis has only considered noise from either minimising electrical noise in the instrumentation circuitry or from light sources independent of the sensor system. In the domain of photoplethysmography there are another two significant contributors to noise: high brightness independent light and motion induced artefacts. The term *artefacts* is used for these noise sources as they are non-stationary in nature - just like the biosignal being measured. In order to address these issues a model for how these artefacts couple in to the sensing system needs to be developed. This chapter presents theory and results of which excerpts have been published in [8] and in [9].

The light received by the photodetector is separated into two classes:

Dependent Light: *DL*, all incident photons originated from the sensor's LED.

Independent Light: *IL*, no incident photons originated from the LED.

As IL is irrelevant to the biosignal it is desirable to remove it completely. Equation 5.1, which was originally posed in chapter 2 as the signal-chain model of a PPG sensing system, shows how DL and IL are linearly separable.

$$y_P(\lambda,t) = A_P S_P(\lambda)\left(\underbrace{e^{\overbrace{-A(\lambda,t)}^{\text{absorption signal}}}\overbrace{S_L(\lambda)A_L x_L(t)}^{\text{LED signal}}}_{\text{Dependent Light}} + \underbrace{\sum_{\forall j}\left[e^{-B_j(\lambda,t)}I_{N,j}(\lambda,t)\right]}_{\text{Independent Light}}\right) \tag{5.1}$$

This gives us the first artefact classification: *Additive Artefacts* are any form of noise detected by the sensor where the light source is independent from the sensor's LED.

5.1 Additive Artefacts

The previous chapter introduced the concept of homodyne sensing for the linear separation of the absorption signal from IL sources in the frequency domain. The main limitation of this technique is that homodyne detection is bounded by the dynamic range of the photodetector's amplifier. This presents

a problem for a wearable PPG sensor as to keep power consumption low it is necessary to minimise the LED current and maximise the gain of the photodetector which makes the amplifier prone to saturating when exposed to large photocurrent induced by ambient light sources, such as daylight, which can be orders of magnitude greater in illumination intensity than the sensor's LED.

In applications where the bandwidth of the IL that causes photodetector saturation is lower than that of the DL being detected it is possible to use a low-pass filter and closed-loop control to stabilise the mean current input to the transimpedance amplifier such that its full dynamic range is made available for the dynamic DL signal - from now on described as *DC stabilisation*. This is relatively common in PPG applications as the heart rate can be derived from the dynamic portion of the absorption signal, but is not necessarily applicable to pulse oximetry as the DC component of the absorption signal is required to normalise the dynamic component. When applying DC stabilisation to PPG it is desirable to set the cut-off frequency of the control loop very low such that it does not reject low bandwidth signal content such as respiration and vasomodulation. Such high time constants (>1s) are not trivial to implement in analogue circuitry [84, 1] or in the digital domain require a high memory to sampling rate ratio, and in both cases the control system output takes a long time to settle from ambient light level disturbances. When considered in the context of a wearable sensor solution, it is unlikely that there will be suitable separation of the biosignal from the IL sources in the frequency domain for DC stabilisation to provide robust protection against high dynamic range IL artefacts.

5.1.1 DC Current Stabilisation

This shortcoming can be solved by combining the closed-loop current control with homodyne sensing. Shifting the biosignal up to a higher frequency band means that the remaining baseband signal can be used exclusively for keeping the mean current through the transimpedance amplifier at a fixed operating point such that the rest of the dynamic range can be used by homodyne detection to remove the remaining IL noise components. Additionally, the low-pass filter used in the DC stabilisation control loop only needs to be first order if the carrier frequency is chosen such that there is significant distance between the IL bandwidth and homodyne passband in the frequency domain. Restricting the feedback network to a first order system minimises the execution time of the filter and makes designing the system for stability trivial. Figure 5.1 shows how homodyne detection is combined with DC stabilisation to create a dual-mode additive artefact rejection system.

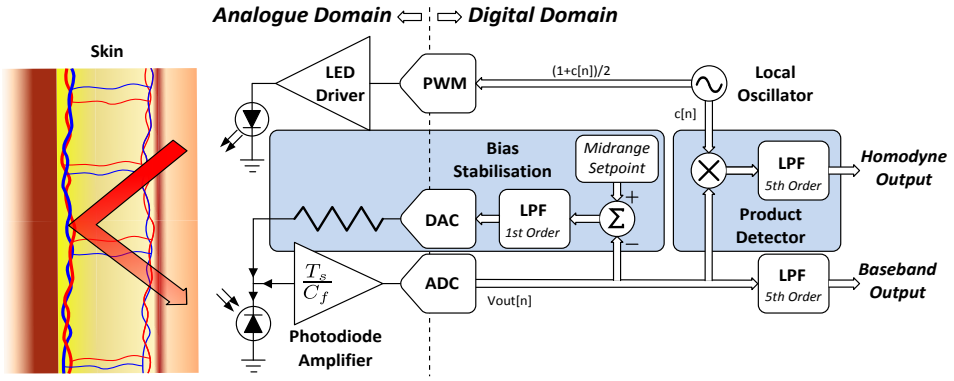


Figure 5.1: A single channel PPG homodyne detection system with DC stabilisation of the input current.

5.1.2 Continuous Time Analysis

Figure 5.2 shows the analogue section of the DC stabilised photodetector. Note that a pair of voltage mode DACs provide both voltage bias of the photodiode as well as implementing the current mode negative feedback of the mixed-signal control loop. The difference in this case is that it integrates the difference between the input photocurrent and the feedback current. Note that an active anti-aliasing filter (AAF) is included to separate the integration capacitor from the input capacitance of the ADC.

$$v_{out}(t) = \frac{1}{C_F}\int_0^t i_P(\tau) - i_{FB}(\tau)\,d\tau - v_{B0}(t) \tag{5.2}$$

Transforming this into the Laplace domain simplifies the design of the feedback loop. Substituting in the variables representing the bias voltage and current gives the following equation:

$$V_{out}(s) = \frac{1}{sC_F}\left(I_P(s) - \frac{V_{B1}(s) - V_{B0}}{R_{FB}}\right) + V_{B0} \tag{5.3}$$

Note that V_{B0} is intended to be constant for the lifetime of any given sensing application. The feedback network is modelled in the analogue domain and has to take in to account the feedback resistance required to translate from the voltage-mode output of the DACs to the current-mode input of the photodetector. The part of the control loop performed by the MCU is modelled with the following transfer function:

$$V_{B1}(s) = A_{DAC}\frac{A_{FB}}{1 + \frac{s}{\omega_{FB}}}(sT_sA_{ADC}V_{out}(s) - SP) \tag{5.4}$$

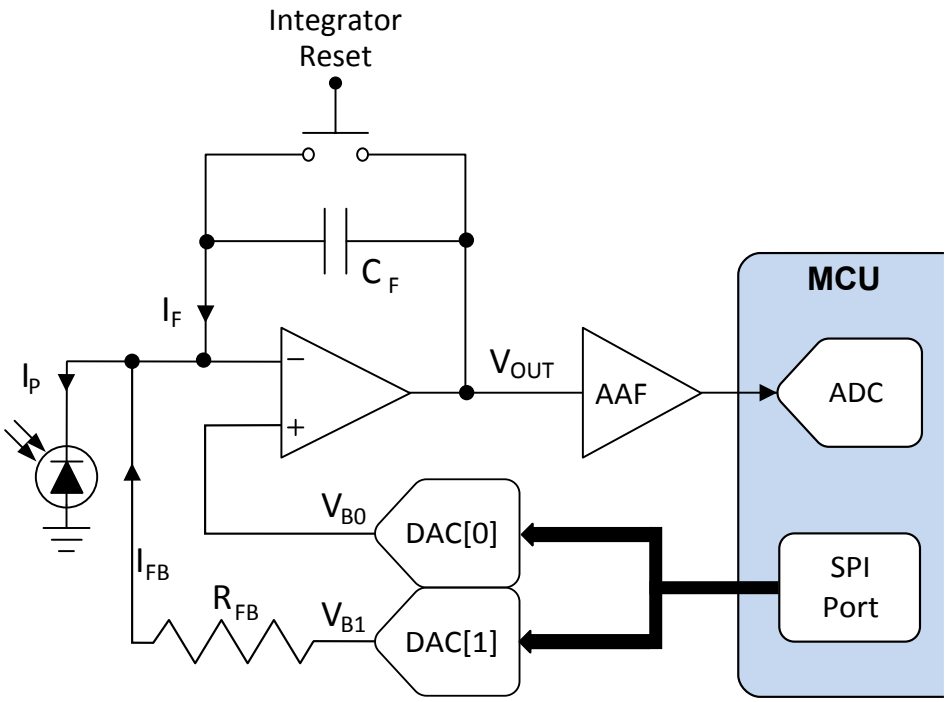


Figure 5.2: Integrating transimpedance amplifier with mixed-signal input current and voltage biasing..

To preserve the rank of the system $v_{out}(t)$ has to be differentiated over the sampling period, T_s. This equation includes terms for the gains of both the ADC and the DAC (A_{ADC} and A_{DAC}), programmable gain and cut-off frequency of the feedback LWDF (A_{FB} and ω_{FB}) as well as the desired steady state current of system (SP). The steady state value is always chosen as the mid point between zero and full scale on the ADC to ensure that dynamic range is maximised for the carrier when homodyne detection is enabled. The full system response is given below.

$$\begin{aligned} &sC_F\left(\frac{s}{\omega_{FB}}+\frac{T_s}{C_F}\frac{A_{DAC}A_{FB}A_{ADC}}{R_{FB}}+1\right)V_{out}(s) \\ &=I_P(s)\left(1+\frac{s}{\omega_{FB}}\right)+\frac{A_{DAC}A_{FB}SP}{R_{FB}}+V_{B0}\left(sC_F+\frac{1}{R_{FB}}\right)\left(1+\frac{s}{\omega_{FB}}\right) \end{aligned} \tag{5.5}$$

5.1.3 Differentiated Solution

The resettable ITIA described in chapter 3 essentially performs the required differentiation on the output of the integrator as well as translating the design in to the discrete-time domain. It is possible to model the amplifier with a

fixed gain of T_s/C_F instead of the original integrator relationship. Dynamic analysis of discrete time systems is conventionally performed in the Z-domain, but as the operating frequency range of the control loop is so much lower than the sampling frequency the remaining analysis is still performed in the Laplace domain. This gives the following approximated system response:

$$\begin{aligned}&\frac{C_F}{T_s}\left(\frac{s}{\omega_{FB}}+\frac{T_s}{C_F}\frac{A_{DAC}A_{FB}A_{ADC}}{R_{FB}}+1\right)V_{out}(s)\\&\equiv I_P(s)\left(1+\frac{s}{\omega_{FB}}\right)+\frac{A_{DAC}A_{FB}SP}{R_{FB}}+V_{B0}\left(sC_F+\frac{1}{R_{FB}}\right)\left(1+\frac{s}{\omega_{FB}}\right)\end{aligned} \tag{5.6}$$

The point for this system is to ensure that the gain of the system is small at low frequencies so it is necessary to check the steady-state response by setting $s = 0$. To simplify the analysis the gain needs to be large enough to satisfy the following condition:

$$\frac{C_F}{T_s} \ll \frac{A_{DAC}A_{FB}A_{ADC}}{R_{FB}} \tag{5.7}$$

Then the steady-state response at the output of the ADC can be approximated as follows:

$$A_{ADC}V_{out}(0) \approx \frac{R_{FB}}{A_{DAC}A_{FB}}I_P(0) + SP + \frac{V_{B0}}{A_{DAC}A_{FB}} \tag{5.8}$$

To further simplify the analysis V_{B0} and SP are both set to zero as only the response of the system in the passband is now of interest now that the steady-state response is robust against saturation.

$$V_{out}(s) \approx \frac{T_s}{C_F}\left[\frac{s+\omega_{FB}}{s+\omega_{FB}\left(\frac{T_s}{C_F}\frac{A_{DAC}A_{FB}A_{ADC}}{R_{FB}}+1\right)}\right]I_P(s) \tag{5.9}$$

It can be seen from equation 5.9 that the system is now has a high-pass filter response with the system zero located at $f_z = 2\pi\omega_{FB}$ and the start of the passband is set by the pole located at:

$$f_p = \frac{\omega_{FB}}{2\pi}\left(\frac{T_s}{C_F}\frac{A_{DAC}A_{FB}A_{ADC}}{R_{FB}}+1\right) \tag{5.10}$$

Note that it is necessary to ensure that this frequency is low enough to ensure that there is minimal attenuation of the carrier and surrounding biosignal image in the homodyne passband.

The DAC has a limited voltage range and hence sets the boundary on the

Table 5.1: System Design Parameters

Supply Voltage	3.0V
Sampling Period, T_s	1.0ms
Integrator Capacitance, C_F	100pF
Reset Period, T_{rst}	62μs
Transimpedance Gain	4.38x10^{6}V/A
Saturating Input Current	**685nA**
DAC Output Range	1.25V
Feedback Resistance, R_{FB}	330kΩ
Photodiode Bias Voltage, V_B	0.1V
Available Subtraction Current	**3.48μA**
ADC gain, A_{ADC}	2^{10}/3V
DAC gain, A_{DAC}	1.25V/2^{12}
Filter gain, A_{FB}	2^7
Filter cut-off frequency, $2\pi\omega_{FB}$	0.25Hz
System pole frequency, f_p	**101Hz**
DC attenuation	**52.1dB**

maximum magnitude of current that can be subtracted from the input in conjunction with the choice of R_{FB}. A further limitation on the feedback current range is the value chosen for the bias voltage. Typically V_{B0} is chosen to be as close to 0V as possible as under ambient lighting conditions it is generally only necessary to inject current in to the input node. However, in low light conditions it is often necessary to draw some current from the ITIA input node as systematic circuit offsets can behave as "negative" light. These constraints suggest choosing a low value for R_{FB}, but the feedback resistor contributes to the input noise current with a spectral density of $\sqrt{4kT/R_{FB}}$. For the purposes of minimising current, R_{FB} should be chosen as the maximum possible value that satisfies the condition in equation 5.7 and achieves the necessary maximum current subtraction level. Without *a priori* knowledge of what the maximum light level is going to be an alternative approach is to choose a current limit that represents a multiple of what current would saturate the ITIA without using any feedback. Table 5.1 lists a set of design parameters that enables the input current stabilisation system to reject low frequency light up to 5 times larger than the expected incident light levels from the sensor's LED. These are the values that were used with prototype version 2 for generating the results presented in [9].

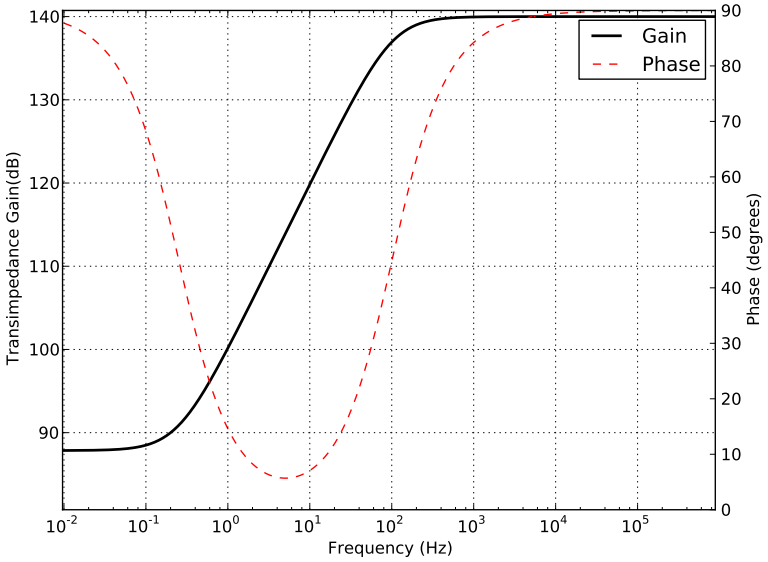


Figure 5.3: Frequency response of the DC stabilised photodetector. Passband start frequency tuned to 101Hz. DC gain is 52.1dB less than pass band gain.

The frequency response of this circuit is presented in figure 5.3. This configuration has up to 52dB of DC current rejection which is set by the distance between f_z and f_p; a first order system has a roll-off of 20dB per decade. To increase the attenuation level at DC requires lowering f_z (via the internal LWDF cut-off frequency) and keeping f_p at the same frequency by increasing the feedback gain, A_{FB}. Note that there is a limit to how high A_{FB} can be set before the integer limit of the DAC's input is reached. In the case that it is observed that the DAC is regularly saturating (i.e. not enough current is available to compensate for the input photocurrent) it is necessary to reduce the value of R_{FB} and accept an increase in current noise.

5.1.4 Validation

Both the text and the results presented in this section are taken directly from "Dual-Mode Additive Noise Rejection in Wearable Photoplethysmography", presented at BSN2012 [9]. In the case that mistakes were found in the original publication they have been corrected here.

To test the ability of any given system configuration to cope with high DC input light levels an experiment was designed where a 400 candela white LED light source was used to illuminate the PPG sensor and the surrounding tissue in order to saturate the sensor's photodetection amplifier. The external light source was physically moved from side to side in order to show how the

photodetector saturates and how quickly it recovers from a step change in light level. Three different system configurations were tested with the same experiment:

1) No offset subtraction.

2) 10Hz offset current compensation

3) 100Hz offset current compensation

For each experiment both the baseband signal and the passband signal are plotted to show any differences between them. The transducer was a reflection mode device worn at the author's fingertip where is was easiest to accurately induce the high optical intensity offsets.

Figure 5.4 shows the output of the PPG system with the offset compensation functionality disabled. In this case the photodetector saturates quickly and neither the baseband or homodyne output have any biosignal content. With the low frequency control loop enabled the base-band signal still retains biosignal content and is only momentarily saturated by bright light while the feedback loop brings the average input current back down to the set level (figure 5.5). In this case both baseband and homodyne outputs experience some corruption although not as bad as without any feedback as it is only during the transitions. However, when the cut-off frequency of the control loop is increased to 100Hz it is able to respond fast enough such that the photodetection amplifier is never saturated. The homodyne output is now completely free of artefacts from high input light level and as such it does not matter that there is no longer any meaningful output from the baseband signal chain.

Figure 5.7 shows the discrete Fourier transform (DFT) of the output of the ADC with current compensation enabled. This proves that the DC stabilisation system does not interfere with spectrum of the homodyne detection system. The spectrum shows the carrier 72dB above the noise floor. Using the 5th order LWDF specified in chapter 4 provides a further 96dB of out-of-band rejection and the DC stabilisation circuit 52dB of rejection at low frequency including 14dB of rejection beyond the range of the transimpedance amplifier.

5.1.5 Implementation Details

The complete system compiled to 4kB of ROM on the MSP430F2122 MCU used in prototype version 2. The algorithm executes in <5700 clock cycles which at 8MHz is only enough to support a single channel at the required sample interval of 1ms. Even reducing the sample rate to the point that 2

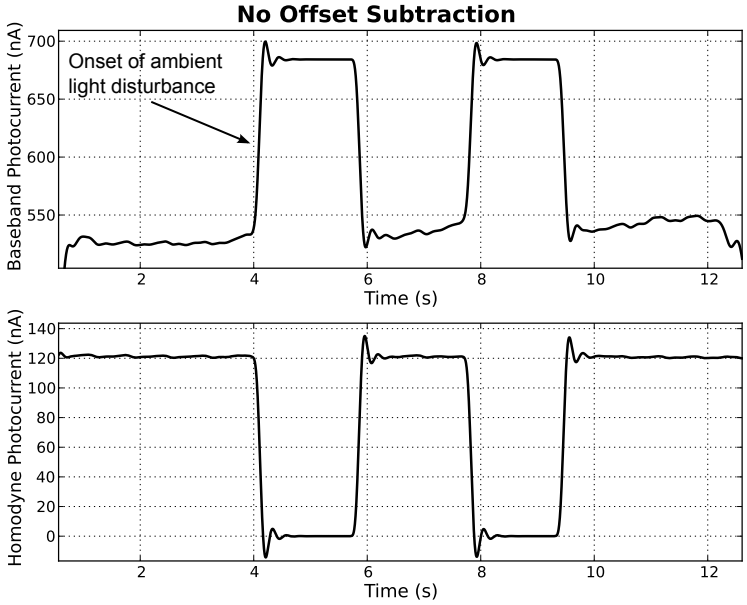


Figure 5.4: Without any input current stabilisation the detector saturates and no PPG biosignal is detectable in either the baseband or passband.

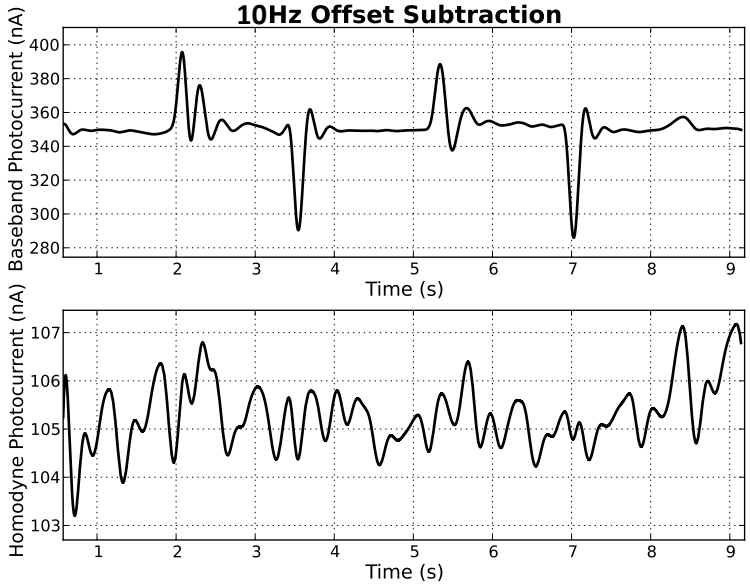


Figure 5.5: Enabling the current stabilisation control loop allows the system to operate with high ambient light levels. However, both baseband and homodyne outputs are corrupted while the input current level stabilises.



Figure 5.6: The high bandwidth input current compensation loop removes most PPG biosignal content from the baseband signal, but the homodyne output is now robust against ambient light disturbances.

channels could be supported would still mean that the MCU is operating 100% of the time which incurs an average current consumption of 21mA, including 2mA for average LED drive current.

The first step in reducing power consumption in the MCU is to increase the percentage of time that the processor core is in a sleep or power-down state. Processing time can be reduced by choosing an alternative MCU core which supports hardware multiplication which is not available on the MSP430F2122. Ported to the MSP430F1611 (the MCU used in the BSN development platform), which has a 16x16 bit multiplier, the algorithm executes in <2700 clock cycles which brings the full-power duty-cycle to down to 67.5% for a 2kHz throughput which still means the system consumes >10mA. The biggest gains are to be made by swapping to a native 32 bit processor core which allows for quicker execution of the LWDF which has been built using 32 bit internal registers. For example, porting the algorithm to an Energy Micro Tiny Gecko MCU (ARM Cortex-M3 core) reduces the execution time to 570 clock cycles and only requires 3.6mA during execution at 14MHz. In conjunction with using a low frequency oscillator to generate the sampling interval the estimated average core consumption falls to <1mA with an average system power consumption of 3mA.

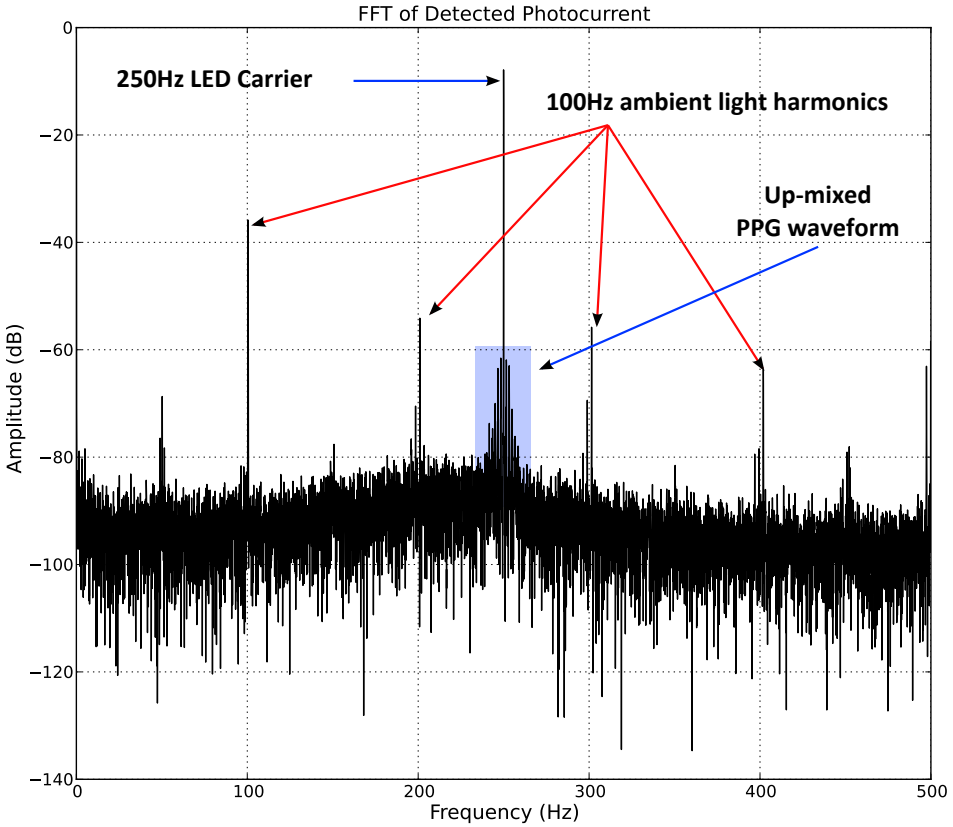


Figure 5.7: A DFT of the discrete-time photodetector output with offset compensation.

5.1.6 Optimised Current Stabilisation

This DC current stabilisation system has two aspects which can be improved upon. The first is the fact that a low-pass filter is used in the feedback path as opposed to using a pure integrator which is the more conventional approach. Using a low-pass filter makes the system zero positive which means the DC system gain at the ITIA output is still infinite which leaves a finite DC gain at the output of the differentiation stage. Using an integrator would place the zero at the origin leaving finite DC gain at the ITIA output and zero DC gain at the differentiation output. There is also the additional computational cost of executing the low-pass filter which requires multiplication whereas the integrator requires only additions.

The second aspect to be improved upon is even with using a pure integrator instead of a low-pass filter the finite gain at the ITIA output leaves it prone to saturation which is why the reset switch is required. To make the reset switch redundant requires making the DC gain zero at the ITIA output zero. This can be achieved by adding another integrator in the feedback path between the ITIA output (non-differentiated) and the input node. This feedback network was described back in chapter 3 as the *biased integrator* and is illustrated in figure 5.8. It was originally discarded as a possible solution to controlling the DC gain of an ITIA due to the loss of low frequency signal components which is no longer a concern thanks to the use of homodyne detection. The voltage at the output of the integrator with optimised feedback system is given by the following equation:

$$V_{out}(s) = \frac{s}{s^2 C_F + A_{FB}\frac{A_{DAC}A_{ADC}}{R_{FB}} + sT_s B_{FB}\frac{A_{DAC}A_{ADC}}{R_{FB}}} I_P(s) \tag{5.11}$$

The gain of the two feedback paths and the time-constant of the differentiation sets the frequency response of the system. Careful choice of these coefficients creates a system band-pass filter response which provides the equivalent DC signal attenuation as the reset-based integrator approach but without sacrificing any gain due to the duration of the reset period. This increase in system gain is at the expense of one additional multiplication and one addition which on an MCU with hardware multiplication can be performed in significantly less time than a safe reset period. In fact the response shown in figure 5.9 is achieved using only feedback coefficients which can be implemented as arithmetic bit shifts leaving only the low pass filter requiring multiplications. An improvement in phase response is achieved by making the A_{FB} coefficient negative. The narrow bandwidth of the 5th order low pass

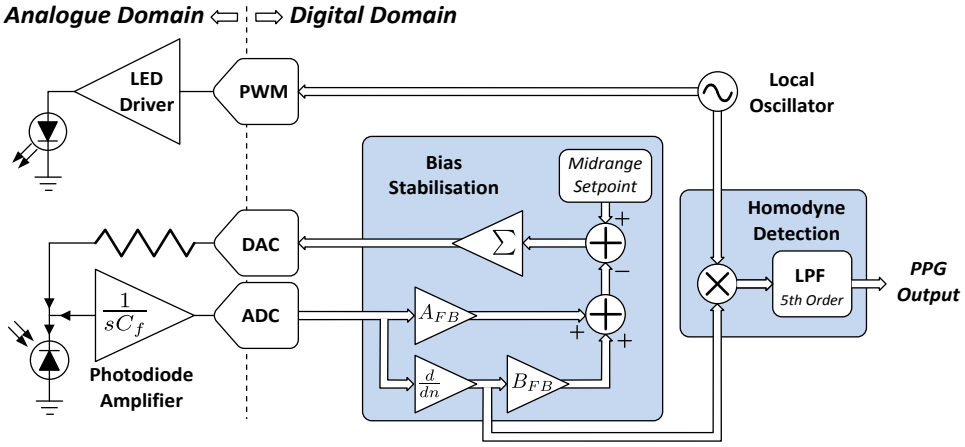


Figure 5.8: The optimised DC current rejection system

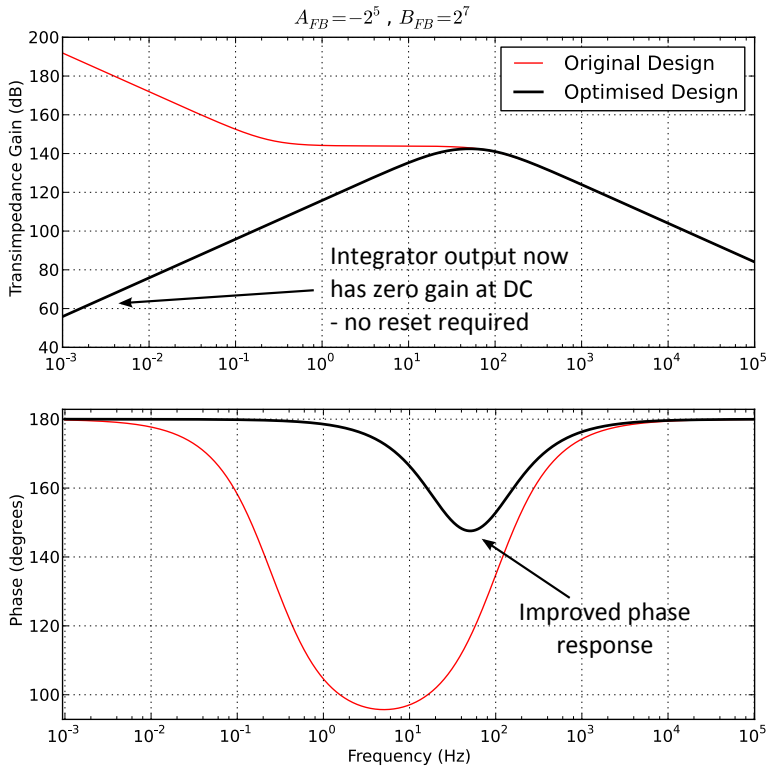


Figure 5.9: Frequency response of the optimised DC current rejection system

.

filter makes this phase response irrelevant, but were it possible to relax the filtering requirements the improvement in stability could be noticed.

5.2 Multiplicative Artefacts

5.2.1 Theory

The second artefact classification comes from any mechanical manipulation on the sensor, or on the tissue being observed by the sensor, causing variations in light path between the LED and the photodiode. Revisiting the transmission mode PPG model allows us to describe the absorption coefficient in term of the constituent components. Figure 5.10 illustrates the model and table 5.2 describes each label. The system is simplified into three physical volumes: arterial blood, venous blood and non-blood related tissues. Both arterial and venous blood volumes are constituted of varying concentrations of haemoglobin (Hb) and oxyhaemoglobin (HbO_2). Refer back to equations 2.5 and 2.6 in chapter 2 to see how the arterial oxygen saturation is derived from these coefficients. It should be reiterated here that this is only an approximate model used to investigate the presence of non-linear noise components - this research does not aim to comment on the accuracy of said model or the choice of wavelengths used in the experiments. Equation 5.12 describes the relationship between these coefficients and the accumulated absorption coefficient, $A(\lambda, t)$, in equation 5.1.

$$\begin{aligned} A(\lambda,t) = {} & d_A(t)[\varepsilon_{AO}(\lambda)c_{AO}(t) + \varepsilon_{AH}(\lambda)c_{AH}(t)] \\ & + d_V(t)[\varepsilon_{VO}(\lambda)c_{VO}(t) + \varepsilon_{VH}(\lambda)c_{VH}(t)] \\ & + d_T(t)\varepsilon_T(\lambda)c_T \end{aligned} \tag{5.12}$$

In this equation the variables of primary interest in a PPG system are $d_A(t)$ which is related to the arterial blood volume, $c_{AO}(t)$ and $c_{AH}(t)$ which allow the calculation of S_PO_2. The conventional method of removing the unwanted terms is to make the assumption that the only dynamic term in the model is $d_A(t)$. However, when a PPG sensor is worn by a physically active patient it is no longer safe to make this assumption as it is likely both $d_V(t)$ and $d_T(t)$ are dynamic. We can rearrange the absorption equation to focus on splitting the depth variable into static component and two dynamic components, one for variations that can be considered to be signals of interest and one for noise induced by motion of the subject; $\overline{d_i}$, $d_i^*(t)$ and $n_i(t)$ respectively for an

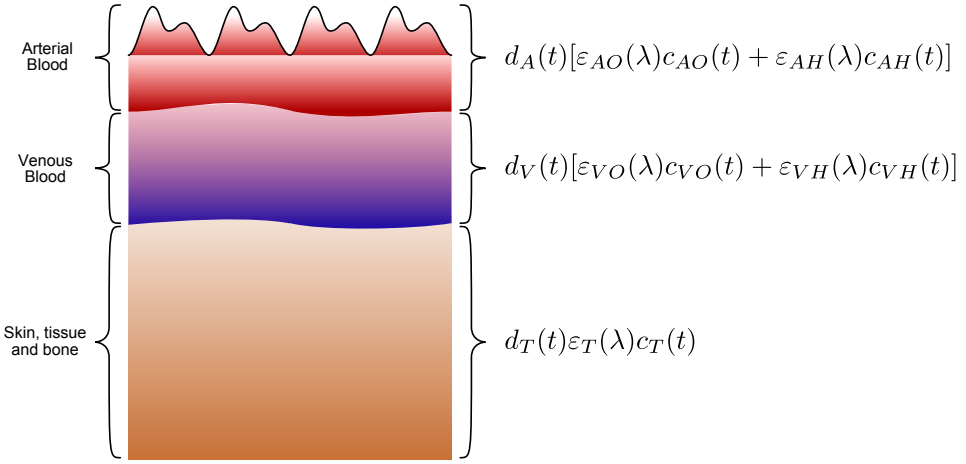


Figure 5.10: Simplified absorption model of transmission mode PPG.

Table 5.2: Absorption Glossary

Substance	**Extinction coefficient**	**Concentration**	**Physical depth**
Arterial HbO_2	$\varepsilon_{AO}(\lambda)$	$c_{AO}(t)$	$d_A(t)$
Arterial Hb	$\varepsilon_{AH}(\lambda)$	$c_{AH}(t)$	$d_A(t)$
Venous HbO_2	$\varepsilon_{VO}(\lambda)$	$c_{VO}(t)$	$d_V(t)$
Venous Hb	$\varepsilon_{VH}(\lambda)$	$c_{VH}(t)$	$d_V(t)$
Tissue, bone, etc.	$\varepsilon_T(\lambda)$	$c_T(t)$	$d_T(t)$

absorbing substance i.

$$\begin{aligned} A(\lambda,t) = &\left[\overline{d_A} + d_A^*(t) + n_A(t)\right] A_A(\lambda,t) \\ &+ \left[\overline{d_V} + d_V^*(t) + n_V(t)\right] A_V(\lambda,t) \\ &+ \left[\overline{d_T} + d_T^*(t) + n_T(t)\right] A_T(\lambda,t) \end{aligned} \tag{5.13}$$

In the context of pulse oximetry the presence of motion term $n_A(t)$ is of no significance as the only term of interest is $A_A(\lambda,t)$ and static components can be removed. Pulse oximetry is still corrupted by any variations in the venous blood or tissue as it relies on their absorptions to be static in order to cancel them out in the ratio-of-ratios calculation.

Linear separation of $d_A^*(t)$ from $n_A(t)$ is not trivial as it cannot be assumed that they are separated in the frequency domain - both being non-stationary signals. Techniques such as Masimo's SET [54] can be applied as mentioned in chapter 2, but are not applicable to wearable sensors due to the computational complexity. One popular solution to this issue is to use motion sensors and a mechanical model to implement an adaptive noise cancellation (ANC) filter. A

more far-fetched solution is to induce a mechanical vibration in the tissue with an actuator and use homodyne detection to lock on to the known vibration - just as is done to linearly separate additive light sources in the optical domain. This, however, is unlikely to be an attractive solution to prospective wearers of the sensor.

Hayes and Smith [41], using diffusion theory, propose a model for the system where the physical affect of motion is independent of both optical wavelength and substance. Expressed in the terms used in this chapter such a model would be expressed as follows:

$$\begin{aligned} A(\lambda,t) = &\left[\overline{d_A} + d_A^*(t)\right] A_A(\lambda,t) \\ &+ \left[\overline{d_V} + d_V^*(t)\right] A_V(\lambda,t) \\ &+ \left[\overline{d_T} + d_T^*(t)\right] A_T(\lambda,t) \\ &+ n(t) \end{aligned} \tag{5.14}$$

This relationship between motion and optical attenuation is considered to be *multiplicative*. As such, it is possible to remove the multiplicative noise term by *ratiometric* comparison between two different optical wavelengths. A single optical channel measurement, separated by TDM of different wavelength LED, can be expressed as a simplification of 2.3:

$$y_P(\lambda,t) = K(\lambda)e^{-A(\lambda,t)}x_L(\lambda,t) \tag{5.15}$$

Here K is the sensitivity of the photodetector system and x_L is the LED drive strength of the particular channel λ. Taking the ratio of two channels, red and infrared say, provides the following relationship:

$$y_R(t) = \frac{y_P(\lambda_0,t)}{y_P(\lambda_1,t)} = \frac{K(\lambda_0)}{K(\lambda_1)}\,\frac{x_L(\lambda_0,t)}{x_L(\lambda_1,t)}\,e^{-[A(\lambda_0,t)-A(\lambda_1,t)]} \tag{5.16}$$

The calculation of the ratio leads to a differencing of the absorption terms. Note that it is necessary that there are no systematic or ambient light offsets in the optical measurement for this ratio calculation to be accurate.

$$\begin{aligned} A(\lambda_0,t) - A(\lambda_1,t) \;= &\left[\overline{d_A} + d_A^*(t)\right] \left[A_A(\lambda_0,t) - A_A(\lambda_1,t)\right] \\ &+ \left[\overline{d_V} + d_V^*(t)\right] \left[A_V(\lambda_0,t) - A_A(\lambda_1,t)\right] \\ &+ \left[\overline{d_T} + d_T^*(t)\right] \left[A_T(\lambda_0,t) - A_A(\lambda_1,t)\right] \end{aligned} \tag{5.17}$$

Now the wavelength independent noise term has been cancelled out, but it is still necessary to choose wavelengths which allow the signals of interest to be measured while minimising terms which are to be ignored. Ideally, the

following conditions would leave only arterial blood signals in the output.

$$\begin{aligned} A_A(\lambda_0, t) - A_A(\lambda_1, t) > 0 \Rightarrow A_{A*} \\ A_V(\lambda_0, t) - A_V(\lambda_1, t) = 0 \\ A_T(\lambda_0, t) - A_T(\lambda_1, t) = 0 \end{aligned} \tag{5.18}$$

These, unfortunately, are not feasible to achieve so a condition is generally imposed that dynamic venous and tissue signals must be negligible. Otherwise, additional wavelengths of light are required to spectrally separate venous blood [27] or the use of adaptive noise cancellation techniques as per Masimo SET devices [54], both of which are currently beyond the scope of a purely wearable PPG system. Although the optical homodyne sensing technique does scale to using more than just two wavelengths LEDs are not readily available in small enough packages to justify their inclusion.

5.2.2 Validation

Both the text and the results presented in this section are taken directly from "Ratiometric Artifact Reduction in Low Power Reflective Photoplethysmography", published in the IEEE Transaction on Biomedical Circuits and System in 2011 [8].

Notes on experiment limitations

At the time that this paper was published homodyne sensing was still in its early phase of development and used envelope detection as well as a low frequency (31.25Hz) LED carrier frequency. The need to increase the carrier frequency to 250Hz, use product detection and higher performance signal filtering, as proposed in the previous chapter, are all outcomes from the experimental data presented in [8]. Furthermore, it was discovered that insufficient communications termination led to significant noise injection into the IR optical channel.

Assessment of Induced Artefact

To test the feasibility of the combined additive and multiplicative artefact reduction a number of experiments were performed on a single subject with the PPG transducer located above the left eyebrow. It should be noted that all of these experiments were essentially qualitative in nature. This is primarily due to the complexity of how motion couples into tissue means that even with a motion sensor to quantify the magnitude of induced motion any model for how

this translates into tissue volume variations will only ever be subjective. Furthermore, the ratiometric method has only been applied to transmission mode PPG in fingertips so these experiments are to analyse the feasibility of translating the technique to reflection mode PPG for use in a wearable platform. As such the method was only tested on a single subject with qualitative motion inputs although an accelerometer measure was included to enable frequency domain separation of the biosignal from the motion signal.

The following 5 different types of physical disturbance were imparted directly to the PPG transducer:

1) Shaking head left-to-right.

2) Nodding head up-and-down.

3) Twitching forehead muscles.

4) Tapping transducer.

5) Pressing transducer on-and-off against the forehead.

By choosing these tests the intention was to have a range of motions that affect both the coupling of the transducer to the sensing location and that may physically induce changes in volume of the blood within the sensing location. The metric used to quantify the rejection of motion based noise is the fundamental pulse-to-motion ratio (fPMR) as defined by the ratio of the magnitude of the first harmonic of the heart rate waveform to the magnitude of the spectral peak of the motion induced noise. To measure this requires that the motion be induced at a frequency independent of the subject's heart rate which may not be realistic in a real-world application of PPG, but serves to quantify the artefact rejection capability.

Figure 5.11 illustrates the three signal processing methods used for a comparative analysis of the performance of the proposed artefact reduction method on the data captured from the experiments listed previously. The three algorithms compared were: a) low-pass filtering of the IR waveform to reject carrier signal, b) IR to red ratio calculated without additive noise rejection (now called *baseband ratiometric method*), and c) the ratio as calculated using the amplitude modulation method (now called *AM ratiometric method*). The final low-pass filter in each method is to remove any high frequency aliasing artefacts caused by poor out-of-band rejection by the bandpass filter and it is included in each algorithm to ensure that there is no bias towards any particular method. Figure 5.12 illustrates how the spectral peak of the noise can be identified from the DFT of the accelerometer data. In the frequency domain

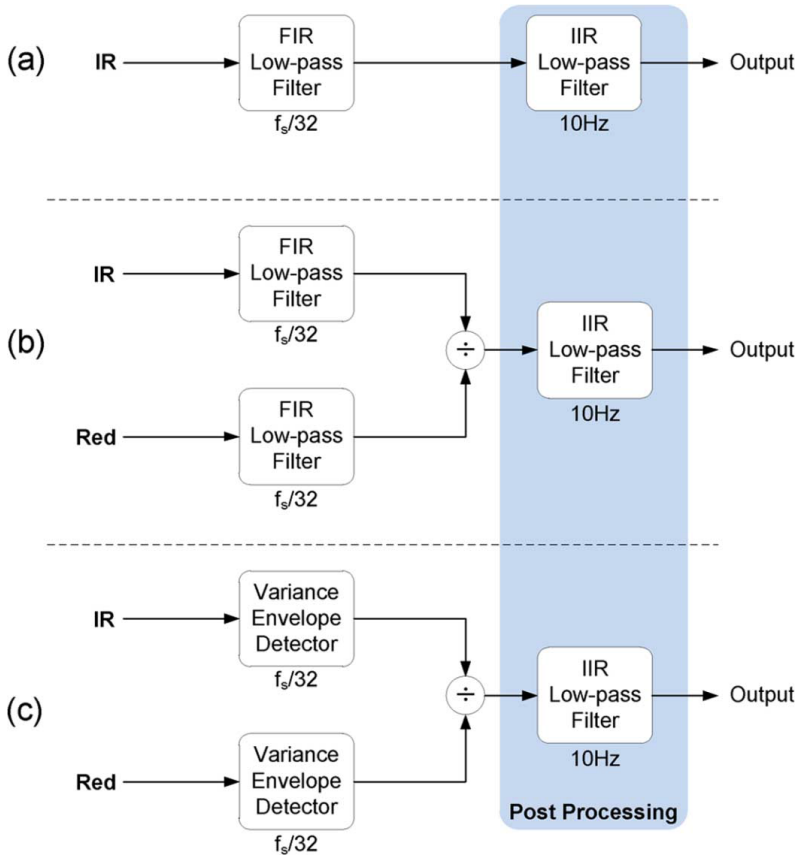


Figure 5.11: The three signal processing techniques used to compare artefact rejection performance: (a) simple low-pass filtering, (b) ratiometric measurement calculated from baseband signals, and (c) ratiometric measurement calculated from amplitude modulated signals.

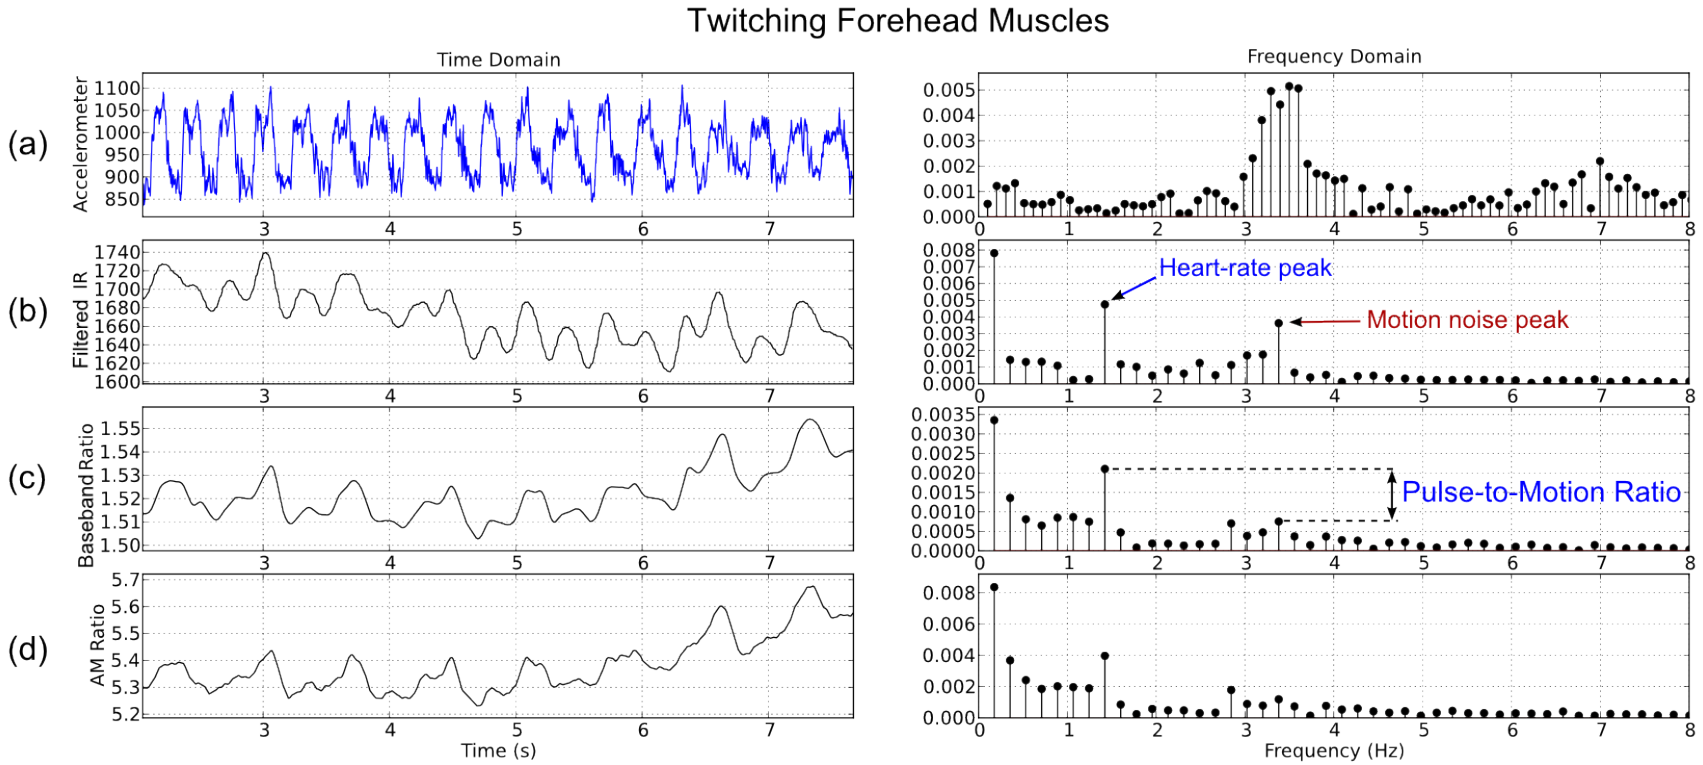


Figure 5.12: Sensor system output from when twitching forehead muscles: (a) is the X-axis output of accelerometer, (b) is the output from the IR ADC channel after the carrier has been removed with low-pass filtering, (c) is the output of the baseband ratiometric algorithm, and (d) is the output of amplitude modulated ratiometric algorithm.

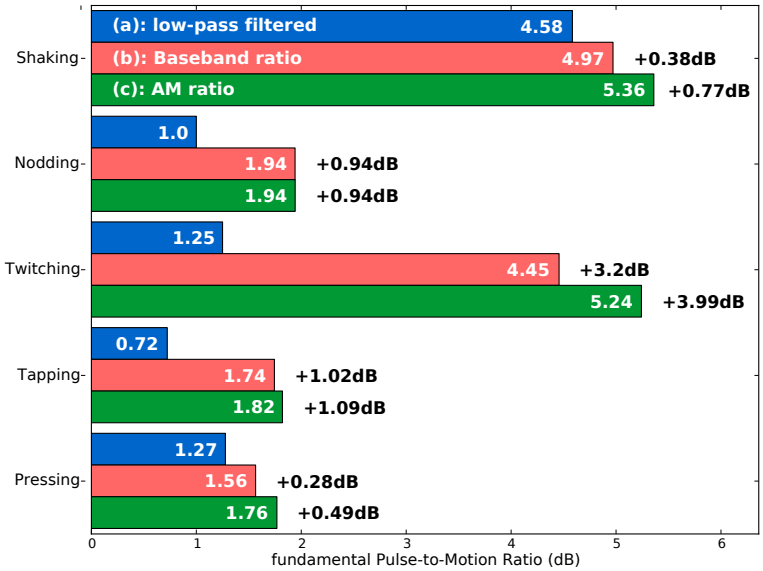


Figure 5.13: Comparison of the fPMR of each algorithm for a given type of motion induced noise.

the separation of the heart rate signal and motion noise is clear and allows a direct measurement of the fPMR for a given algorithm over the duration of a particular experiment.

Figure 5.13 compares fPMR of each of the three signal processing methods for each of the 5 motion tests. In each case the AM ratiometric method provided an increase in signal to noise ratio over both the filtered IR signal and the baseband ratiometric method.

When comparing the baseband ratiometric method to the AM ratiometric method there appears to be little gained for the extra processing effort required in calculating the variance compared to a rectangular finite-impulse-response(FIR) filter although it should be noted that these experiments were performed indoors under incandescent lighting so the photocurrent offsets are likely to be small. What is most clear from this data is that for some forms of motion the ratiometric method, while always providing a nominal increase in fPMR, still cannot stop the motion induced signal from dominating the physiological signal. However, in one of the motion experiments (twitching of forehead muscles) the physiological signal was recovered from an unreadable signal by the AM ratiometric algorithm providing a 4dB boost in fPMR. This supports the part of the model (equation 5.13) where there is a multiplicative artefact, $n(t)$, that is relatively independent of optical wavelength. The four other experiments, where the recovery of a physiological signal failed, suggests significant presence of the $n_A(t)$ term as shown in equation 5.13. This also

supports the model - even though a method has not yet been provided to compensate for this noise term.

Assessment of Natural Motion

To test the system as a whole the subject performed a series of simple activities while the sensor collected PPG data. The protocol involved the subject sitting, standing, walking on a treadmill at 2.5km/h, 3.5km/h, 4.5km/h, jogging at 7.5km/h, 8.5km/h, back to walking at 3.5km/h then finally standing again, each for 30 second intervals. This protocol was chosen to include a range of motion intensities and induce variation in the subject's heart rate. This particular experiment was conducted indoors, but with incandescent lighting turned on and the treadmill located by an open window to expose the subject to direct sunlight.

The spectrograph in Figure 5.14 illustrates the spectral output of the red channel of the PPG system during the natural motion tests. Over the 270 seconds of the experiment the subject's heart rate can be seen to increase from around 80 beats per minute (BPM) while sitting to around 140BPM when jogging, but between 150s and 210s the motion artefact is the dominant spectral component in the baseband signal. The amplitude modulated red PPG signal shows removal of the additive component of the motion noise during this experiment can significantly increase the signal to noise ratio, but there is still a residual motion component on the plethysmogram. Note that for the natural motion experiments it is not possible to provide an fPMR measure due to the motion component not always being spectrally separable from the cardiac signal. Figure 5.15 shows the same natural motion experiment performed again, but now compares just AM additive noise rejection vs. full ratiometric noise rejection. In both cases the same cardiac signal trend as before is visible, as is the motion induced noise during the jogging phase. There is a small reduction in the relative intensity of the motion noise component, which is consistent with the synthetically induced motion experiment (figure 5.13). Unfortunately, there is significant noise coupling into the IR channel which corrupts the time domain signal. This noise source was eventually traced down to poorly terminated communication transmission line reflections coupling into the analogue power supply voltage. The noise only affected the IR channel as all communications happened to be synchronised to this channel.

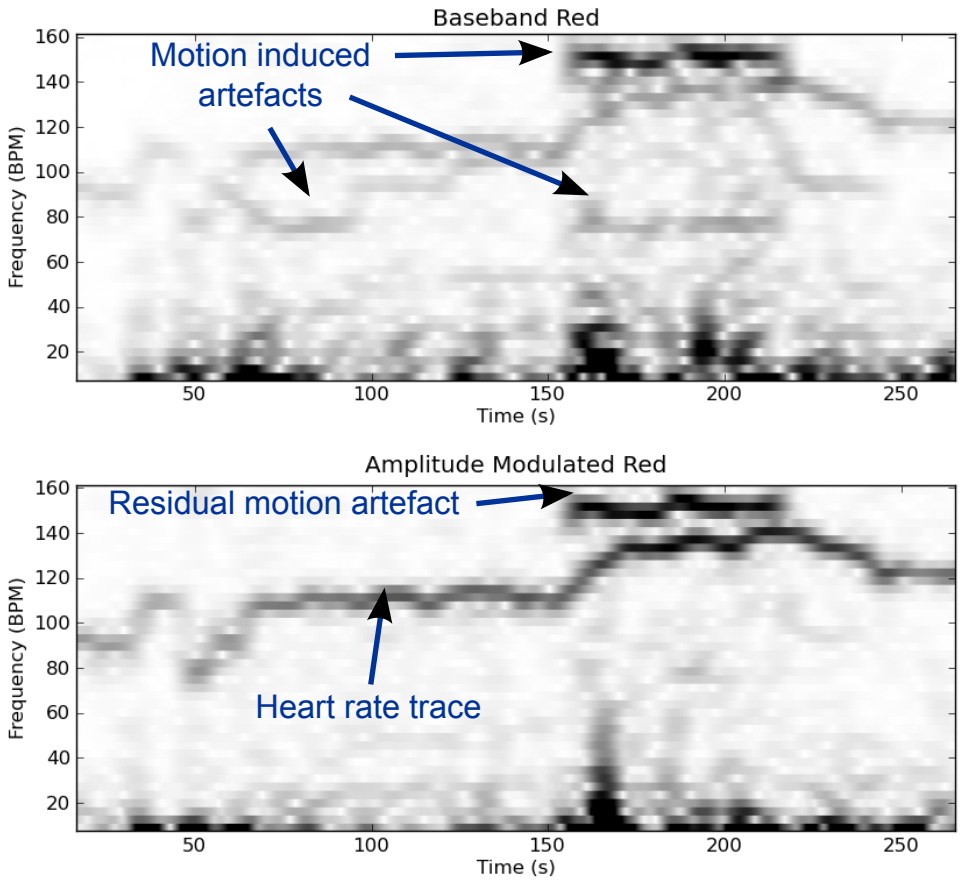


Figure 5.14: Spectrograms from the baseband and amplitude modulated filter outputs of the red PPG channel from the natural motion experiment. Removal of additive optical signal components has significantly reduced presence of motion induced noise during jogging although there is still some residual artefact.

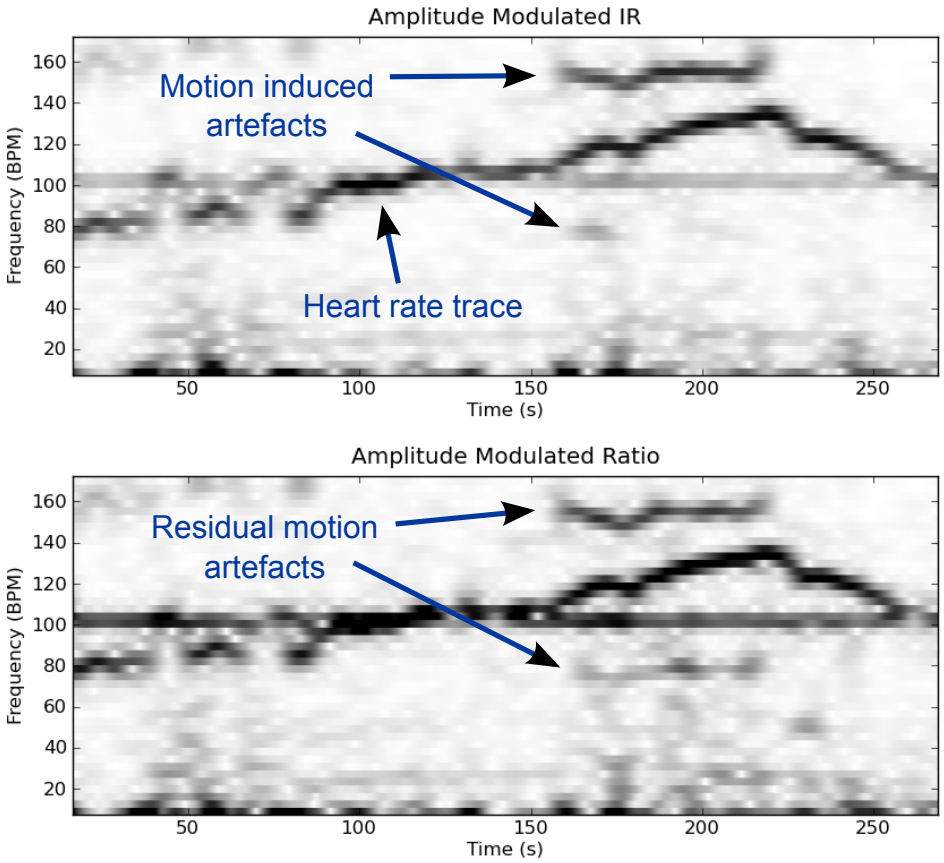


Figure 5.15: Spectrograms comparing non-ratiometric with ratiometric noise rejection. There is a small reduction in the relative intensity of motion noise from jogging. Note that the constant signal at 102BPM is the coupling of power supply noise in to the IR PPG channel.

5.3 Conclusions

5.3.1 Additive Artefacts

In this chapter a PPG signal processing configuration has been proposed that has additive noise rejection of >70dB (depending on LED brightness and sensor location) by using homodyne detection to place the biosignal in an available frequency band which has the least interference. Furthermore, mixed-signal negative current feedback stabilises the operating range analogue front-end amplifier which provides a further 14dB of noise cancellation at low frequencies without affecting the desired signal. Experimental results show that a cardiovascular waveform can be detected in the optically harsh environment of a typical home and that the system can operate even when exposed to high static light levels. Having proven the efficacy of the system the next step is to optimise the system in terms of operating range, size and power consumption.

To further simplify the PCB implementation an optimised form of the current feedback system has been proposed which gives the system zero gain at DC throughout the signal chain allowing the integrating transimpedance amplifier to be implemented without the need for reset circuitry.

The MSP430F2122 MCU on prototype PCB v2 can only execute the entire algorithm for single channel in real time at 8MHz and the processor core spends >70% of the time in full power processing mode which leads to an average current consumption of 21mA (including 2mA for average LED drive current). The first step in reducing power consumption in the MCU is to increase the percentage of time that the processor core is in a sleep or power-down state. Processing time can be reduced by choosing an alternative MCU core which supports hardware multiplication which is not available on the MSP430F2122. Ported to the MSP430F1611, which has a 16x16 bit multiplier, the algorithm executes in <2700 clock cycles which brings the awake duty-cycle to nearer 30%. Furthermore, swapping to a native 32 bit processor core allows for quicker execution of our implementation of the LWDF which is internally 32 bits itself. For example, porting the algorithm to an Energy Tiny Gecko MCU (ARM Cortex-M3 core) reduces the execution time to 570 clock cycles and only requires 3.6mA during execution. This leads to an estimated average core consumption of <1mA and an average system current consumption of 3mA. These power and processor considerations will be addressed in chapter 7 when all proposed techniques are drawn together in a minimised system implementation.

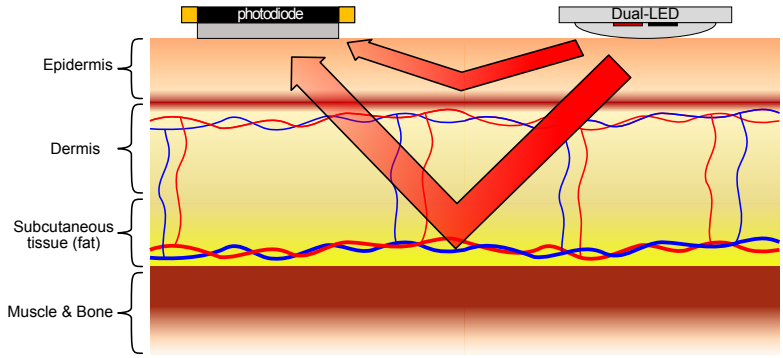


Figure 5.16: LED light bypassing arterial blood in reflection mode PPG.

5.3.2 Multiplicative Artefacts

The combination of ratiometric cancellation of multiplicative noise with amplitude modulation to reject additive noise has been shown to provide up to 4dB increase in signal-to-noise ratio in experimental data although some of this is likely to be due to removal of additive motion induced artefacts. In all of the data sets the ratiometric technique provided an increase on SNR over the basic PPG measurement method - even in the absence of homodyne sensing. The experimental data has shown clearly that reflective PPG sensors are prone to motion induced modulation of the blood volume which neither ratiometric or additive rejection techniques can separate from true physiological contributors and hence an additional artefact rejection technique is required to perform this separation. Although the results suggest that the combination of the simple additive and ratiometric techniques may not be enough on there own to cancel out all motion artefacts they can act as a useful pre-filtering stage in a more advanced motion artefact rejection algorithm - not all of which has to necessarily be performed on the sensor node itself if a method of compressing the amount of data to be transmitted can be found.

Another aspect that needs to be considered for multiplicative artefacts is the possibility that when PPG is configured in reflection mode there is the possibility that there are multiple paths that light can travel from the LED to the photodiode - not all of which necessarily pass through arterial blood (figure 5.16). If present, these paths, when dynamic in nature, would render the proposed artefact rejection techniques ineffectual.

Even if the ratiometric cancellation method's relatively poor performance does not currently justify its inclusion in the wearable PPG system, the high performance of the additive noise removal system means that it is now feasible to perform some on-node signal analysis - even if it is only possible while the subject is stationary.

6 On-Node Signal Processing

A wearable PPG system must have some way of minimising the amount of data transmitted wirelessly in order to minimise the system's power consumption. The obvious approach would be to implement some form of naive lossless data compression algorithm, such as one based on Lempel–Ziv compression algorithms [85, 86, 87], but these are not necessarily applicable to low latency data streaming as sampled data has to be segmented into blocks which are monolithically compressed. The achievable compression ratios are also rarely better than 50% which is not a significant saving against the extra processor cycles and memory required in order to perform the compression.

There are many techniques for analysing raw PPG signals and extracting the key features such as the Variable Frequency Complex Demodulation (VFCDM, [88, 89]), Hilbert-Huang Transform (HHT, [90]), as well as more classical frequency domain techniques based on the Fast Fourier Transform (FFT) [91]. However, the intrinsic mathematical nature of these methods means they are not natural choices for implementation on ultra-low-power MCUs.

Compressed Sensing (CS, [92]) provides a method of reconstructing a signal from a sparsely sampled data set. CS has been applied to pulse oximetry [93] with an under sampling ratio of up to 30. While CS can greatly reduce the sample set size there is still significant complexity in determining the sparse sample set that will allow for the lowest error in the reconstruction phase.

Alternatively, data transmission rate can be reduced by only reporting the morphological features of the waveform being measured. By using morphological features one is not only reducing the required data rate, but also directly providing information about the subject's cardiovascular function. These features are often extracted in the time domain via some form of signal *delineation*. Delineation, the process splitting a signal into a set of vectors that describe its shape, is a much more attractive solution for embedded compression due to its computational simplicity and the scalability of its compression ratio against desired accuracy. Back in 2003 Zong et al. used the dominant slope of the plethysmogram to detect the onset of "arterial blood pressure pulses" (ABP) [94]. Since then both Linder [95] and Cox [96] have used morphological features of the plethysmogram to assess hypovolemia via envelope

detection and extrema location respectively. Farooq et al. has specifically applied a very simple delineation algorithm to ubiquitous health monitoring [97]. Li et al. focus on justifying the use of delineation for ABP pulses and dicrotic notch detection by extensive simulation of the algorithm on open databases for cardiovascular data [98]. Most recently Sukor et al. use extract morphological features in order to quantify the quality of a received pulse-oximeter waveform in order to inform an observer if the data can be trusted or not [99]. With this recent research justifying that delineation is a valid approach for morphological analysis then it makes sense that it could also be used for the on-node compression of data for wireless transmission.

The choice of delineation algorithm for a PPG system is dependent on a number performance specifications which are discussed below. Numbers are chosen, where possible, to be competitive with the Nonin OEM III pulse-oximeter module.

Temporal Resolution: Must be able to report the period of a cardiac cycle to ±1ms. This provides a single cycle heart rate accuracy of ±1 beat per minute (BPM) at a heart rate of 245BPM. A high single cycle accuracy also allows high resolution measurement of heart rate variability as opposed to improving frequency resolution by averaging over multiple cycles.

Heart Rate Range: Must be able to measure heart rate in the range of 10 to 300BPM. This large range is to support all physiological possibilities. For a sample period of 1ms this translates to a maximum sample count of 6,000 which is easily contained in a 16 bit register. The range also provides boundaries for measuring the validity of a detected cardiac cycle.

DC Accuracy: The mean value of the plethysmogram over the cardiac cycle must be conserved in the delineation output. This is to enable S_PO_2 calculation and allow for subsequent low frequency data analysis, e.g. sinus rhythm detection, at the other end of the wireless data link.

Computational Complexity: In a prototyped system the algorithm will need to be implemented on a low power MCU. This means avoiding floating point calculations and recursive algorithms. It is also desirable that the algorithm is transferable to both digital and analogue IC implementations such that translation to an application specific IC (ASIC) design requires minimal extra effort.

Data Throughput: To allow the use of an ultra-low-power wireless communication solution (such as ANT+™, DASH7®, Bluetooth® Low Energy) it is necessary to keep the data rate low as well as keeping the messaging rate low to allow the radio to spend >90% of the time in a low power "sleep" mode. Hence the message rate should be kept to ≤4 per heart beat and the peak data rate to <20kbit/s

Adaptiveness: PPG waveforms are highly non-stationary. Amplitude and frequency are variable from subject to subject, sensor location to location. Even if fixed in location on a single subject the perfusion level in the arteries being sensed will drift over time. As such the delineation process must be able to adapt to variable base lines, cycle amplitudes and cycle frequencies.

As photoplethysmography is fundamentally the measurement of a cyclic waveform is was decided to chose a delineation process that segments PPG data into individual cardiac cycles. Each cardiac cycle is further segmented into four phases which provide additional morphological information about each cardiac event in order to allow physiological analysis beyond mere heart rate tracking as well as quantification of the quality of a delineated cycle.

6.1 Quadrature Phase Segmentation

The term *quadrature* can refer to both a complex representation of a periodic time-domain signal as well as to numerical integration used to calculate the one dimensional definite integral of a signal over a given time period. In the context of this thesis quadrature exclusively refers to the complex representation of a cyclic signal.

A pure cosine wave has 4 symmetrical phases as shown in figure 6.1. The Hermitian symmetry of the wave is best described by the complex function, C, that traces a unit circle about the origin. The quadrature representation of the function $C(t)$ in the time domain can be written as Euler's formula:

$$C(t) = e^{-j\omega_0 t} = \underbrace{\cos(\omega_0 t)}_{\text{real}} + \underbrace{j\sin(\omega_0 t)}_{\text{imaginary}} \tag{6.1}$$

Segmenting the 4 phases of the cosine wave can be found by placing boundaries at the extrema of the real wave and at the extrema of the imaginary wave (which are essentially the zero crossing points of the real wave). With the vectors for each of the 4 phases it is possible to use sinc interpolation to perfectly reconstruct the original signal. In fact, if all frequency content is below the

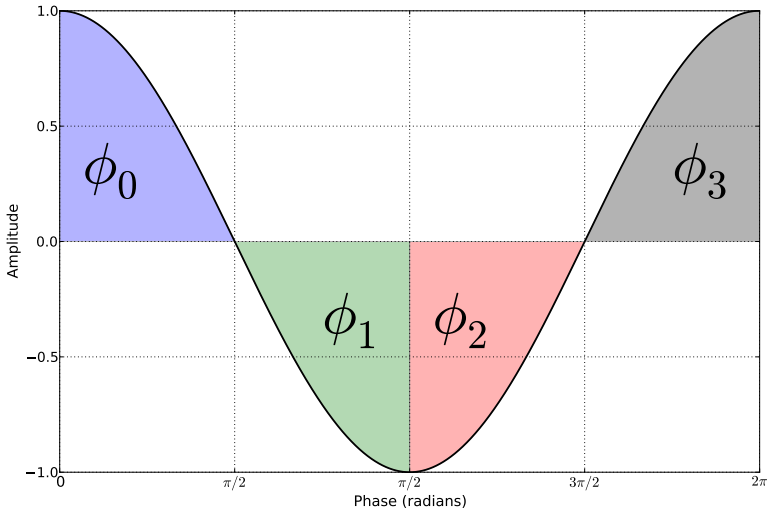


Figure 6.1: The four phases of a cosine wave.

Nyquist limit then only 2 segments, equally spaced in the time domain, are required to perfectly reconstruct the original signal. However, when analysing non-stationary signals it will not be sufficient to rely on Nyquist sampling to preserve the signal content.

$C(t)$ is also the *analytic representation* of a cosine wave with a time-invariant frequency, ω_0. Analytic signals are useful tools in analysing non-stationary and non-linear data as they allow a signal to be described in polar format with time varying amplitude and phase components, $A(t)$ and $\phi(t)$ respectively.

$$C(t) = A(t)e^{-j\phi(t)} \tag{6.2}$$

These components describe the input signal without any loss of information, but a Hilbert transform is required to generate the complex term in $C(t)$ which is effectively the input signal phase shifted back by 90° at all constituent frequencies. Having to perform a Hilbert transform in real time on an MCU would be unrealistic so instead of generating a true analytic representation we define a pseudo-analytic signal which is the complex combination of an input signal, $f(t)$, with its derivative scaled by the estimated primary frequency of the signal.

$$C(t) = f(t) - j\frac{f'(t)}{w_0^*} \tag{6.3}$$

A cardiac cycle is labelled with a sequential index i and occurs in a time window, $\mathbf{T}_i$, which has the range $t_i[0] \rightarrow t_{i+1}[0]$ where $t_i[0]$ is the first peak of cycle i. Then the four phases within the cycle can be segmented by the temporal locations of the extrema. The ordering of phases is fixed by the fact

Table 6.1: Quadrature Phases of a Non-Stationary Signal

Phase Label	$\mathbf{f(t)}$	$\mathbf{f'(t)}$	**Final Value**
ϕ_0	$<f_{\text{MAX}}[i-1]$	$<f'_{\text{MAX}}[i-1]$	$f'_{\text{MIN}}[i]$
ϕ_1	$<f_{\text{MAX}}[i-1]$	$>f'_{\text{MIN}}[i]$	$f_{\text{MIN}}[i]$
ϕ_2	$>f_{\text{MIN}}[i]$	$>f'_{\text{MIN}}[i]$	$f'_{\text{MAX}}[i]$
ϕ_3	$>f_{\text{MIN}}[i]$	$<f'_{\text{MAX}}[i]$	$f_{\text{MAX}}[i]$

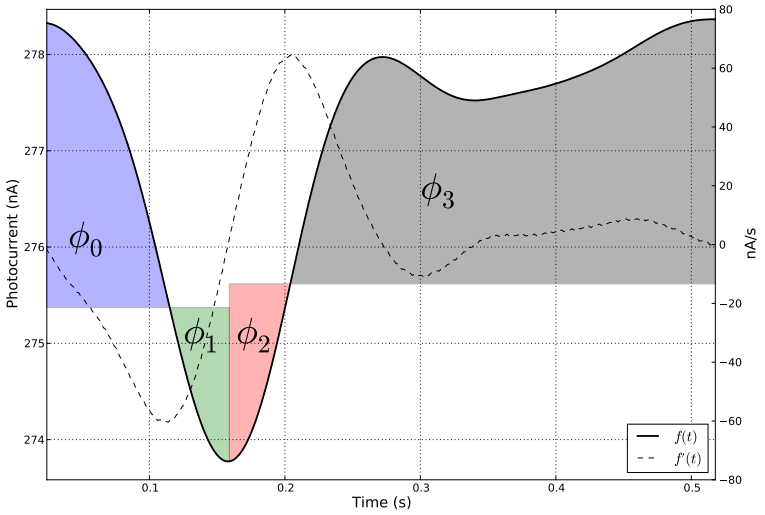


Figure 6.2: Quadrature phases of a typical PPG cycle.

that the sign of the signal's derivative determines whether it is heading towards a maximum or a minimum. For the cosine wave these points would refer to phase locations $\pi/2$, π, $3\pi/2$ and 2π. For a non-stationary signal, such as a PPG waveform, the phases are not necessarily distributed uniformly in time as illustrated in figure 6.2. Note that the data used for analysis throughout this chapter is the same data as was used in chapter 4 which was captured from a reflection mode transducer worn at the author's fingertip.

For each phase it is only necessary to report the relevant extrema which just occurred, but in discrete systems it is beneficial to report the mid points as well as one cannot assume the true zero crossing will occur at a quantisation boundary. So for each phase final values of $f(t)$ and $f'(t)$, and the time will be compiled into a report for transmission; each message consisting of 12 values. To generate a conservative estimate of required bandwidth allocate 32 bits for each sample and assume that two optical channels are required for pulse oximetry. For a maximum heart rate of 300BPM this requires a data throughput of only **3840**bps which is well within the operating parameters of

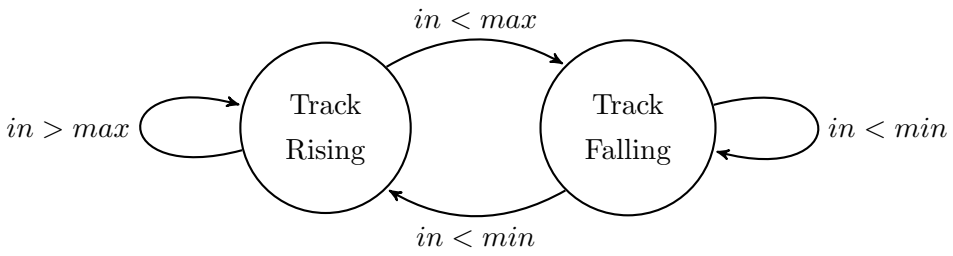


Figure 6.3: Simple Binary State Extrema Detector

low power digital radios.

Each cycle of a detected PPG signal has now been segmented into quadrature phases for transmission over a wireless communications channel and subsequent data analysis. The analytical techniques that are available for use on the delineated signal will be discussed later in this chapter. First a robust method for detecting the extrema needs to be implemented.

6.2 Extrema Detection

The extrema are the turning points of a one dimensional signal: also called the peaks and troughs or local maxima and minima. For a data set in the absence of any noise the first derivative test is the traditional method of detecting the turning points. However, differentiation amplifies higher frequency noise and creates more turning points than are generated by the principle function which is being investigated. A more common implementation is to store the value of an input signal if, and only if, it is greater than the previous largest value; or lesser in the case of minimum detection. To extend this to detecting local turning points the algorithm needs a way of deciding when to report the current turning and resetting the stored maximum value used for comparison.

6.2.1 Finite State Machines

For signals that are predominantly cyclic in nature it is possible to model them as having one of two internal states: rising or falling. The correct time to report an extrema value and reset the tracker is on any transition between states. Figure 6.3 describes this simple detector. This state machine detects all turning points, which includes noise, and will create too much data to be transmitted over a low power radio. To reject noise a combination of linear filtering and thresholding is used. A high sampling rate is used to provide precise temporal resolution, but only a small portion of the available spectrum is occupied by biosignals so filtering is very effective for removing spectrally

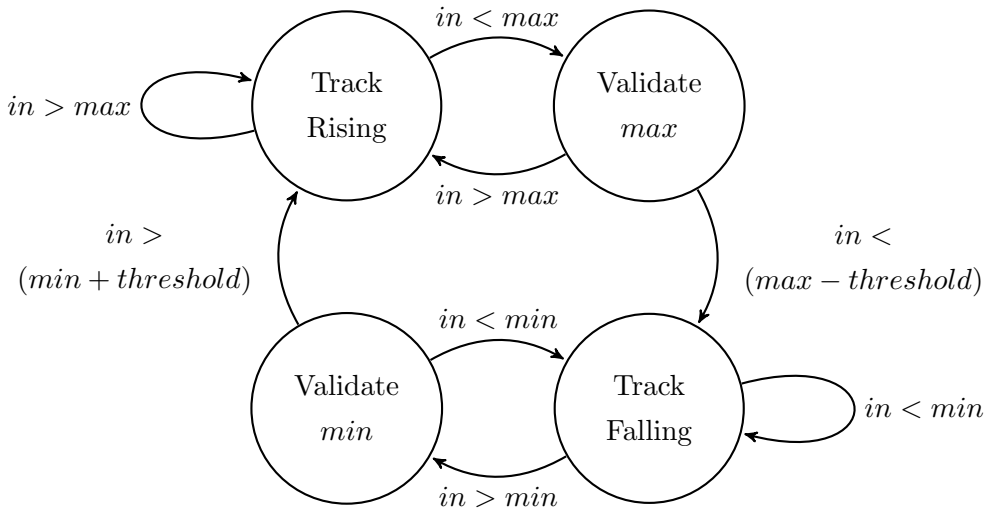


Figure 6.4: Relative Thresholding Extrema Detector

white noise.

6.2.2 Relative Thresholding

For in-band noise thresholding is used to select only the extrema that have a magnitude larger than a set level. For a cyclic waveform it is possible to apply relative thresholding - where the magnitude of one extrema is measured with respect to the preceding opposite extrema. Relative thresholding allows the baseline of a signal to drift significantly over time without saturating the threshold comparison. Validation of any proposed extrema with a threshold comparison requires extra internal states as shown in figure 6.4. The extrema value is reported and extrema tracker reset on a transition from a validation state to a tracking state. An important feature of this system is that it can return from a validation state to a tracking state if the signal changes direction before the threshold is crossed. This is particularly useful for PPG signals in the presence of a dicrotic wave which is caused by the closure of the aortic valve which induces a reflected wave of blood. There are some analytical purposes for detecting and measuring the dicrotic wave, but for the purposes of heart rate detection it is of no interest so can be disregarded. As the dicrotic wave is generated by the cardiac cycle it occupies the same region of the spectrum so linear filtering cannot be used to reject it. Thresholding provides the simplest method of classifying the source of a turning point. Figure 6.5 shows how a dicrotic wave can confuse a simple extrema detector by detecting every change is direction. In this case the effect is that the dicrotic wave would be classified as a complete cardiac cycle and lead to over estimation of the heart rate. The

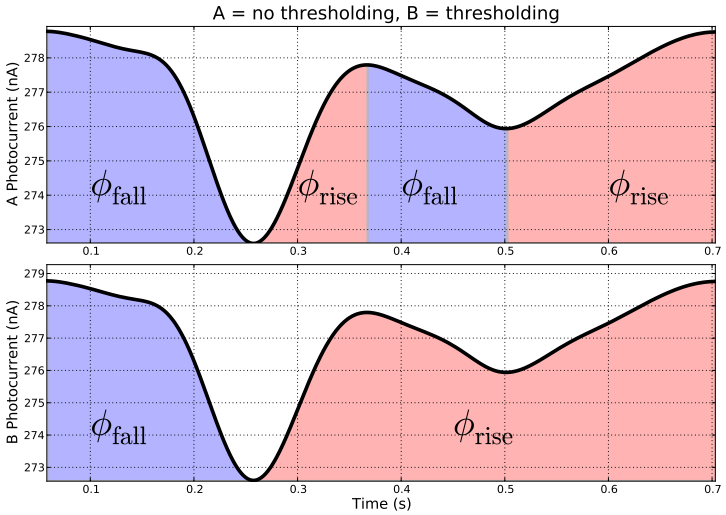


Figure 6.5: Without thresholding the dicrotic wave is classified as its own waveform rather than a subset of the cardiac cycle.

bottom plot in figure 6.5 shows how using thresholding selects the largest peak.

6.2.3 Extension to Quadrature

For full quadrature segmentation one of these state machines is required both the raw input signal, $f(t)$, and its derivative, $f'(t)$, and each requires a different threshold value. The interdependence of a value and its derivative for a cyclic signal means that the state transition conditions in one state machine are dependent of the current state of the other machine which complicates the implementation of quadrature segmentation with relative thresholding as the added validation states necessarily have to overlap with tracking states from the derivative state machine. And as the threshold for each machine can be independently chosen, it is possible that one machine will validate a phase transition while the other machine rejects the transition. To avoid this it is necessary to constrain the state transition such that they have to follow the order prescribed in table 6.1. This is done by modifying the state machine of figure 6.4 such that forward transitions from validation to tracking states are dependent on the state of the other machine as well as the threshold which is illustrated in figure 6.6.

Timing the duration of each quadrature phase is a task of timing from the occurrence of an extrema in one state machine until the next extrema in the other state machine. The ability to return from a validation state to a tracking state means that the system must continue to simultaneously track time during

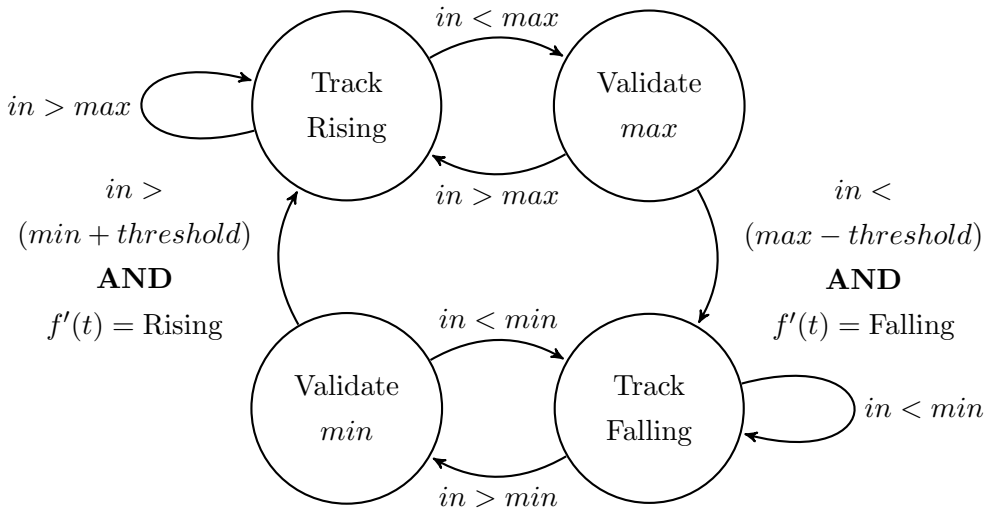


Figure 6.6: FSM for the $f(t)$ half of a quadrature extrema detector. Note the extra condition on the state of the $f'(t)$ half of the detector required to transition from a validation phase into a tracking phase.

the validation state with an separate timer for the next phase. It is possible for three timers to be tracking simultaneously making it necessary to have a separate timer for each quadrature phase and for each phase the duration is latched at the same time as a new extrema value is recorded.

6.2.4 Choosing a Threshold

Scalable Accuracy and Data Rate

The threshold value can be scaled depending on how much information is to be preserved by the extrema detector. In the extreme case, setting the threshold to zero allows the preservation of all noise content in the signal at the expense of increasing the require data throughput. Consider the case of the dicrotic wave in figure 6.5. If doubling the data rate is acceptable then the threshold can be reduced to include the turning points of the dicrotic wave in the cardiac cycle report. However, reducing the threshold makes the detector more susceptible to noise that is of comparable magnitude to the dicrotic notch.

To investigate the relationship between the threshold and the amount of data points detected the quadrature delineation algorithm was run multiple times on a single data set with a different threshold for each run. The results are presented in table 6.2. The number of bits is based on generously allocating 32 bits for each sample of the signal and is derivative for 2 channels, 16 bits for the duration of the phase (in samples), and 2 bits of phase information

Table 6.2: Effect of varying threshold on data rate on a 17s data set.

Threshold *(nA)*	**Extrema**	**Total Bits**	**Data Rate** *(bps)*
0.0001	328	37392	2216.7
0.001	316	36024	2135.6
0.01	312	35568	2108.6
0.1	276	31464	1865.3
0.25	232	26448	1567.9
0.5	184	20976	1243.5
1.0	140	15960	946.2
1.5	120	13680	811.0
2.0	**112**	**12768**	**756.9**
2.5	112	12768	756.9
3.0	108	12312	729.9
5.0	56	6384	413.8

making 114 bits per extrema. In each case the threshold for the derivative half of the detector is fixed at the lowest level possible that still rejects any residual signal content from the mixing process. If the signal were truly noiseless then it would not be necessary at all to use thresholding on the derivative signal due to the fact that its state machine transitions are locked by the state of the main primary extrema detector and as such is indirectly thresholded. The row highlighted in bold font (2.0nA) is the threshold setting that perfectly delineates the data set by the cardiac cycle validated by visual inspection of the choice of peaks. It this case it is possible to estimate the heart rate directly from the number of extrema detected and the duration of the experiment (i.e 95 BPM for the subject in this experiment). As a consequence this means that the required data rate is also dependent on the heart rate although at less than 2.5kbps for the lowest threshold setting this is still within the bounds of what is achievable by the ultra low power digital radio solutions mentioned earlier.

Comparison to Nyquist Sampling

Chapter 4 designed the detection system to have the signal bandwidth limited to 5Hz by linear filtering. By using Nyquist sampling the output sample rate could be reduced to 10Hz which for 2 optical channels would result in 640bps while still allowing perfect reconstruction at the receiving end of the com-

munications channel by using a sinc interpolator. This suggests that Nyquist sampling is immediately superior in performance to the quadrature delineation in compression performance while being simpler to implement on the sensor node. However, Nyquist sampling requires the input signal to be very strictly bandwidth limited in order to prevent aliasing which is not necessarily a safe assumption to make due to the finite attenuation of the filter used in the homodyne detection system beyond the 5Hz cut-off frequency. Also, note that the data rates quoted above include extrema values for both the signal and its derivative at all extrema which includes the redundant samples of the derivative when the signal is at either its maximum or minimum.

The primary reason for choosing quadrature delineation over Nyquist sampling is that it combines compression with the ability to measure the instantaneous frequency (e.g. heart rate) of period signals in the PPG waveform. A further benefit of quadrature delineation is the two levels of scalability it provides. Firstly, the data rate is proportional to the frequency of dominant periodic component of the input signal. For example a 60BPM heart rate can be compressed to 4 extrema per second and hence transmitted in 456bps - better than Nyquist sampling by 28%. Secondly, thresholding allows data rate to be scaled back were it to be determined, either by the data recipient or the sensor node itself, that the signal-to-noise ratio has fallen - due to either a reduction in detected optical light level or to a high level of artefacts. This means that data is transmitted, and hence extra power consumed, when there is a high level of confidence in the quality of data that the sensor can provide.

Comments on Adaptive Thresholding

The choice of value for a threshold can, in many cases, be a static value - determined from *a priori* knowledge of the amplitude range of the primary oscillating component in a signal. On the contrary, photoplethysmography needs to measure non-stationary signals where the amplitudes vary significantly over time. For example ambient temperature variations can modulate the baseline perfusion of blood in the sensing location which scales both the static blood level and the amplitude of the peak absorption. There are three available approaches to locking on to variable amplitude photometric signals in PPG: 1) vary *LED brightness*, 2) vary *photodetector gain*, or 3) vary the extrema detector *threshold*. Of these the least intrusive into the operation of a PPG system is varying the threshold although the threshold can only be lowered to a certain level at which point baseline system noise will be the signal with the dominant amplitude.

The threshold is simple to change and can be done either locally, i.e. by

the MCU implementing the PPG system, or remotely, i.e. by an authorised device that is receiving data from the sensor. The most attractive of these two solutions is remote control of the threshold as it reduces the computational load on the PPG sensor's processor. However, it cannot always be assumed that there will be a master processor device to aid with setting of the threshold and the PPG sensor may have to operate autonomously for some period of time using its own threshold adaptation rule. One such solution is to scale the threshold to the relative amplitude of the preceding extrema. In keeping with the system philosophy of simplicity for low power, the adaptation rule needs to be compact enough to be performed by an MCU. This rules out using statistical methods to generate an expectation for the scale of the amplitude. Fortunately, it is effective enough to implement an adaptation rule that is calculated solely from the previous peak-to-peak value and the time since the last extrema was detected. There are two parts to the rule:

Amplitude Scaling: For each validated extrema the threshold value, δ, is set to fraction, α, of the difference between the last two extrema

Decay: At each new sample the threshold value is reduced by a small amount, β, to cope with rapidly reducing peak-to-peak amplitudes.

$$\delta(t) = \alpha \cdot |f_{MAX}[i] - f_{MIN}[i]| - \beta\,(t - \tau) \tag{6.4}$$

Here τ is the time at which the previous extrema was detected. The choice of the α is dependent on the expected maximum amplitude of dicrotic notches and the choice of β is dependent on the expected maximum cardiac cycle period. Boundaries are also placed in the value of δ to protect against it falling to zero under low perfusion and also to protect against input signal saturation from momentarily setting infeasibly large threshold values, i.e. threshold boundaries allow the extrema detector to recover from system errors much more quickly than if the threshold was allowed to adapt linearly to all inputs.

6.3 Signal Reconstruction

Where it is necessary to analyse a PPG signal based on a samples uniformly spaced in time the signal has to be reconstructed from the delineation data using interpolation. There are many interpolation algorithms and the choice of algorithm is dependent on the mechanism by which the original data set was generated.

The cubic Hermite interpolation (CHI) algorithm is a particularly suitable choice for quadrature delineation as it creates a polynomial function to interpolate between an interval between two adjacent samples using both the interval extrema values and the tangent of the curve measured at the extrema. For a non-stationary data set the polynomial can be created piecewise from each pair of samples while still being differentiable due to the preservation of the gradient at each sample.

The polynomial is generated by first creating an interpolation time base between the two extrema which occur at sample indices k_i and k_i (where i in the sequential index of the extrema as before).

$$t = \frac{(k - k_i)}{k_{i+1} - k_i} \tag{6.5}$$

Then the CHI polynomial for the sample interval i to $i + 1$ can be expressed as:

$$\begin{aligned} f_i[k] = &\left(2t^3 - 3t^2 + 1\right) f[k_i] \\ &+ \left(t^3 - 2t^2 + t\right) (k_{i+1} - k_i) f'[k_i] \\ &+ \left(-2t^3 + 3t^2\right) f[k_{i+1}] \\ &+ \left(t^3 - t^2\right) (k_{i+1} - k_i) f'[k_{i+1}] \end{aligned} \tag{6.6}$$

One problem with the standard CHI is that as it is a third degree polynomial it does not necessarily preserve monotonicity, i.e. it is possible for the interpolated curve to overshoot passed the extrema as is illustrated in figure 6.7 leading to significant error in estimation of the PPG waveform. Fritsch and Carlson proposed a method for preserving monotonicity [100] that adjusts the gradients at each sample point to prevent overshoot. Changing the gradients means that the interpolated signal is no longer differentiable, but this is not necessarily a problem if the accuracy of the curve fitting is improved.

For PPG it is possible to use the asymmetric nature of a waveform to implement a simpler method for improving the interpolation. During the systolic phase of the cardiac cycle the CHI is an excellent match - it is only during the low pressure phase that overshoot occurs. So during phases ϕ_0, ϕ_1 and ϕ_2 the CHI polynomial is used unchanged, but during phase ϕ_3 the peak gradients are scaled down by a factor empirically chosen to provide the best match. There is still a small amount of overshoot, particularly in phase ϕ_0, but asymmetric CHI (ACHI) still significantly improves the interpolation accuracy over normal CHI at large threshold values as can be seen in table 6.3. Note that at small threshold values the normal CHI algorithm has the lowest error as

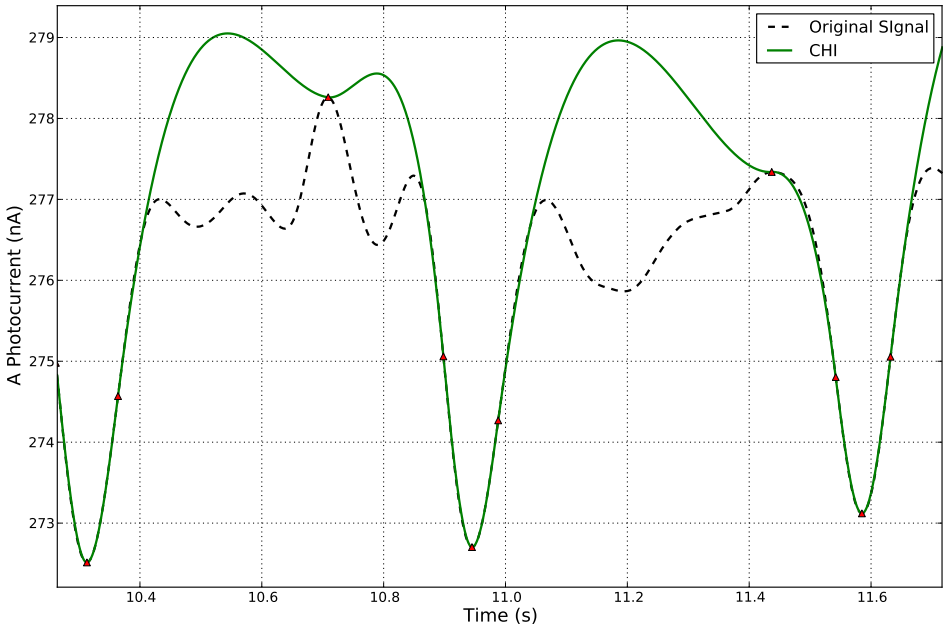


Figure 6.7: Overshoot caused by the cubic Hermite interpolation polynomial with a 2nA threshold.

the all of the smaller peaks during the diastole phase are now included in the sample set and there is less requirement for intelligent interpolation. The linear interpolation (LI) algorithm is also included as a baseline for comparison. The fact that LI consistently outperforms CHI above a threshold of 0.25nA is solely down to the poor matching of CHI during phase ϕ_3. If the RMS error were constrained to just phases ϕ_1 and ϕ_2 then CHI would significantly outperform linear interpolation. Furthermore, RMS might not be the best metric of accuracy as the mean error includes no concept of the different phases of the cycle and hence does not reflect how well the CHI based algorithms work with the higher energy phases To put these error values in to context observe that the peak-to-peak current measurement in a typical cycle for the data set used in these tests is above 4nA.

6.4 Extension of Delineation to Pulse Oximetry

A useful alternative to reporting raw photocurrent measurements would be to perform a logarithmic conversion on node which would allow separation of system dependent parameters, such as LED brightness and transimpedance gain as was discussed in chapter 2. Performing a logarithmic conversion in the analogue domain is relatively simple as most semiconductor devices exhibit some form of exponential relationship between inputs and outputs. However, per-

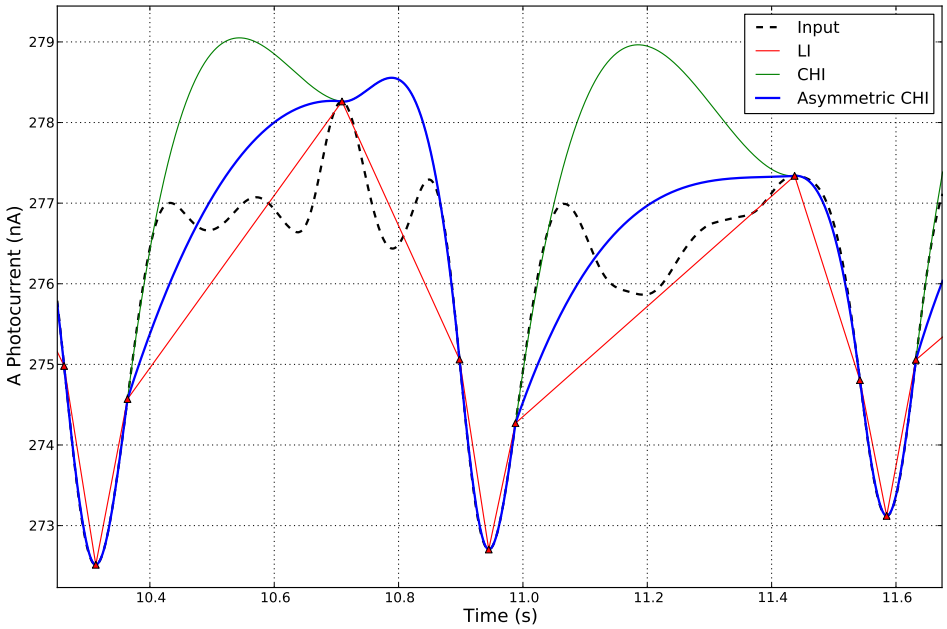


Figure 6.8: Asymmetric cubic Hermite interpolation polynomial with a 2nA threshold compared. The gradient is scaled down by 2.5 in phase ϕ_3.

Table 6.3: Accuracy of three simple interpolation algorithms for a range of thresholds.

Threshold *(nA)*	**LI Error** *(nA_{RMS})*	**CHI Error** *(nA_{RMS})*	**ACHI Error** *(nA_{RMS})*
0.0001	0.46	0.11	0.17
0.001	0.45	0.10	0.17
0.01	0.45	0.10	0.17
0.1	0.49	0.29	0.26
0.25	0.57	0.59	0.40
0.5	0.69	0.90	0.46
1.0	0.79	1.15	0.54
1.5	0.86	1.27	0.60
2.0	0.89	1.34	0.63
2.5	0.89	1.34	0.63
3.0	1.00	1.72	1.25
5.0	1.86	7.76	7.38

forming a logarithmic conversion digitally is not so simple and often requires the use of floating point data types which have been ruled out for implementation in a wearable PPG system. A popular alternative approach is to calculate the logarithmic derivative which is the ratio of the derivative of a signal to itself.

$$\frac{d}{dt}\ln\left(f(t)\right) = \frac{1}{f(t)}\frac{df(t)}{dt} \tag{6.7}$$

Quadrature segmentation already uses the derivative so the only extra calculation is a division. A division is a non-trivial calculation to perform in a MCU, but modern cores such as the ARM Cortex®-M3 core have hardware dividers that operate in 12 or less cycles so this is a feasible operation to perform. The discrete-time derivative is calculated as follows where T_s is the sampling period such that $t_{n+1} = t_n + T_s$:

$$\frac{d}{dt}f[n] = \frac{f[n+1] - f[n-1]}{2T_s} \tag{6.8}$$

It is not necessary to normalise the derivative to the sampling time - instead is can remain as a static "gain" parameter in the signal chain. Care needs to be taken to preserve fixed-point dynamic range when performing a discrete division. For example, if the desired output is to have at least 16 bits of dynamic range and the calculation is to be performed in a single divide instruction then it is necessary to keep the numerator to between 16 and 32 bits and the denominator to less than 16 bits. To keep the signal in such a range scaling factors are applied to the input signal and its derivative which are represented as further static "gain" parameters P and Q. These scaling factors are implemented as bit shifts and their values are chosen based on *a priori* data captured by the sensor system. Each discrete log derivative calculation has an intrinsic gain K which is set by the sample period and the division scaling factors. If the output of the scaled ratio calculation, $g[n]$, is equal to a scaled version of the natural logarithmic derivative.

$$g[n] = \frac{P(f[n+1] - f[n-1])}{f[n]Q} = \frac{2T_sP}{Q}\frac{d}{dt}\left(\ln f(t)\right) \tag{6.9}$$

To show how this applies to pulse-oximetry revisit the signal chain model from equation 2.3 except remove the additive noise terms and assume that the emitted LED brightness is static.

$$f(t) = y_P(\lambda, t) = G_P S_P(\lambda) e^{-A(\lambda,t)} S_L(\lambda) G_L x_L \tag{6.10}$$

Calculating the logarithmic derivative of this signal yields:

$$g[\lambda, n] = \frac{2T_s P}{Q} \cdot \frac{d}{dt}\left(-A(\lambda, t)\right) \tag{6.11}$$

All that remains in the signal is the gain of the log derivative operation and the derivative of the absorption signal. This can now be used to directly calculate R - the ratio-of-ratios used in equation 2.5 to calculate S_PO_2.

$$R = \frac{g[\lambda_{IR}, n]}{g[\lambda_{Red}, n]} = \frac{\varepsilon_{Hb}(\lambda_{\mathrm{IR}})c_{Hb} + \varepsilon_{HbO_2}(\lambda_{\mathrm{IR}})c_{HbO_2}}{\varepsilon_{Hb}(\lambda_{\mathrm{Red}})c_{Hb} + \varepsilon_{HbO_2}(\lambda_{\mathrm{Red}})c_{HbO_2}} \tag{6.12}$$

This metric is now independent of the pulsatile component of absorption, $d(t)$ and the gain of the log derivative calculation. For maximum SNR R should be calculated when the signal derivative is at its maximum which ties in well with signal delineation. Also, the calculation of S_PO_2 does not actually need to be performed on-node, only when it is required to display the value, so as long as the wireless transceiver sends $g[\lambda_{IR}, n]$ at least once per cardiac cycle the oxygen saturation can be recovered. Reddy et. al. published a similar method in 2009 that uses the "slope" of the PPG waveform after logarithmic conversion [101] which they claim is both "novel" and "calibration free". The novelty is questionable as it is simply performing the calculations in a different order to how it has been done in the past - the exemplar case being US patent 5,934,277 issued to Datex-Ohmeda Inc. in 1999 which uses regression analysis on the slope to improve signal quality. It can also be argued that it is not calibration free either due to the fact that the calculation still relies on the extinction coefficients of haemoglobin and oxy-haemoglobin at fixed wavelengths which requires the chosen LEDs to have peak emission at those chosen wavelengths both at time of manufacture and throughout the operating life of the sensor which is not a safe assumption to make.

Using the derivative directly, without any DC component signal, removes from the calculation of R any effect of static side-channel light as was introduced in figure 5.16 in chapter 5 making this the most suitable method of calculating S_PO_2 for reflection mode PPG.

As the logarithmically converted signal contains both absorption information as well as temporal information there is no need to transmit the raw photocurrent signal. In some circumstances it is desirable to infer the changes in blood volume for each phase of the cardiac cycle which requires integrating the derivative signal. The volume couldn't be recovered if only the derivative extrema were transmitted. Fortunately, the delineator design already has phase synchronised integrators in use - the individual phase timers are essen-

tially integrators with fixed inputs. Adding an extra integrator per phase is all that is needed to allow changes in blood volume to be inferred from the quadrature delineated signal. In essence, integrators enable the conservation of photonic information between sampling events and their conservative nature will find them applied to electrical photodetection as well in a subsequent chapter.

It must be noted that equation 6.12 only holds true if it can be assumed that arterial blood is the only dynamic component in the absorption of light. The previous chapter highlighted situation where it might not be possible to make this assumption. Furthermore, if there are non-arterial dynamic signals then peak detection is prone to magnify these artefacts when calculating S_PO_2. To illustrate this it is necessary to convert the calculate value of R in to S_PO_2. This is done either by using the Beer-Lambert law based relationship as described in equation 2.5 with estimated values for the extinction coefficient, or by calibration from known S_PO_2 values. For the purposes of illustration the Beer-Lambert law method is used and example coefficient values are chosen from the chart in [101] using 660nm and 905nm as the transducer wavelengths as an OSI-Optoelectronics DLED-660/905-LLS-3 dual LED was used in the collection of the example data. The chosen values are shown in equation 6.13 in the context of how they are used to calculated S_PO_2. Only approximate values are used as this exercise was not intended to provide usable S_PO_2 values, but only to illustrate the potential errors that peak-derived data points can create.

$$S_pO_2 = \frac{3.2 - 0.8R}{3.2 - 0.3 + [1.2 - 0.8]R} \tag{6.13}$$

Figure 6.9 shows how noisy the S_PO_2 signal is when calculated directly from the peak values of each phase of the signal $g[\lambda, n]$. The samples at the end of phases ϕ_0 and ϕ_2, of course, provide the most accurate estimation of S_PO_2 as the calculations based on the other samples produce impossible oxygenation values suggesting that these regions are predominantly noise. Note that the S_PO_2 are not in the expected range of 90-96% which strongly suggests the chosen coefficients are not a good match for the properties of the sensing system and proper calibration is required in order to provide reliable oxygen saturation readings.

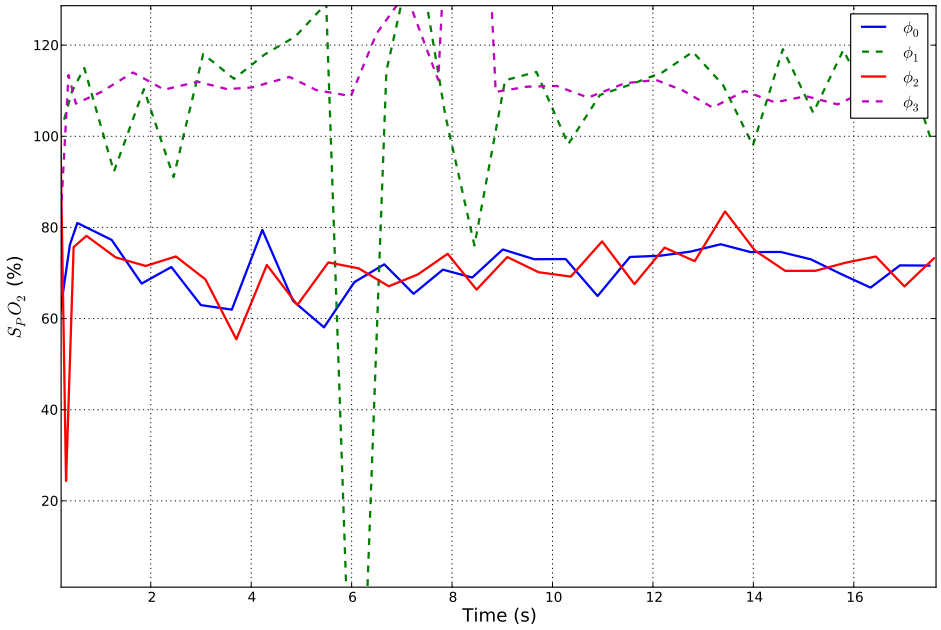


Figure 6.9: S_PO_2 calculated at each of the four quadrature extrema.

6.5 Conclusions

This chapter has presented a computationally efficient method of delineating a PPG waveform into vectors which reduce the required communication bandwidth to well within the operating range of modern ultra low power radio solutions. Delineating both the signal and its first derivative preserves sufficient temporal and amplitude information to allow the waveform to be reconstructed at the other end of a communications channel using simple interpolation polynomials.

Simple adjustments have been made to the signal reconstruction algorithm's interpolation model in order to better match specific aspects of the PPG waveform and its cyclic nature. There is still considerable scope for improvements in the choice of interpolation model, but the presented solution achieves an RMS error of less than 20% of the typical peak-to-peak signal amplitude which provides significant margin for estimating the heart rate of the subject wearing the sensor.

Integrated into the signal delineation is the conversion of raw photometric data into values relevant to physiology and independent of the instrumentation system's configuration. However, caution must be taken when calculating S_PO_2 directly from values generated by an extrema detector as any additive noise that is presence in the delineated data gets amplified when calculating ratios.

Delineation by extrema is an attractive choice over standard information compression methods due to its simplicity and high compression ratio. The example presented in this chapter reduced the required data rate from 32kbps down to the range of 0.7–2.5kbps depending on the desired signal accuracy. To estimate what this might mean in terms of radio power consumption assume that a radio is available with 30mA power consumption when transmitting at 1Mbps and 1μA consumption during sleep mode - these numbers are a conservative estimate based on the Bluegiga BLE112 Bluetooth Low Energy module [102]. To support a 300BPM cardiac signal with support for the dicrotic notch would require the radio to be transmitting at a duty cycle of about 0.45% of its maximum bit rate which leads to an average current consumption estimate of $<$140μA. This would mean that the LEDs would definitely become the dominant energy consumers in the system.

7 A Complete Wearable Photoplethysmography System

Having researched each component of the PPG system individually it is possible to integrate them into a single platform. The first platform in this chapter is another PCB implementation that is a physically minimised embodiment of the system that uses a state-of-the-art mixed-signal MCU to allow the footprint to be shrunk to less than 250mm^2 (compared to ∼254mm^2 for a UK 5p coin) for integration into a wearable patch enclosure. The second platform is a full-custom integrated circuit solution designed to have a general purpose architecture for sensing applications where one of the first target applications is a PPG system.

7.1 Smart-Patch Compatible PCB

One of the great benefits of reflection mode PPG is the range of locations on which it can be placed, but this creates an additional dimension when choosing how the opto-electronics are to integrate with the sensor enclosure. The "smart-patch" is an attractive solution because it is as location agnostic as reflection mode PPG.

7.1.1 Modular Platform

The Pervasive Sensing group at the Hamlyn Centre is also actively researching wearable ECG, glucose monitoring and other wearable sensor solutions so it was decided to target fitting the PPG sensor into a similar enclosure as had been proposed for those other sensor systems. Figure 7.1 shows a 3D printed prototype of a silicone patch that has a cavity for containing the electronics as well as tabs for secure adhesion. While this is only a prototype, it was chosen by user survey as the preferred solution of several designs so is a good target platform for the PPG system to integrate into.

Given the range of low power radio solutions, batteries and sensor options it was decided to keep the system modular such that each component can be developed independently. In order to ensure that size is not compromised

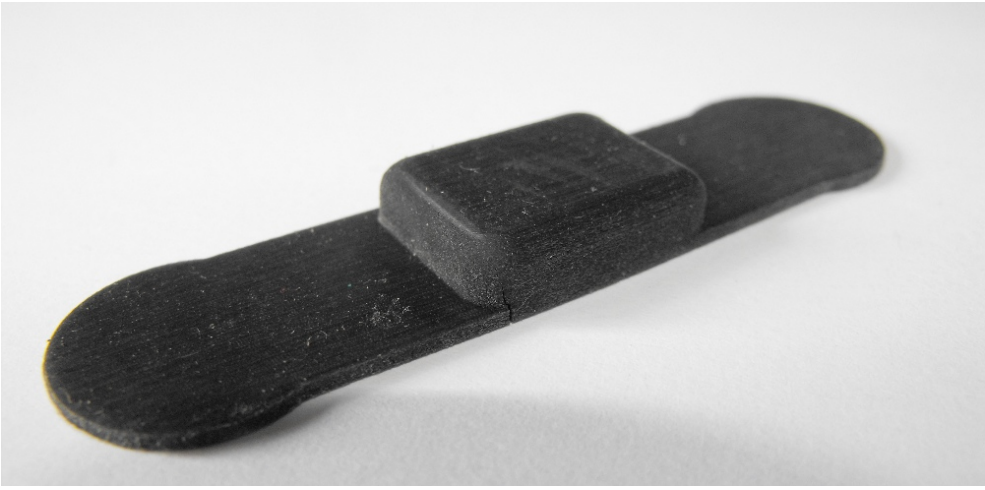


Figure 7.1: A prototype a smart-patch with a cavity for electronics.

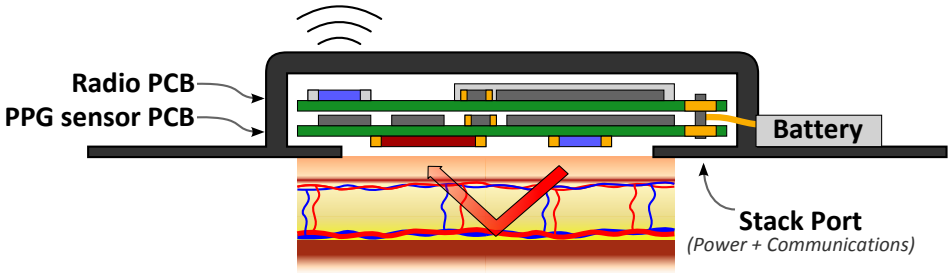


Figure 7.2: An example PCB stack for use in a wireless smart-patch sensor.

the PCB form factor was fixed to a maximum of 13x20mm (similar in dimensions to the Bluegiga BLE112 Bluetooth Low Energy module mentioned in the previous chapter[102]). Separate PCBs can be stacked either vertically or laid out horizontally. Unlike conventional PCB stacks there are no connectors in the design so the system configuration is fixed at build time, but there is neither the size nor the financial cost of including modular board-to-board connectors in the design. It is also feasible to adopt the patch style of Haahr et. al. [63] and make the battery and enclosure a separate sub-assembly from the electronics such that the expensive sensor components can be reused with single-use patch casings. This modular approach is intended to be the optimal balance between speed of development, patient adoption and build cost.

7.1.2 PPG Module

The final version of the PCB is a two channel (IR and red) PPG system which includes a single integrating photocurrent amplifier, homodyne detection, bias current stabilisation, and signal delineation as illustrated in figure 7.3. Note

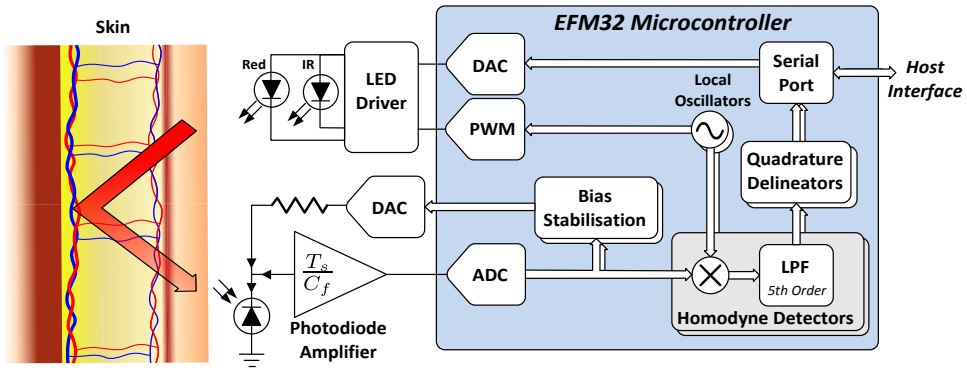


Figure 7.3: Architecture of the complete PPG sensor system.

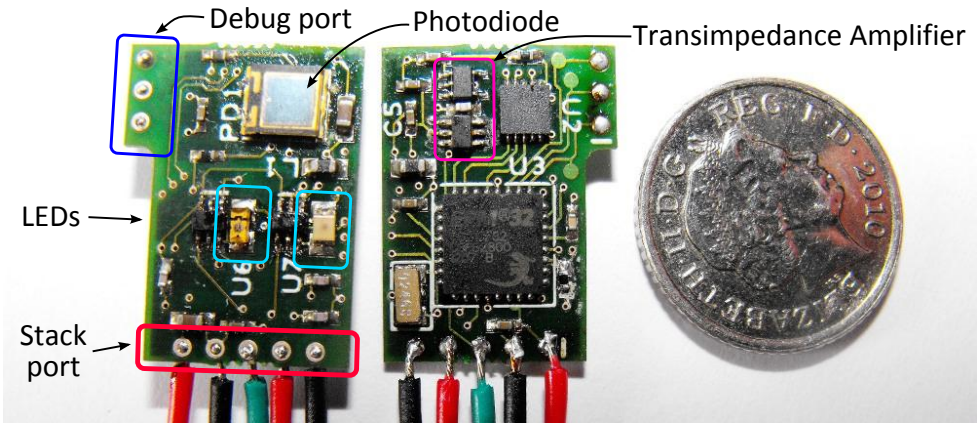


Figure 7.4: Version 3 of the wearable PPG system PCB.

that each optical channel requires its own set of bias stabilisation, homodyne detector and delineator. The microcontroller used is the Energy Micro® Tiny Gecko EFM32TG210F32 which has the ARM Cortex-M3 processor core, 32kB of flash memory, and crucially includes an ADC, DAC and 3x op-amps which are required to implement the mixed signal portion of the system. Aside from passives the only other PCB components are the photocurrent amplifier's op-amp, 3x analogue switches, the photodiode and 2x 0805 footprint LEDS. The opto-electronic components are mounted on the underside of the PCB, but the photodiode can be attached via wires to implement a more flexible solution if required. The total footprint of the device is 12.5x19.5mm after removal of the debug connector which is visible in the photos of the device in figure 7.4. Full schematics of the circuit can be found in Appendix C on page 213.

Performance

The firmware was written in the C programming language and compiled using Keil μVision v4.23 to 10,192 bytes of flash memory with only cross-source optimisation enabled, and 8,008 bytes with level 2 optimisations enabled which would allow the use of the 8kB version of the MCU which saves 10s of US cents at high volume (>1000 units).

The inclusion of a 32kHz oscillator on the PCB allows the MCU to run in energy mode 2 when not actively processing data. When in this low power mode the core only draws 1μA of current but can transition to full operating mode in 2μs so there is little overhead associated with swapping between operating modes. While developing the system the serial port must run all of the time in asynchronous mode to support the full data throughput, but a more realistic implementation would be the sensor acting either as an I2C slave or a low speed UART (<9600 baud) both of which can operate in energy mode 2. The serial port was disabled to quickly get an accurate estimate of the real power consumption when operating as an I2C slave device. Under these conditions the algorithm, timed from starting an ADC capture to being ready to transmit, executed in <85μs which achieved a total average current consumption of 2.4mA at 3V giving the headline power consumption figure of 7.2mW, 5.1mW of which is consumed by the LEDs.

Output

Figure 7.5 shows data captured by the device from the author's fingertip to prove that it is possible to capture real blood volume pulsation in such a system with so few individual components. This data was generated with an average red LED current of 1.5mA and an average IR LED current of 0.2mA. At these illumination levels the cardiac signal has a typical peak-to-peak signal of 3.5nA for the IR channel against a background noise of <0.5nA (1.5nA peak signal for the red channel).

Battery Options

In chapter 2 a 110mAh battery was provided as an example cell that could be used in BSN applications. While the quoted capacity of batteries cannot necessarily be used to accurately estimate a system's operating life the 2.4mA consumption of this system suggests an operating life (excluding radio) of over 40 hours. However, at 15x28mm this battery has a considerably larger area (420mm^2) than the sensor itself (254mm^2). A battery equivalent in area to the sensor PCB can be sourced from manufacturers such as China BAK Bat-

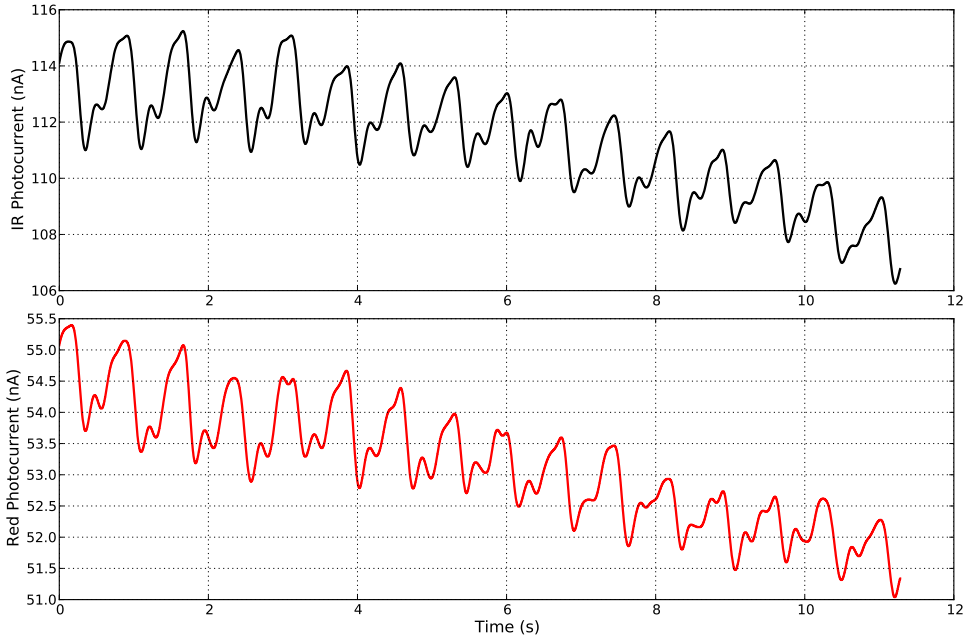


Figure 7.5: Data output by the sensor patch PCB to validate operation. Measured at the finger tip with total LED current consumption <2mA.

tery, Inc. (`http://www.bak.com.cn`). One such battery is the BAK381221P 55mAh cell which measures 3.8x12x21mm and, in theory, could operate the PCB (again, excluding radio) for over 20 hours which would allow the sensor to be operated from waking till sleeping for most patients. Having to recharge everyday means that the sensor now has a usage profile like that of a smart phone. The necessity to interact with the sensor at least once per day may be considered a significant barrier to adoption by many users, especially considering how many more functions a smart phone provides for the same level of maintenance. Also, this 55mAh battery is equivalent in size to the PCB making it the dominant consumer of space. The combined influence of size and operating life time of the battery highlight it as a major barrier to the adoption of wearable sensing technology as it currently stands.

7.1.3 PCB Summary

Significant effort was required to transfer all of the advancements developed throughout this thesis to the new MCU platform so no research studies have yet been performed with the complete system. Before taking the sensor to trials it still needs to be integrated with a radio and battery into a functional patch-type enclosure. However, when it does come to trials this platform will provide more precise cardiac signal capture, for longer periods of time, and

with more robustness against the environment outside of the research lab.

The bill-of-materials for this design comes to about $8.20 (excluding passives and PCB) of which 29% is from the MCU and 40% is from the external 12bit dual DAC used to bias the transimpedance amplifier (see figure 5.2). Having one component that performs such little functionality for such a large share of the cost does not translate well to the high volume manufacturing scale that would be needed to support pervasive cardiac monitoring on a global scale. Of course there is scope for further optimisation, e.g. fixing the voltage bias and moving to the physically larger EFM32TG230 MCU which has 2 internal DACs, but for such a specific application at such a large production scale the costs justify moving towards a fully integrated circuit solution.

7.2 Towards an ASIC Implementation

The Pervasive Sensing group at the Hamlyn Centre are collaborating with Hong Kong firm ASTRI (*Applied Science and Technology Research Institute*) to produce a mixed-signal system-on-chip ASIC for BSN applications, both wearable and implantable, based on the ARM Cortex-M0 MCU IP. The work presented in this chapter reflects part of the Hamlyn Centre's contribution to the project. At the time of publishing this thesis the chip had not yet been produced so only simulation results will be presented. An acknowledgement must be made to Dr. Ching-Mei Chen who joined the group late in the design phase and with her expertise in analogue IC design translated the designs from research concepts into truly synthesisable circuits.

Up to this point significant research has gone in to finding a way to move most of the PPG system functionality into the digital domain to reduce the number of discrete components in the system. This design procedure does not necessarily carry over to IC design where it is possible to implement the entire system on a single chip so the choice of where to divide the analogue and digital domains can be made based on power and sensitivity parameters as opposed to size. The simplicity of all components in the PPG system means that they could all be translated to the analogue domain, but that would constrict the ASIC to just performing PPG functions. The two functions which could conceivably be transferable to other sensing modalities are the amplification and the homodyne detection modules. The rest of the functionality can be trivially implemented in the MCU core of the ASIC.

7.3 Photocurrent Amplification

Directly translating the PCB-based integrator design into an IC-based design will definitely reduce the power consumption, but to realise the full benefits of an IC-based implementation it is worth reconsidering the circuit architecture with regards to the fundamental requirements. The first question to answer is do we still want to use the integrating amplifier? Kim et. al. have compared the capacitive feedback amplifier to the conventional resistive feedback amplifier in IC form for biomedical applications [103, 104] and their findings back up the findings presented in chapter 3; that the integration of small currents can be advantageous from a signal to noise ratio perspective as long as correlated double sampling is used to reject low frequency noise.

Back in section 2.3.4 it was stated that the op-amp in the transimpedance amplifier performed two functions: input voltage stabilisation and current-to-voltage conversion/amplification. The combination of functionality using a general purpose amplifier leads to certain design complications. For example, the gain-bandwidth limitations caused by large parasitic input capacitance are entirely due to output of the amplifier being connected directly to the input. Additionally, the resetting of the integration capacitor is complicated by the act of shorting the input node to the amplifier output. Due to these observations it was decided to separate the tasks of voltage stabilisation and current amplification for implementing an IC transimpedance amplifier. There already exists a circuit that can perform these two functions independently: *the current conveyor*.

7.3.1 Current Conveyor

First proposed by Sedra and Smith back in 1968 [105], the current conveyor is an elementary circuit building block with 3 ground referenced terminals: X, Y, Z. This first generation current conveyor (CCI) implemented the following rules:

1. The voltage at X is to be the same as the voltage applied to Y.
2. The current into Y is to be the same as the current into X.
3. The current into Z is to be the same as the current into X.

Rule 3 gives the circuit its name: current is *conveyed* from node X to node Z. Only rules 1 and 3 are required in order to implement a transimpedance amplifier. Figure 7.6 shows a CMOS implementation of a CCI connected as

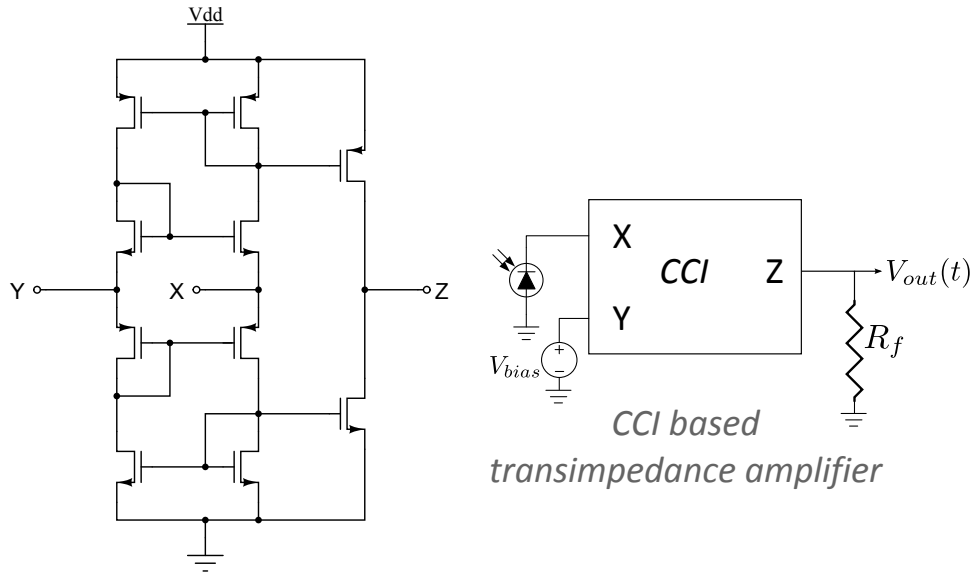


Figure 7.6: A first generation current conveyor circuit and its application as a conventional transimpedance amplifier.

a conventional transimpedance amplifier. The resistor which sets the transimpedance gain, R_f, is no longer connected to the input node so can now be considered independent of the input capacitance. The amplifier bandwidth is now a function of the internal bandwidth of the CCI between nodes X and Z as well as the combination of R_f with any parasitic capacitance between node Z and ground. Node Y can then be used to set a static bias voltage across the photodiode without it being added to the output voltage as is the case with the voltage-mode op-amp based transimpedance amplifier.

The current into node Y is superfluous to the requirements of a transimpedance amplifier and hence so is rule 2 of the CCI. This ties in perfectly with the second generation current conveyor (CCII) as proposed by Sedra and Smith in 1970 [106] which is often described by the following matrix equation:

$$\begin{bmatrix} v_X \\ i_Y \\ i_Z \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ \pm 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} i_X \\ v_Y \\ v_Z \end{bmatrix} \tag{7.1}$$

Now no current flows through node Y and it is possible to have both non-inverting and inverting versions of the conveyor (CCII+ and CCII- respectively). Figure 7.7 shows a CMOS implementation of a CCII± built around a voltage-mode op-amp where the current conveying is performed by using current mirrors in line with op-amp supply nodes. The second generation current

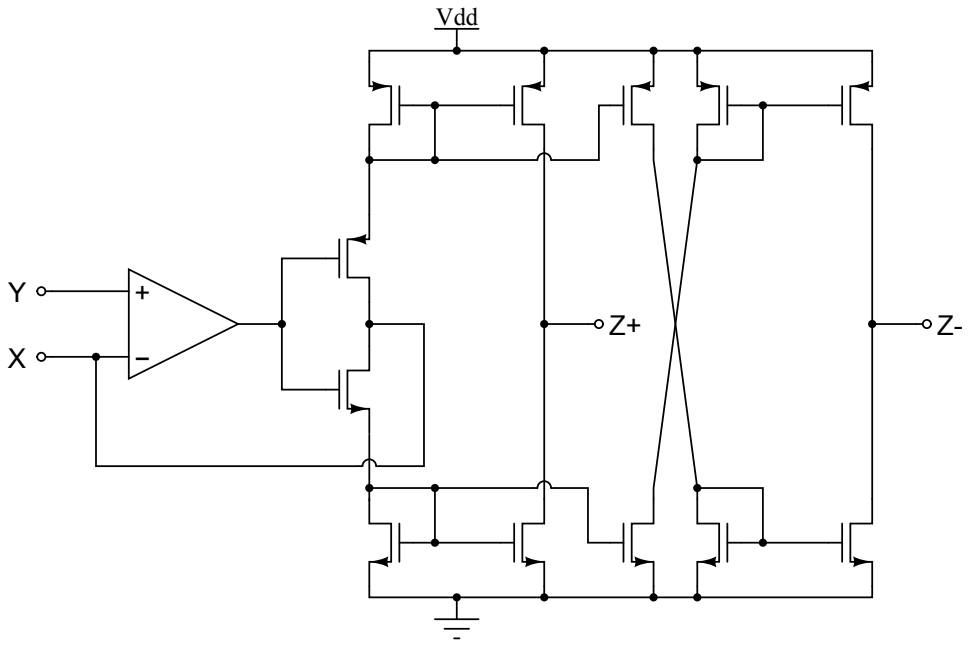


Figure 7.7: A second generation current conveyor circuit implemented using an op-amp to servo the voltage on node X. Complementary output nodes are included.

conveyor can be used as a general purpose element for the implementation of a wide range of circuit and computational functions [107]. However, the current feedback operational amplifier is the most common usage of the current conveyor due to the very high bandwidth they can provide for a given current consumption independently of the amplifier gain and input capacitance when compared to conventional voltage feedback amplifiers. A commercially available example of a current feedback op-amp is Analog Devices' AD844, but with a quiescent current of 1mA, a full power bandwidth of 300MHz and $5\text{pA}/\sqrt{\text{Hz}}$ input referred noise current is over-specified for a low power PPG system. This does however suggest that much lower current consumptions can be achieved for a much lower bandwidth implementation of the current conveyor if the noise current can be kept low enough. Along with high input current noise, the improved bandwidth comes at the cost of poor input current and voltage offsets, but as the PPG system has to be designed independently of input offsets anyway these disadvantages are relatively inconsequential in the design of the photocurrent amplifier so the CCII+ is a good match to the system requirements. In fact, Kim et. al. also compared the current conveyor to [103, 104] the transimpedance amplifiers although their results suggested that the input referred noise was an order of magnitude worse than for the capacitive feedback transimpedance amplifier. This is liekly due to the use of the

current conveyor for current amplification rather than using transimpedance gain at the output to boost the signal to noise ratio.

A further benefit of the current conveyor is that it is a current-mode circuit and hence is suitable for low-voltage operation. This allows for implementation in deep sub-micron CMOS processes as are used for manufacturing contemporary MCU devices. Using a process that is suitable for implementing low power MCUs means that the design can be relatively transferable to a range of system-on-chip solutions. The ultimate financial reward for an ASIC implementation of a pulse oximeter module would be the ability to licence the design to manufacturers that already have target markets with existing manufacturing and delivery infrastructure. Throughout this thesis the TSMC CM018 180nm process will be used. The 180nm process can be considered to be a relatively mature process and well suited to both low power digital IC implementation and for analogue cell design. Its suitability to mixed-signal MCU is highlighted by Energy Micro® choosing to use a 180nm process to implement their low power MCUs. The very latest MSP430 devices from Texas Instruments have moved to an ultra low leakage 130nm process in order to minimise the current draw during sleep modes, but 130nm processes are still relatively expensive for prototyping predominantly analogue designs.

The specific choice of CM018 was due to its availability on multi-project wafer (MPW) services at EUROPRACTICE, MOSIS and at the Hong Kong Science and Technology Parks which gives more financial options by having flexibility in the location of fabrication.

7.3.2 CCII Implementation

There are a number of possible ways of implementing a CCII, but to avoid going into very fine level analogue design the CCII based transimpedance amplifier is based on the implementation which uses a voltage feedback op-amp to stabilise the voltage and current-mirrors to convey the current to the output. Sticking to a more system level design might not lead to an optimum implementation of a CCII based transimpedance amplifier, but allows for quicker implementation and analysis.

The primary requirement of whatever topology is chosen is that it has a class-AB output stage to make the cell as general purpose as possible while still being applicable to PPG. A class-AB output stage allows rail-to-rail swing in the output voltage and the ability to both sink and source current though node Z. One possible implementation is proposed by Elwan and Soliman [108]. The stabilisation amplifier is implemented as a balanced operational transconductance amplifier (OTA) [67]. The balanced OTA is desirable for its simplic-

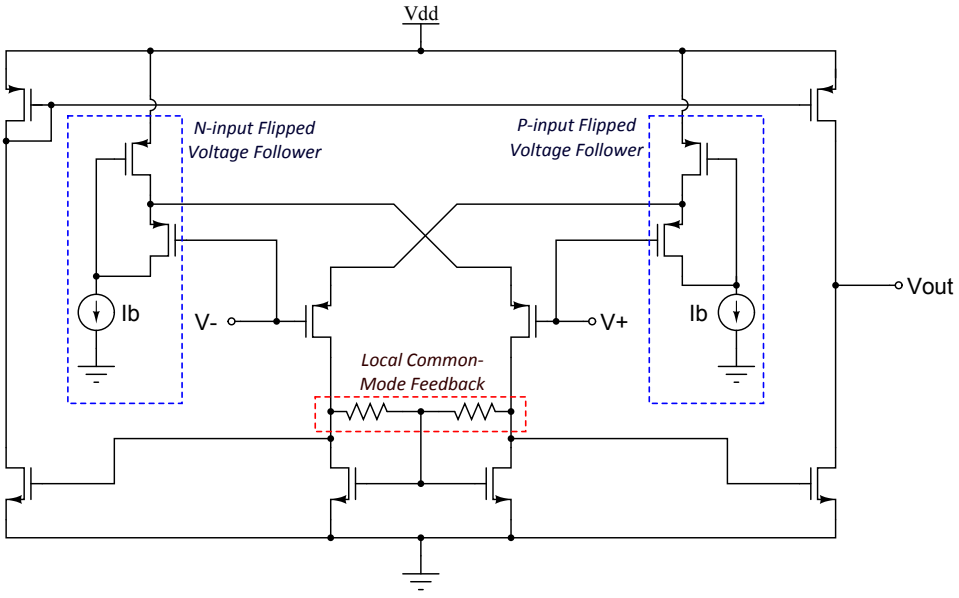


Figure 7.8: The *super-OTA*: a class-AB balanced OTA built using Flipped Voltage Followers and local common-mode feedback.

ity and low voltage operation. A further benefit for our application is that the construction of a CCII from a balanced OTA is done simply by creating copies of the output stage transistors and sharing the gate connections to create current mirrors. The main consideration with applying a balanced OTA to a transimpedance amplifier is that the current capability of the output stage is limited by the bias current of the differential pair. Setting the bias current to match the maximum expected input current of the transimpedance amplifier means wasting energy when the input current is low. To overcome this issue adaptive biasing can be added to allow the current consumption to scale with the magnitude of the input current of the cell. However, from a standby current of 56μA the maximum current available at the output node is only 300μA - a range factor of less than 6. A more usable current output would be up to 1mA which would be capable of driving the IR LED of the PPG system - even better would be up to 6mA for driving the red LED.

To improve the ratio of quiescent to maximum output current ratio the current conveyor will be built using a balanced OTA with a flipped voltage follower (FVF, [109]) in the input stage as proposed by Lopez-Martin as a "super" class-AB OTA [110]. The FVF has also been proposed for implementing a current conveyor without the super-OTA [111], but that design does not have a push-pull output which is desirable to make the current conveyor as flexible a building-block as possible. The intention of the FVF in the super-OTA is

to keep the voltage at the common node of the differential pair constant while supplying as much current as is required to satisfy the individual demands of each transistor in the differential pair. The super-OTA also uses local common mode feedback to boost gain and decrease the influence of common mode input signals on circuit operation. The circuit architecture of the super-OTA is shown in figure 7.8 where the two FVFs are highlighted in blue boxes. The cross-coupled FVFs work such that with zero differential input voltage each input transistor draws a quiescent current set by the intrinsic level-translation voltage in the FVF. Typically this is designed such that the input transistors receive the same bias current as the FVFs. As the differential input voltage increases one input transistor will turn off while the current through the other one will keep increasing until the upper FVF transistor is no longer operating in strong inversion and its current supply saturates.

7.3.3 Simulation

The CCII system has been designed and simulated in the Cadence IC5.1.41 environment and the results presented here are plots from the output of that tool. It is not within the scope of this thesis to present in depth circuit analysis of the proposed BSN ASIC solution - only to justify the system within the context of a PPG application. As such the details will be kept to a minimum.

The desired performance of the CCII is based on the specification in table 3.1 in chapter 3 with the parameters adapted for the current-mode nature of current conveyors. The primary requirements are that the CCII has a -3dB bandwidth greater than 400kHz in current follower mode and that the input referred noise current is less than $100\text{fA}/\sqrt{\text{Hz}}$ in the 250Hz region used for homodyne sensing. The 400kHz bandwidth is an order of magnitude larger than specified for the PCB solution such that the BSN ASIC can support much narrower current pulses should an application arise that requires it.

Gain

The gain performance is tested by simulating the current gain of the circuit from the X node to the Z node. A 100nF capacitor is connected to the output to make sure the amplifier is capable of slewing into a significantly large load. Figure 7.9 shows that the gain extends to the 400kHz requirement. The gain peak after 1MHz is due to some gain peaking in the FVF cells whose bandwidth does not necessarily line up with the resonant points in other parts of the circuit. The gain is also slightly less than unity with 3.15% attenuation which is insignificant in comparison with the intended transimpedance gain of $\geq$

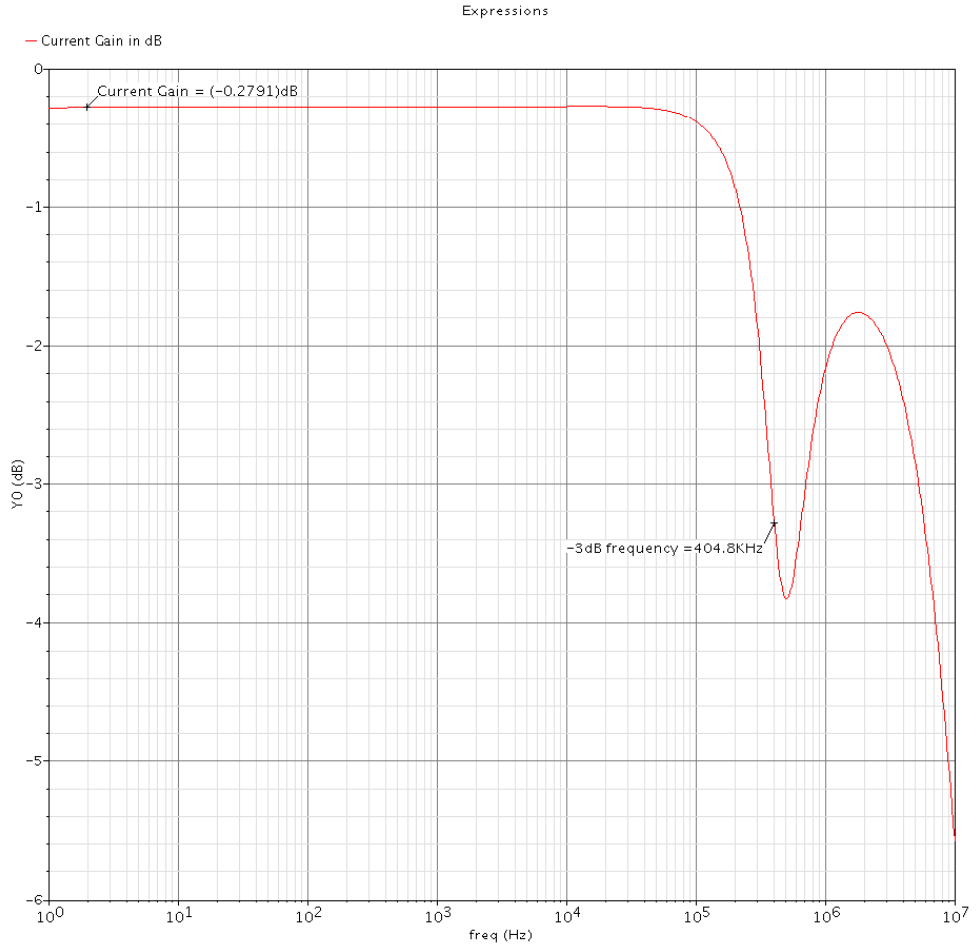


Figure 7.9: Simulated gain response of the proposed CCII based transimpedance amplifier.

1MV/A. Also, PPG has no quantified outputs except in the time domain so some gain inaccuracy is quite acceptable.

Transient

The phase has not been included here which is generally used to infer the stability of the amplifier. Instead, figure 7.10 shows the transient response of the CCII configured as a transimpedance amplifier with a 1μA input square-wave at 250Hz and a 1MΩ resistor connected to the Z node to implement the transimpedance gain. This plot shows two things: there is no overshoot from a current edge with a rise time of 2μs, and there is a significant offset of 47mV in the output. The lack of any overshoot proves that the phase response in current follower mode is acceptable. The presence of overshoot is not an issue for the PPG application due to the system level resilience to additive signal

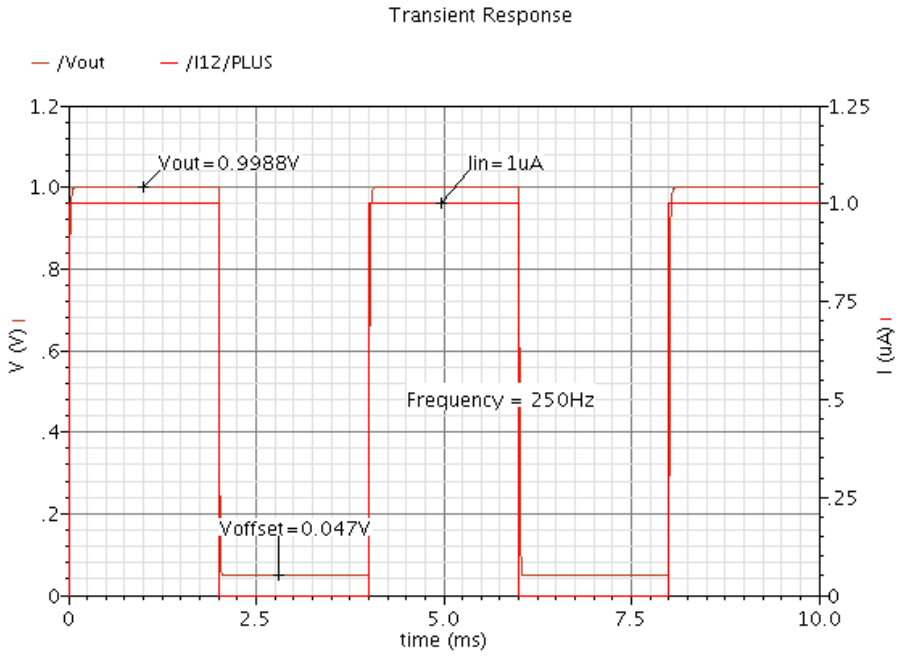


Figure 7.10: Simulated transient response of the proposed CCII based transimpedance amplifier.

artefacts in the photodetector stage. In this simulation the transimpedance gain is only 0.95MV/A suggesting the presence of finite output resistance although, as stated previously, the accuracy of the gain is not crucial to the performance of the system as a whole.

Noise

The circuit simulator outputs the noise voltage spectrum at a specified node so to estimate the input-referred noise current of the CCII it is necessary to configure it as a conventional transimpedance amplifier by connecting node Z to 0V via an ideal (i.e. noiseless) 1MΩ resistor and inferring the noise current. Figure 7.11 shows the noise voltage spectrum at node Z. In the low frequency part of the spectrum flicker noise is clearly visible with a corner frequency of about 1kHz which justifies the need for homodyne sensing. The vertical axis is plotted in decibels so, for reference, the noise floor at -164dB translates to 6.3nV/$\sqrt{\text{Hz}}$. If we assume a transimpedance gain of 0.95MV/A from the transient simulation this suggests an input referred noise current of 6.6fA/$\sqrt{\text{Hz}}$ which is well within the requirements of the system. There is some gain peaking at about 800kHz, but this is below 10nV/$\sqrt{\text{Hz}}$. Even down at 1Hz the noise density is still below the target of 100fA/$\sqrt{\text{Hz}}$ so it can be stated that the circuit noise performance is acceptable throughout the entire

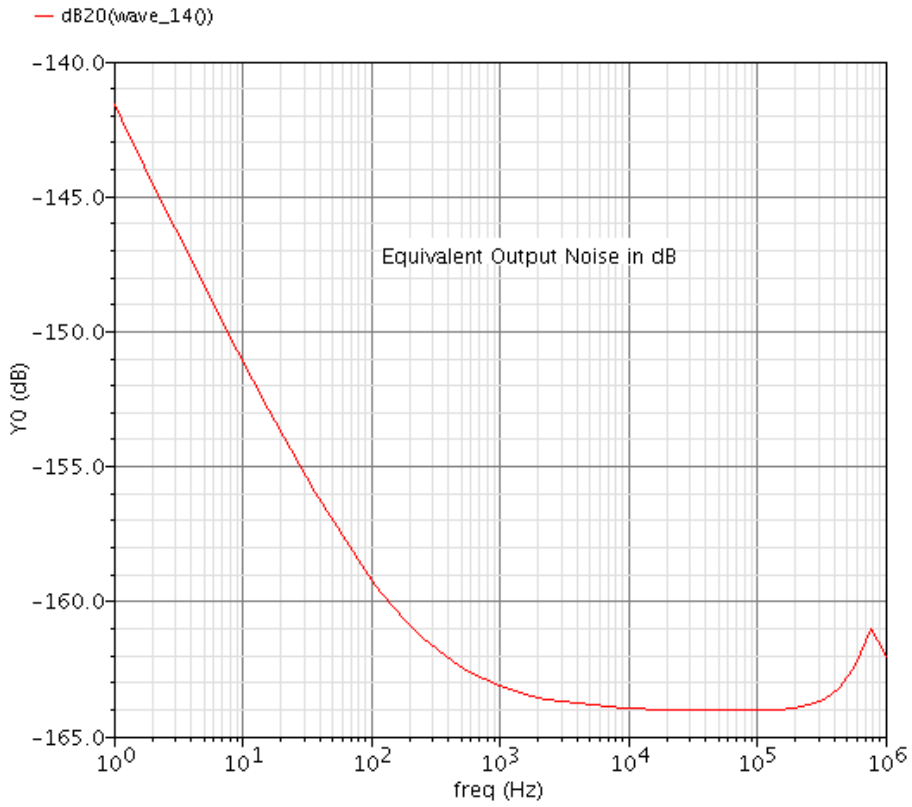


Figure 7.11: Simulated noise voltage spectrum at the output of the proposed CCII configured as a resistive transimpedance amplifier.

frequency range.

Power

As the ASIC is implemented with a 180nm process the voltage supply is set to 1.8V. Low supply voltage does constrain the sensing range of a transimpedance amplifier, but the as the noise is so low it is possible to sacrifice some gain to maintain dynamic range. The only significant constraint is that the forward voltage of RED LED used in version 3 of the PPG PCB has a forward voltage of 1.9V so some external circuitry will be required so that is can be driven from a higher supply voltage.

Each FVF is biased with a 200nA current source which leads to a quiescent current draw of 1.8μA for a CCII with complimentary outputs. There are only 6 current paths which are linked to the current flowing through node X and with internal scaling the total added current draw is equal to 5x the input current. For example, an input current of 1μA (the top end of the expected

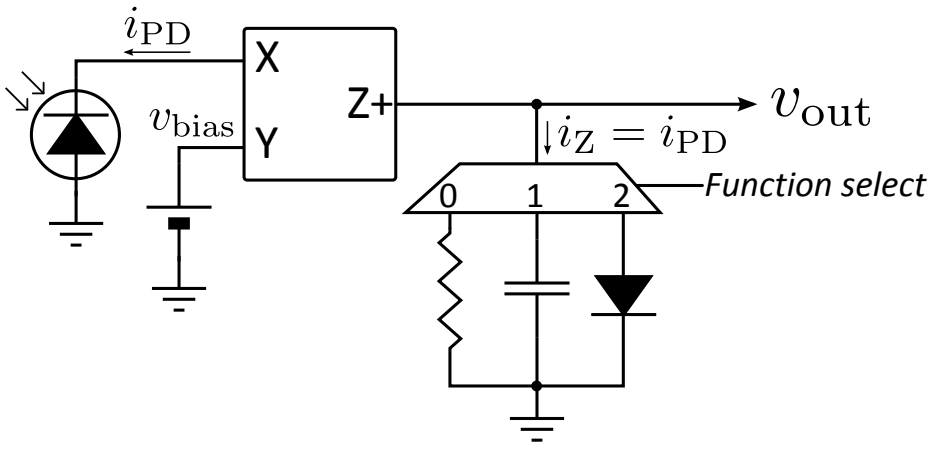


Figure 7.12: A transimpedance amplifier with a choice of linear, integrating and logarithmic transfer functions.

photocurrent range) would draw 6.8μA and consume 12.24μW of power. This can be compared to the static current consumption of the MCP6231 op-amp used in the PCB design which at 20μA would consume 36μW if operated at an equivalent 1.8V supply.

7.3.4 Alternative Transfer Functions

As with PCB based transimpedance amplifiers, it is possible to implement alternative transfer functions with the CCII+ element by changing the element which converts current to voltage, but with the added benefit that the element does not need to be used in the feedback path. For example, an integrating transimpedance amplifier can be constructed by replacing R_f of figure 7.6 with a capacitor C_f. Alternatively, R_f can be replaced with a diode to create a logarithmic response. For the purposes of flexibility the transimpedance element can either be an internal IC component or located off-chip, selectable with an analogue multiplexer.

7.3.5 Differential Amplification

Chapter 2 mentioned the issue of electrostatic interference in high gain transimpedance amplifiers. Differential input transimpedance amplifiers are possible in PCB implementations [32], but were discounted for this research due to the extra PCB area and power consumption required by the two additional op-amps. Furthermore, implementing any form of DC input current stabilisation is complicated by the presence of two input current nodes. This leaves electromagnetic screening as the only suitable solution to preventing charge

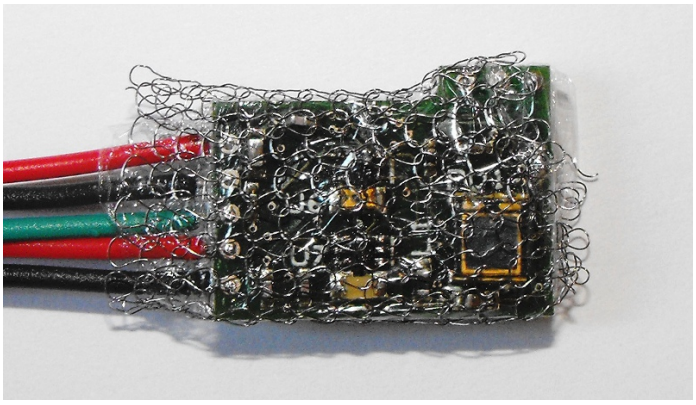


Figure 7.13: Wire mesh electromagnetic shielding required to protect singled ended transimpedance amplifier.

injection. Figure 7.13 shows the use of wire mesh to shield the pulse oximeter circuit board. The conductive mesh is bonded to the ground potential of the circuit with a solder joint on the reverse of the board. The gap between the wires is intended to allow the short wavelength optical EM waves to pass, but block the longer wavelength noise sources as well as providing a path for the wearer to electrostatically couple to the ground net instead of the input node of the photocurrent amplifier.

Moving to the ASIC domain means that size and power are much less of an issue so it is possible to revisit a differential input architecture. The current conveyor is again a suitable solution as it separates the bias voltage from the input current node. This allows two independent voltage sources to define the bias voltage across the photodiode without affecting the differential measurement of current through the photodiode. The circuit diagram for the differential version of the CCII based transimpedance amplifier is shown in figure 7.14.

7.3.6 Homodyne Detector

Translating the PCB-based mixed-signal homodyne detector to an integrated circuit design should not be a direct copy as this would make the IC architecture very specific to sensing systems that exhibit multiplicative relationships which restricts the number of sensing applications that the ASIC could be used for. For a more general purpose circuit that still supports homodyne detection it is possible to combine the multiplication stage of product detection into the current conveyor used for the transimpedance amplification stage. To match the fact that the 4 sample carrier only has the values or +1, -1 and 0 it is necessary to design the current conveyor with complimentary outputs.

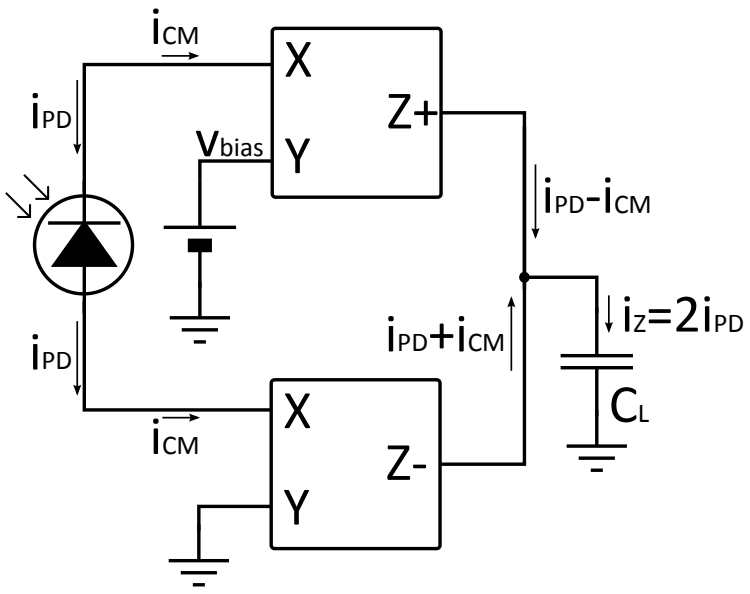


Figure 7.14: A differential transimpedance amplifier using a pair of current conveyor circuits. A capacitor is connected to the output to perform integrating transimpedance functionality.

For a second generation current conveyor this cell is labelled as CCII±. The system then only needs to select which of the CCII± outputs is used as the input to the integration capacitor according to the phase of the carrier. Figure 7.15 shows the reconfigurable CCII building block used in the BSN ASIC. To implement the discrete-time mixer node *CASC_OUT* would be used as the output node with the *Z_POLARITY* switch used to select either +1 or -1, and the *CASC_OUT_Z_EN* switch would override the normal *Z_OUT* signal to set the output current to 0.

One possible alternative is to implement a multiplication mode current conveyor as proposed by Hwang et. al. [112] which integrates a Gilbert multiplication cell into the current conveyor allowing an external voltage to scale the gain between the input current and the output current. This would only be needed if carrier had to be more complex than the proposed 4 sample sinusoid.

Filtering

The low-pass filter of the product detector could be implemented purely in analogue circuitry on the IC, but to make the system configurable requires increasing the complexity of the design. Also, the analogue filter would either need to operate constantly at a very low current or would need to be able to be turned on and off at the sampling events associated with the discrete-time

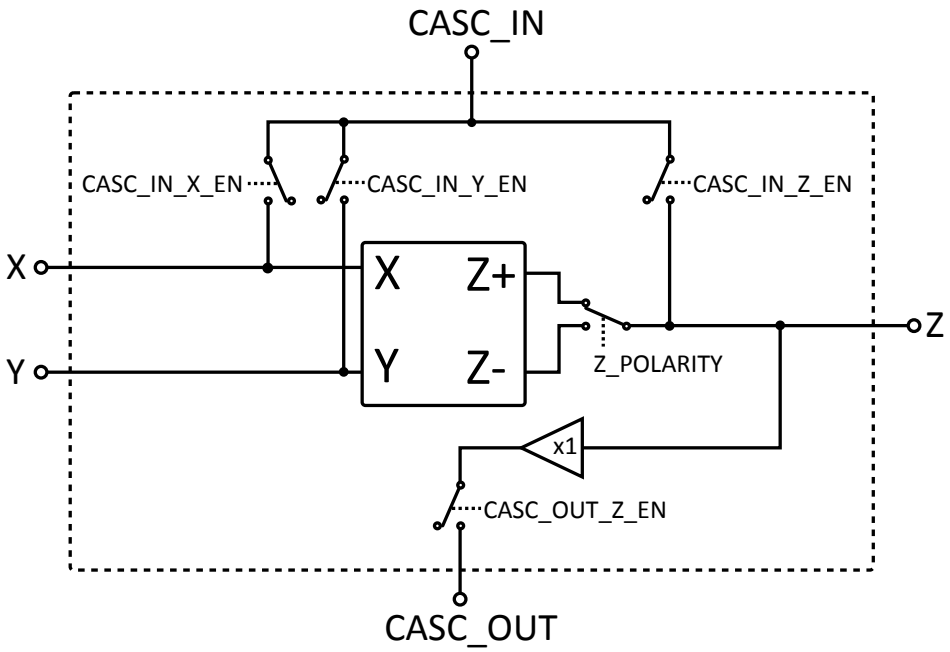


Figure 7.15: Reconfigurable current conveyor cell. Complimentary outputs enable the implementation of the coarse product detection required for homodyne sensing.

nature of TDM used for multiple optical channels. The simpler and more generic approach would be to perform an analogue to digital conversion on the output of the integration stage and implement the filter as an LWDF in either re-configurable digital logic or in software as in the previous section. This allows the system to take advantage of the low current consumption of CMOS circuitry while no changes are occurring as well as being much more configurable than a purely analogue implementation. This trade-off is dependent on the ability to convert from analogue to digital using as little energy as possible and the ability of the processor code to enter a low power sleep mode and wake up again in a short period of time. Fortunately these are both achievable features with the BSN ASIC ADC specified to work at 180μW and the processor specified to wake-up in less than 10μs.

Carrier Synthesis

Using an ITIA for photocurrent amplification means that it is only necessary to provide PWM outputs from a logic core in order to modulate the average LED drive current. PWM is a very common peripheral for MCU cores so carrier synthesis does not need specific design consideration for IC implementation. However, the LED still needs to have its peak current controlled and for this

it is acceptable to reuse the OTA core of the CCII to implement the voltage to current convertor originally proposed for homodyne detection in chapter 4 (figure 4.4).

7.4 IC Summary

This section has been more of a commentary of the methods that are being used to translate the PPG system from the PCB domain to the IC design domain. The content is intended to show that it is possible to build the PPG system using only an MCU core and current conveyors. Simulation results have been presented to show that the core cells used in the ASIC meet the gain, bandwidth and noise requirements while consuming less than half the power of the equivalent op-amp based PCB design.

The primary outcome from researching analogue circuit cells for implementing a BSN ASIC has been the selection of the second generation current-conveyor as a general purpose processing element that can be used as both a sensor interface as well as signal processing elements. Being a current-mode circuit element means that it is naturally suited to low-voltage, high dynamic range applications, as well as having enhanced immunity to parasitic capacitance when compared to equivalent voltage-mode cells. This chapter has focussed on using the CCII as a photodiode interface, and has set the bias conditions of the circuit to match those requirements, but it can also be used as an interface to voltage mode sensors as well, e.g. for implementing instrumentation amplifiers for ECG applications [113, 114], as well as for chemical sensing applications [115] and microelectrode arrays [116].

What still remains to be addressed in the BSN ASIC project is the application of a CCII as an information processing element for the purposes of enhancing on-node data analysis. The ability to use a current conveyor for implementing a wide range of functions, e.g. filters and rectifiers, has been known for many years [117], but more recently they have been adopted into artificial neural network designs [118] and the addition of a multiplication function [112] opens up many interesting opportunities for the future without having to add significantly to the current library of ASIC cells.

The *super-OTA* has been chosen as the core element of the CCII in order to enhance the flexibility of the circuit. On its own the super-OTA can be used as a rudimentary op-amp, but where it excels is its ability to adapt its biasing currents in order to match the requirements of the output by using flipped voltage followers. The output stage is also trivial to mirror meaning that very little extra circuitry is required to implement a CCII with complementary

outputs. As such we have implemented a cell that has >400kHz of bandwidth while being able to sink and source over 1mA of current while requiring $<10\mu$A of current under quiescent conditions.

Overall, while still in an early development phase, the CCII shows promise as a very useful, low power processing element for the implementation of pervasive sensing systems while minimising the amount of circuit elements that need to be designed.

8 Conclusions

This thesis has presented work which has addressed issues of size, power consumption, artefact rejection and data compression as a coherent system for implementing a wearable PPG sensing for pervasive cardiac health monitoring on a manufacturable format that is accessible on a global scale by both high and low income countries. The focus of this thesis has been that a cardiovascular monitoring solution cannot be created from individual functional components developed in isolation; each aspect of the system must be developed from the bottom up with a uniform methodology. In the context of this thesis the PPG system is implemented exclusively in the discrete-time domain. Even the analogue components of the LED driver and photocurrent amplifier are discrete-time although the circuit analysis is performed in the continuous time domain such that it is familiar to readers experienced in analogue circuit design.

The single time domain methodology has simplified the physical implementation of the system allowing most of it to be implemented in firmware. To the knowledge of the author no other researchers have managed to implement a complete PPG system in so few components without resorting to designing an ASIC.

8.1 Contributions

The research performed in the field of wearable PPG systems has resulted in the original contributions listed below:

- Analysis and design methodology for implementing an integrating transimpedance amplifier in PCB format with particular focus on the effect of parasitic input capacitance and micro-power op-amps with low gain-bandwidth products. A amplifier design is presented which has a transimpedance gain of 3MV/A using a single 20μA op-amp and analogue switch which has 600kHz bandwidth and $<$1nA$_{\mathrm{RMS}}$ of input referred current noise.

- A homodyne detection system for the rejection of light from sources independent from the sensor. The signal processing is performed completely

in the discrete-time domain such that it can be performed in a low power MCU. A 5th order filter is used to provide 96dB of stop band rejection and the entire signal chain executes in <30μs on an ARM Cortex-M3 equipped MCU. Placing the carrier frequency at 250Hz avoids a number of ambient noise sources as well as allowing the integrating transimpedance amplifier to operate in a low noise region of its spectrum.

- A low frequency/DC current rejection control loop that uses the spectrum remaining from the homodyne process which can have a faster response time to the current state-of-the-art DC rejection system while also preserving DC PPG information. The circuit is configured to subtract up to 8x as much current as would normally saturate the transimpedance amplifier to achieve a very high dynamic range system.

- An embedded signal analysis algorithm that delivers a scalable relationship between desired signal precision and data transmission bandwidth while using negligible MCU resources by delineating the cyclic nature of the PPG signal. Included is an interpolation procedure for accurate reconstruction of the delineated signal in the time domain at the receiver end of the communication channel. The method is capable of compressing a 14kbps data stream to 3kbps with a peak detection granularity of 1pA.

- A complete PPG sensor for integration into a patch-type product using standard PCB technology and COTS components. All of the previously mentioned functionality is implemented on the node while consuming <7.5mW (excluding communications). From a small 55mAh battery this allows the sensor to be operated for one waking day without intervention.

- An implementation of the current-conveyor analogue IC building block using a 180nm CMOS fabrication process that can implement the integrating transimpedance amplifier circuit in only 6.8μA with comparable bandwidth and noise performance to the 20μA op-amp based circuit. The current-conveyor has the advantage that the passive component that implements the transimpedance gain is physically separated from the photodiode and any associated parasitic capacitance. The current mode nature of the devices also makes it trivial to implement a differential input structure in order to reject electromagnetic and electrostatic interference from the low impedance input node.

8.2 Future Work

Significant progress has been made towards making PPG an adoptable technique for pervasive cardiac monitoring by addressing the issues of power, additive noise and size. Most significant of the remaining requirements for end-user adoption is reliability, both in terms of the quality of data produced by the sensor and the infrastructure for preserving the privacy and security of the subject's data. In order to validate the proposed system's suitability the following challenges and opportunities for further research remain:

Human Factors and Ergonomics

The transducer and electronics needs to be packaged in a format that is both desirable to the end-user but functional in providing quality PPG signals and being resilient to artefacts. This is where this PPG system still lags behind the systems from RWTH Aachen [60] and DTU [63].

System Validation

Once the system has be adequately packaged then a thorough set of user-based experiments and trials need to be performed to prove that the developed algorithms operate outside of the lab in real world situations. Consideration needs to be given to what type of cohort a given system implementation is tested with as a system optimised for use by athletes will unlikely to meet the requirements of CVD sufferers.

Enhanced Data Analysis

Delineation of the raw data does provide a significant reduction in data rate and heart rate can be derived from the output, but it is still a relatively naive algorithm. A very useful metric that has yet to be well defined is signal quality which can be readily extracted from the morphological features already present in the quadrature of the waveform as is done by Sukor et al. [99]. Other physiological metrics that are theoretically already present in the the delineated signal. For example, as homodyne sensing and extrema detection preserve DC information respiration rate should be recoverable if it is present in the baseline blood volume variations.

Sensor Fusion

Chapter 5 presented results which question the ability of the system to reject motion induced artefacts which are derived from actual changes in blood vol-

ume that did not originate from within the body. In the absence of any extra information it is unlikely that this problem will be adequately solved, but the use of accelerometers has opened up the opportunity to use the extra contextual information to cancel out motion artefacts. To date most accelerometer based solutions assume an additive model for the noise coupling, but the system model strongly suggests there is both additive and multiplicative coupling; additive coupling via motion modulation of non-coherent light sources and multiplicative coupling via modulation of tissue and blood volume.

Optical System Design

The proposed wearable PPG system constrained itself to the standard red and IR optical wavelengths as these are the standard choices for transmission mode pulse oximeters. The validity of these wavelengths have been questioned as far back as 1992 [119]. In addition to measuring oxygen the development of CO-oximetry by Masimo [120] has highlighted the opportunities presented by measuring at multiple wavelengths to which the homodyne sensing technique scales well thanks to its ability to distinguish channels in both time and frequency domains.

Power Source

Having to either recharge or replace a wearable sensor's battery once per day will not be a very attractive feature to users who want to adopt the technology for its long-term monitoring functionality. One solution would be to further reduce the power consumption of the system to the point where the sensor can operate without intervention for days at a time. There is, however, likely to be a diminishing return on reducing LED brightness and increasing photodetector gain as designs approach limits in signal to noise ratio due to fundamental optical attentuation in either transmission or reflection mode PPG. An attractive alternative would be to replace the use of batteries with a power source which is pervasive in its power delivery such as harvesting energy from the user's motion[10]. As it stands, the pulse oximeter requires more power than miniature energy scavenging systems can provide, but as power requirements and deliverable power capabilities converge wearable sensor solutions should emerge that are more attractive to end users.

Complete ASIC Solution

If user trials prove successful then it is likely that the expense of developing an ASIC can be justified as suitably low risk. Many of the decisions made early on

in this research were based on the assumption that a multichannel had to be implemented using time division modulation to minimise component count. This is not a constraint in IC design so it would be possible to design the system from the bottom up in the continuous-time domain while still making use of the integrating transimpedance amplification, homodyne sensing and DC current stabilistion methods used in the discrete-time domain. Even the delineator can be implemented in continuous time and the output from the circuit would be completely event-driven and take advantage of that exciting field of digital circuit research [121].

Bibliography

[1] A. K. Y. Wong, K.-P. Pun, Y.-T. Zhang, and K. N. Leung, "A Low-Power CMOS Front-End for Photoplethysmographic Signal Acquisition With Robust DC Photocurrent Rejection," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 2, pp. 280–288, Dec. 2008.

[2] United Nations Dept. of Economic and Social Affairs. Population Division, "World Population Ageing: 1950-2050," tech. rep., United Nations Department of Economic and Social Affairs, Population Division, 2002.

[3] G.-Z. Yang, ed., *Body Sensor Networks*. London: Springer London, 2006.

[4] World Health Organization, "WHO Patient Safety Pulse Oximetry Project," tech. rep.

[5] Y. Yih, ed., *Handbook of Healthcare Delivery Systems*. CRC Press, 2010.

[6] J. A. Patterson, D. C. McIlwraith, and G.-Z. Yang, "A Flexible, Low Noise Reflective PPG Sensor Platform for Ear-Worn Heart Rate Monitoring," in *2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, pp. 286–291, IEEE, June 2009.

[7] J. A. C. Patterson and G.-Z. Yang, "Ratiometric Artefact Reduction in Low Power, Discrete-Time, Reflective Photoplethysmography," in *2010 International Conference on Body Sensor Networks*, pp. 174–179, IEEE, June 2010.

[8] J. A. C. Patterson and G.-Z. Yang, "Ratiometric Artifact Reduction in Low Power Reflective Photoplethysmography," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 5, pp. 330–338, Aug. 2011.

[9] J. A. Patterson and G.-Z. Yang, "Dual-Mode Additive Noise Rejection in Wearable Photoplethysmography," in *2012 Ninth International Conference on Wearable and Implantable Body Sensor Networks*, pp. 97–102, IEEE, May 2012.

[10] P. Mitcheson, E. Yeatman, G. Rao, A. Holmes, and T. Green, "Energy Harvesting From Human and Machine Motion for Wireless Electronic Devices," *Proceedings of the IEEE*, vol. 96, pp. 1457–1486, Sept. 2008.

[11] G. J. Pottie, "Wireless sensor networks," in *Information Theory Workshop*, pp. 139–140, 1998.

[12] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Next Century Challenges: Mobile Networking for "Smart Dust"," 1999.

[13] SparkFun Electronics, Inc., "110mAh Lithium Polymer Battery." http://www.sparkfun.com/datasheets/Batteries/UnionBattery-110mAh.pdf.

[14] L. Seon-Woo and K. Mase, "Activity and location recognition using wearable sensors," *Pervasive Computing, IEEE*, vol. 1, no. 3, pp. 24–32, 2002.

[15] O. L. Macsorley, "High-Speed Arithmetic in Binary Computers," *Proceedings of the IRE*, vol. 49, no. 1, pp. 67–91, 1961.

[16] B. Gilbert, "A precise four-quadrant multiplier with subnanosecond response," *Solid-State Circuits, IEEE Journal of*, vol. 3, pp. 365–373, Dec. 1968.

[17] L. Chua and L. Yang, "Cellular neural networks: theory," *IEEE Transactions on Circuits and Systems*, vol. 35, no. 10, pp. 1257–1272, 1988.

[18] B. Sheu, J. Choi, and C.-F. Chang, "An analog neural network processor for self-organizing mapping," in *1992 IEEE International Solid-State Circuits Conference Digest of Technical Papers*, IEEE.

[19] C. Mead, *Analog VLSI and neural systems*. Addison-Wesley Longman Publishing Co., Inc., 1989.

[20] L. Wang, S. Thiemjarus, B. Lo, and G.-Z. Yang, "Toward a mixed-signal reconfigurable ASIC for real-time activity recognition," in *Medical Devices and Biosensors, 2008. ISSS-MDBS 2008. 5th International Summer School and Symposium on*, pp. 227–230, 2008.

[21] P. Scarborough, K. Wickramasinghe, P. Bhatnagar, and M. Rayner, "Trends in coronary heart disease, 1961-2011," tech. rep., British Heart Foundation, Oct. 2011.

[22] R. F. Yazicioglu, T. Torfs, J. Penders, I. Romero, H. Kim, P. Merken, B. Gyselinckx, H. J. Yoo, and C. Van Hoof, "Ultra-low-power wearable biopotential sensor nodes.," *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 2009, pp. 3205–8, Jan. 2009.

[23] J. Yoo and H.-J. Yoo, "Emerging low energy Wearable Body Sensor Networks using patch sensors for continuous healthcare applications.," *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 2010, pp. 6381–4, Jan. 2010.

[24] S. S. Lobodzinski and M. M. Laks, "New devices for very long-term ECG monitoring.," *Cardiology journal*, vol. 19, pp. 210–4, Jan. 2012.

[25] T. Schermer, J. Leenders, H. in 't Veen, W. van den Bosch, A. Wissink, I. Smeele, and N. Chavannes, "Pulse oximetry in family practice: indications and clinical observations in patients with COPD.," *Family practice*, vol. 26, pp. 524–31, Dec. 2009.

[26] Y. Mendelson and C. Pujary, "Measurement site and photodetector size considerations in optimizing power consumption of a wearable reflectance pulse oximeter," in *Proceedings of the 25th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE Cat. No.03CH37439)*, vol. 4, pp. 3016–3019, IEEE, 2003.

[27] T. Aoyagi, "Pulse oximetry: its invention, theory, and future.," *Journal of anesthesia*, vol. 17, pp. 259–266, Jan. 2003.

[28] J. E. Sinex, "Pulse oximetry: Principles and limitations," *The American Journal of Emergency Medicine*, vol. 17, pp. 59–66, Jan. 1999.

[29] P. D. Mannheimer, "The light-tissue interaction of pulse oximetry.," *Anesthesia and analgesia*, vol. 105, pp. S10–7, Dec. 2007.

[30] J. Allen, "Photoplethysmography and its application in clinical physiological measurement.," *Physiological measurement*, vol. 28, pp. R1–39, Mar. 2007.

[31] V. König, R. Huch, and A. Huch, "Reflectance Pulse Oximetry Principles and Obstetric Application in the Zurich System," *Journal of Clinical Monitoring and Computing*, vol. 14, pp. 403–412, 1998.

[32] J. G. Graeme, *Photodiode Amplifiers: OP AMP Solutions*. McGraw-Hill Professional, 1995.

[33] B. Pease, "What's All This Transimpedance Amplifier Stuff, Anyhow? (Part 1)," 2001.

[34] M. Hayes, "A nonlinear optical preamplifier for sensing applications," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 49, no. 1, pp. 1–9, 2002.

[35] D. Hoe and D. Ribner, "An auto-ranging photodiode preamplifier with 114 dB dynamic range," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 2, pp. 187–194, 1996.

[36] C. Ciofi, F. Crupi, C. Pace, and G. Scandurra, "How to Enlarge the Bandwidth Without Increasing the Noise in OP-AMP-Based Transimpedance Amplifier," *IEEE Transactions on Instrumentation and Measurement*, vol. 55, pp. 814–819, June 2006.

[37] V. M. Joyner and Y. Zhang, "A CMOS Analog Front-end Receiver with Desensitization to Input Capacitance for Broadband Optical Wireless Communication," *Journal of Communications*, vol. 4, pp. 572–581, Sept. 2009.

[38] L. Rovati, P. Lago, and L. Priori, "High performance single stage shaper-preamplifier for optical gated signals," *Electronics Letters*, vol. 30, no. 20, p. 1709, 1994.

[39] Texas Instruments, "Burr-Brown IVC102: Precision Switch Integrator Transimpedance Amplifier." http://www.ti.com/lit/ds/symlink/ivc102.pdf, 1996.

[40] Texas Instruments, "Burr-Brown DDC112: Dual Current Input 20-Bit analog-to-digital converter." http://www.ti.com/lit/ds/symlink/ddc112.pdf, 2004.

[41] M. J. Hayes and P. R. Smith, "Artifact Reduction in Photoplethysmography," *Applied Optics*, vol. 37, p. 7437, Nov. 1998.

[42] M. Tavakoli, L. Turicchia, and R. Sarpeshkar, "An Ultra-Low-Power Pulse Oximeter Implemented With an Energy-Efficient Transimpedance Amplifier," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 4, pp. 27–38, Feb. 2010.

[43] T. Delbruck and C. Mead, "Adaptive photoreceptor with wide dynamic range," in *Proceedings of IEEE International Symposium on Circuits and Systems - ISCAS '94*, vol. 4, pp. 339–342, IEEE, 1994.

[44] A. A. Dorrington and R. Künnemeyer, "A Simple Microcontroller Based Digital Lock-in Amplifier for the Detection of Low Level Optical Signals," p. 486, Jan. 2002.

[45] K. Gan, E. Zahedi, and M. A. M. Ali, "Design and Implementation of A PC-Based Digital Synchronous Detection System for Biological Signal Measurement," in *2007 5th Student Conference on Research and Development*, pp. 1–4, IEEE, 2007.

[46] M. R. Grubb, B. R. Hayes-Gill, J. A. Crowe, D. Sharkey, N. Marlow, and N. J. Miles, "Towards Forehead Reflectance Photoplethysmography to Aid Delivery Room Resuscitation in Newborns," in *4th European Conference of the International Federation for Medical and Biological Engineering* (J. Sloten, P. Verdonck, M. Nyssen, J. Haueisen, and R. Magjarevic, eds.), vol. 22 of *IFMBE Proceedings*, pp. 1053–1056, Springer Berlin Heidelberg, 2009.

[47] S. J. Barker and S. Morgan, "A Laboratory Comparison of the Newest "Motion-Resistant" Pulse Oximeters During Motion and Hypoxemia," *Anesthesia and Analgesia*, vol. 98, no. S2, 2004.

[48] B. S. Kim and S. K. Yoo, "Motion artifact reduction in photoplethysmography using independent component analysis.," *IEEE transactions on bio-medical engineering*, vol. 53, pp. 566–8, Mar. 2006.

[49] A. R. Relente and L. G. Sison, "Characterization and adaptive filtering of motion artifacts in pulse oximetry using accelerometers," in *[Engineering in Medicine and Biology, 2002. 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society] EMBS/BMES Conference, 2002. Proceedings of the Second Joint*, vol. 2, pp. 1769–1770 vol.2, 2002.

[50] H. H. Asada, J. Hong-Hui, and P. Gibbs, "Active noise cancellation using MEMS accelerometers for motion-tolerant wearable bio-sensors," in *Engineering in Medicine and Biology Society, 2004. IEMBS '04. 26th Annual International Conference of the IEEE*, vol. 1, pp. 2157–2160, 2004.

[51] P. T. Gibbs, "Active motion artifact cancellation for wearable health monitoring sensors using collocated MEMS accelerometers," in *Proceedings of SPIE*, vol. 5765, (San Diego, CA, USA), pp. 811–819, SPIE, 2005.

[52] J. M. Graybeal and M. T. Petterson, "Adaptive filtering and alternative calculations revolutionizes pulse oximetry sensitivity and specificity during motion and low perfusion.," *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 7, pp. 5363–6, Jan. 2004.

[53] M. J. Hayes and P. R. Smith, "A new method for pulse oximetry possessing inherent insensitivity to artifact," *Biomedical Engineering, IEEE Transactions on*, vol. 48, no. 4, pp. 452–461, 2001.

[54] P. S. Addison and J. N. Watson, "Oxygen saturation determined using a novel wavelet ratio surface," *Medical Engineering & Physics*, vol. 27, no. 3, pp. 245–248, 2004.

[55] J. Allen and A. Murray, "Variability of photoplethysmography peripheral pulse measurements at the ears, thumbs and toes," *IEE Proceedings - Science, Measurement and Technology*, vol. 147, no. 6, p. 403, 2000.

[56] S. Rhee, B.-H. H. Yang, and H. H. Asada, "Artifact-resistant power-efficient design of finger-ring plethysmographic sensors.," *IEEE transactions on bio-medical engineering*, vol. 48, pp. 795–805, July 2001.

[57] P. Branche, W. Johnston, C. Pujary, and Y. Mendelson, "Measurement reproducibility and sensor placement considerations in designing a wearable pulse oximeter for military applications," in *IEEE 30th Annual Northeast Bioengineering Conference, 2004. Proceedings of the*, pp. 216–217, IEEE, 2004.

[58] U. Anliker, J. A. Ward, P. Lukowicz, G. Tröster, F. Dolveck, M. Baer, F. Keita, E. B. Schenker, F. Catarsi, L. Coluccini, A. Belardinelli, D. Shklarski, M. Alon, E. Hirt, R. Schmid, and M. Vuskovic, "AMON: a wearable multiparameter medical monitoring and alert system.," *IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 8, pp. 415–27, Dec. 2004.

[59] D. B. McCombie, A. T. Reisner, and H. H. Asada, "Adaptive blood pressure estimation from wearable PPG sensors using peripheral artery pulse wave velocity measurements and multi-channel blind identification of local arterial dynamics.," *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 1, pp. 3521–4, Jan. 2006.

[60] S. Vogel, M. Hülsbusch, T. Hennig, V. Blazek, and S. Leonhardt, "In-ear vital signs monitoring using a novel microoptic reflective sensor.," *IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 13, pp. 882–9, Nov. 2009.

[61] M.-Z. Poh, N. C. Swenson, and R. W. Picard, "Motion-tolerant magnetic earring sensor and wireless earpiece for wearable photoplethysmography.," *IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 14, pp. 786–94, May 2010.

[62] Y. Mendelson, R. J. Duckworth, and G. Comtois, "A wearable reflectance pulse oximeter for remote physiological monitoring.," *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 1, pp. 912–5, Jan. 2006.

[63] R. G. Haahr, S. B. Duun, M. H. Toft, B. Belhage, J. Larsen, K. Birkelund, and E. V. Thomsen, "An Electronic Patch for Wearable Health Monitoring by Reflectance Pulse Oximetry," *IEEE Transactions on Biomedical Circuits and Systems*, no. 99, pp. 1–1, 2011.

[64] H. Asada, P. Shaltis, A. Reisner, and R. Hutchinson, "Mobile monitoring with wearable photoplethysmographic biosensors," *IEEE Engineering in Medicine and Biology Magazine*, vol. 22, pp. 28–40, May 2003.

[65] A. Dhukaram, C. Baber, L. Elloumi, B.-J. van Beijnum, and P. De Stefanis, "End-User perception towards pervasive cardiac healthcare services: Benefits, acceptance, adoption, risks, security, privacy and trust," in *2011 5th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth) and Workshops*, 2011.

[66] Y. Mendelson and B. D. Ochs, "Noninvasive pulse oximetry utilizing skin

reflectance photoplethysmography.," *IEEE transactions on bio-medical engineering*, vol. 35, pp. 798–805, Oct. 1988.

[67] P. E. Allen and D. R. Holberg, *CMOS Analog Circuit Design*, vol. 269. Oxford University Press, 2002.

[68] E. A. Peláez and E. R. Villegas, "LED power reduction trade-offs for ambulatory pulse oximetry.," *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference*, vol. 2007, pp. 2296–9, Jan. 2007.

[69] J. Fischer, "Noise sources and calculation techniques for switched capacitor filters," *IEEE Journal of Solid-State Circuits*, vol. 17, pp. 742–752, Aug. 1982.

[70] C.-A. Gobet and A. Knob, "Noise analysis of switched capacitor networks," *IEEE Transactions on Circuits and Systems*, vol. 30, pp. 37–43, Jan. 1983.

[71] L. Toth, I. Yusim, and K. Suyama, "Noise analysis of ideal switched-capacitor networks," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 46, pp. 349–363, Mar. 1999.

[72] O. Oliaei, "Thermal noise analysis of multi-input SC-integrators for delta-sigma modulator design," in *2000 IEEE International Symposium on Circuits and Systems. Emerging Technologies for the 21st Century. Proceedings (IEEE Cat No.00CH36353)*, vol. 4, pp. 425–428, Presses Polytech. Univ. Romandes, 2000.

[73] O. Oliaei, "Noise analysis of correlated double sampling sc integrators with a hold capacitor," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 50, pp. 1198–1202, Sept. 2003.

[74] N. Stuban and M. Niwayama, "Optimal filter bandwidth for pulse oximetry.," *The Review of scientific instruments*, vol. 83, p. 104708, Oct. 2012.

[75] B. Lo, "BSN Development Kit v3.." http://ubimon.doc.ic.ac.uk/bsn/a1892.html.

[76] K. N. Glaros and E. Drakakis, "Trade-offs for low power integrated pulse oximeters," in *2009 IEEE Biomedical Circuits and Systems Conference*, pp. 245–248, IEEE, Nov. 2009.

[77] C. Turner, "Recursive discrete-time sinusoidal oscillators," *IEEE Signal Processing Magazine*, vol. 20, pp. 103–111, May 2003.

[78] J. E. Volder, "The CORDIC Trigonometric Computing Technique," *IEEE Transactions on Electronic Computers*, vol. EC-8, pp. 330–334, Sept. 1959.

[79] Y. Hu, "CORDIC-based VLSI architectures for digital signal processing," *IEEE Signal Processing Magazine*, vol. 9, pp. 16–35, July 1992.

[80] L. Gazsi, "Explicit formulas for lattice wave digital filters," *IEEE Transactions on Circuits and Systems*, vol. 32, pp. 68–88, Jan. 1985.

[81] A. Fettweis, "Wave digital filters: Theory and practice," *Proceedings of the IEEE*, vol. 74, no. 2, pp. 270–327, 1986.

[82] A. B. Williams and F. J. Taylor, *Electronic filter design handbook*. McGraw-Hill, 2006.

[83] M. Mandal and A. Asif, *Continuous and Discrete Time Signals and Systems*. Cambridge University Press, 2007.

[84] A. Wong and K. Hung, "A near-infrared heart rate measurement IC with very low cutoff frequency using current steering technique," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 52, pp. 2642–2647, Dec. 2005.

[85] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," *IEEE Transactions on Information Theory*, vol. 23, pp. 337–343, May 1977.

[86] J. Ziv and A. Lempel, "Compression of individual sequences via variable-rate coding," *IEEE Transactions on Information Theory*, vol. 24, pp. 530–536, Sept. 1978.

[87] T. Welch, "A Technique for High-Performance Data Compression," *Computer*, vol. 17, pp. 8–19, June 1984.

[88] H. Wang, K. Siu, K. Ju, and K. H. Chon, "A high resolution approach to estimating time-frequency spectra and their amplitudes.," *Annals of biomedical engineering*, vol. 34, pp. 326–38, Feb. 2006.

[89] S. Dash, K. H. Shelley, D. G. Silverman, and K. H. Chon, "Estimation of respiratory rate from ECG, photoplethysmogram, and piezoelectric pulse transducer signals: a comparative study of time-frequency methods.,"

IEEE transactions on bio-medical engineering, vol. 57, pp. 1099–107, May 2010.

[90] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung, and H. H. Liu, "The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 454, pp. 903–995, Mar. 1998.

[91] K. A. Reddy, B. George, and V. J. Kumar, "Use of Fourier Series Analysis for Motion Artifact Reduction and Data Compression of Photoplethysmographic Signals," *Instrumentation and Measurement, IEEE Transactions on*, vol. 58, pp. 1706–1711, May 2009.

[92] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, pp. 1289–1306, Apr. 2006.

[93] P. K. Baheti and H. Garudadri, "An Ultra Low Power Pulse Oximeter Sensor Based on Compressed Sensing," in *2009 Sixth International Workshop on Wearable and Implantable Body Sensor Networks*, pp. 144–148, IEEE, June 2009.

[94] W. Zong, T. Heldt, G. Moody, and R. Mark, "An open-source algorithm to detect onset of arterial blood pressure pulses," in *Computers in Cardiology, 2003*, pp. 259–262, IEEE, 2003.

[95] S. Linder and S. Wendelken, "Using the morphology of the photoplethysmogram envelope to automatically detect hypovolemia," in *Computer-Based Medical Systems, 2006. CBMS 2006. 19th IEEE International Symposium on*, pp. 371–371, IEEE, 2006.

[96] P. Cox, C. Madsen, K. L. Ryan, V. a. Convertino, and E. Jovanov, "Investigation of photoplethysmogram morphology for the detection of hypovolemic states.," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society.*, vol. 2008, pp. 5486–9, Jan. 2008.

[97] U. Farooq, D. Jang, J. Park, and S. Park, "PPG delineator for real-time ubiquitous applications," in *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, pp. 4582–4585, IEEE, 2010.

[98] B. N. Li, M. C. Dong, and M. I. Vai, "On an automatic delineator for arterial blood pressure waveforms," *Biomedical Signal Processing and Control*, vol. 5, pp. 76–81, Jan. 2010.

[99] J. A. Sukor, S. J. Redmond, and N. H. Lovell, "Signal quality measures for pulse oximetry through waveform morphology analysis.," *Physiological measurement*, vol. 32, pp. 369–84, Mar. 2011.

[100] F. N. Fritsch and R. E. Carlson, "Monotone Piecewise Cubic Interpolation," *SIAM Journal on Numerical Analysis*, vol. 17, pp. 238–246, Apr. 1980.

[101] K. Reddy, B. George, N. Mohan, and V. Kumar, "A Novel Calibration-Free Method of Measurement of Oxygen Saturation in Arterial Blood," *IEEE Transactions on Instrumentation and Measurement*, vol. 58, pp. 1699–1705, May 2009.

[102] Bluegiga Technologies, "BLE112 Bluetooth 4.0 Module." http://www.bluegiga.com/BLE112_Bluetooth_low_energy_module.

[103] D. Kim, W. Tang, B. Goldstein, P. Weerakoon, H. Montanaro, B. Martini, and E. Culurciello, "Performance comparison of low current measurement systems for biomedical applications," in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, pp. 3469–3472, IEEE, May 2010.

[104] D. Kim, B. Goldstein, W. Tang, F. J. Sigworth, and E. Culurciello, "Noise Analysis and Performance Comparison of Low Current Measurement Systems for Biomedical Applications," 2012.

[105] K. Smith and A. Sedra, "The current conveyor; A new circuit building block," *Proceedings of the IEEE*, vol. 56, no. 8, pp. 1368–1369, 1968.

[106] A. Sedra and K. Smith, "A second-generation current conveyor and its applications," *IEEE Transactions on Circuit Theory*, vol. 17, no. 1, pp. 132–134, 1970.

[107] C. Toumazou, F. Lidgey, and D. Haigh, eds., *Analogue IC Design: The Current-Mode Approach*. IET, 1992.

[108] H. Elwan and A. Soliman, "Low-voltage low-power CMOS current conveyors," *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 44, no. 9, pp. 828–835, 1997.

[109] R. Carvajal, J. Ramirez-Angulo, A. Lopez-Martin, A. Torralba, J. Galan, A. Carlosena, and F. Chavero, "The flipped voltage follower: a useful cell for low-voltage low-power circuit design," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 52, pp. 1276–1291, July 2005.

[110] A. Lopez-Martin, S. Baswa, and R. Carvajal, "Low-Voltage Super class AB CMOS OTA cells with very high slew rate and power efficiency," *IEEE Journal of Solid-State Circuits*, vol. 40, pp. 1068–1077, May 2005.

[111] A. Lopez-Martin, J. Ramirez-Angulo, and R. Carvajal, "Low-voltage low-power wideband CMOS current conveyors based on the flipped voltage follower," in *Proceedings of the 2003 International Symposium on Circuits and Systems, 2003. ISCAS '03.*, vol. 1, pp. I–801–I–804, IEEE.

[112] Y.-S. Hwang, W.-H. Liu, S.-H. Tu, and J.-J. Chen, "New building block: multiplication-mode current conveyor," *IET Circuits, Devices & Systems*, vol. 3, p. 41, Feb. 2009.

[113] M. Ayachi, L. Zhou, and J.-P. Blonde, "A current mode CMOS IC for biological signals measurement in noisy environment," in *2008 International Conference on Microelectronics*, pp. 248–251, IEEE, Dec. 2008.

[114] G. Ferri, V. Stornelli, and A. Di Simone, "A CCII-based High Impedance Input Stage for Biomedical Applications," *Journal of Circuits, Systems and Computers*, vol. 20, pp. 1441–1447, Dec. 2011.

[115] P. Pookaiyaudom, F. Lidgey, K. Hayatleh, P. Seelanan, and C. Toumazou, "Chemical current conveyor (CCCII+): System design and verification for buffer index/capacity measurement," *Sensors and Actuators B: Chemical*, vol. 147, pp. 228–233, May 2010.

[116] P. Livi, F. Heer, U. Frey, D. J. Bakkum, and A. Hierlemann, "Compact Voltage and Current Stimulation Buffer for High-Density Microelectrode Arrays," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 4, pp. 372–378, Dec. 2010.

[117] B. Wilson, "Recent developments in current conveyors and current-mode circuits," vol. 137, 1990.

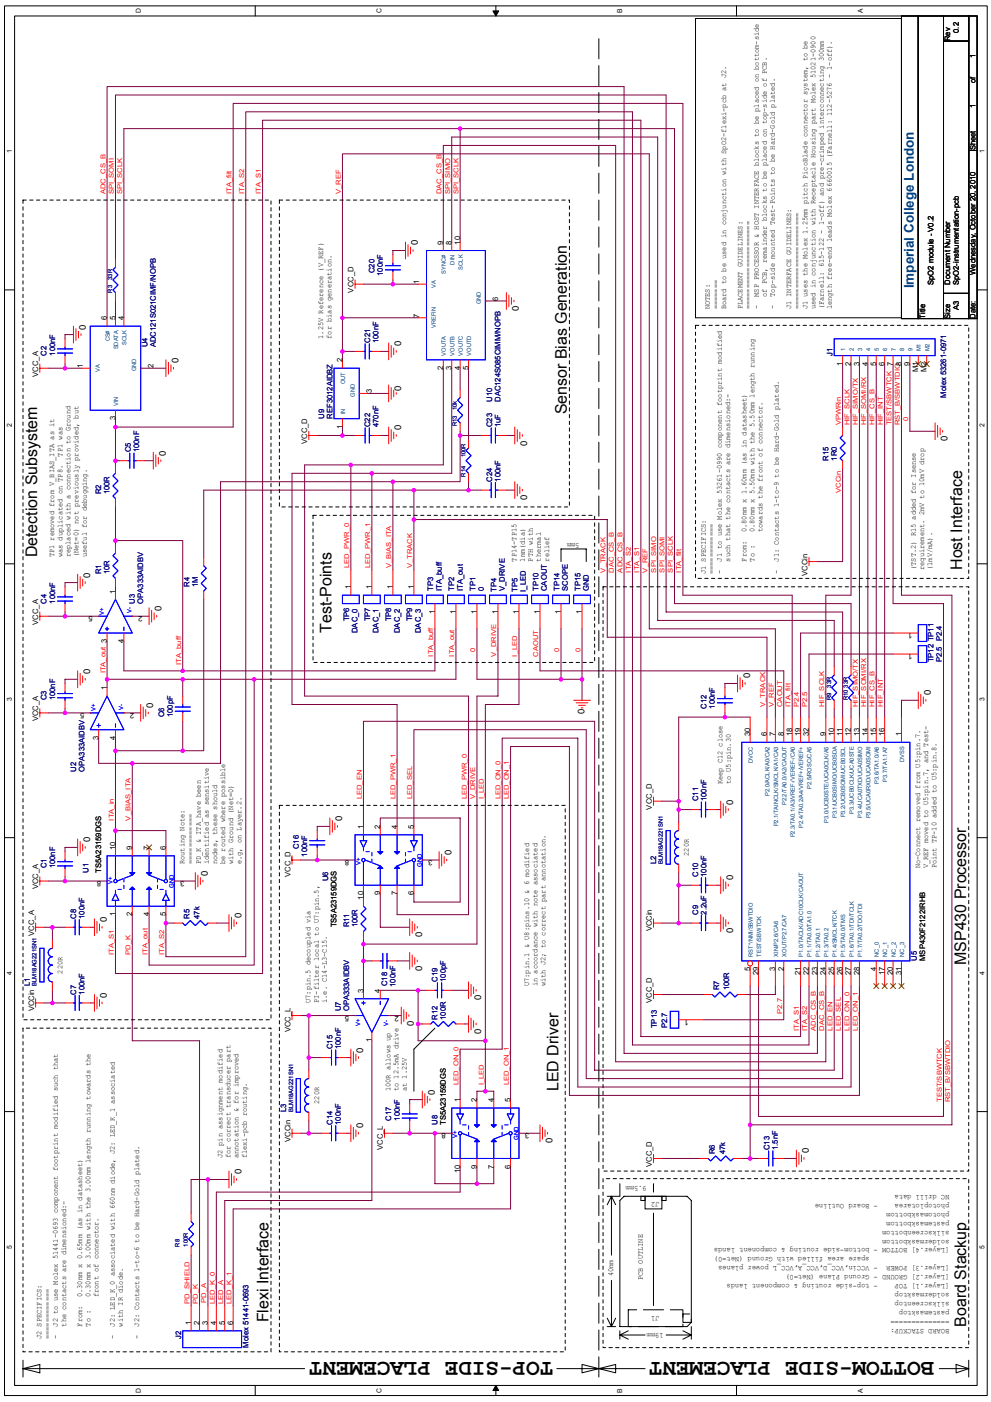
[118] G. Zatorre-Navarro, N. Medrano-Marqués, and S. Celma-Pueyo, "Analysis and simulation of a mixed-mode neuron architecture for sensor conditioning.," *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, vol. 17, pp. 1332–5, Sept. 2006.

[119] D. Damianou, "The wavelength dependence of pulse oximetry," in *IEE Colloquium on Pulse Oximetry: A Critical Appraisal*, vol. 1996, pp. 7–7, IEE, 1996.

[120] M. Sayer, "Implications of a New Medical Technology in the Hospital: CO-OXIMETRY," in *Proceedings of the IEEE 32nd Annual Northeast Bioengineering Conference*, pp. 45–46, IEEE, 2006.

[121] Y. Tsividis, "Event-Driven Data Acquisition and Digital Signal ProcessingA Tutorial," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 57, pp. 577–581, Aug. 2010.

Appendix A: Prototype Version 2


TOP-SIDE PLACEMENT
BOTTOM-SIDE PLACEMENT
Flexi Interface
LED Driver
Detection Subsystem
Test-Points
Sensor Bias Generation
Host Interface
MSP430 Processor
Board Stackup
Imperial College London
SpO2 module - V0.2
SpO2instrumentation.pcb
Wednesday, October 20, 2010

Appendix B: Lattice Wave Digital Filter Prototype in Python

Full code available at `https://github.com/jimurai/lwdf`

```
#! /usr/bin/python
from scipy import signal
import numpy as np

# An adaptor class for use in Lattice Wave Digital Filters
class Adaptor():
  type = 0
  dtype = np.float
  alpha = np.zeros((1,),dtype=np.float)
  coeff_scale = 0
  def __init__(self, gamma=0, coeff_scale=1, dtype=np.float):
    self.set_params(gamma, coeff_scale, dtype)

  def set_params(self, gamma=0, coeff_scale=1, dtype=np.float):
    # Define the adaptor
    self.dtype = dtype
    self.alpha = np.zeros((1,),self.dtype)
    self.coeff_scale = coeff_scale
    # Choose the adaptor type
    if (gamma > 0.5) & (gamma < 1):
      self.type = 0
      temp = (1<<self.coeff_scale)*(1 - gamma)
    elif (gamma > 0) & (gamma <= 0.5):
      self.type = 1
      temp = (1<<self.coeff_scale)*gamma
    elif (gamma >= -0.5) & (gamma < 0):
      self.type = 2
      temp = (1<<self.coeff_scale)*abs(gamma)
    elif (gamma > -1) & (gamma < -0.5):
      self.type = 3
      temp = (1<<self.coeff_scale)*(1 + gamma)
    else:
      print('Invalid gamma.')
    if self.dtype != np.float:
      temp = temp.round()
    self.alpha[0] = temp
```

```
    print gamma, self.alpha, self.type

  def process(self, wave_in):
    _diff = np.zeros((1,),dtype=self.dtype)
    _waveout = np.zeros((2,),self.dtype)
    _wavein = np.array((wave_in),self.dtype)
    if self.type == 0:
      _diff = _wavein[0] - _wavein[1]
      _waveout[1] = self.multiply_by_alpha(_diff) + _wavein[1]
      _waveout[0] = _waveout[1] - _diff
    elif self.type == 1:
      _diff = self.multiply_by_alpha(_wavein[1] - _wavein[0])
      _waveout[1] = _diff + _wavein[0]
      _waveout[0] = _diff + _wavein[1]
    elif self.type == 2:
      _diff = self.multiply_by_alpha(_wavein[0] - _wavein[1])
      _waveout[1] = _diff - _wavein[0]
      _waveout[0] = _diff - _wavein[1]
    elif self.type == 3:
      _diff = _wavein[1] - _wavein[0]
      _waveout[1] = self.multiply_by_alpha(_diff) - _wavein[1]
      _waveout[0] = _waveout[1] - _diff
    else:
      print('Invalid type.')
    return _waveout

  def multiply_by_alpha(self, input):
    if self.dtype == np.float:
      return (self.alpha[0]*input)/(1<<self.coeff_scale)
    else:
      return ((np.cast[np.int64](self.alpha[0])*input)+
          (1<<(self.coeff_scale-1)))>>self.coeff_scale

# Class for constructing and building a Lattice Wave Digital Filter
class Filter():
  order = 0
  adaptors = []
  registers = []
  output = [0,0]
  dtype = np.float
  coeff_scale = 0
  def __init__( self, order=3, wn=0.01, coeff_scale=1,
        rp=1, rs=60, ftype='butter', dtype=np.float):
    if order%2 == 0:
      print 'No even order filters allowed.'
      return
    # Filter settings
    self.order = order
```

```
self.wn = wn
self.rp = rp
self.rs = rs
self.dtype = dtype
self.ftype = ftype
self.coeff_scale = coeff_scale
gammas = []
psi = np.tan(np.pi*self.wn/2.0)
psi2 = psi*psi
if self.ftype == 'butter':
  (z,p,k) = signal.iirfilter( self.order, psi, btype='lowpass',
                  analog=1, ftype='butter', output='zpk')
  filter_roots = np.sort(p)
elif self.ftype == 'bessel':
  print 'Please avoid using Bessel filters'
  print 'They don\'t translate well to LWDFs.'
  (z,p,k) = signal.iirfilter( self.order, psi, btype='lowpass',
                  analog=1, ftype='bessel', output='zpk')
  filter_roots = np.sort(p)
elif self.ftype == 'cheby1':
  (z,p,k) = signal.iirfilter( self.order, psi, rp=1, btype='lowpass',
                  analog=1, ftype='cheby1', output='zpk')
  filter_roots = np.sort(p)
elif self.ftype == 'cheby2':
  (z,p,k) = signal.iirfilter( self.order, psi, rs=self.rs,
                  btype='lowpass', analog=1,
                  ftype='cheby2', output='zpk')
  filter_roots = np.sort(p)
else:
  print 'Invalid filter type.'
  return
# Separate the real pole from the complex poles
real_index = 0;
for i in range(self.order):
  if abs(filter_roots[i].imag) <= 1e-16:
    real_index = i
    break
complex_roots = np.concatenate((filter_roots[0:real_index],
                  filter_roots[real_index+1:]))
# Put in the real pole's gamma value
h_B = -1.0*filter_roots[real_index].real
gammas.append((1.0 - h_B) / (1.0 + h_B))
# Calculate coefficients of the individual Hurwitz polynomials
for i in (range((order-1)/2)):
  h_A = -2.0*(complex_roots[2*i].real)
  h_B = abs(complex_roots[2*i])**2
  gammas.append((h_A - h_B - 1.0)/(h_A + h_B + 1.0))
  gammas.append((1.0 - h_B) / (1.0 + h_B))
```

```
    # Construct filter
    for i in range(self.order):
      self.adaptors.append(Adaptor(gammas[i],self.coeff_scale,self.dtype))
      self.registers.append(0)

  def push(self, value):
    # Insert the new value into both branches
    if self.dtype != np.float:
      output = [value<<self.coeff_scale,value<<self.coeff_scale]
    else:
      output = [value*(1<<self.coeff_scale),value*(1<<self.coeff_scale)]

    # Process real pole in the top path
    wave_in = [output[0], self.registers[0]]
    wave_out = self.adaptors[0].process(wave_in)
    self.registers[0] = wave_out[1]
    output[0] = wave_out[0]

    # Process top branch
    for i in range(int(self.order/4)):
      x = [4*i+3, 4*i+4]

      wave_in = [self.registers[x[0]], self.registers[x[1]]]
      wave_out = self.adaptors[x[1]].process(wave_in)
      self.registers[x[1]] = wave_out[1]

      wave_in = [output[0], wave_out[0]]
      wave_out = self.adaptors[x[0]].process(wave_in)
      self.registers[x[0]] = wave_out[1]

      output[0] = wave_out[0]

    # Process bottom branch
    for i in range(int((self.order+1)/4)):
      x = [4*i+1, 4*i+2]

      wave_in = [self.registers[x[0]], self.registers[x[1]]]
      wave_out = self.adaptors[x[1]].process(wave_in)
      self.registers[x[1]] = wave_out[1]

      wave_in = [output[1], wave_out[0]]
      wave_out = self.adaptors[x[0]].process(wave_in)
      self.registers[x[0]] = wave_out[1]

      output[1] = wave_out[0]

    # Save the output and return it
    self.output = output
```

```
        return output

if __name__ == "__main__":
    from scipy import fftpack, signal
    import matplotlib.pyplot as plt
    # Fetch example data
    [ir,red,amb] = np.load('test_data.npy')

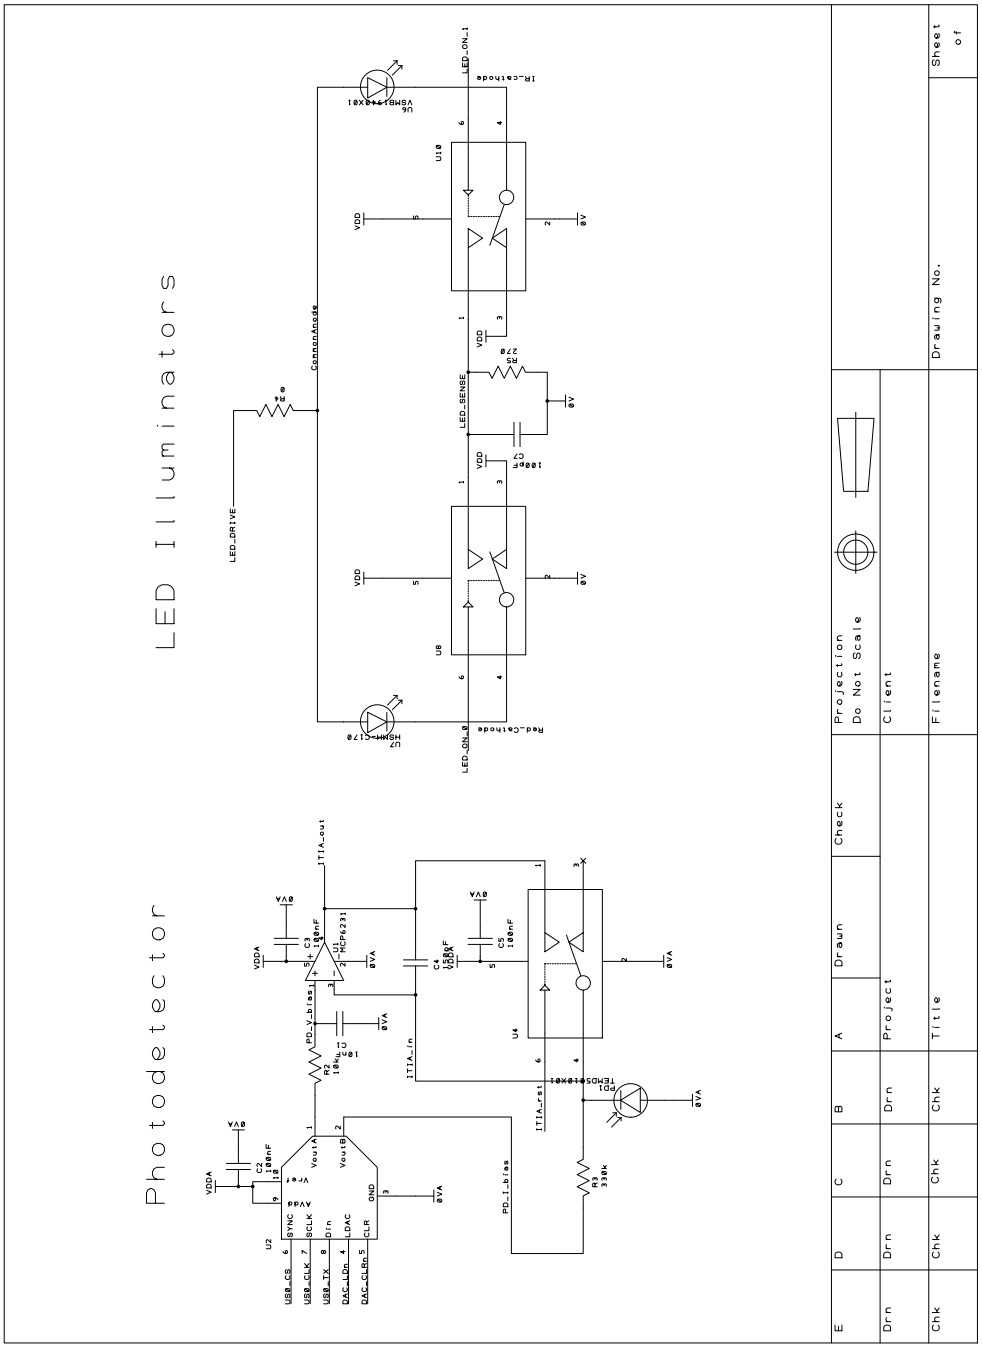
    # Define filter
    order = 5
    wn = 50.0/500.0
    coeff_scale = 16
    dtype = np.int32
    rs = 96
    dtype=np.int32
    ftype = 'cheby2'
    lwdfilter = Filter(order,wn,coeff_scale,rs=70,ftype=ftype,dtype=dtype)

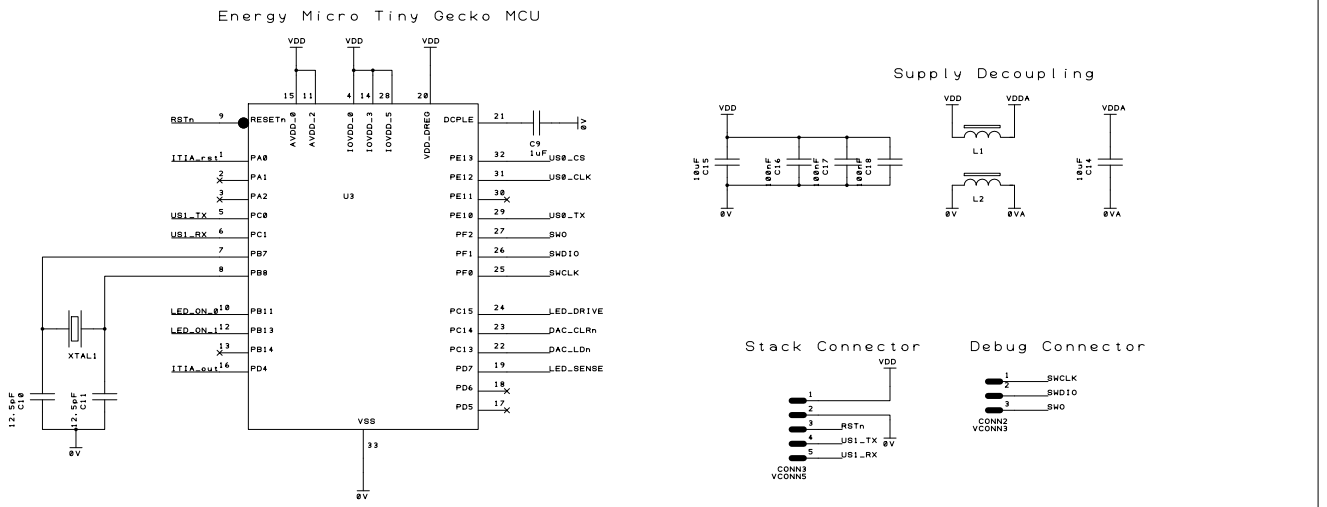
    # Perform sample-by-sample filtering
    ir_out = []
    for x in ir:
        temp = lwdfilter.push(x)
        if dtype==np.float:
            ir_out.append((temp[0]+temp[1])/(2.0*(1<<coeff_scale)))
        else:
            ir_out.append((temp[0]+temp[1]+(1<<coeff_scale))>>(coeff_scale+1))

    # Compare with SciPy equivalent
    (B,A) = signal.iirfilter( order, wn, rs=rs,btype='lowpass', analog=0,
                    ftype='cheby2', output='ba')
    zf = signal.lfiltic(B,A,np.zeros(len(A)-1), np.zeros(len(B)-1))
    ir_comp, zf = signal.lfilter(B,A,ir,zi=zf)

    # Plot the data
    plt.plot(ir, 'k-', linewidth=2.0)
    plt.plot(ir_comp, 'r-', linewidth=2.0)
    plt.plot(ir_out, 'b-', linewidth=2.0)
    plt.grid()
    plt.legend( ('Raw data', 'SciPy filtered - float32', 'LWDF - int32'),
          loc='upper left')
    plt.show()
```

Appendix C: Prototype Version 3


LED Illuminators
Photodetector
LED_DRIVE
Common_node
LED_SENSE
Red_Cathode
IR_Cathode
LED_ON_0
LED_ON_1
ITIA_out
ITIA_in
ITIA_rst
PD_V_bias
PD_I_bias
VDDA
0VA
VDD
0V
Projection
Do Not Scale
Client
Project
Title
Filename
Drawing No.
Sheet
of



E	D	C	B	A	Drawn	Check	Projection Do Not Scale		
Drn	Drn	Drn	Drn	Project			Client		
Chk	Chk	Chk	Chk	Title			Filename	Drawing No.	Sheet of