LINEAR ARRAY CMOS DETECTORS FOR LASER DOPPLER BLOOD FLOW IMAGING

John Michael Himsworth, M.Eng (Hons)

Thesis submitted to the University of Nottingham

for

the degree of Doctor of Philosophy

December 2011

Abstract

Laser Doppler blood flow imaging is well established as a tool for clinical research. The technique has considerable potential as an aid to diagnosis and as a treatment aid in a number of situations. However, to make widespread clinical use of a blood flow imager feasible a number of refinements are required to make the device easy to use, accurate and safe.

Existing LDBF systems consist of 2D imaging systems, and single point scanning systems. Single point scanning systems can offer fast image acquisition time, and hence high frame rate. However, these require high laser power to illuminate the entire target area with sufficient power. Single point scanning systems allow lower laser power to be used, but building up an image of flow in skin requires mechanical scanning of the laser, which results in a high image acquisition time, making the system awkward to use.

A new approach developed here involves scanning a line along a target, and imaging the line with a 1D sensor array. This means that only one axis of mechanical scanning is required, reducing the scanning speed, and the laser power is vastly reduced from that required for a 2D system.

This approach lends itself well to the use of integrated CMOS detectors, as the smaller pixel number means that a linear sensor array can be implemented on an IC which has integrated processing while keeping overall IC size, and hence cost, lower than equivalent 2D imaging systems.

A number of front-end and processing circuits are investigated in terms of their suitability for this application. This is done by simulating a range of possible designs, including several logarithmic pixels, active pixel sensors and opamp-based linear front-ends. Where possible previously fabricated ICs using similar sensors were tested in a laser Doppler flowmetry system to verify simulation results.

A first prototype IC (known as BVIPS1) implements a 64x1 array of buffered logarithmic pixels, chosen for their combination of sufficient gain and bandwidth and compact size. The IC makes use of the space available to include two front-end circuits per pixel, allowing other circuits to be prototyped. This allows a linear front-end based on opamps to be tested. It

is found that both designs can detect changes in blood flow despite significant discrepancies between simulated and measured IC performance. However, the signal-noise ratio for flux readings is high, and the logarithmic pixel array suffers from high fixed pattern noise, and noise and distortion that makes vein location impossible.

A second prototype IC (BVIPS2) consists of dual 64x1 arrays, and integrated processing. The sensor arrays are a logarithmic array, which addresses the problems of the first IC and uses alternative, individually selectable front-ends for each pixel to reduce fixed-pattern noise, and an array of opamp-based linear detectors. Simulation and initial testing is performed to show that this design operates as intended, and partially overcomes the problems found on the previous IC - the IC shows reduced fixed pattern noise and better spatial detection of blood flow changes, although there is still significant noise.

Acknowledgements

I would like to thank my supervisors Professor Barrie Hayes-Gill and Professor Stephen Morgan for their advice and support during this work. I would also like to thank my colleague Xianhong Xu, who designed the digital parts of the second prototype IC (BVIPS2), Hoang Nguyen and Diwei He for sharing their knowledge of Laser Doppler Flowmetry, Xiabing Huang and Martin Potter from Moor Instruments Ltd for their help with experimental equipment and LDF devices, and Roger Light for his guidance and assistance with all things chip related.

This work has been financially supported by the Engineering and Physical Sciences Research Council, and also through the UK Technology Strategy Board as part of the Blood Vessel Imaging for Venipuncture, Phlebotomy and Surgery project.

Finally I would like to thank my friends, family, and particularly Emma, for preserving my sanity throughout this project.

List of Abbreviations

ADC - Analog to Digital Converter

AMS - Austria Micro Systems

BVIPS - Blood Vessel Imaging for Phlebotomy, Veinipuncture and Surgery

CCD - Charge Coupled Device

CMOS - Complementary Metal Oxide Semiconductor

CMRR - Common Mode Rejection Ratio

LDF - Laser Doppler Flowmetry

LDBF - Laser Doppler Blood Flowmetry

FFT - Fast Fourier Transform

FPN - Fixed Pattern Noise

FPGA - Field Programmable Gate Array

GMC - Transimpedance-Capacitance Filter ($g_m - C$)

HDA - Hysteretic Differentiator Amplifier

IC - Integrated Circuit

LASCA - Laser Speckle Contrast Analysis

PCB - Printed Circuit Board

PSRR - Power Supply Rejection Ratio

SNR - Signal-to-Noise Ratio

SRAM - Static Random Access Memory

THD - Total Harmonic Distortion

TIA - Transimpedance Amplifier

VLSI - Very Large Scale Integration

Contents

Chapter 1

Introduction

This chapter will introduce the type of device developed during this work, along with the context of this work within a wider DTI research project, and the novelty of the work presented here. A summary of the thesis structure is also given.

The basic principle of Laser Doppler Flowmetry will then be discussed, along with related detection electronics. The physics of the Doppler effect will be introduced, and the way the effect applies to light reflecting from blood cells is discussed. The information available from this process is explained, and a generic electronic system for processing these signals is introduced. Possible processing methods are introduced, and the signal expected from the photosensors is estimated. Finally, the structure of the thesis is described.

1.1 Context of this Work

This section introduces the wider DTI research project into Blood Flow Imaging, of which the work in this thesis forms a part. The intended work to be done as part of this project will be explained, including a description of exactly what parts of the presented work were performed by the author.

1.1.1 DTI Project - Blood Vessel Imaging for Veinipuncture, Phlebotomy and Surgery (BVIPS)

The aim of the project was to develop a line scanning blood flow imager for vein location, for use in a clinical setting. The project partners included Moor Instruments Ltd., a manufacturer of LDBF systems, Sifam Ltd (now Gooch & Housego), a manufacturer of Laser packages, Peninsular Medical School in Exeter, Nottingham City Hospital, and the University of Nottingham. The work presented here is based on the development of an integrated sensor array on a CMOS IC, intended to allow an improved imaging device to be produced.

1.1.2 Focus of this Work Within the BVIPS Project

This section of the BVIPS project described within this thesis focused on the research and development of a custom integrated CMOS sensor for the imager. It was the intention at the outset of the project to fabricate two ICs during this project. This section describes the work required for this project and some of the constraints placed on this work by the wider project.

The first IC (referred to as 'BVIPS1'), was to include a 64x1 sensor array, which could be integrated into an existing imager. This IC was not intended to include digital processing, being focused on developing the analogue signal acquisition and conditioning stages. As no space on the IC was consumed by digital processing, a series of extra circuits could be included, allowing smaller arrays of alternative front-end and filtering circuits to be tested, with a view to inclusion on a second prototype IC. The additional circuits covered within this work include a 32x1 array which duplicated the pixel design used within the main array, with the addition of extra visibility of intermediate signals which was not available in the main array, and also an additional analogue processing stage which was not used in testing. Additionally, an opamp-based linear front-end circuit was fabricated, investigating the feasibility of using opamps as photodetector front-ends in 1D integrated arrays. The design and testing of this IC is shown in Chapter 4.

The second IC, BVIPS2, was to develop on the first IC by adding on-chip digital processing and analogue-digital conversion. This was intended to allow further refinements to an imaging instrument by minimising the size of the electronics elements. The IC consists of dual 64x1 arrays of two types of sensor (a technique developed for the first IC), and links these to

on-chip analogue-digital conversion and digital processing. The design and testing of the analogue parts of this IC is shown in Chapter 5.

One significant constraint placed on the design of both ICs by other aspects of the project, and the line scanning technique, is the physical size of the photodiodes. As the ICs produced were to be built into an existing optical/mechanical system (the Moor LDLS imager), the IC had to address concerns about the line imaged onto the sensor 'wandering' across the sensors as the line is mechanically scanned. This is partly caused by imperfections in the mechanical components, and partly by the changing optical path as the scanning mirror moves causing the expected line location to move. This movement is not a problem in the imager using discrete components, as the discrete photodiodes used are considerable larger, at $3.6\,\mathrm{mm} \times 0.5\,\mathrm{mm}$.

Intuitively, larger photodetectors would seem to be an advantage, with the larger area resulting in more light being gathered and hence higher photocurrent and easier current to voltage conversion. However, for LDBF applications the important photocurrent is the AC component rather than the overall intensity, and this does not increase linearly with area, and the modulation depth can actually decrease with larger photodiodes, as the AC component is linked to the size of the speckles produced by the light imaged onto the detector. This issue is discussed in more detail in Section 1.8.2.

1.1.3 Work Done by the Author

As this work builds on previous research at the University of Nottingham, and involves external partners, this section clarifies which work was done by the author, and also where the work of others was used.

Overall IC design of both prototypes used was performed by the author. However, individual analogue circuits (the current to voltage converters, filters and amplifiers described in Chapter 3) were developed during previous work. Additionally, the first prototype includes an opamp-based front-end, which makes use of an opamp developed for use as an off-chip buffer for previous ICs developed at the University of Nottingham, although the front-end circuit itself is the authors work. The second version of this circuit, used on the second prototype IC, uses an opamp that was designed by the author for this project. The first IC also included two other prototype arrays designed by other researchers, but these circuits are

not described within this thesis.

On the second prototype IC, described in Chapter 5, all analogue parts of the IC were designed by the author (excepting the individual circuit elements as described above). The digital parts of the IC were designed by other researchers, and are described in this thesis in order to demonstrate the suitability of these circuits as components for the overall system. Top level design of the IC including the analogue and digital elements was performed by the author.

All test results shown are from testing performed by the author.

1.1.4 Other Aspects of the BVIPS project

Other parts of the project, not performed by the author, included:

- Work on processing hardware and algorithms, using FPGA based processing to increase frame rates, performed at the University of Nottingham by other researchers[Hoang et al., 2010]. The later stage of this work included developing processing techniques suitable for use on an ASIC. The digital processing block produced was included on the second prototype IC, integrating the processing techniques developed with the analogue sensors developed as part of this work. (However, this thesis is concerned with the analogue parts of the IC rather than the use of the digital processing also included.)

- Device development using discrete component front-ends, with the improved data acquisition/FPGA processing system developed here, by Moor Instruments Ltd. It was then intended to replace the discrete sensors here with the integrated CMOS sensors developed. Initial integration was performed to investigate the suitability of the IC produced, but this did not progress to producing a full imager with a CMOS array front-end.

- Production of either a stable visible red, fibre-coupled laser diode package of the required power level for this application (~40mW), or a combination of a stable, near infra-red laser with the required power along with a low power visible red laser as a 'guide' beam, allowing users to see where the laser is illuminating. The visible beam also improves laser safety, as it allows the blink response of the human eye to restrict

exposure time of the eye to the laser in the event that the laser shines into a user's or patient's eyes. This work was carried out by Sifam Ltd, up until this company was bought by Gooch & Housego.

- Clinical use of a prototype device, Peninsular Medical School, Exeter and City Hospital, Nottingham. This work was intended to investigate the clinical usefulness of an imager using the line scanning technique, and the user needs of such a device.

1.2 Novelty of this Work

The work presented in this thesis has a number of novel aspects. The ICs produced are the first 1D array integrated sensors designed for use in laser Doppler blood flow imaging, with the second IC in particular showing the possibilities of CMOS integration, combining the photodetectors, analogue signal conditioning, data acquisition and blood flow processing on one IC.

Some of the novel aspects have wider application than blood flow imaging. The 1D array/line scanning imaging technique has two basic advantages:

- Reduced laser power - as only one line of the imaging target has to be illuminated, the total illumination power required in a 1D system to give equivalent optical power density to a 2D system is reduced, to approximate the square root of the 2D illuminating power.

- Reduced area required for the 1D sensor array, due to reduced pixel count. This additional area can be used for any purpose - the second IC fabricated here makes use of this area to implement an alternative, full-size, array of front-end circuits, as well as integrated ADCs and digital flow processing.

The reduction in illumination power has technical advantages, in terms of only requiring a lower power laser, and also safety advantages because of the reduced laser safety risks. For a system such as a blood flow imager, intended for use by non-specialists in an open setting (i.e. a normal hospital ward), this represents a significant advantage - a 2D imager could potentially be designed that has superior frame rate, larger imaging area and is simple to use, however if the high laser power required means the device is only considered safe when

used by a specialist operator in a specific, private room, then the utility of the imaging device is greatly reduced.

The reduction in IC physical size applies to both pixel level and system level size. At the system level, the reduction in size from 2D to 1D sensors allows space for system-level components such as output-buffers, ADCs and the processing block, while keeping overall size and hence cost down. However, in the context of a full system the cost saving may not be a major advantage. However, the size advantage also applies at the pixel level - as the pixel count is reduced to the square root of the 2D equivalent pixel count, each pixel can be made much larger without having the severe impact that such changes would have on a 2D system. On the ICs produced here, this advantage is used by designing pixels that use large circuits such as opamps (where the front-end is generally too large for 2D integrated sensors), include higher order filters (which cannot be implemented at column level of a 2D array due to the settling time of the filters), and alternative signal paths (either alternative amplifier designs, or duplicate circuits to give redundancy). Similarly, the photodiode itself can be made larger if required by the application for increased light gathering, although at the expense of high input capacitance. The ICs produced here have large photodiodes ($1\,\mathrm{mm} \times 50\,\mu\mathrm{m}$), although this is due to concerns about the mechanical line scanning causing the line imaged onto the sensor to wander, rather than for improved detection reasons.

1.3 Layout of the Thesis

Chapter 1 introduces the type of device to be developed as part of this project. The Doppler effect and its application to measuring blood flow using reflected laser light is described. A generic Laser Doppler Flowmetry system architecture is presented, with the function of its various components discussed. Variations to this generic system due to the type of detection circuits and the processing methods that might be used here are considered.

The optical mixing process that produces an electronic signal in a photosensor is described, and based on previous experimental measurements, the signals that might be expected from the photodiodes to be used are calculated, allowing simulations on specific circuits to be performed using a known input current. The frequency component of this signal is also considered, allowing the bandwidth requirements of the system to be set. However, the light level measurements and calculations in that chapter were shown to be incorrect by later

measurements (shown in chapter 4), potentially reducing the accuracy of some simulations for this specific application.

Chapter 2 reviews previous laser Doppler flowmetry systems, from initial development of the techniques through to modern systems in clinical use and current research trends in LDBF imagers. Applications of LDBF other than basic imaging are also introduced to show the potential of LDBF as a clinical tool.

As the system designed for this project is intended for vein location rather than general flow imaging, alternative technological methods of vein location are considered. This includes low-cost, simple devices intended as a minor aid to traditional clinical skills, as well as advanced vein imaging methods.

Chapter 3 investigates the circuits that can be used on ICs for laser Doppler detection, investigating each block of a generic LDBF system in turn. The general principles of the current-voltage converter circuits are introduced, along with several variations on the basic circuit that could be used in this application. Simulations are performed on these circuits to select the most suitable for this application. Various processing circuits (AC amplifiers and filters) are described to select other parts of the system, with simulations performed to show typical behaviour and the effect of manufacturing variations where necessary.

Chapter 4 describes the first IC designed for this project, known as BVIPS1. The system structure and pixel design are shown. Characterisation and simulation results from the IC are compared. Testing results from Doppler imaging with this IC are shown, and the implications of these results for vein imaging are discussed. Problems with the IC behaviour and their likely causes are described. Design options made possible by the use of a linear array (rather than a 2D array) are shown, such as dual front-end options, and physically larger front-end designs including opamp based linear detectors.

Chapter 5 shows the changes to be made to the BVIPS1 design for the 2nd IC fabricated for this project, known as BVIPS2. These changes are partly to solve issues found in testing the first IC, and also to develop the system to add on-chip analog-digital conversion and digital processing to find flux. A linear front-end using opamp based front-ends similar to those tested on BVIPS1 is also used as an alternative main circuit on this IC, which has two optional 64x1 detector arrays. Initial testing results from this IC in terms of characterisation and Doppler imaging are shown, although testing was limited by time. The tests show some improved performance, athough not all issues are fully resolved. Further work using this IC

is considered.

The possibilities of a linear array sensor are further demonstrated on the BVIPS2 IC by the use of two full size sensor arrays along with the use of integrated processing circuitry on one IC. In particular, the linear array layout allows both sensor arrays to use novel features. One of the arrays is composed of opamp based front-using opamps ends, which has not been previously demonstrated on an IC for laser Doppler flowmetry, partly due to the large size of this type of pixel not being feasible on a 2D array sensor. The other array is composed of logarithmic pixels which combine the larger pixel area available with a compact front-end circuit design to allow redundant front-end circuits to be made available.

Finally, Chapter 6 presents conclusions from the project, summarising the work undertaken and its potential for use in clinical systems. Further work using the ICs and circuits developed is suggested, as well as developments to the systems designed. Possible limitations on the further development of this work are also considered.

1.4 System to be Developed

This thesis discusses the research into and development of a line scanning blood flow imager intended for vein location, or more specifically the detector for such a system. The detector uses a 64x1 linear array of CMOS photodetectors. The use of CMOS technology for the detectors themselves allows processing circuitry to be integrated onto the same device. The work covers the selection and prototyping of detector circuits and analogue signal conditioning suited to this application, and also the integration of the analogue system produced into a system-on-chip with integrated data acquisition and digital processing.

Vein location will be performed by producing a 2D flow image of a target (tissue containing veins such as a hand or forearm), from which the high flow in a vein allows these structures to be identified. The 2D image will be formed by projecting a line onto the target, which will be imaged back onto the sensor array, allowing measurement of blood flow along this line. The line will then be mechanically scanned along the target, allowing a 2D image to be built up from successive lines.

1.5 The Doppler Effect

The Doppler effected was first described by Christian Doppler in 1842 and demonstrated for sound waves shortly after [Houdas, 1991, Serway and Jewett, 2009]. It describes the nature of waves when the source is moving relative to the observer. This can most clearly be demonstrated for sound waves, such as the change in pitch of sound from a moving vehicle as it passes the observer.

As the source travels towards the observer, the apparent wavelength is shorter, as later wavefronts are produced nearer to the observer than the earlier fronts. As the wave velocity is constant, this means that frequency appears to be higher.

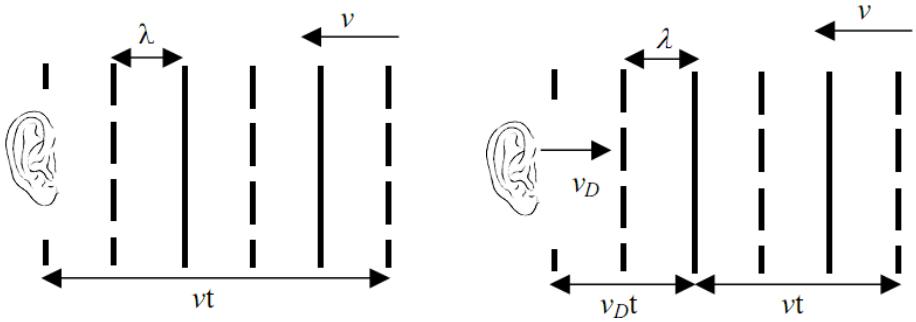

(a) Initial situation (b) Situation after movement

Figure 1.1: Wavefronts with wavelength λ travelling at velocity v, with an observer moving at v_D over time t

Figure 1.1 shows wavefronts travelling from a source to an observer, and the effect of movement on this situation. The Doppler effect can be considered in terms of wavelength or frequency, but as in electronics it is easier to measure frequency, this will be considered the main term. For the situation without movement (Figure 1.1a), the frequency, f, can be found as:

$$f = \frac{\text{oscillations received by observer}}{\text{total time}} = \frac{vt/\lambda}{t} = \frac{v}{\lambda} \tag{1.1}$$

In the case of a moving observer, the situation changes, as the observer passes a larger number of wavefronts than in the static case:

$$f' = \frac{\text{oscillations received by observer}}{\text{total time}} = \frac{(vt + v_D t)/\lambda}{t} = \frac{v + v_D}{\lambda} \tag{1.2}$$

The difference between these two frequencies is termed the Doppler shift frequency (f_d), and is found by subtracting 1.1 from 1.2:

$$f_d = f' - f = \frac{v + v_D}{\lambda} - \frac{v}{\lambda} = \frac{v_D}{\lambda} = \frac{v_D f}{v} \tag{1.3}$$

$$v_D = \frac{f_d v}{f} \tag{1.4}$$

This means that if the Doppler frequency can be measured, for example by emitting a known wave (known f, v and λ) and then measuring it in a different location, then v_D can be found as in equation 1.4, giving the relative velocity between the source and observer. This technique can be used for various situations including measuring speed of moving vehicles [Halliday et al., 2000] and astronomical measurements such as motion of stars, using the emitted light from the star as a source [Serway and Jewett, 2009, Cutnell and Johnson, 2006].

1.6 Laser Doppler Flowmetry

1.6.1 Doppler Shift from Moving Cells

Laser Doppler Blood Flowmetry involves illuminating the skin with coherent, monochrome light of constant intensity. As this light penetrates into the skin it interacts with cells by scattering and absorption. The absorbed light is lost, reducing the total reflected light. Some light scatters off static tissue, resulting in reflected light at the same wavelength as the incident light.

The Doppler shift applies to photons that are scattered by moving red blood cells. In this case the light is subject to a Doppler shift as described above. In this situation, where light is reflected back to the original source, the shift occurs twice, affecting the incident and reflected light. The first shift is the same as the situation described above - the light interacts with a blood cell that is moving relative to the source, so the light incident on the cell appears to have undergone a shift in frequency. When this light is scattered by the cell, the cell itself acts as a moving source relative to the static sensor, so undergoes another shift of the same frequency. This is referred to as a double Doppler shift [Shepherd and Oberg, 1990].

The situation is complicated by the complex movement of the cells. The flow being measured is that of a large number of blood cells flowing in a network of vessels. This means that cells are moving at a variety of speeds in a variety of directions, so the relative speed between the cell and the source is widely varying.

Similarly, the nature of the tissue means that incident photons can follow a variety of different paths, with variation in the number of scattering incidents with static tissue and with moving cells.

Figure 1.2 illustrates the interaction of incident light with tissue and the variety of paths that photons can take.

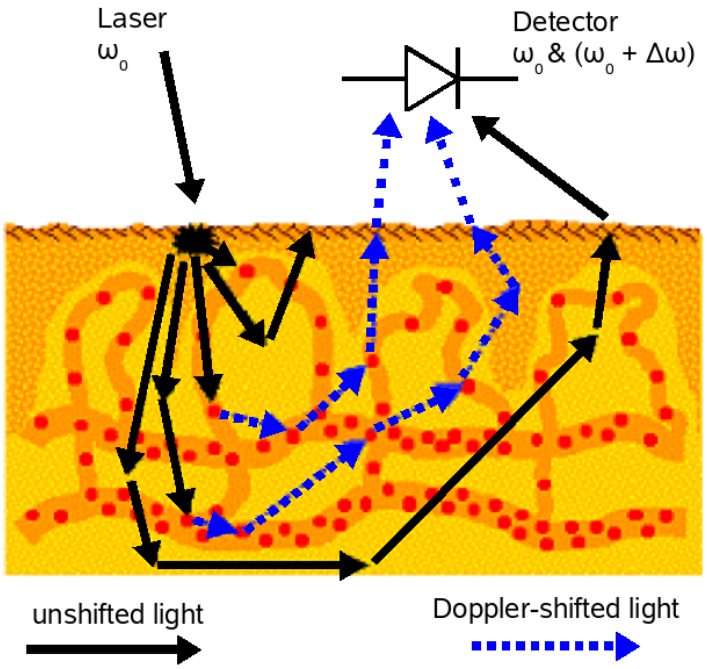

Figure 1.2: Path of incident light through the skin

For a single interaction, the Doppler frequency can be calculated as:

$$\triangle f = \frac{2v f_0 cos\,\theta}{c_{tissue}} \tag{1.5}$$

where f_0 is the frequency of transmitted light, v is the speed of a blood cell, c_{tissue} is the speed of light in tissue, and $cos\,\theta$ is the angle between the direction of cell movement and incident light (shift is highest if the cell is moving directly towards or directly away from the source). The factor of two represents the double Doppler shift.

For example, if the speed of light in tissue is $2 \times 10^8 \mathrm{ms}^{-1}$ (assuming refractive index of ~1.5 [Tuchin, 2007]), and f_0 is that of 650 nm laser light then:

$$f = \frac{v}{\lambda} = \frac{2 \times 10^8}{650 \times 10^{-9}} \approx 3 \times 10^{14}\,\mathrm{Hz} \tag{1.6}$$

If we apply this to equation 1.5, with blood cells moving at roughly $2\,\mathrm{cms}^{-1}$ [Tanaka et al., 1974], we can find the Doppler shift frequency expected:

$$\triangle f = \frac{2 v f_0 cos\theta}{c_{tissue}} = \frac{2 \times 2 \times 10^{-2} \times 3 \times 10^{14}}{2 \times 10^8} = 60\,\mathrm{kHz} \tag{1.7}$$

In the real case this frequency would be considerably lower, as θ will generally be less than 0, and the flow value used here is that in retinal blood vessels rather than capillaries.

This gives the Doppler shift frequency for a single interaction. The overall received signal will be composed of many such interactions, including photons which have interacted with more than one moving cell and have therefore undergone several shifts. This results in the spectrum of the reflected light being spread around the original single wavelength, as shown in Figure 1.3. Section 1.8.3 shows the frequency spectrum from Doppler signals produced from moving blood cells in more detail.

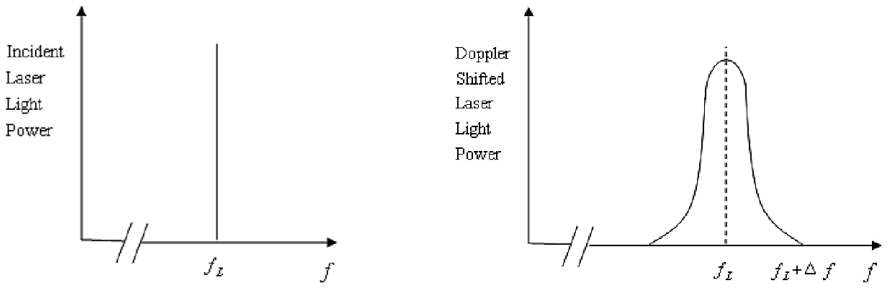

(a) Incident light (b) Reflected light

Figure 1.3: (a) Spectrum of incident light and (b) reflected light after Doppler shifting and static reflection

1.6.2 Optical Mixing

The frequency of the light cannot be directly measured by electronics, being of the order of 10^{14} Hz [Albrecht et al., 2002]. The Doppler shift from equation 1.5 is considerably lower in magnitude than this, so the reflected light is still at a frequency of a similar order of magnitude.

However, the reflected light consists of static and shifted wavelengths. These components

mix in the detector giving an electrical signal at the difference frequency on a DC background. This mixing process can be expressed by the addition of the shifted and un-shifted components in equations 1.8 and 1.9.

$$E_1 = E_0 cos(\omega_1 t + \phi) \tag{1.8}$$

$$E_2 = E_0 cos(\omega_2 t + \phi) \tag{1.9}$$

where E_0 is the amplitude of the wave, ω_1 and ω_2 are the angular frequencies of each wave and ϕ is the phase of the wave. When both of these fields are incident upon a detector, the overall field is the sum of the individual fields:

$$E = E_0 \left[cos(\omega_1 t + \phi) + cos(\omega_2 t + \phi)\right] \tag{1.10}$$

Using a trigonometric sum-to-product identity as in equation 1.11[Abramowitz, 1970], this can be seen as a mixing (i.e. multiplication) of signals at half of the sum and difference frequencies, shown in equation 1.12.

$$cos(A) + cos(B) = 2cos(\frac{A+B}{2})cos(\frac{A-B}{2}) \tag{1.11}$$

$$E = 2E_0 cos\frac{1}{2}((\omega_1 + \omega_2)t + 2\phi)cos\frac{1}{2}((\omega_1 - \omega_2)t + 2\phi) \tag{1.12}$$

Re-writing this as equation 1.13 shows the detected field, E, as a travelling wave at frequency $\omega_1 + \omega_2$ and amplitude E_m:

$$E = E_m cos\frac{1}{2}((\omega_1 + \omega_2)t + 2\phi) \tag{1.13}$$

where:

$$E_m = 2E_0 cos\left(\frac{1}{2}\left((\omega_1 - \omega_2)t + 2\phi\right)\right) \tag{1.14}$$

The sum frequency of the travelling wave, $\omega_1 + \omega_2$ is still in the optical frequency range, so is undetectable by the electronics. The modulated amplitude E_m will be detected by the photosensor. The amount of the light illuminating the detector is known as the irradiance [Hecht, 2001], and is given by:

$$E_m^2 = 4E_0^2 cos^2\left(\frac{1}{2}\left((\omega_1 - \omega_2)t + 2\phi\right)\right) = 2E_0^2\left[1 + cos\left((\omega_1 - \omega_2)t + 2\phi\right)\right] \quad (1.15)$$

It can be seen from equation 1.15 that the DC level of the irradiance is $2E_0^2$, giving the DC photocurrent detected. The modulation frequency of the irradiance is $\omega_1 - \omega_2$, which is equal to the Doppler shift frequency. This means that measuring the frequency of the modulated signal gives the Doppler shift frequency, and hence the speed of the moving cells.

The amplitude in equation 1.15 assumes that the two frequency components have the same amplitude. In this application this is unlikely to be the case as more than half of the light incident on the skin will reflect off static tissue, mainly the top surface of the skin. This means the amplitude of the modulated component will be reduced. However, this gives us more information about the total blood flow. An increase in moving cells means that more of the reflected light will be from moving cells, and hence the amplitude of the Doppler shifted component will be higher. This will result in a higher modulation depth. From this, it is possible to measure the amount of moving cells, as well as the speed of cells. Combining this information gives an overall figure for flow.

1.6.3 Concentration and Flow Processing

The two values most commonly used as output from LDBF systems are concentration and flow. Concentration refers to blood volume rather than speed. Flow is found by weighting the concentration result according to frequency, as faster moving cells contribute to more overall movement of blood than a similar but slower moving blood volume.

Analytically, concentration is found from the total power in the Doppler signal [Belcaro et al., 1994]:

$$Concentration = \int_{\omega_1}^{\omega_2} P(\omega)d\omega \quad (1.16)$$

Where $P(\omega)$ represents the power of the signal from the photosensor as a function of frequency, and ω_1 and ω_2 are the upper and lower limits of the frequency range of interest. ω_1 is sufficiently above DC to remove flicker noise and movement artifacts (i.e. Doppler shifts caused by movement of the target object itself, rather than blood within the target object), while ω_2 is set above the highest expected Doppler frequency, ignoring noise signals above this. In the ideal case, ω_1 could be anything greater than DC, and ω_2 could be infinite.

Flow (often referred to as flux) is found by weighting the power spectrum by frequency, and finding the total power in the weighted spectrum. This is equivalent to multiplying the total concentration found in equation 1.16 by the mean blood cell velocity [Belcaro et al., 1994]:

$$Flow = \int_{\omega_1}^{\omega_2} \omega.P(\omega)d\omega \tag{1.17}$$

However, this can be complicated if the processing is to be performed on a voltage signal ($V(\omega)$) rather than a power spectrum ($P(\omega)$). In this case the flux equation becomes:

$$Flow = \int_{\omega_1}^{\omega_2} \omega.P(\omega)d\omega = \int_{\omega_1}^{\omega_2} \left(\omega^{\frac{1}{2}}.|V(\omega)|\right)^2 d\omega \tag{1.18}$$

From equation 1.18, it can be seen that as well as calculating the flux from frequency domain processing of the signal, time domain processing is also possible. Weighting the voltage signal from the system by the square root of the angular frequency, before squaring this voltage (e.g. using a mixing circuit) and averaging over time would give an equivalent result to integrating the power spectrum over all frequencies.

1.7 Typical Laser Doppler Flowmetry System

This section discusses the basic structure of a typical laser Doppler flowmetry system, and the techniques that may be used to apply the principles shown in Section 1.6.

1.7.1 Typical LDBF System using Time-Domain Processing

1.7.1.1 System Diagram

The system shown in Figure 1.4 is a generic LDBF system, using time domain methods to extract final concentration and flow values. The system developed for this project uses frequency domain processing in the form of fast Fourier transform processing on a field programmable gate array (FPGA) to extract the frequency spectrum of the signal. However, the frequency domain method still requires an analogue front-end to detect and amplify the signal prior to digitisation, which uses the photodetector and filter from the early stages of Figure 1.4.

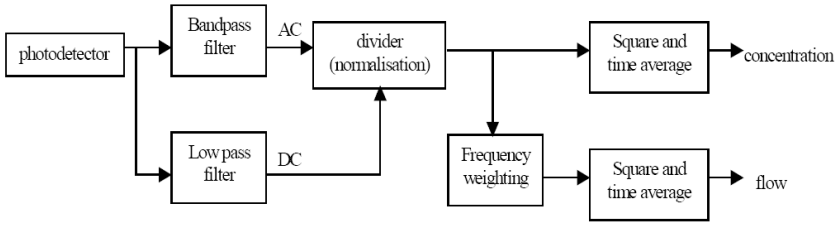

Figure 1.4: Block diagram of a laser Doppler flowmetry system [Belcaro et al., 1994]

1.7.1.2 System Components

The first element of the system in Figure 1.4 is the photodetector, which is generally a photodiode connected to a current-voltage converter. This produces a voltage signal proportional to the incident light power on the detector. This will have a fairly large DC value with the Doppler signal superimposed upon it. There are a variety of circuits that can perform this function, which will be discussed in more detail in Chapter 3.

The bandpass filter is used to remove the frequency components of the signal that are not part of the required Doppler signal (see Section 1.8.3), removing movement artifacts, the DC level and removing noise frequencies above the Doppler bandwidth. The filter may also amplify the required frequencies, otherwise sampling of the signal may require an ADC with very high resolution or a range very closely matched to the output of the photodetector.

For normalisation purposes the signals are divided by the DC value of the photodetector output as given by a low pass filter. This is done to correct for variations in illuminating optical power or skin reflectivity. While this is a separate stage of the common LDBF architecture given in Figure 1.4, in this system the normalisation may be provided as an inherent feature of the logarithmic current to voltage converter to be used. This principle is described in Chapter 3.

1.7.1.3 Processing

After normalisation, processing can be done to give a voltage directly proportional to the concentration and flow of blood. The ideal behaviour of the system here implements the functions described in Section 1.6.3. Concentration is found by squaring and time averaging of the signal, giving a signal with a DC level proportional to blood concentration. This signal is then averaged to give a voltage directly proportional to blood concentration.

For flow calculations, the signal is first passed through a frequency weighted filter. This increases the magnitude of the higher frequency components of the Doppler signal relative to the lower frequency signals. The output of this signal is then squared and time averaged in the same way as for the concentration signal.

The important characteristic for this stage of the system is the frequency response of the frequency weighted filter. The ideal response performs $\omega^{\frac{1}{2}}$ weighting as in equation 1.18, however this can be difficult to implement. If non-ideal weighting is to be used, then the effect of this non-ideal behaviour must be considered. Too little weighting could result in the increased speed of some cells being ignored, giving a flow signal lower than reality. Too much weighting would mean that a small number of fast moving cells, or a high noise floor, would give a large flow signal when actual blood flow may be fairly low.

The complexity of the micro-circulation, conventionally the target of LDBF, means that a definite answer is difficult to find. The precise mathematical definitions of flow do not directly translate to the level of perfusion of tissue due to this flow. There is also little standardisation between different types of LDBF instrument. Instead, LDBF systems can be used to show relative differences in flow, such as between different areas of skin, or on the same area of skin over time in response to some form of stimulus. This means that an approximation to the $\omega^{\frac{1}{2}}$ response may be acceptable [Belcaro et al., 1994, Gu, 2007, Hoang et al., 2010].

1.7.2 Considerations for Frequency Domain Processing

Alternatively if using the frequency domain method, the time domain frequency weighted filter and the square and time average blocks are not required. Instead, the signal will be digitised and an FFT will be performed after the low pass filter. The separate concentration and flow signals can then be found in a similar way, albeit working in the frequency domain rather than the time domain. Taking an average of the spectrum given by the FFT will give a value proportional to blood concentration. To give a similar value for flux, the spectrum must first be amplified by the same transfer function as the time-domain frequency weighted filter above. Taking an average of the new, weighted spectrum will give the flow value.

1.8 Expected Doppler Signal

This section describes the input signal to be expected from the photodiodes in the system. These details will then be used as the specification for the system design.

1.8.1 Light Levels and Photocurrents

The design of the IC depends on the expected incident light levels, as various circuit characteristics can depend on input current DC level, as this sets the DC operating point.

Previous measurements of incident light power for laser Doppler blood flowmetry used a laser power of 7 mW on to the skin, which gave a measured power density at 5 cm of 50 μWcm^{-2} or 0.5 Wm^{-1}[Kongsavatsak, 2005]. Over the area of the 1000 x 50 μm photodiode used here this gives an optical power of:

$$P_{optical} = 50 \times 10^{-6} \times 1 \times 10^{-3} \times 0.5 = 2.5 \times 10^{-8}\,\mathrm{W} = 25\,\mathrm{nW} \tag{1.19}$$

This photodiode size is chosen based on two factors. The pixel width of 1 mm gives a wide enough array to reliably image a laser line onto the array, without the mechanical scanning of the beam causing alignment or beam wandering issues (the beam moving off the array as the line scans). At the same time, it is not so large as to leave insufficient silicon area for the signal detection and conditioning circuits. The pixel pitch of 50 μm for 64 pixels gives sufficient space for pixel level circuitry while not increasing the size and cost of the IC unnecessarily. There is an additional factor that requires consideration, that of diode capacitance. The photodiode has a large area and periphery, and the capacitance of the diode is proportional to the area of the photodiode depletion region. If this capacitance is too large, it will reduce the bandwidth of the detector by creating a high RC constant at the front-end. This is investigated by simulation in Chapter 3, Section 3.2.1.3.

The responsivity of the photodiodes used here at red and near infrared wavelengths is approximately 0.3 A/W[Kongsavatsak, 2005]. This gives a total photocurrent of:

$$I_{DC\,photo} = 2.5 \times 10^{-8} \times 0.3 = 7.5 \times 10^{-9}\,A = 7.5\,\mathrm{nA} \tag{1.20}$$

However this is for illumination of a single spot with roughly 1 mm diameter. For a 64 pixel system it can be assumed that the area on to which light is projected will increase by 64.

The laser to be used by this system is also more powerful, at around 45 mW, increasing by a factor of 6.43. This results in a drop in incident power of $\frac{64}{6.43} = 9.96$. This will give a DC photocurrent of:

$$I_{photo} = \frac{7.5 \times 10^{-9}}{9.96} = 753\,\text{pA} \tag{1.21}$$

The 45 mW laser is also an infrared laser. This gives higher skin penetration depth due to the longer wavelength [Belcaro et al., 1994], which will marginally increase the photodiode responsivity above that for the light wavelength used for previous readings, increasing the DC photocurrent. The effect of light penetrating deeper into the tissue may not significantly affect the DC photocurrent, as the reflected light depends on scattering rather than absorption, although again this adds some margin of error to the above photocurrent figure. However, the greater penetration depth should mean more light interacts with moving blood cells rather than static tissue near the surface. This may lead to an increase in the Doppler ratio, i.e. the ratio of modulated light to constant intensity light at the detector, and hence to an increase in the AC signal being measured.

The typical DC photocurrent used here in simulations with a $1000\,\text{x}\,50\,\mu$m photodiode is therefore $765\,p$A, with an AC peak-peak current of 75 pA. This assumes a Doppler ratio of approximately 10%. This ratio is independent of illuminating power - if the incident optical power rises, the amount of directly reflected light will increase, however the number of photons scattering of moving blood cells will increase by the same proportion. [Belcaro et al., 1994]

To ensure the system works over a range of light levels, simulations will also be performed at a minimum and maximum light level, with DC photocurrents of 75 pA and $1,500$ pA respectively.

1.8.2 Large Photodiodes and Speckle Size

The ICs produced here were required to have larger photodiodes than usual in CMOS sensors, due to concerns about the beam 'wandering' off the sensor during mechanical scanning. The 1D array topology means this photodiode size is feasible, however the increase in detector area also affects the photocurrent input to the detection circuits. Conventional thought would suggest that the large photodiodes used here are an advantage for signal detection, resulting in higher light levels. However, changes in lens focal lengths to focus light to the

size of the sensor means that smaller sensors should receive the same light level.

It is possible that the large sensor actually reduces the Doppler ratio, hence reducing the SNR. This is because the larger pixels may result in more speckles being imaged onto each photodiode.

Speckle size for LDBF is given by [Steenbergen, 2004]:

$$Size_{speckle} = \frac{1.22 \times \lambda_{laser}}{NA} \tag{1.22}$$

Where NA is the numerical aperture of the system. For the imaging setup used here, the NA can be found based on the size of lens used to focus light onto the IC and the distance from the lens to the target.

$$NA = \frac{r_{lens}}{\sqrt{D_{target}^2 + r_{lens}^2}} \tag{1.23}$$

To give an estimate of speckle size, the lens diameter for this system will be approximately 25 mm, while the distance from the lens to the target will be approximate 250 mm. If a near infra-red laser is used, the wavelength will be approximately 750 nm, This gives a speckle size of:

$$Size_{speckle} = \frac{1.22 \times 750\,\text{nm}}{0.1} = 9.1\,\mu\text{m} \tag{1.24}$$

The photodiode size is $1000 \times 50\,\mu$m, giving an estimate of ~600 speckles per pixel. Compared to previous CMOS detectors with diode size around $50 \times 50\,\mu$m [He et al., 2009, Gu et al., 2008], giving approximately 30 speckles per pixel. The AC photocurrent is due to changes in intensity at each pixel caused by the speckles appearing and disappearing. If the number of speckles on a pixel is sufficiently large, the change in overall intensity as the speckles could average out to a constant intensity. However, the existing Moor system to be modified for this project has larger sensors ($3.6\,\text{mm} \times 0.5\,\text{mm}$) than the prototype ICs, while having similar lens size and distance to target, suggesting the the speckle size is not a significant problem. However, this reduction in modulation depth would result in a smaller AC photocurrent to that expected from the calculations shown in Section 1.8.1.

1.8.3 Spectrum of the Doppler Signal

As well as the magnitude of the photocurrent, it is important to know the spectrum of the Doppler signal. As not all blood cells are moving in the same direction or at the same speed, the modulated signal will have a continuous spectrum from DC up to several kHz. Figure 1.5 [Belcaro et al., 1994] shows the type of signal expected from two different targets.

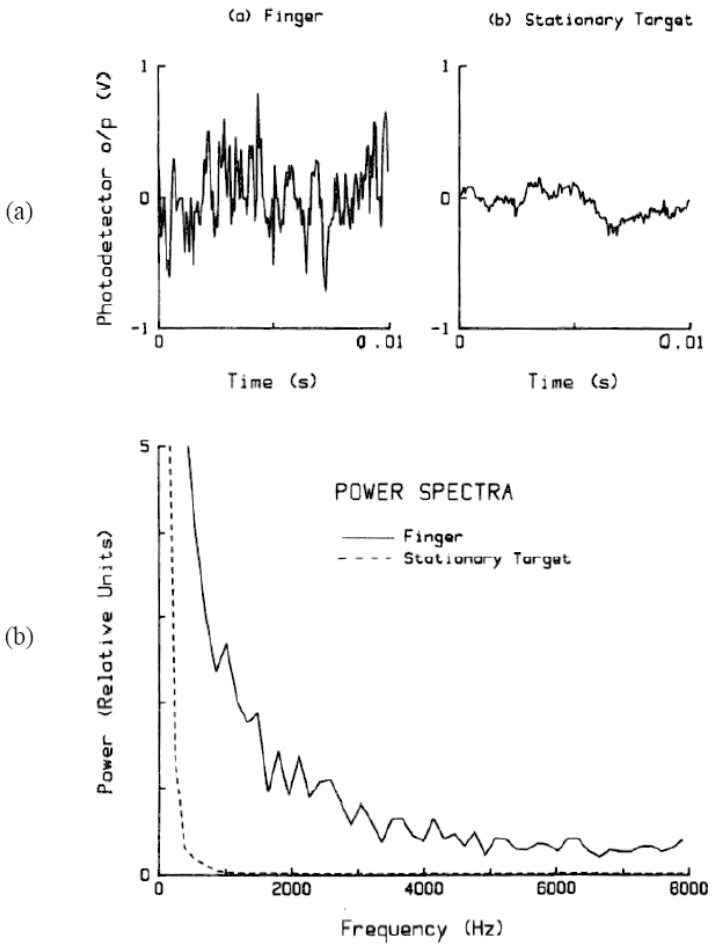

Figure 1.5: Time and frequency domain plots of Doppler signals from a finger and a static target [Belcaro et al., 1994]

It can be seen from the power spectra for the finger that the majority of the signal is at the lower frequencies, below 2 kHz, with a tail extending towards 10 kHz. The shape of this curve could also be seen as a histogram of red blood cell velocity, with most blood cells moving at lower velocities, with a smaller number of cells moving faster. The spectrum is mostly Gaussian, except for the tail extending to higher frequencies due to a small number of cells with high velocity relative to the detector.

While most of the signal is at lower frequencies, the signal must not be excessively truncated by low pass filtering. While the majority of the cells are moving fairly slowly, a single fast moving cell is more significant to the total blood flow than a slower moving cell. The overall flow value for any one point on the skin is therefore a frequency weighted average of the power spectrum.

For the application considered in this thesis the bandwidth to be used is 20 kHz. This is higher than apparently necessary from the spectrum shown in Figure 1.5 for a number of reasons. The first is that in some circumstances the flux will increase above that shown - for example when blood flow in a limb or finger is occluded by a tourniquet, there is a brief but considerable increase in flow when the tourniquet is first released. The magnitude and duration of this increase can contain clinically useful information, and as such the higher blood velocity must not be ignored due to filtering. Secondly, as this project is intended to image flow in veins rather than micro-circulation in capillaries, there is the possibility of faster moving blood cells. This second factor is less certain, as while the total flow in a vein is larger, the larger diameter of veins means that cell speed may be lower. The veins intended to be imaged here also tend to run along the skin. This makes them suitable for inserting needles and cannulae, but also means that the majority of the velocity component of each cell is perpendicular to the incident light. This means that the velocity of the cells relative to the incident photons - the relevant portion for causing Doppler shifts - is fairly low.

Given the above, the bandwidth requirement has been set fairly high. However the use of a tuneable anti-aliasing filter means that if a high cut-off frequency of 20 kHz is deemed excessive, the cut-off could be reduced. This should reduce noise by filtering out noise between the original and reduced cut-off frequency, and would also allow a lower sampling speed, possibly facilitating the digital processing and read-out. Conversely, the cut-off of the filter could be set higher if faster flow is to be observed, although if the digital back-end of the system is designed around a 20 kHz sampling rate this would require a more significant change to the system.

There is also a low cut-off frequency, as while the Doppler spectrum is continuous, at lower frequencies the signal from the photodiode includes movement artifacts which will detract from the actual blood flow signal. Removing lower frequencies also vastly reduces $1/f$ noise in the system. The low cut off to be used here is around 100 Hz.

Chapter 2

Review of Existing Laser Doppler and Vein Location Systems

2.1 Introduction

Laser Doppler has been used for blood flow measurement since the 1960s. This chapter discusses the development of Laser Doppler Blood Flow (LDBF) systems from early prototypes to modern systems in clinical and research use. This will include looking at a range of systems that employ the LDBF technique in different ways, and the various ways that such systems are used or could be used.

While this chapter focuses on LDBF systems, alternative technologies for vein location are also considered. If LDBF is to be used for a specific application rather than general blood flow imaging, it has to offer significant advantages in terms of either performance, useability or cost over any existing alternatives.

Finally, issues that apply to the widespread clinical use of LDBF imagers are considered. Existing LDBF devices are mostly used for research purposes, so may not be subject to all the constraints placed on a more widely used device.

2.2 Development of Laser Doppler Blood Flow Measurement Systems

2.2.1 Early Development of Laser Doppler Blood Flowmetry

Laser Doppler Flowmetry (LDF, when not used for blood flow assessment) was first used for the measurement of fluid flow in 1964 by Yeh and Cummins [Yeh and Cummins, 1964]. This was measuring the flow of fluid in a tube in a test rig, with a separate reference arm as in a Mach-Zehnder interferometer. The technique was first applied to blood flow by Riva et al in 1972, on retinal arteries of rabbits [Riva et al., 1972]. Riva also applied the technique to 200 μm capillaries.

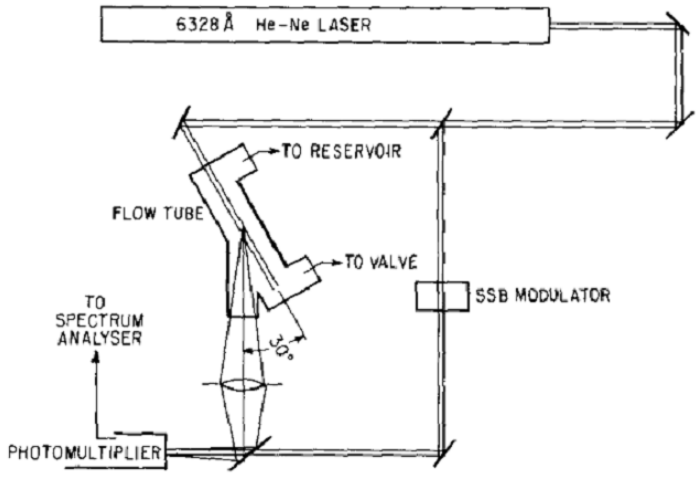

Figure 2.1: Interferometer used for Laser Doppler Velocimetry in retinal vessels [Yeh and Cummins, 1964]

The technique was first used on human retinal blood vessels in 1974 by Tanaka and Riva [Tanaka et al., 1974]. This required a reduction in laser power for safety, which in turn required a photomultiplier tube and a photon counting detector.

The technique used initially was Laser Doppler velocimetry (LDV) rather than flowmetry (LDF), as the measured quantity was the speed of blood flow rather than an overall figure for flow itself.

Stern used the technique on flow in capillaries, using statistical processing of the raw signal to determine overall flow rather than flow velocity [Stern, 1975]. This is more useful as a clinical measure, as it is overall flow which is relevant for tissue receiving sufficient oxygen,

nutrients etc from the circulatory system. Stern also compared LDBF to existing techniques, mainly xenon washout, and found good correlation.

2.2.2 Single Point Measurement Systems

A clinical instrument making use of LDBF was built in 1978 [Watkins and Holloway, 1978], but this had poor reproducibility and was difficult to use. At the same time, Powers and Frayer developed a Laser Doppler velocimetry instrument for use during surgery [Powers and Frayer, 1978].

One of the problems with early LDBF systems was instability of the lasers used. Mode hopping of the lasers used meant the Doppler shift of the incident light was masked by the shift in wavelength of the laser [Sargent and Scully, 1972]. One technique used to overcome this was developed by Nilsson, that of using a differential sensor. As a result of this, common mode sources of noise such as laser mode hopping are removed [Nilsson et al., 1980]. Differential detection is made possible as the flow of blood causes a randomly changing speckle pattern, where the signal from an individual speckle has random phase. This means the Doppler signals from two detectors have components with similar frequency and magnitude, but different phase. Measuring the signal differentially therefore removes common mode noise, but the Doppler signals themselves do not cancel each other out.

2.2.3 Scanning Systems

The single point technique was used in a clinical setting shortly after its initial development for wound assessment [Oberg et al., 1979, Wunderlich et al., 1980]. However, the single point nature of the technique presents obvious limitations.

Flow assessment over an area was developed by combining a single point measurement system with a pair of mirrors to implement a raster scanning system. Essex and Byrne developed a system that scanned a $500 \times 700\,\mathrm{mm}$ area, with $2\,\mathrm{mm}$ resolution. However, the scanning requires repeated re-positioning of the mirrors, which limits the frame rate. For example, The system developed by Essex and Byrne took 6 minutes to acquire an image [Essex and Byrne, 1991, Essex, 1994]. Figure 2.2 shows the structure of this type of system, as used in an imager produced by Moor Instruments Ltd.

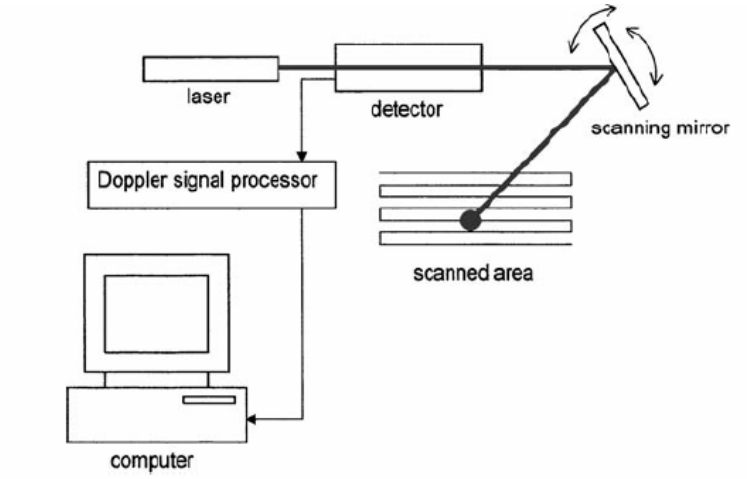

Figure 2.2: Main optical and electronic elements of a scanning LDBF system (Moor Instruments, UK)

2.2.4 2D Doppler Flow Imaging Systems

The scanning systems previously described proved the usefulness of obtaining flow images. To improve the acquisition time for such images, Serov and Steenbergen used a 2D sensor array to measure flow over an area without mechanical scanning [Serov et al., 2002].

Figure 2.3 shows the setup used for imaging with a 2D commercial CMOS sensor [Serov et al., 2005]. Figure 2.4 shows images produced by this system. Image acquisition time for this system for a 256×256 image, captured as a series of 64×8 sub-frames, is 90 seconds. The majority of this time is consumed by performing FFT calculations on a DSP to calculate flux from the raw signal. This system used a sampling frequency at each pixel of 16.8 kHz, which is lower than the preferred figure of 20 kHz [Belcaro et al., 1994]. The system can randomly address pixels, such that it can work as a single point system, which allows a sampling frequency of 40 kHz to be used if single point measurement rather than imaging is sufficient.

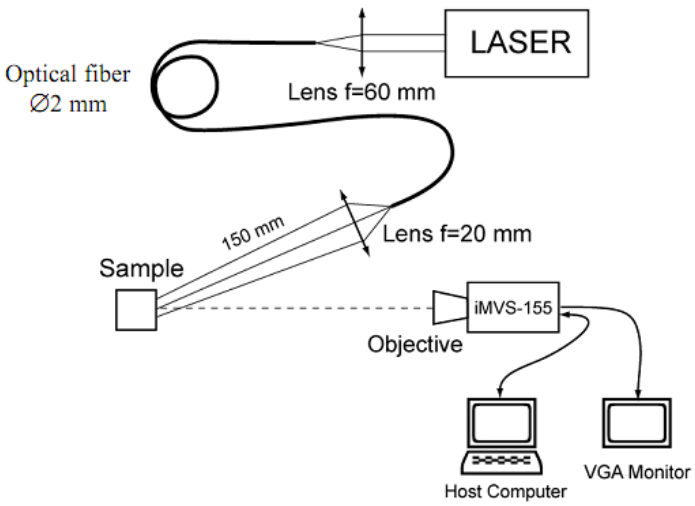

Figure 2.3: 2D imaging setup used by Serov et al.[Serov et al., 2005]

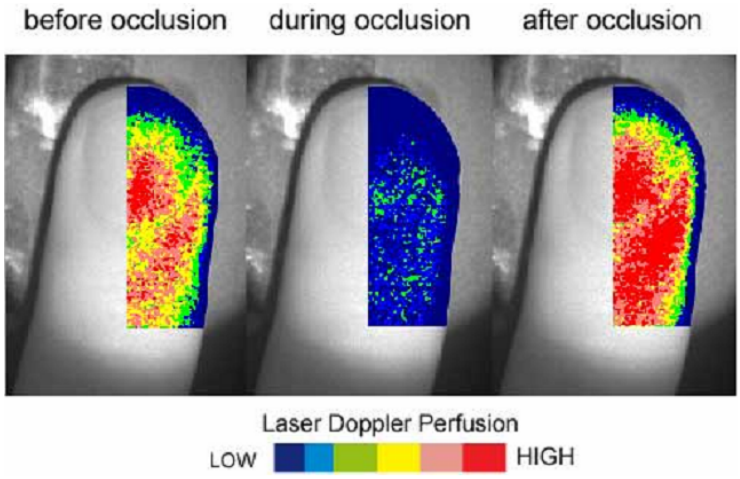

Figure 2.4: 2D images of a finger from equipment in Figure 2.3 [Serov et al., 2005]

Improvements to the data acquisition system for a similar system architecture have improved the speed of image acquisition to 1.2 s, making the system much more clinically useful, although the pixel bandwidth is still lower than ideal, at around 4 kHz [Serov et al., 2006b]. Figure 2.5 shows a sequence of images produced by this newer imager during occlusion and release of blood flow in a finger. This imager also combined the LDBF technique with laser speckle contrast analysis (see Section 2.4.1), which allows a higher frame rate (10 frames/second) to be achieved using a less accurate flow assessment method.

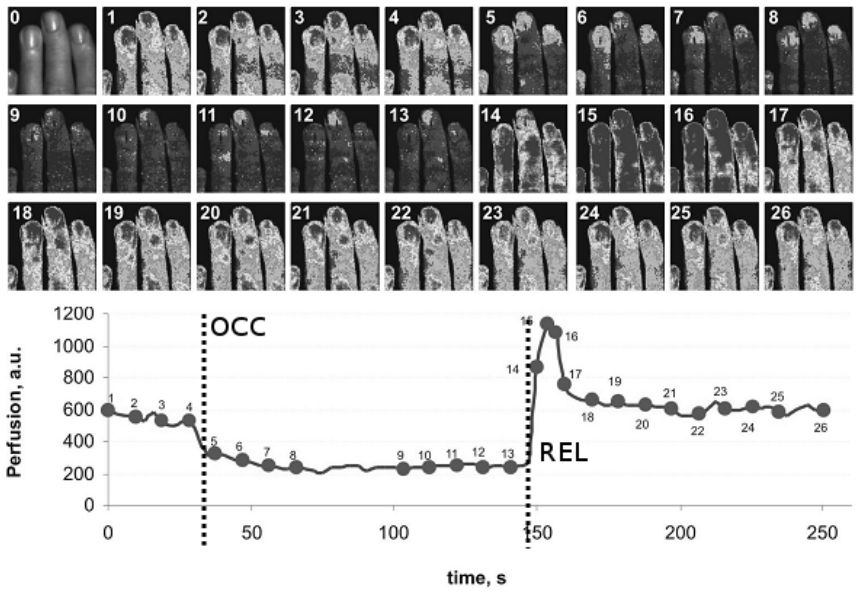

Figure 2.5: Images and single pixel plot of flux in a hand during occlusion and release, recorded by a 2D sensor [Serov et al., 2006b]

The system developed by Serov et al is also capable of locating veins. Figure 2.6 shows a flow map produced by averaging 10 individual frames, in which veins are clearly visible. The speed of this system means that the averaging can be performed without making acquisition time too long for viable use, with this image taking 12 s (10×1.2 s) to acquire.

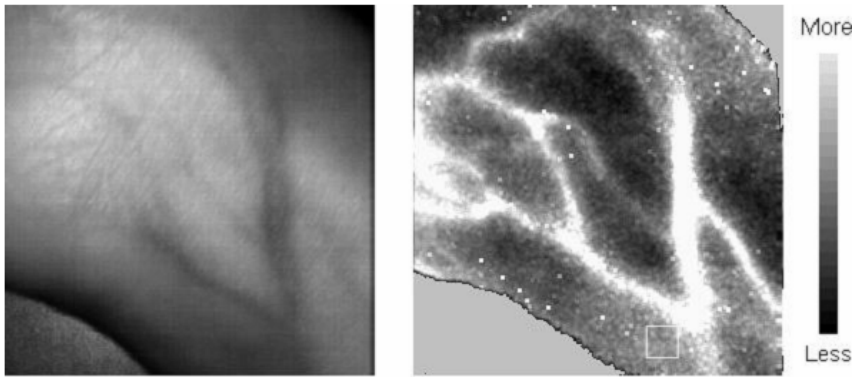

Figure 2.6: Image of veins (right) and visible image (left) of a hand produced by 2D sensor [Serov et al., 2006b]

2.2.5 2D integrated CMOS LDBF Sensors

A major limitation in imaging systems is data bottlenecks. If data is to be processed on a PC, then every pixel has to be sampled with high bandwidth, ideally 40 kHz, for a 20 kHz

bandwidth [Belcaro et al., 1994]. This can be mitigated by increasing data acquisition speed or number of channels and by using a lower sampling speed, as with the systems made by Serov shown in Section 2.2.4.

An alternative technique, given the use of CMOS detectors, is to integrate processing electronics onto the sensor IC. By performing the processing on-chip, the amount of data to be acquired by a PC is vastly reduced.

Kongsavatsak developed a 16x1 IC which included sensors, amplifiers and digital flux processing on one IC [Kongsavatsak, 2005]. Gu developed a pixel design that integrated the photodetector with analog flux processing [Gu, 2007].

One difficulty of this approach is the increased time and cost required to prototype designs, particularly when the full potential of combining digital processing circuits with analogue sensors and signal conditioning is used. This can be mitigated by using FPGAs to develop processing techniques [Zhu et al., 2006, He et al., 2009]. Figure 2.7 shows an image captured using an FPGA, with an acquisition time of around 4 seconds. This system uses line illumination and a 1D array of pixels, with mechanical scanning of the line. This is intended to combine some of the advantages of full-field imaging systems without requiring the same laser power as these [Hoang et al., 2010]. As only one axis of scanning is required, the time required for mechanical movement is greatly reduced, allowing faster image acquisition than single point scanning systems. Some delays are still caused by the mechanical scanning, and therefore the sensors and processing circuits used are capable of faster image acquisition if other limitations of the system can be reduced.

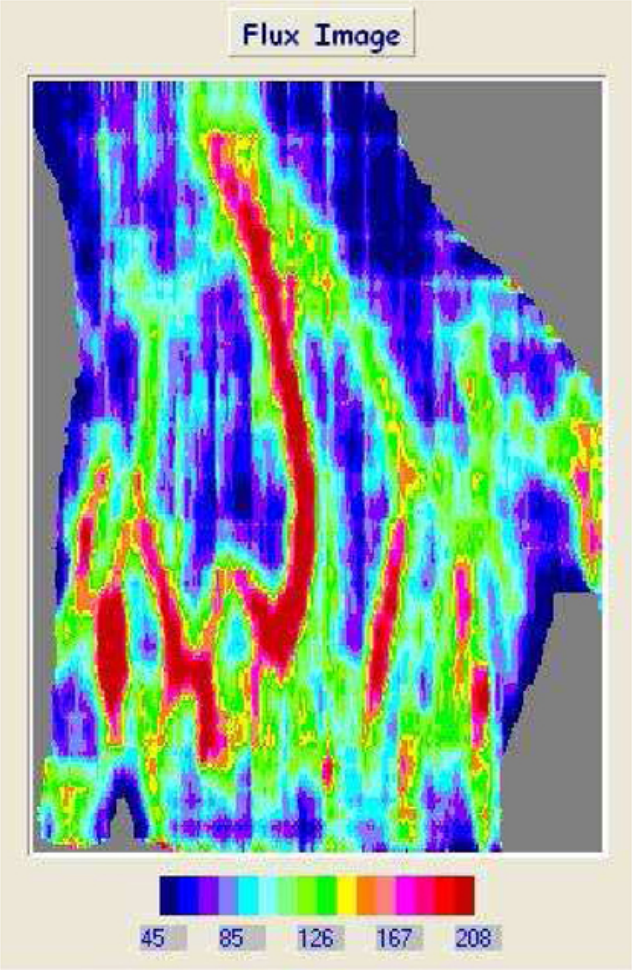

Figure 2.7: Image of blood flow showing veins in a hand, processed by an FPGA system with an algorithm suitable for implementation on-chip [Morgan and Hayes-Gill, 2009]

2.2.6 Choice of Sensor technology - CCD or CMOS

Traditionally charge coupled devices (CCDs) have been the sensor of choice for conventional imaging applications. Technologies used to fabricate CCDs have been optimised for optical performance, enabling low-noise, high contrast and low cross-talk between pixels [Theuwissen, 1995]. However, CMOS sensors have become more widely used for imaging applications [Fossum, 1997], and are becoming more common in consumer devices such as digital cameras [Theuwissen, 2001]. There are various reasons for this; the specialisation of CCD processed means that it is not feasible to integrate any processing circuits onto the same IC, requiring additional components and interconnections; CCDs require serial readout, so selecting areas of interest or single pixels is not possible (or at least does not allow a speed in-

crease), whereas CMOS sensors are generally randomly addressable; CMOS pixels can work at lower voltages, and only access pixels when required, resulting in lower power consumption than CCDs, which require require larger voltages (8-15 V), switched more frequently for the charge transfer process [Holst and Lomheim, 2007]. The reduction in size of CMOS features as the processes have developed has improved the fill factors and resolutions that can be achieved, with single chip sensor arrays having up to 14M pixels available [Collins, 2009].

For LDBF applications, the major advantage is frame rate - the serial readout of CCDs means that it is very difficult to achieve a high enough sampling rate for Doppler signals. Randomly addressable pixels are also very significant, as it means that it is easy to trade off image size for acquisition time, so an area of interest can be viewed at improved frame rates with only software settings needing to be changed.

However, faster CCDs are also available, such as frame transfer CCDs. These interleave rows of sensors with rows of buffers and acquisition circuits, allowing the entire collected charge to be transferred to a buffer very quickly, and the next image exposure to be started while the previous image is sampled and converted to a digital output. However, this makes the sensor area considerably larger (roughly double the size), and hence the cost is higher. As the fabrication process used is the same as for conventional CCDs, it is still not possible to integrate standard CMOS circuitry onto the same IC as a frame transfer CCD.

2.3 Development of LDBF for Medical Use

2.3.1 Research Use of Laser Doppler Flowmetry Systems

LDBF is a well established technique in medical research. It has been used for pressure sore diagnosis in wound healing [Nixon et al., 1999], ischemic ulcers [Gschwandtner et al., 1999], burns assessment [Brown et al., 1998], joint inflammation [Ferrell et al., 1997], allergic reactions [Clough et al., 1998], dermatology [Quinn et al., 1991], wound assessment [Khan and Newton, 2003] and physiology of diabetes [Morris et al., 1995].

Some of these studies (e.g. burns and wound assessment) may lead to clinical use of LDBF devices by demonstrating the utility of LDBF in clinical applications, or an improvement in care made possible. However, the use of the technique for research purposes does not mean that it is suitable for routine clinical use. Research use is likely to be under much more

controlled conditions, in terms of selection of patients, location of testing (for laser safety reasons, the use of high power lasers in some systems is sometimes confined to separate rooms, rather than general wards), and device useability may be a secondary consideration to the information provided by the device.

However, the later stages of development of LDBF imagers introduced in Sections 2.2.4 and 2.2.5 show obvious improvements in useability and information provided that make routine clinical use more feasible - for example, a rapid acquisition time means that the subject (or patient) does not have to remain in a fixed position for an uncomfortable time, and movement artifacts are reduced as there is less time in which movement can occur. A higher frame rate also allows changes over time to be seen, making assessment of blood flow in response to external stimuli easier.

2.3.2 Non-Imaging Applications of LDBF Techniques

As well as development of LDBF techniques for general imaging, there are applications where simpler systems, based on single point sensors, can be useful. Koelink developed a sensor that used a combination of two wavelengths for illumination [Koelink et al., 1994]. The two wavelengths have different penetration depths due to differing absorption of the light by tissue, so separate flow measurements in superficial capillaries and deeper blood vessels can be made.

This technique could potentially be combined with an imaging technique, giving two images of surface and deeper tissue blood flow. The integrated 2D sensors introduced in Section 2.2.5 are particularly well suited for this application, as the combination of processing on chip, coupled with the low size of each sensor and cost per unit (in high volume) means that systems using multiple sensors are more feasible. For commercial systems, the use of two sensors would further increase the demands on the data acquisition system, which is already a limitation on the system.

An alternative to imaging systems to increase clinical useability is to develop a single point system that is compact and easy to use. To this end Serov developed an integrated probe that combined sensor and laser into one small package (6 mm diameter x 6 mm height), with data acquisition performed by the sound system of a laptop PC [Serov et al., 2006c].

Such sensors have various applications. Afshari et al used a small LDBF system built into the

handle of a golf club [Afshari et al., 2002]. This demonstrates a method of monitoring blood flow of a subject performing various activities - sports, using power tools etc. The compact size of the system used means that the measurements can be taken in a way that does not interfere with the activity itself. In addition, the use of a compact Laser Doppler sensor built into other equipment could have non-medical applications, for example monitoring flow or vibration in industrial processes.

2.4 Alternative Flow Assessment Methods

While this thesis is mainly investigating laser Doppler blood flow imaging, there are other methods available that use the coherent nature of laser light to measure blood flow. This section considers two of these techniques with regard to their use for a vein location system.

2.4.1 Laser Speckle Contrast Analysis

An alternative technique to direct measurement of Doppler shift from moving blood cells is that of laser speckle contrast analysis (LASCA). This technique makes use of the speckle pattern produced by a coherent light source illuminating a scattering medium, such as skin. This pattern is produced in the same way as the interference fringes produced by two intersecting beams from coherent sources - each individual scatterer in the medium acts as a separate source, producing a complex interference pattern of light and dark spots.

The movement of red blood cells in capillaries in the skin means that the scattering medium changes, changing the speckle pattern produced. LASCA uses a standard CCD camera to capture this speckle pattern. If there is no flow, the pattern does not change so the captured image has high contrast. With higher flow, the change of the pattern during image capture causes blurring of the image produced. Measuring the contrast of the subsections of the image (e.g. 5x5 pixel blocks) allows a 2D flux map to be built up. As this can be done using a single image captured with a commercial CCD sensor, achieved frame rates can be very high. The technique was developed by Briers, Fercher and Richards [Fercher and Briers, 1981, Richards and Briers, 1997], and a commercial device was developed by Moor Instruments. The Moor Instruments FLPI system can operate at up to 25 frames/second, giving images of the sort shown in Figure 2.8.

This obviously has considerable advantages in terms of speed, however the major limitation is penetration depth, as the technique only detects flow in the top surface of the skin. This is partly due to the lower power used in the Moor FLPI system, which is sold as a class 1 laser product to make clinical use more feasible.

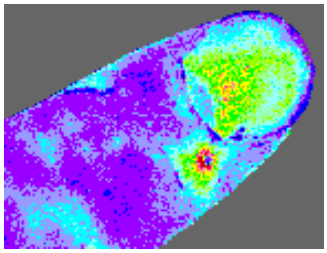

Figure 2.8: Image of flow in a finger using laser speckle contrast analysis [MoorInstruments, 2007a]

Briers also compared LASCA and LDBF, investigating the similarities between the techniques [Briers, 1996]. Both techniques were developed separately, but both have similar roots. Both Doppler shifts and speckle pattern changes are related to the changes in optical path length caused by movement of the scatterers, i.e. red blood cells. Stewart compared speckle analysis and LDBF and found that the speed of acquisition of LASCA was a considerable advantage in clinical use [Stewart et al., 2005].

2.4.2 Laser Velocimetry

As well as measuring the Doppler shift of light reflected from a single laser beam, it is possible to detect the speed of moving particles using an arrangement of two intersecting beams. This produces an interference pattern, and the movement of particles through this pattern produces a signal that can be measured to determine velocity. A diagram of a system using this technique is shown in Figure 2.9.

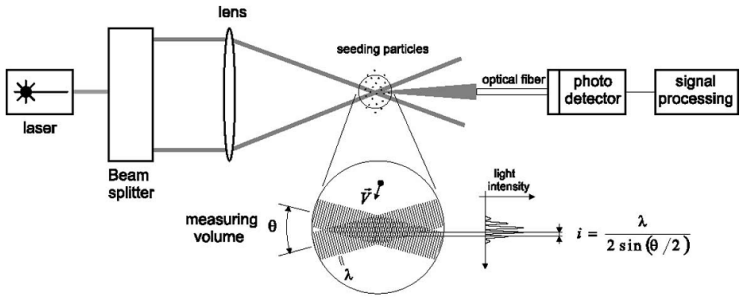

Figure 2.9: Two beam technique for Laser velocimetry [Le Duff et al., 2004]

The two beam technique was developed at a similar time to the early LDBF systems [Rudd, 1969, Stein and Pfeifer, 1969], but did not become widely used for this application. This technique is best suited to a smaller number of particles - a single particle crossing a series of fringes will clearly produce a wide variation in reflected intensity, whereas the light reflected from a stream of particles will tend to average out, giving less variation. This technique is not ideally suited to blood flow assessment, as blood flow consists of a high number of particles in a continuous stream, meaning that the variation in reflected intensity from each particle is averaged out by surrounding particles.

Despite this, as the technique avoids some of the complexities of LDBF systems (such as requiring a stable single mode laser) it has been used to develop a low-cost LDV sensor [Le Duff et al., 2004]. Reduced costs of the sensor, and hence costs of a device built using the sensors, could increase the feasibility of widespread adoption of such devices for clinical use.

2.5 Alternative Vein Location Technologies

So far the techniques discussed have all been based on laser imaging. However, as the system to be developed here is intended for vein location rather than purely flow imaging, alternative methods for vein location should be considered.

2.5.1 Trans-Illumination

The simplest of these methods is using back illumination. The principle here is to illuminate tissue beneath a vein,such that the vein appears as a shadow on the skin. The most crude form of this involves shining a torch on the back of smaller limbs (e.g. through a hand) in a darkened room. However, the scattering of light through tissue means that the width of tissue across the whole limb results in very poor resolution .

A device that applies the back illumination technique in a more advanced way is the vein lite. This uses a ring of LEDs such that tissue is illuminated through the skin next to the region of interest. Light then scatters down into the tissue and back to the surface, except where blocked by a vein. This variation of the back-illumination technique is known as trans-illumination. Figure 2.10 shows this device in use.

Figure 2.10: A veinlite being used to highlight a vein in an arm [Translite, 2001]

While this is a very basic device compared to the LDBF systems shown, it has been shown to be useful in clinical practice [Kussman, 2001, Lindsey, 2005, Zimerman, 1991]. The main advantage of the system is it's simplicity - it is easy to use, and cheap to buy. An LED version is also available, such that it is very portable, making it feasible to carry around for use whenever required.

However, its simplicity also causes some of it's limitations. The veinlite does not detect flow, instead showing the physical structure of a vein. This means it cannot assess the flow within the vein itself, which may be an important indicator of the suitability of a vein for canulation. It also requires contact of the light head with the skin, which may not be suitable in some situations (for example, where skin is damaged), whereas LDBF can be a totally non-invasive and non-contact technique.

2.5.2 Infra-Red Image Processing

Infra-red light is often used in vein location and flow imaging, as the longer wavelength of this light penetrates further into the skin [Belcaro et al., 1994]. This means that a single image of a body part such as a hand taken using infra-red illumination shows some information about the features under the skin, such as veins. This can be done using commonly available sensors and light sources, with processing being performed by a standard PC. This technique has been used to identify veins using an imaging process [Shuwang et al., 2005, Chen et al., 2007]. Figure 2.11 shows a raw image taken in this way, and a processed binary image from an

automated process of locating veins within the image.

This technique is in some ways a more advanced application of the trans-illumination technique, using the scattering of light through tissue to produce shadows where the light is blocked by veins. Here the illumination is applied over a wider area, with image processing used to locate the veins. The veinlite uses this technique with more controlled illumination such that the processing is not required.

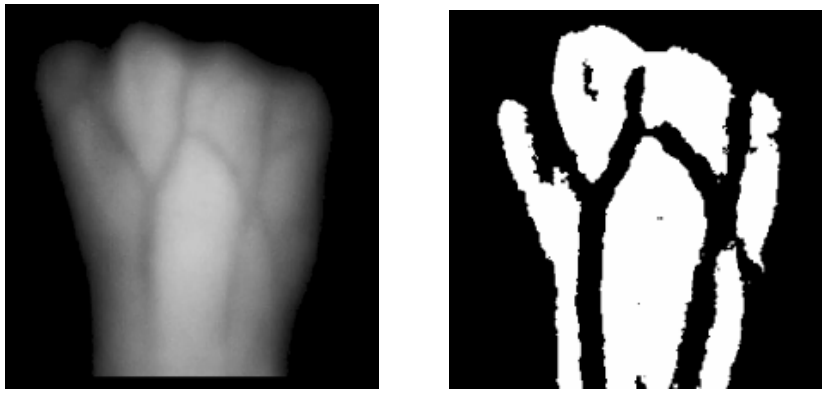

(a) original IR image (b) processed image

Figure 2.11: Processed image of veins and original IR image [Chen et al., 2007]

An extra step that can be used with this technique, which is applicable to other vein location systems, is to project the vein image produced back on to the skin. This approach was developed by Zeman [Zeman et al., 2005], and a commercial device has been developed by Luminetx. This technique avoids the issue of relating a point on a vein image to a physical location. The limitation of this technique is the additional complexity and cost of the projection equipment. The projection must also be well aligned, even on a non-uniform surface such as tissue

2.5.3 Haptics

Haptics refers to recognition of objects through touch [Klatzky et al., 1985]. This would generally refer to a person identifying objects through touch, but this approach to vein location could potentially be automated. A system has been developed by Zivanovic that uses a movable probe to palpate veins in the same way as a human would locate veins by touch [Zivanovic and Davies, 2000]. The system scans the probe across an arm, building up a force-position plot for each point. A fluid-filled vein has a force-position response that can be

recognised, locating the vein. This system was intended to be used as part of a robotic blood sampling system, so is intended to be fully autonomous rather than an aid for a clinician.

As a vein location method, the system is slower than existing methods due to the extra mechanical scanning required (the probe must be scanned across a line, and at each point must be moved towards the arm to build up the force-pressure curve). The system developed took roughly 1 minute to measure along a 15mm line. However, the technique could potentially be developed to improve speed. For example if an array of pressure sensors could be developed, e.g. using MEMS processes, it would be possible to produce a force-position map across an area. This could be combined with processing techniques similar to image processing to locate veins.

Zivanovic's system also used a similar force-position sensing technique to automate the needle insertion procedure. While Zivanovic did not expect the automated needle insertion technique to be trialled on humans, or to enter routine use, in the foreseeable future, this does raise the issue of the level of automation possible if vein location and flow mapping systems become sufficiently reliable.

2.5.4 Ultrasound

Ultrasound is a common technique in medical diagnosis, allowing images of tissue structure at any depth. This could be used for vein location by looking for the structure of veins rather than the flow within them. Ultrasound devices generally require significant user training, as well as the use of a gel to couple ultrasound from the transducer into tissue, which is a significant disadvantage over a non-contact flow imager. However, the increasing use of ultrasound in hospitals may mean that the hardware and user training becomes commonplace for other reasons, allowing its adaptation for vein location [Mbamalu and Banerjee, 1999]. Figure 2.12 shows an image from an ultrasound imager while being used for vein location. The complexity of this image, compared with vein images from 2D LDBF systems (such as those in Figure 2.7) demonstrates the additional challenges in using current ultrasound devices. The ultrasound technique does offer information on depth of the vein within the tissue, which could be an advantage if visualised in an accessible manner. However, unlike for LDBF systems the images produced do not display flow, meaning that the actual flow of blood within located veins cannot be known.

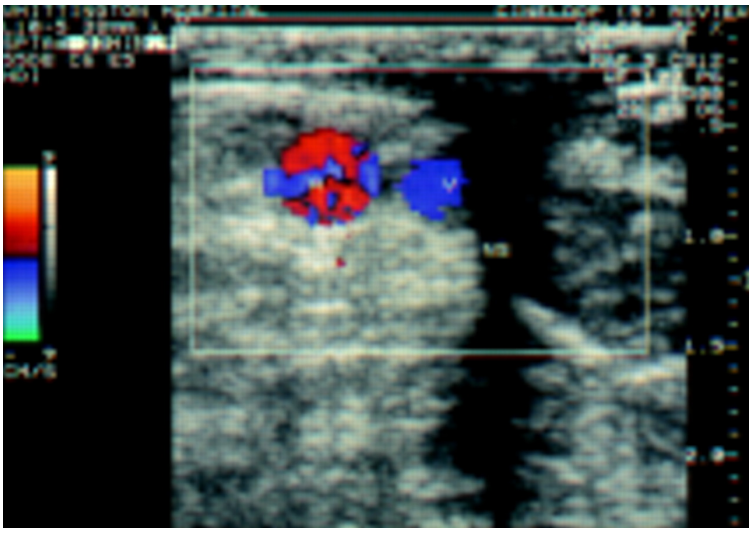

Figure 2.12: Image of a vein and artery from an ultrasound imager [Mbamalu and Banerjee, 1999]

2.6 Clinical Use of Vein Location Systems

This section considers aspects of vein location systems that must be addressed to make a device useful and practical in a clinical setting.

2.6.1 Laser Power against Acquisition Speed

As shown in Section 2.2.4, 2D imagers have considerable advantages in terms of acquisition time, as mechanical scanning delays are removed. However, the drawback here is that to illuminate all areas of the tissue and hence all pixels of the sensor with sufficient light requires very high overall illumination power. For example, the system developed by Serov and Lasser used a 250 mW laser, compared with 1-2 mW in the early single point and scanning systems [Serov et al., 2006a, Oberg et al., 1979, Nilsson, 1996].

For use in a hospital this has potential problems with laser safety. Even though the beam is diffused over a wide area, it is possible that it could be reflected so as to cause a laser hazard, particularly if widely used in an environment where common laser safeguards and training cannot be assured.

A possible compromise here is to use a 1D array, illuminating tissue with a line that can be scanned in one dimension, while the array is electronically scanned in another. Removing one

dimension of mechanical scanning improves image acquisition time over raster scan systems, while the reduced area of illumination allows lower laser power to achieve sufficient power density. This technique is used by the Moor Instruments Laser Doppler Line Scanner (LDLS), which uses a 40 mW laser with a commercial 64×1 pixel sensor and discrete electronics for signal conditioning, with digital processing performed by an FPGA [Hoang et al., 2010].

This is the architecture to be investigated in this thesis, as the combination of linear arrays and on-chip processing present a variety of possible IC designs. Only requiring one line of pixels means that the IC can include significant processing circuitry in space that would otherwise be filled with sensors. Combining the processing circuits into a 2D imager can mean that the design requires a large silicon area, and is therefore very expensive. Alternatively, the space available means that more than one type of front-end circuitry can be included on the IC for each pixel, allowing a choice of detector circuit for different applications, or allowing prototyping of new detector designs.

2.6.2 Location Marking on Patients

The LDBF systems shown have all focused on producing a flow map as an output. While this is clinically useful, if the flow map is to be used to direct clinical procedures such as veinipuncture, it is necessary to link points on the flow map to their physical location. An example of this would be marking a point for needle insertion.

The best way to do this would be to project the flow image onto the skin as done by Zeman [Zeman et al., 2005], giving a direct link between flow map features and their location. However, this is also the most complex option, requiring projection hardware, as well as a way of aligning the projected image with the imaged tissue such that the mapping is direct. This could be achieved by using references such as the edge of a limb, which could be taken using a visible light image. Alternatively, if visible light images can be taken, the system could record an image of the projected image, and use this feedback to adjust its position until alignment is achieved.

A simpler method is to identify a point on the skin by projecting a single point, arrow or cross-hair. While this adds an extra step for the user of selecting a point to be marked, the hardware is simpler, and therefore cheaper and more reliable. The projection could potentially use the same laser as used for the main imaging, which also makes verification

of the marker positioning using the main imager a possibility. A drawback of this technique is that it involves placing a needle, likely to be highly reflective, in the path of the laser. This risk could be reduced by lowering the laser power for this process, as this stage may be separate from the flow measurement itself.

An even more simple method may be to use a pen to mark a point on the tissue. The mark and the pen itself will be visible on the flow image (due to changes in reflectivity and an absence of flow), allowing a mark to remain after the imaging equipment has been removed.

2.6.3 Confidence in Results

Section 2.6.2 raises a potential barrier to the widespread adoption of LDBF systems in clinical use, that of user confidence. For all systems such as this, the final decision on diagnosis or treatment will be made by a clinician. It is therefore important that clinicians can have confidence in the results produced by automated measurement systems. This is part of the advantage of systems such as the veinlite, as this is a technological device that can be used as an aid to normal clinical practice, rather than requiring significant changes in techniques.

Development of clinical devices must therefore be undertaken in partnership with clinicians to ensure that devices produced are suitable in terms of useability and user acceptance. For example the blood flow/vein images produced by the 2D imagers (see Figures 2.7 and 2.2.5) show the location of veins very clearly. From this, it would be possible to develop algorithms to asses each vein and decide where canulation should occur. The marking method described in Section 2.6.2 could then automatically mark the selected location. A system such as the haptics system developed by Zivanovic (see Section 2.5.3) could even automatically insert the needle and draw blood or administer drugs. A system developed in this way would be unlikely to be easily accepted by clinicians or patients. This demonstrates the need for user confidence in developed systems, and for user input into the device development process.

2.7 Commercially Available Imagers and Figures of Merit

To be able to assess the performance of any system that is built during this work, the performance and typical figures of merit of available imager systems should be considered. Table 2.1 shows the specifications of two commercial LDBF imagers, each of which is a

scanning beam type. The main issue shown in this table is that neither imager gives quantified units of flow, instead giving arbitrary 'perfusion units'. This can be calibrated against a known motility target in both cases, although changes to the reflected light level mean that comparing readings between two applications could still be problematic. Part of the reason for not giving quantified flow is that there is no simple way to quantify flow in a complex system such as capillaries - flow would be measured as volume of blood moving through a given area of tissue over time, but given the random direction of blood flow and the non-uniform nature of the circulatory system, giving quantified values for flow is likely to be very difficult, and cannot be guaranteed to be be more clinically useful.

	Perimed PIM3	Moor LDI2-IR
Wavelength	670-690 nm	785 nm (+660 nm target beam)
Laser power (max)	1 mW	2.25 mW (+200 μW target beam)
Camera	CMOS 1280x1024	CCD 752 x 582
Scan area	50x50 cm (approx. max.)	50x50 cm (max, at 100 cm from target)
Scan time	4s (2.7x2.9cm, 10x10 points, 25cm from object) 4m29s (29x29cm, 85x85 points)	20 s (15x15 cm, 64x64 points) 5 m (50x50 cm, 256x256 points)
Measuring depth	0.5-1 mm, depending on tissue properties	not given - depends on tissue properties
Measuring units	Perfusion Units (arbitrary units)	Perfusion Units (arbitrary units)
Measurement resolution	Not given, but calibration is at no flow (0 $\pm$ 1 PU) and high flow (250 $\pm$ 15 PU)	0-5000 PU $\pm$ 10% relative to moorLDI2 standard

Table 2.1: Specifications for two commercially available LDBF imaging (scanning) systems - Perimed PIM3 [PerimedAB, 2011] and Moor LDI2-IR [MoorInstruments, 2007b]

For this reason, the specifications given for measurement accuracy refer to a range of arbitrary units, rather than quantified flow units, with the accuracy figure giving some indication of what level of flow change can be detected. The main differentiator between systems is then then the area that can be scanned, the spatial resolution of the flow image, and the time taken to produce an image.

2.8 Summary

The development of Laser Doppler Flowmetry for blood flow measurement has been introduced, along with the type of devices being developed using this technique. A range of alternative technologies for vein location has also been considered, including experimental

techniques and devices already being sold. A number of other factors affecting the use of LDBF in a clinical setting have also been considered, with a view to ensuring that developed instruments are suitable for widespread use.

Based on this review, a 1D scanning array would represent a new type of device, which could combine high imaging speed, compact size, safe laser power and high accuracy. The use of a fully integrated sensor in such a device will allow high performance while keeping overall system size and cost low. While alternative systems may be superior in some aspects, a line scanner represents a good compromise of all factors, allowing a vein location device to be produced that is suitable for clinical use.

Chapter 3

Investigation of Current to Voltage Converter Circuits and Processing Methods

3.1 Introduction

This chapter will describe circuits that can be used on an IC to detect laser Doppler flowmetry signals. The chapter addresses each element of the basic pixel structure shown in Figure 3.1 in turn. For each element, the design and operating principles of the circuit in question is given, simulations are used to show the performance of each circuit and compare alternative designs. Finally, design decisions taken on the basis of these simulations are described.

3.1.1 Pixel Elements for LDBF

Figure 3.1 shows the blocks used in a general LDBF pixel. The first part of this pixel is the photodiode, in which photons from the incident light generate electron-hole pairs, resulting in a photo-generated current proportional to the light intensity. While several types of photodiode can be implemented on a CMOS IC, the wavelength of light used here (visible red - near infra red, approximately 630-750 nm) means that only one available design is appropriate, and hence this is not discussed in depth here. This design uses the pn junction

formed between the n-well and the p-substrate of the IC to produce a relatively wide depletion region, hence increasing sensitivity to light with a longer wavelength of light.

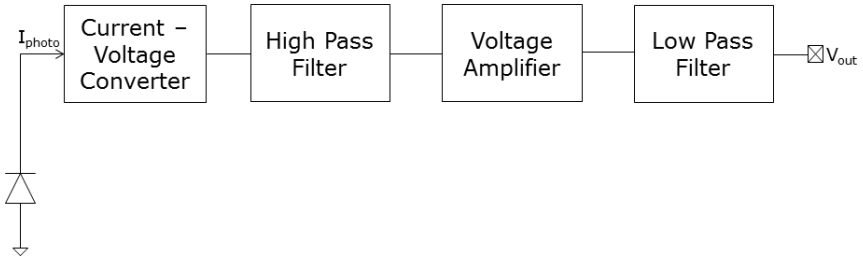

Figure 3.1: Main elements of the pixel

The current to voltage converter is a transimpedance amplifier, taking the photocurrent as an input and producing a voltage waveform as an output. While it is possible to perform signal conditioning on a current signal, it is more common for circuits to operate with voltage signal inputs and outputs.

The high pass filter is used to remove the DC and low frequencies from the input signal. The DC is removed at this stage to prevent amplification of the signal causing saturation, as the DC signal is generally higher than the modulated AC part. Removing AC signal components lower than the Doppler bandwidth also removes some or all of the 1/f noise in the signal, as well as removing movement artefacts, caused by relatively slow (e.g. < 5 Hz) movement of the target object relative to the imager.

The remaining AC signal is then amplified, increasing the amplitude of the frequencies of interest relative to other sources of noise in the system (read out noise, power supply noise), and allowing the use of lower-spec analogue to digital converters. This stage could be omitted if an ADC with sufficiently low noise and sufficient precision over the required voltage range was available, but this generally requires a fixed DC level which is not always possible.

Finally, the low-pass filter removes AC components above the Doppler bandwidth, reducing the overall noise level. This filter also acts as an anti-aliasing filter, removing signals above the nyquist frequency of the ADC.

3.1.2 Structure of the Chapter and Simulations Performed

The chapter consists of three main sections - current-voltage converters, amplifier/filter circuits, and low-pass filters. This is followed by a summary of the design choices made from

this chapter.

Section 3.2, looking at current-voltage converters, begins with logarithmic detectors. This section (3.2.1) includes a discussion of basic log pixels, followed by an introduction to two variations of the logarithmic pixel. Simulations are then performed in Section 3.2.4 on these three variants to show their various strengths and weaknesses, with the simulations including nominal DC and AC response, bandwidth over a range of photocurrents, and noise response including input- and output-referred noise.

A linear pixel design is then discussed in Section 3.2.5, including the operating principles of the circuit, its strengths and weaknesses for use on an IC, followed by simulations of a design that could be fabricated in an array on chip (nominal DC, AC, bandwidth, noise).

The final type of front-end circuit considered is the active pixel sensor, shown in Section 3.2.6 The method used to measure light levels used by these circuits is discussed, and their advantages and disadvantages are considered. Simulations are used to demonstrate why this type of front-end circuit is not used on this IC (using transient simulations to show the operating of the pixel during sampling).

Section 3.3 discusses the circuit used for the high-pass filter and amplifier elements of the pixel, both being variants of the 'hysteretic differentiator amplifier' (HDA). The two circuit designs are shown, and simulations are used to compare the two. The simulations include transient simulations to show the characteristics of the output signal of each design, followed by comparison of nominal DC response, AC/frequency response, and noise levels. To address concerns of manufacturing variations, monte-carlo and corner simulations are used to show the susceptibility of these circuits to these variations.

Section 3.4 introduces the circuit to be used as a low-pass/anti-aliasing filter. As the signal has been amplified at this point, the behaviour of this component of the pixel is less critical, and therefore only one design is considered here. The filter design is shown, and simulations of the nominal filter bandwidth, transient response (checking for distortion), DC and noise responses are performed to show that the filter is suitable for this application.

Finally, Section 3.5 summarises the design decisions taken due to the simulations shown in this Chapter.

3.1.3 Simulator and Device Models

All simulations shown in this thesis were performed using Cadence Design Systems IC5 or IC6, using the Virtuoso Spectre simulator, a spice based simulator. Transistor level device models were provided as part of the austriamicrosystems design kit for the c35 (0.35 μm) process used to fabricate the prototype ICs. Models include sections for typical mean conditions, as well as sections for Monte Carlo analysis and corner analysis.

3.2 Current to Voltage Converter Circuits

This section will discuss the various types of circuits used to convert the current from the photodiode into a voltage signal which can be amplified, filtered and sampled. The main categories that will be discussed here are logarithmic pixels, linear pixels based on opamps, and active pixels. The first two of these categories are 'continuous time' circuits, in that their output is a continuous voltage waveform proportional to the input current, while active pixels make a series of discrete samples of the photocurrent by measuring the rate at which the photodiode capacitance is discharged by the photocurrent.

3.2.1 Logarithmic Current to Voltage Converters

This section describes the structure and operating principles of logarithmic pixels. Advantages of the logarithmic pixel relevant to LDBF include provision of natural normalisation of the AC signal, and its smaller physical footprint compared to the standard linear operational amplifier pixel, which requires a compensation capacitor requiring a large silicon area [Allen and Holberg, 2002]. It also gives a continuous analogue output waveform, unlike active pixel CMOS sensors (also known as integrating pixels) discussed in Section 3.2.6 which give a series of output samples. This means the circuit can be used with a range of common continuous time signal conditioning circuits.

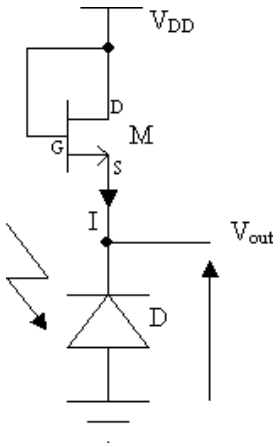

Figure 3.2: Basic Normalising logarithmic front-end pixel

Analysis of the logarithmic circuit will be carried out by first considering its DC operating point. It is this operating point that determines the AC behaviour, which requires a separate analysis based on the AC small signal characteristics of the MOSFETs used.

3.2.1.1 DC Operation of Logarithmic Pixels

In Figure 3.2 the NMOS transistor has its gate connected to the drain (labelled 'G' and 'D' in Figure 3.2) and acts simply as a resistive load for the photodiode. At the very low DC photocurrent levels typically seen in laser Doppler imaging (with small photodiodes), the MOS transistor operates below its normal saturation region and is in a region called "subthreshold" [Moini, 2000].

In this subthreshold region the DC current (I_{DC}) through the MOS transistor (equal to the DC photocurrent in the photodiode in Figure 3.2) is given by [Allen and Holberg, 2002]:

$$I_{DC} = I_0 e^{\left(\frac{qV_{GS}}{n'KT}\right)} \tag{3.1}$$

Where:

I_0 is related to the saturation current of the drain to substrate diode of the MOS transistor, found from the current where $V_{GS} = V_T$ (A)

n' is the subthreshold slope factor of the drain - substrate diode (~1.1 - 1.5)

V_{GS} is the DC voltage between the MOS gate and source terminals (V)

k is Boltzmann's constant

T is the temperature (^{o}K)

q is the electron charge, 1.6×10^{-19} C

For the above and following equations, upper case subscripts denote DC signals, lower case denotes AC signals.

From Figure 3.2 the DC voltages can be observed to be:

$$V_{GS} = V_{DS} = V_{DD} - V_{OUT}$$

Re-arranging:

$$V_{OUT} = V_{DD} - V_{GS}$$

Substituting from equation 3.1 this becomes:

$$V_{OUT} = V_{DD} - n'U_T ln(I_{DC}) + n'U_T ln(I_0) \tag{3.2}$$

where $U_T = \frac{kT}{q}$ is referred to as the thermal voltage

From equation 3.2 the DC output voltage of the pixel is shown to be related to the natural logarithm of the DC photocurrent, hence the term logarithmic pixel.

3.2.1.2 AC Operation of Logarithmic Pixels

In transistor circuits such as the logarithmic pixel shown in Figure 3.2, the DC operating point establishes the AC behaviour. The AC transimpedance (i.e. gain, R_{ac}) of the log pixel can be derived by differentiating equation 3.2 with respect to I_{DC}:

$$R_{ac} = \frac{dV_{OUT}}{dI_{DC}}$$

Hence:

$$R_{ac} = \frac{d\left(V_{DD} - n'U_T ln(I_{DC}) + n'U_T ln(I_0)\right)}{dI_{DC}}$$

resulting in:

$$R_{ac} = \frac{n'U_T}{I_{DC}} \tag{3.3}$$

So the AC transimpedance is proportional to the reciprocal of the DC photocurrent. Note that the transistor dimensions (gate width and length) are not present in these equations, as these are not dominant factors in device behaviour in the sub-threshold region. There is some impact on the sub-threshold slope factor, n', but this is not simple behaviour, being dependent on the device oxide capacitance and depletion layer capacitance which are affected by device geometry.

Assuming that the AC Doppler photocurrent (i_{ac}) is directly proportional to the DC photocurrent (I_{DC}) caused by the reflected light intensity [Belcaro et al., 1994] then:

$$i_{ac} = \frac{I_{DC}}{m} \tag{3.4}$$

where m is the Doppler ratio whose typical value in blood flowmetry ranges from 10 to 100.

The output AC voltage of the pixel is given as:

$$v_{ac} = i_{ac} \cdot R_{ac}$$

and if substitutions are made for i_{ac} from equation 3.4 and R_{ac} from equation 3.3 and respectively we obtain:

$$v_{ac} = \frac{I_{DC}}{m} \cdot \frac{n'U_T}{I_{DC}} = \frac{n'U_T}{m} \tag{3.5}$$

Equation 3.5 shows that the AC output voltage is independent of the DC photocurrent. Instead, it is determined by the subthreshold slope factor (n'), thermal voltage (U_T) and Doppler ratio (m). Therefore, for a given Doppler ratio, any fluctuations in the laser source power output (i.e. in I_{DC}) or variations in skin remittance (also causing a change of I_{DC}) will not affect the AC output voltage. In other words this logarithmic pixel performs normalisation in the detector, making it unnecessary to implement a separate normalisation function.

The behaviour described above can be demonstrated by simulation. The simulation plot in Figure 3.3, which shows how the DC output voltage and its derivative varies against DC photocurrent, illustrates the logarithmic response and the normalisation principle. The orange line is the DC output voltage of the pixel and follows Equation 3.2. The gradient of this plot, shown in orange, represents the AC transimpedance (R_{ac}) of the pixel. At low photocurrents the gradient of the DC response is steep. A small change in photocurrent

(such as that caused by the AC photocurrent) therefore causes a relatively large change in output voltage, giving a high AC transimpedance. As DC photocurrent rises, the gradient of the line becomes smaller. This means that an equal change in photocurrent causes a smaller change in output voltage, giving a low AC transimpedance.

The above analysis assumes that the MOSFETS remain in the sub-threshold region (i.e. equation 3.1 applies throughout). As the photocurrent rises the devices will enter the saturation region (as the diode connected layout means that at all times $V_{DS} = V_{GS}$, and therefore $V_{DS} > V_{GS} - V_T$), where the relationship between R_{ac} and I_{DC} changes and the normalisation no longer occurs as described here. The simulation performed for Figure 3.3 showed that the devices were operating in the subthreshold region over the full range of photocurrents used.

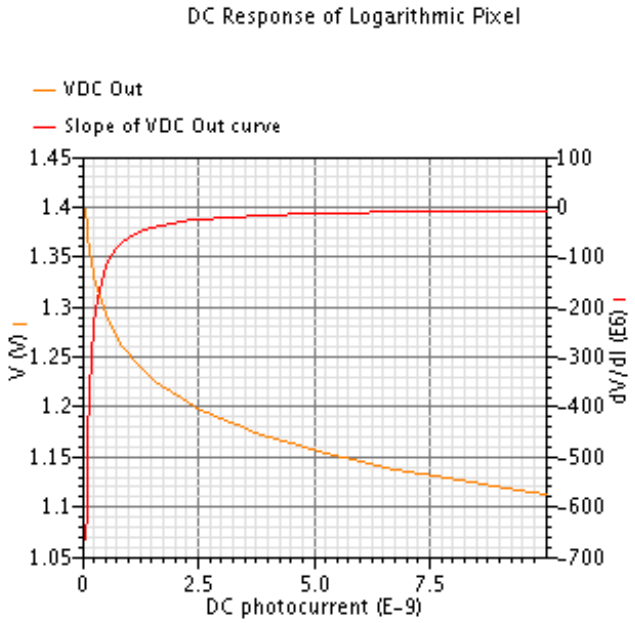

Figure 3.3: DC output voltage (left y-axis) and AC gain (right y-axis) plotted against DC photocurrent for log pixel

3.2.1.3 Bandwidth of Logarithmic Pixels

This section describes the bandwidth of the logarithmic pixel, including its dependence on DC photocurrent and hence light power. Modifications made to the basic design to improve bandwidth are shown, along with simulations to demonstrate both the improvement and the variation with I_{DC} of the modified design.

The basic log pixel was shown in Figure 3.2 and consists of a photodiode (D) and an NMOS transistor (M). The main aspects of this in terms of bandwidth are the AC resistance of the transistor and the capacitance of the photodiode. The latter assumes that the capacitance of the photodiode is much greater than the capacitance of the transistor - this is a valid assumption given the large photodiodes to be used here (which have to be large to collect sufficient light and to allow alignment of the imaging optics), and the small size of the transistors to be used in the front-end circuit (the compact size of which is a major advantage of its use in integrated detectors).

Knowing the AC resistance from equation 3.3 allows the bandwidth of the detector to be found:

$$f_c = \frac{1}{2\pi RC} = \frac{1}{2\pi \left(\frac{n' U_T}{I_{DC}}\right) C}$$

re-arranging:

$$f_c = \frac{I_{DC}}{2\pi n' U_T C} \tag{3.6}$$

The capacitance of the diode is a function of voltage across the diode, which here is equal to output voltage. For simplicity, the bandwidth will first be derived assuming a constant diode capacitance. The change in capacitance is small enough that this is a useful simplification to make even if it is not totally accurate. From equation 3.6 it can be seen that the lower half of the equation is constant, therefore:

$$f_c \propto I_{DC} \tag{3.7}$$

Hence it is important for all I_{DC} values to ensure that the bandwidth is above 20 kHz if the device is to be used for LDBF measurements. However, the diode capacitance is not constant with V_{DC} and has the following relationship:

$$C = \frac{WLC_J}{\left(1 + \frac{V}{PB}\right)^{MJ}} + \frac{2\left(W + L\right) C_{JSW}}{\left(1 + \frac{V}{PB}\right)^{MJSW}} \tag{3.8}$$

Where W,L are the diode dimensions ($50 \times 1000\,\mu$m), C_J is the junction area capacitance ($0.08\,\mathrm{fF}/\mu\mathrm{m}^2$), PB is the junction potential (0.53 V), MJ is the area junction grading coefficient (0.39), C_{JSW} is the junction side-wall capacitance ($0.51\,\mathrm{fF}/\mu$m) and $MJSW$ is the side-wall junction grading coefficient (0.27). This is taken from the process parameters of the $0.35\,\mu$m CMOS technology used by the work in this thesis [austriamicrosystems, 2007].

This change in C is relatively small over the DC voltage range of interest, but should still be included in the modelling. To consider the effect of the change in capacitance on bandwidth, the dependency of the photodiode voltage on photocurrent must be included. This was given in equation 3.2, and is dependent on DC photocurrent and transistor saturation current I_0. Equations 3.2, 3.6 and 3.8 can then be combined to describe more accurately the relationship between bandwidth and I_{DC}:

$$fc = \frac{I_{DC}}{2\pi n' U_T \left(\frac{WLC_J}{\left(1+\frac{V_{DIODE}}{PB}\right)^{MJ}} + \frac{2(W+L)C_{JSW}}{\left(1+\frac{V_{DIODE}}{PB}\right)^{MJSW}} \right)} \tag{3.9}$$

where:

$$V_{DIODE} = V_{DD} - n' U_T ln(I_{DC}) + n' U_T ln(I_0)$$

Rather than solving the above analytically, the relationship of capacitance against voltage can be displayed graphically after modelling using Matlab. This is shown in Figure 3.4, from which it can be seen that diode capacitance is inversely proportional to output voltage. Figure 3.5 shows the bandwidth increasing linearly with DC photocurrent. The pixel used for this is a basic pixel using a single PMOS load transistor of size $2\,\mu\text{m} \times 0.8\,\mu\text{m}$, and a photodiode of $1000\,\mu\text{m} \times 50\,\mu\text{m}$. This is the same size as those used on the BVIPS1 IC, although the IC uses a more advanced design of pixel as introduced in Section 3.2.2. From Figure 3.5 we can see that for typical light levels (up to 30 nA for this photodiode size), there are two key points:

- The bandwidth is linearly dependent on I_{DC}, despite variation of the diode capacitance.
- The bandwidth is below the required value in the lower range of operating photocurrents (below 6.5 nA)

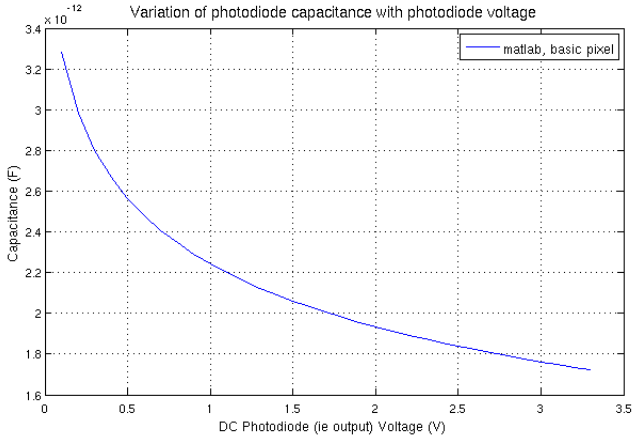

Figure 3.4: Variation of photodiode capacitance with DC output voltage from Matlab modelling

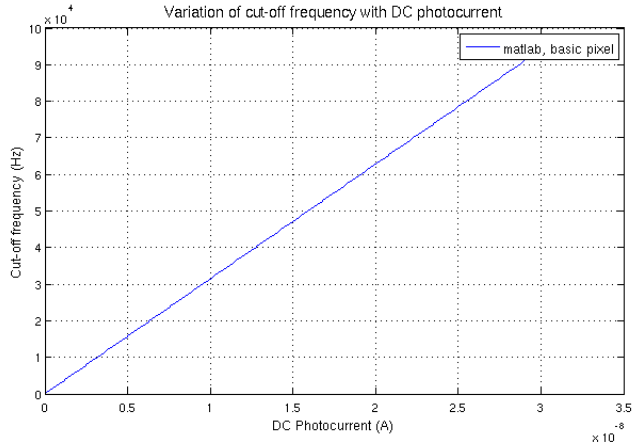

Figure 3.5: Variation of high frequency cut-off with photocurrent of basic log pixel from Matlab modelling

3.2.1.4 Noise in Logarithmic Pixels

This section describes the noise of a logarithmic pixel, including its dependence on DC photocurrent and hence light power. Theoretical noise equations are shown for the basic log pixel structure.

The main sources of noise in a trans-impedance amplifier are thermal and shot noise. $1/f$ noise should be removed by the high-pass filter used to remove the DC component of the signal and movement artefacts caused by Doppler shifts from the moving surface of the skin. However, this is a potential source of additional noise that will not be covered in hand calculations.

As a number of different circuits are used here, the noise values found are input referred. This means that voltage noise at the circuit output is divided by the transimpedance of the circuit, to give a noise current that would give the same output voltage in an ideal (noiseless) circuit. This allows noise in circuits with different gains to be compared.

The total input noise current is found by adding the thermal and shot noise [Horowitz and Hill, 1989]:

$$i_{noise} = \sqrt{i_{shot}^2 + i_{th}^2}$$

$$i_{noise} = \sqrt{i_{shot}^2 + \left(\frac{v_{th}}{R}\right)^2}$$

$$i_{noise} = \sqrt{2qI_{DC} + \left(\frac{\sqrt{4kTR}}{R}\right)^2}$$

$$i_{noise} = \sqrt{2qI_{DC} + \frac{4kT}{R}} \tag{3.10}$$

For the sub-threshold MOSFET the trans-impedance is given by:

$$R_{ac} = \frac{n'U_T}{I_{DC}} \tag{3.11}$$

Hence the input referred noise current is given by:

$$i_{noise} = \sqrt{2qI_{DC} + \frac{4kTI_{DC}}{n'U_T}}$$

which can be simplified to:

$$i_{noise} = \sqrt{\left(2 + \frac{4}{n'}\right) qI_{DC}} \tag{3.12}$$

This shows how the input referred current noise increases with the square root of I_{DC}. However, the voltage noise at the output depends on the trans impedance of the circuit, R_{ac}, shown in equation 3.11 and is therefore:

$$v_{noise} = i_{noise} \cdot R_{ac}$$

Substituting for R_{ac} and i_{noise} gives:

$$v_{noise} = \sqrt{\left(2 + \frac{4}{n'}\right) qI_{DC}} \cdot \frac{n'U_T}{I_{DC}}$$

Rearranging:

$$v_{noise} = \sqrt{\left(2 + \frac{4}{n'}\right) qI_{DC} \frac{n'^2k^2T^2}{q^2I_{DC}^2}}$$

$$v_{noise} = \sqrt{(2n'^2 + 4n')\,\frac{k^2T^2}{qI_{DC}}}$$

$$v_{noise} = \sqrt{(n'^2 + 2n')\,\frac{2U_T^2 q}{I_{DC}}} \tag{3.13}$$

Equation 3.13 shows that while the input referred noise current rises with the square root of I_{DC}, the noise voltage at the output is inversely proportional to the square root of I_{DC}. Given the normalisation performed by the log pixel, the AC signal voltage at the output is constant with I_{DC}. Hence the SNR increases in proportion to the square root of I_{DC}. By contrast, a linear transimpedance amplifier with constant AC resistance would be expected to have a directly proportional relationship between output voltage noise and input referred current noise, so the SNR would fall with increasing I_{DC}.

3.2.2 Buffered Pixel Design for Increased Bandwidth

3.2.2.1 Operating Principle of the Buffered Logarithmic Pixel Design

The Matlab modelling in Section 3.2.1.3 shows that the bandwidth of a basic log pixel with a $50 \times 1000\,\mu$m photodiode is lower than the 20 kHz considered necessary for Doppler flowmetry. Consequently, a modified design is used, as shown in Figure 3.6. This also shows a source follower buffer (to the right of the dotted line) added to the output of the pixel, ensuring that large external capacitances do not reduce the bandwidth of the circuit. The circuit also uses two diode connected transistors (MP1 and MP2), instead of the original single transistor. This increases the AC gain, but reduces the bandwidth (the two transistors in series results in higher R_{ac} due to higher n' in equation 3.11 [Moini, 2000]).

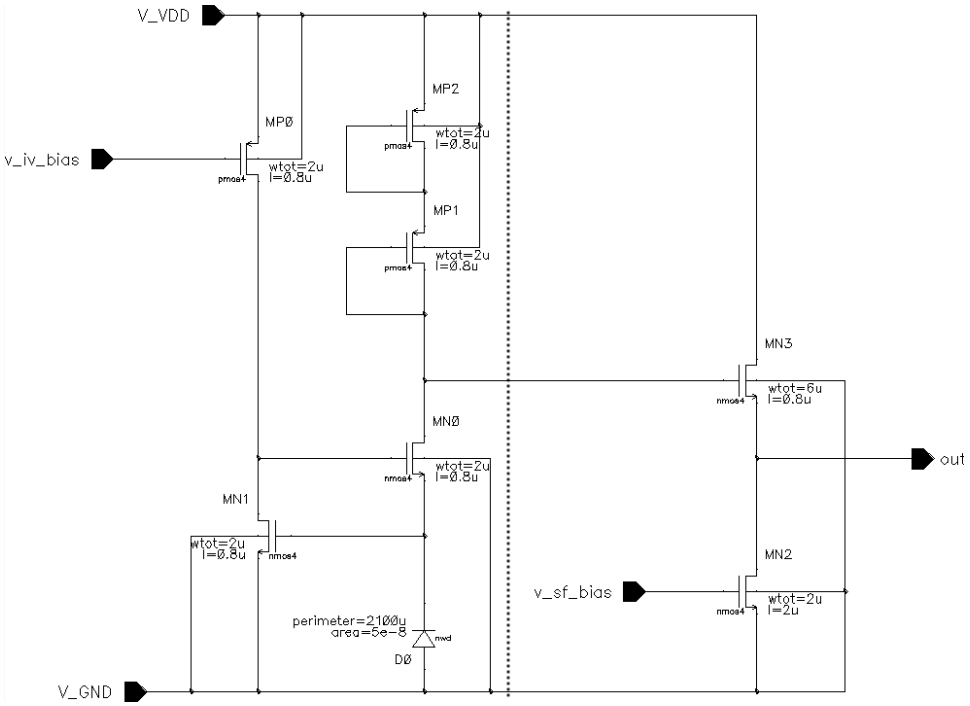

Figure 3.6: Schematic of buffered logarithmic pixel - detector is to the left of the dotted line, right section is a source-follower buffer with unity gain)

The main difference in this circuit is the addition of transistors (MN0 and MN1) between the diode and the output. The two transistors to the right of the main circuit (MN2 and MN3) form a source-follower buffer to allow large output loads to be driven. The circuit also uses two load transistors (MP1 and MP2) rather than one, giving a higher load resistance and hence increased transimpedance (i.e. gain). This change increases the RC time constant, so would by itself be expected to lead to a drop in bandwidth. Any improvement seen in bandwidth is therefore due to the photodiode buffering caused by MN0 and MN1, although the circuit will be simulated with either one or two load transistors to observe the effect of this change.

Transistors MN0 and MN1 buffer the photodiode capacitance from the two diode connected transistors (MP1 and MP2), hence buffering the large diode capacitance from the voltage swing at the output of the circuit. The buffering occurs as a rise in photocurrent causes the diode voltage to fall, as the current through the load transistors increases. This causes MN1 to turn more off (higher resistance), which increases the source-drain voltage of MN1 due to the constant current set by the current source MP0 – biased by an external circuit. This causes the gate voltage of MN0 to increase, which turns more on (decrease in resist-

ance). This decreases the drain-source voltage of MN0, the current through which cannot increase as this is the photocurrent. This means that a small change in diode voltage results in a large change in voltage across MN0, which is what causes the output voltage swing [Kongsavatsak et al., 2008, Gu et al., 2008, Dmochowski et al., 2004].

3.2.2.2 Simulation of Bandwidth of Buffered Logarithmic Pixels

The bandwidth of this modified circuit can be compared with the basic pixel design by simulation, using Cadence Design Systems software, a PSPICE based simulator. This method is used as the more complex circuit cannot be easily modelled in Matlab. The basic pixel is also simulated in Matlab to show the difference between the basic Matlab model used previously, and the more thorough device models and simulation methods used in Cadence. Both circuits are also simulated with one or two load transistors in series with the photodiode, showing the effect on bandwidth of increasing the transimpedance by increasing the AC load resistance.

Figure 3.7 shows the bandwidth of each circuit with increasing DC photocurrent as found by simulation in Cadence, along with the original results from Matlab modelling of the basic pixel. It can be seen that the bandwidth of the buffered pixel is much higher for a given photocurrent. For the single PMOS load variants, the buffered pixel has a bandwidth of 280 kHz at 2.5 nA I_{DC}, where the basic pixel has a bandwidth below 10 kHz. The un-buffered pixel with a single transistor load requires a DC photocurrent of around 10 nA, which cannot be guaranteed for this application - this could lead to variations in flux readings due to changing light levels, and hence changing frequency response, rather than an actual change of flow.

The additional load transistor causes a decrease in bandwidth for both pixel variations. This drop is greater for the basic pixel where the high frequency cut-off is halved with a second load transistor, while for the buffered design the drop is approximately 20%.

There is significant variation between the Cadence and Matlab modelled results for the basic pixel. While this change could be cause for concern, the more advanced simulation performed by Cadence would be expected to give different results, as it includes the effect of all circuit components rather than just diode capacitance, more accurate modelling of the DC operating point and AC characteristics, and more complete transistor models. The results here are sufficient to demonstrate the operating principles of main advantages and disadvantages of

these circuits, and are not used to make key design decisions, so this discrepancy is not a cause for concern.

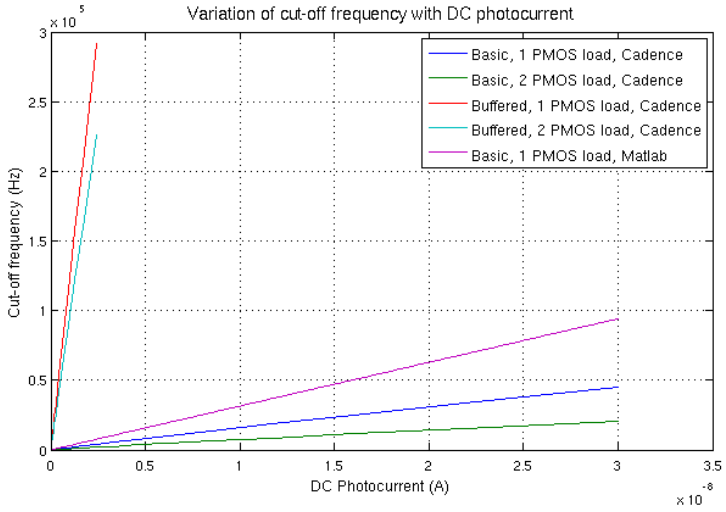

Figure 3.7: Variation of high frequency cut-off with DC photocurrent of buffered log pixel, simulated using Cadence

3.2.2.3 Simulation of Noise of Buffered Log Pixel

As Section 3.2.1.4 assumes the pixel is a simple design using one diode connected transistor in series with the photodiode, simulations must be performed to determine the noise performance of the more complex buffered pixel design described in Section 3.2.2. The buffering is used to increase bandwidth, but should not introduce large amounts of extra noise. The simulations do not show flat noise density over the frequency spectrum, so the noise densities shown are calculated as an average over 50 Hz−30 kHz (50 Hz being a conservative low cut-off frequency of the pixel amplifiers, and 30 kHz being based on the noise bandwidth ($1/4RC$ instead of $1/2\pi RC$) of a system with 20 kHz signal bandwidth). Figure 3.8 and Figure 3.9 compare the noise levels found by Cadence simulation of the various pixel types, along with the theoretical values found from Matlab modelling of equations 3.12 and 3.13.

For input referred current noise (i.e. voltage noise at the output divided by circuit transimpedance, allowing direct comparison of circuits with different gains) in Figure 3.8, the buffered pixel design does have an increase in noise over the basic pixel, but the increase is limited, with a much larger increase seen when going from two to one load transistors on either variant. This decrease in noise with more components seems counter-intuitive, but the increased transimpedance from the additional load means that an equal noise level at

the output corresponds to a lower input-referred current noise, and as the change does not affect DC current, the increase in noise is limited (although there is some increase due to n' in equations 3.12 and 3.13). This effect is shown by the voltage noise results in Figure 3.9, where the noise level increases with the additional load transistor for both variants. These results do show an increased noise level for the buffered pixel, particularly at lower photocurrents. However, as the input referred simulation is better suited to direct comparison of different circuits, these simulations do not show a major drawback of the buffered pixel design. The lower noise levels seen in the basic pixels at low frequencies are partly due to the restricted bandwidth of these circuits in this region, as this effectively filters off noise at higher frequencies.

The simulated results show lower noise levels than the Matlab modelling of basic noise sources. This suggests that the basic model is inaccurate in terms of DC operating point or transimpedance used in these equations, however the basic shape of the response is similar, and the the Cadence simulations would be expected to produce more accurate results due to the more complete circuit and device models.

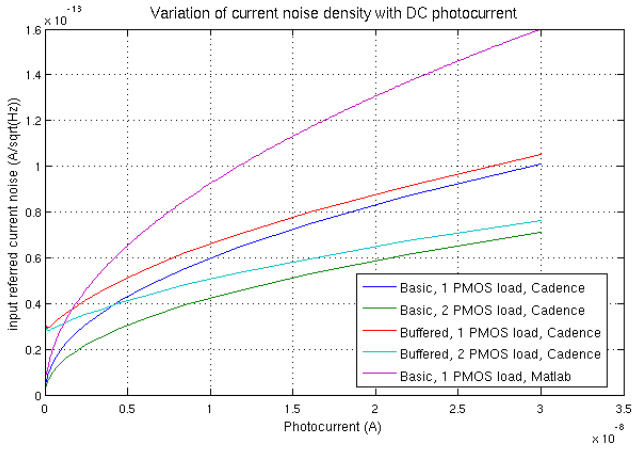

Figure 3.8: Increase in input referred current noise as DC photocurrent increases for basic logarithmic pixel (from Matlab model) and buffered logarithmic pixel (from Cadence simulation)

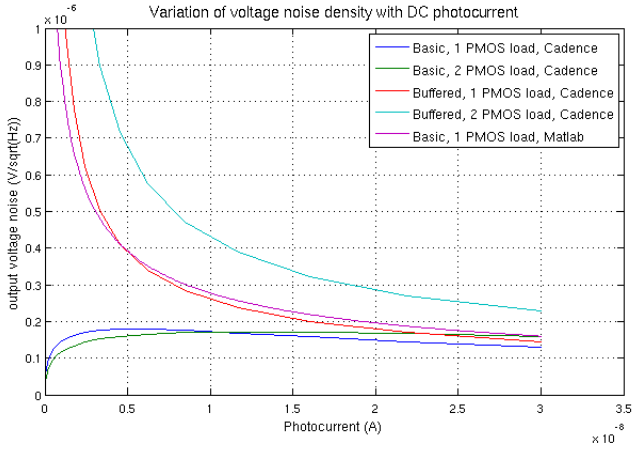

Figure 3.9: Increase in output voltage noise as DC photocurrent increases for basic logarithmic pixel (from Matlab model) and buffered logarithmic pixel (from Cadence simulation)

The Cadence simulations performed on the buffered pixel allow a noise spectrum to be found for input and output noise, allowing more information about noise in this circuit. Figure 3.10 shows the input referred current noise spectrum of the buffered pixel at DC photocurrents of 100 pA, 1 nA, 5 nA and 10 nA. As expected, the noise density at all frequencies rises with DC photocurrent. However, the increase seen is much greater at lower frequencies, suggesting the 1/f noise components are linked to DC photocurrent. The input referred noise density rises at higher frequencies, as the fall in transimpedance at frequencies approaching the high frequency cut-off means that a constant output voltage noise density is equivalent to an increase in input referred noise. The lower bandwidth caused by lower photocurrent (see Section 3.2.1.3) means that this effect is more pronounced at lower DC photocurrents, where the increase in noise density starts at lower frequencies, and causes a greater overall increase.

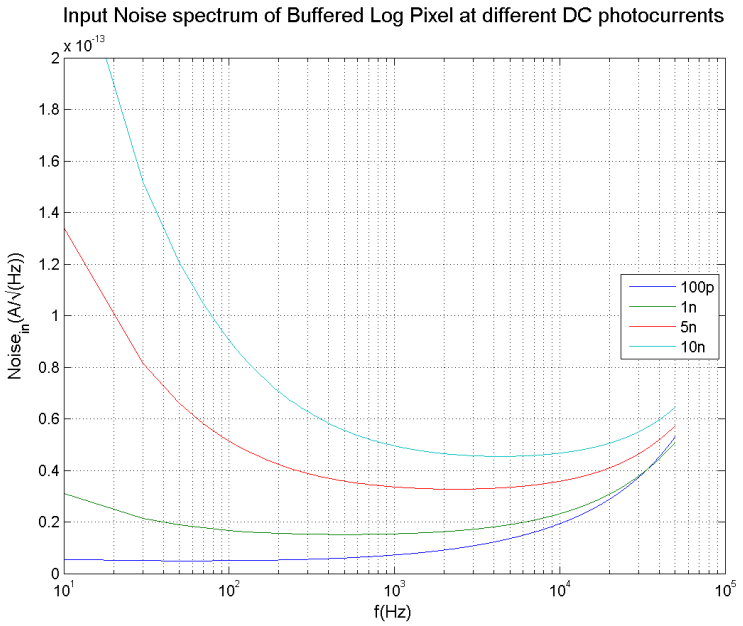

Figure 3.10: Input referred current noise spectrum at different DC photocurrent values, buffered logarithmic pixel (from Cadence simulation)

The output voltage noise spectrum at the four DC photocurrent values is shown in Figure 3.11. This shows the expected fall in noise seen at the logarithmic pixel output as DC photocurrent rises. All spectra show an increase in noise at low frequencies due to 1/f noise, and again this increase as a proportion of the noise density seen in the pass band increases at higher photocurrents.

An unexpected result is that the noise density increases at higher frequencies. This is expected in the input referred case, due to the change in transimpedance at higher frequencies. The noise spectrum at the 100 pA, where the high frequency cut-off is lowest, shows a peak in noise density which then falls off, rather than a continuing increase. It is possible that this effect is occurring in all cases, but the peak of the noise spectrum is above the frequency range simulated. This peak may be caused by a change in the noise contributions of individual elements of the circuit approaching the high frequency cut-off.

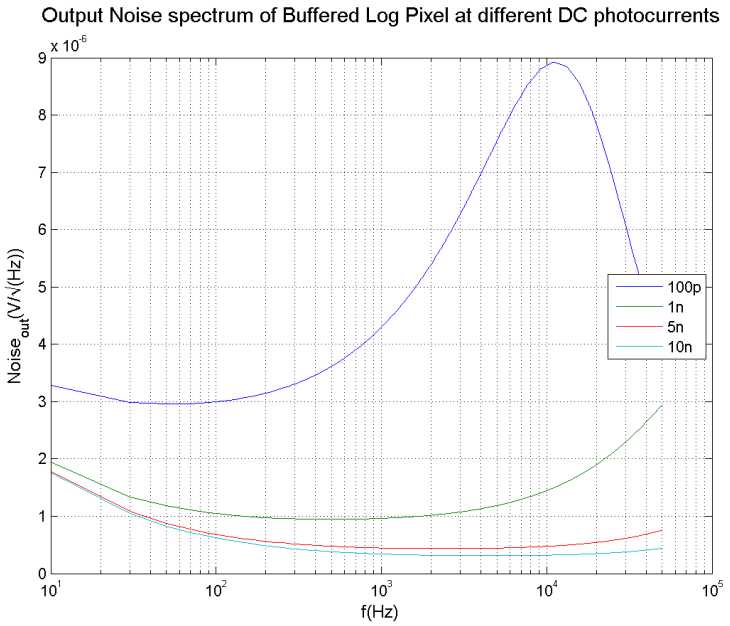

Figure 3.11: Output voltage noise spectrum at different DC photocurrent values, buffered logarithmic pixel (from Cadence simulation)

3.2.2.4 Summary of Buffered Pixel Design

The limited bandwidth of the basic pixel when using a large photodiode was shown in Section 3.2.1.3. The bandwidth is therefore increased by the use of a buffered front-end circuit. Simulations performed on this circuit show that it has the same linear relationship between bandwidth and photocurrent, but with a considerably higher bandwidth (by a factor of ~ 450) for any given photocurrent.

The noise simulations shown here demonstrate the relationship between noise and DC photocurrent (i.e. light level) in a logarithmic pixel. Noise current (input referred) increases proportionally to the square root of photocurrent. Voltage noise at the output is inversely proportional to the square root of DC photocurrent. Simulations were performed on the buffered log pixel that may be used in the BVIPS1 IC. The simulation results were found to follow the same trends as the theory, with a slight reduction in input referred current noise due to the increased trans-impedance of this circuit. All these results are for theoretical or simulated noise, actual noise may be higher due to other aspects such as power supply noise. This will depend on circuit board design (power supply de-coupling/voltage regulation etc), and on the PSRR (power supply rejection ratio) of the design.

3.2.3 Current to Voltage Converter Circuit using CMOS inverter feedback

Another option for a front-end using buffering to increase the bandwidth is shown in Figure 3.12 [Moini, 2000, Johnston et al., 2009]. This uses a CMOS inverter to provide feedback to an NMOS device (MN1) which provides the load impedance. MN0 and MP0 form a high gain amplifier, such that the voltage swing at the I-V output is larger than that at the photodiode cathode. The feedback to the NMOS device ensures that the inverter is always biased in the linear region, rather than in either fully-on or fully-off states as in a standard digital inverter. As photocurrent rises, the voltage at the diode cathode is pulled down as the current though MN1 increases. The drop in voltage turns MN0 more off and MP0 more on, increasing the voltage at the inverter output. This increase in voltage turns MN1 more on, and pulls the voltage at the diode cathode up, reducing the voltage swing in a similar way to the buffered front-end circuit. MN2 and MN3 again form a source follower buffer to drive the next stage.

This circuit provides a potential alternative to the buffered logarithmic pixel presented above. Simulations will be performed on both pixel designs in Section 3.2.4 to determine the superior circuit in terms of noise and gain. The gain of this circuit is likely to be lower, as the feedback arrangement does not allow the use of multiple load transistors in series, although the load transistor is not diode connected so it's AC resistance is harder to predict.

Another drawback is that the feedback arrangement keeps the inverter biased in its transition region (operating as an analogue component), rather than the on/off states of an inverter when used as a digital component, and hence its power consumption is relatively high. Simulation of this circuit shows a DC current through the inverter of ~50 μA, compared to ~11 μA to bias the buffering in the buffered log pixel (bias current for MN1 in Figure 3.6, Section 3.2.2). The basic logarithmic pixel does not have any additional circuitry to increase current consumption, so the current in this circuit is limited to the photocurrent only. This limit also applies to the sections in series with the photodiode in the buffered and inverter-feedback logarithmic pixels. For a 64 pixel array, the 50 μA current in each pixel results in a 3.2 mA current consumption, so while this is higher than for the other circuits, this is not an unreasonable level of current consumption. This represents a further advantage of a 1D detector array, as the lower pixel count means that the maximum current consumption per pixel can be higher, and if the current consumption is large, designing a suitable power

supply network onto the IC is a simpler task than if power had to be routed to a 2D pixel array.

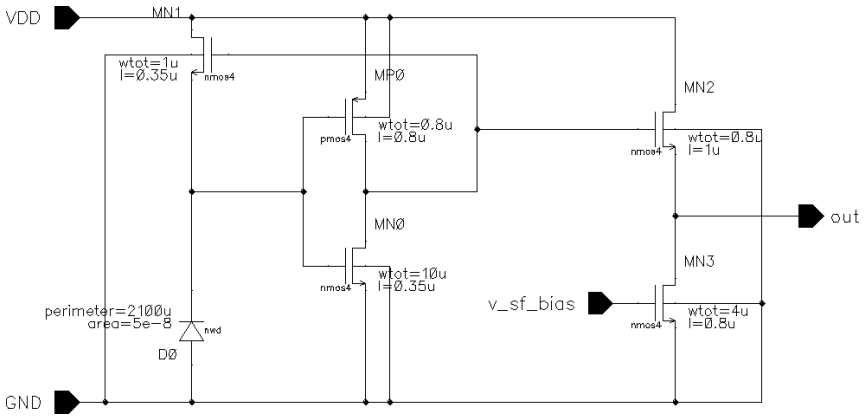

Figure 3.12: Inverter feedback front-end circuit

3.2.4 Simulation of Logarithmic Front-End Circuits

As the logarithmic front-ends using feedback or buffering are fairly complex, analysis and comparison by hand is very difficult. This section shows results of simulations performed in Cadence of the diode connected front-end with no feedback, the CMOS inverter front-end and the cascode front-end circuit, showing why the buffered logarithmic front-end circuit was chosen for this design. Simulations are performed over the range of currents expected (see Section 1.8.1), DC photocurrent of $75\,\text{pA} - 1.5\,\text{nA}$, Doppler ratio of 10%) with a $1000 \times 50\,\mu\text{m}$ photodiode.

3.2.4.1 DC Simulation

Figure 3.13 shows DC simulation results from the diode connected front-end with no feedback, the CMOS inverter front-end and the cascode front-end circuit. It can be seen from the straight line plots on a logarithmic scale that all front-end circuits give a log response, showing that the feedback front-ends give a similar response to the non-buffered version analysed in Section 3.2.1. The buffered feedback circuit is similar to the un-buffered design, although the DC level is offset by roughly $-0.1\,\text{V}$. The CMOS inverter feedback gives a different response due to the inverting buffer, which results in a rising DC voltage as the DC photocurrent rises. The DC level is also considerably lower at around 350 mV. This can be an advantage, as the same AC signal on a smaller DC level means that an analogue-to-digital converter used to sample the signal would require lower dynamic range. For example, if the input range was chosen to be $0 - 500\,\text{mV}$ and with a 14-bit (16384 level) ADC, the accuracy would be $30.5\,\mu\text{V}$. To achieve this resolution with a signal at a higher DC level, say with the range set as $0 - 1.5\,\text{V}$, would require 49152 steps, which would require a 16-bit (65536 level) ADC. This assumes that the ADC range starts at 0 V - if the lower limit of the output voltage (including AC and DC) is known, the ADC range could be offset, for example covering 1-1.5 V, such that the ADC resolution requirement would not depend on DC level. This is not difficult to implement, but can be complicated by the need to set a lower voltage limit, which could mean that in unexpected situations (e.g. lower/higher DC photocurrent than the range expected, or higher AC photocurrent than expected) the ADC saturates while the analogue sections of the circuit are still working correctly.

The lower DC output voltage of the inverted feedback pixel could cause clipping of the AC output signal at voltages near 0 V when a large AC signal is present. For the inverter feedback

pixel, this will occur with a 0.7 V peak-peak signal for the inverted feedback circuit (assuming a separate AC amplifier is used after the front-end, otherwise 0.7 V p-p is far higher than expected), whereas a 2.4 V peak-peak signal could be present on the DC level given by the other circuits without causing clipping at the 0 V or 3.3 V supply rails.

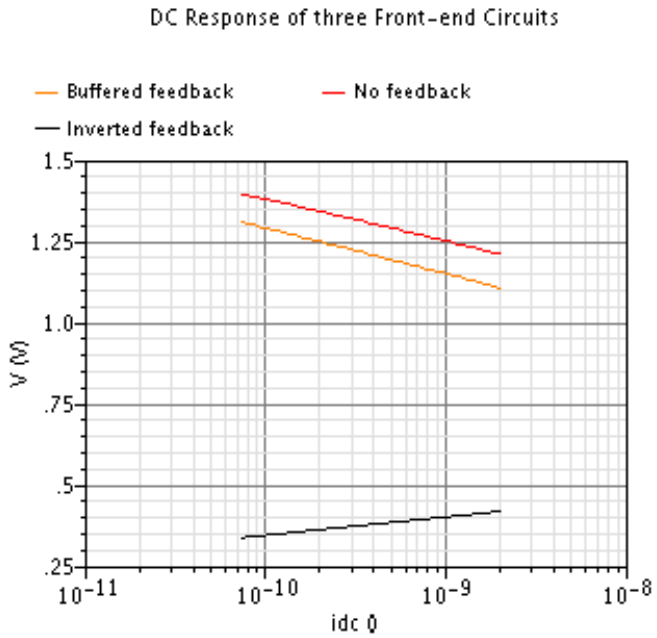

Figure 3.13: DC simulation results for different front-end designs

3.2.4.2 AC Simulation

Figures 3.14 and 3.15 show AC frequency sweep simulation results for the three types of pixels at 765 pA DC photocurrent. Figure 3.14 shows the output magnitude, while Figure 3.15 shows the transimpedance These plots show the output voltage magnitude and transimpedance for each circuit in the pass band (the flat part of the frequency response), as well as the bandwidth (where the roll-off at higher frequencies reaches -3dB). The units used are dBs relative to 1 V. The simulations were performed for a typical DC photocurrent of 765 pA. It can be seen that for the same input current the cascode feedback front-end gives the largest output signal, with the highest trans-impedance of the three. The non-buffered front-end is similar at low frequencies to the buffered IV. This is to be expected as the two circuits use a similar load - two diode connected PMOS transistors. However, the buffered front end has the extra pair of feedback transistors which, as well as buffering the diode voltage to increase bandwidth, also provide some extra gain, as the source-drain AC

resistance of one of the buffering transistors (MN0 in Figure 3.6) is in series with the main load transistors. The CMOS inverter front-end has considerably lower AC gain as well as the lower DC level shown in Figure 3.13. The AC output signal for the cascode front-end is roughly 2.5 times that of the inverter feedback front-end, compared to a DC level roughly 3.5 times larger. The inverter feedback front-end therefore has a larger ratio of AC:DC signal, which can have advantages in terms of dynamic range of the signal when sampled with an ADC (previously discussed in Section 3.2.4.1).

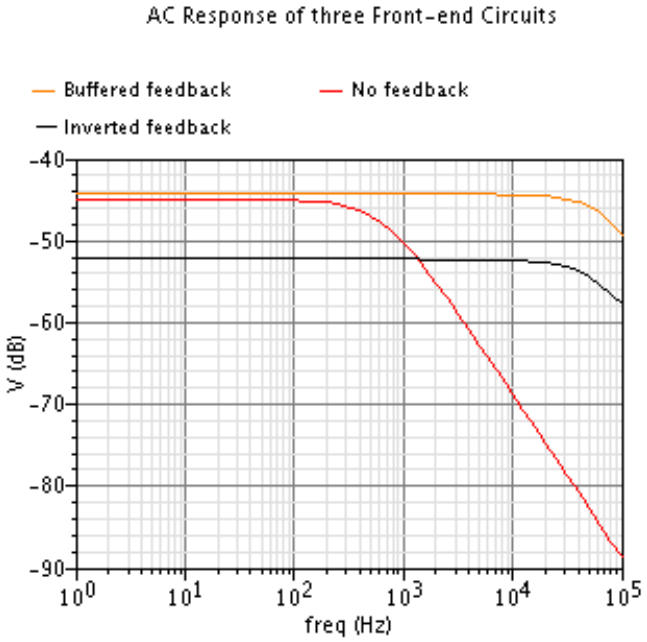

Figure 3.14: AC simulation results for different logarithmic front-end designs, showing output voltage AC magnitude at 765 pA DC photocurrent

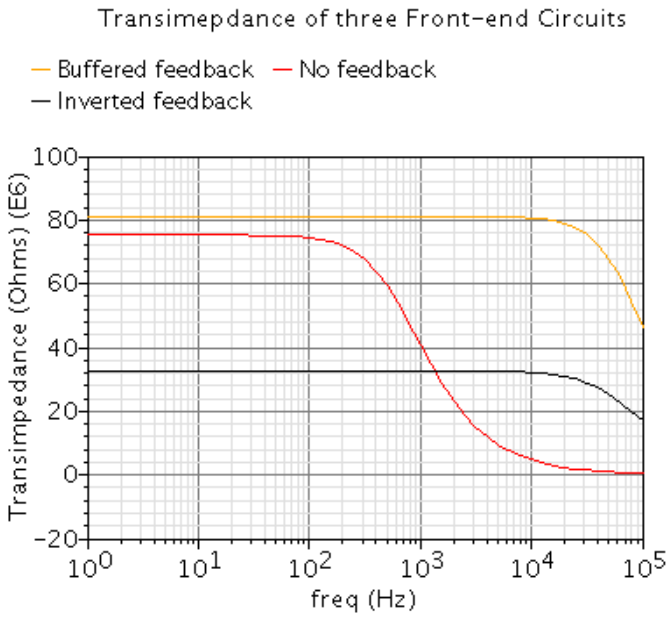

Figure 3.15: AC simulation results for different logarithmic front-end designs, showing transimpedance of each front-end at 765 pA DC photocurrent

Figure 3.14 also shows that as expected the no feedback front-end has considerably lower bandwidth than the pixels using feedback. Of the front-ends using feedback the buffered feedback front-end bandwidth is marginally higher, although at the typical photocurrent used here both front-ends have sufficient bandwidth. However, for all pixels the bandwidth is dependent on DC photocurrent. The logarithmic response of the pixels means that at higher DC currents the AC resistance is lower. This results in a lower RC time constant (and hence increased bandwidth) for higher photocurrents. Figure 3.16 shows this variation of bandwidth with DC photocurrent.

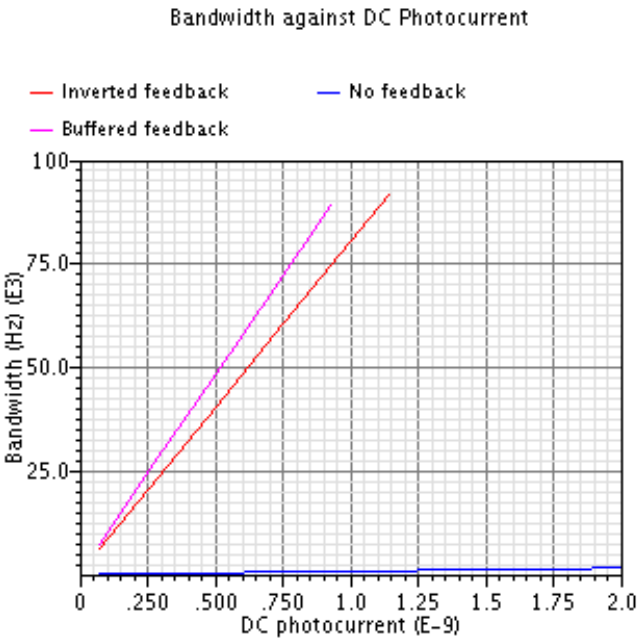

Figure 3.16: AC simulation results showing bandwidth against DC current for different front-end designs

This shows that all three front-ends have an increasing bandwidth with increasing DC photocurrent. The buffered and inverted feedback front-ends have a similar slope on the bandwidth-DC current graph, suggesting that these two front-ends will have a similar frequency response at all photocurrents, although the buffered front-end has slightly higher bandwidth at all DC photocurrents. The front-end without feedback is clearly not a suitable design for this circuit due to its inadequate bandwidth even at high photocurrents. For all front-ends the bandwidth drops below 20 kHz (as required based on estimate in Section 1.8.3) at the lower DC photocurrent limit, although here the buffered front-end has an advantage due to its higher bandwidth- At 75 pA the buffered front-end has a bandwidth of roughly 8 kHz. While this is lower than the specification, it still covers the majority of the Doppler bandwidth. The bandwidth of the buffered front-end rises to 20 kHz at around 200 pA. For the inverter front-end, the bandwidth at 75 pA is roughly 5.5 kHz, with the bandwidth rising to 20 kHz at around 350 pA.

3.2.4.3 Noise Simulation

Given the low typical AC signal levels that are to be measured by this system, it is important that the front-end has very low noise. As the Doppler signal is a continuous range of fre-

quencies it is not possible to remove large amounts of noise by filtering or lock-in techniques. The signal can be averaged to reduce noise, but this obviously reduces the system frame rate as the system is effectively sampling each pixel several times. If FFT processing is used this effect is multiplied by the number of points used in the FFT. The noise simulations shown here give input referred current noise. This considers the total voltage noise at the output and the trans-impedance gain of the front-end, giving a measure of noise which allows direct comparison of different circuits. The noise current values given can be considered as the size of input current that would have to be applied to a noiseless circuit to give a signal at the output equal to the output noise found by simulation.

Figure 3.17 shows input referred noise spectra for all three front-end circuits, while Figure 3.18 shows corresponding output noise spectra. All three input referred spectra have common characteristics, such as the increase in noise density at higher frequencies. This is partly caused by the fall in transimpedance at these frequencies - as the transimpedance falls, an equal level of noise at the output corresponds to a higher input referred noise. The output noise spectra in Figure 3.17 show that there is no major increase in noise density for the un-buffered or inverter feedback pixels, and therefore for these circuits the increase is due to the fall in transimpedance. However, the buffered pixel shows an increase in output noise at high frequencies, although this appears to be a peak rather than a continuous increase, as the noise density begins to fall at around 70 kHz. This suggests that the buffering added to this circuit causes an increase in noise at higher frequencies. This could be a significant disadvantage given the frequency weighting performed as part of LDBF processing, but this is mitigated by the use of a high-pass filter to attenuate signals above 20 kHz, removing the majority of the high frequency noise caused by the buffered front-end.

The noise also rises at lower frequencies, which is partly due to a 1/f noise and partly due to the performance of the circuit at frequencies approaching DC. The circuits used here do not include high-pass filtering, so the effect seen at the high frequency cut-off (where the input-referred noise rises due to falling transimpedance caused by the high-pass filter) is not duplicated at low frequencies.

The noise spectra show that the lowest noise is achieved with the no feedback front-end. This is to be expected as this is the simplest front end, with fewer devices contributing noise to the system. Of the feedback front-ends, the buffered front-end has the lowest noise at lower frequencies, but the spectrum rises more rapidly than the other circuits at higher frequencies.

From around 9 kHz the inverter front-end has lowest noise. This is difficult to balance, as while most of the Doppler signal is at lower frequencies, the higher frequencies are given more significant weighting for flux calculations. As well as this, the noise spectrum does not show which front-end has the highest total noise within the signal bandwidth. Figure 3.17 does not show the effect of changing DC photocurrent on the noise spectrum. Noise is expected to rise with I_{DC}, but depending on the sources of noise this may affect different circuits in different ways. Because of this, simulations were also done to show the effect of I_{DC} on total noise within the signal bandwidth.

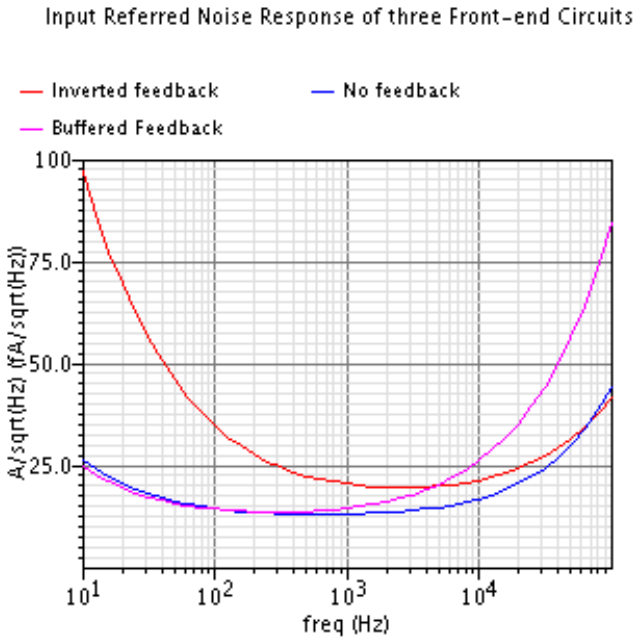

Figure 3.17: Input referred noise spectrum at typical photocurrent for different front-end designs

Figure 3.19 shows the total noise within the noise bandwidth of $100-30$ kHz as the DC photocurrent increases (30 kHz noise bandwidth of $1/4RC$ instead of 20 kHz signal bandwidth of $1/2\pi RC$). This shows that noise for the buffered front-end is higher than for the inverter front-end at lower photocurrents, but it does not rise as quickly with increasing photocurrent as noise in the inverter front-end. Again this gives a crossover part way through the range of signals that may be encountered, making it unclear which circuit is most suited to this application. It should be considered that as DC photocurrent rises the AC signal photocurrent will also rise, given the fixed Doppler ratio of 10%. Over the range here the AC signal may be expected to increase from $7.5-150$ pA peak-peak, or $1.14-22.7$ pA RMS (based

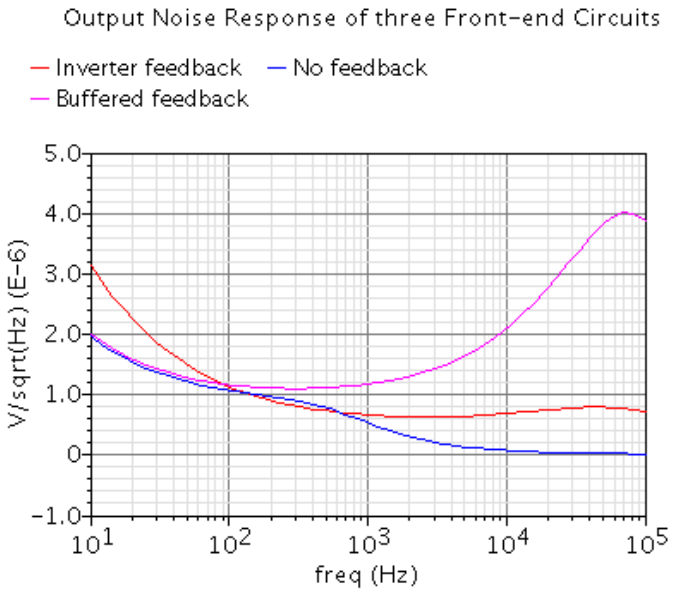

Figure 3.18: Output noise spectrum at typical photocurrent for different front-end designs

on peak-peak = 6.6×RMS value for a Gaussian signal [Jung, 2005]). This means that signal current rises faster than noise current, suggesting that at higher DC currents the SNR will increase considerably regardless of small differences between front-ends. At lower currents, however, it is vital to get the best possible noise performance - as the estimated minimum signal above is below the noise levels shown in Figure 3.19. This would make the inverter front-end the most suitable choice.

However, even with this front-end the SNR at the minimum DC photocurrent is around or slightly lower than 1. This is a considerable problem, and if the calculated light powers and simulated noise levels are correct this would require an increase in laser power or a change to the frequency limits to make the SNR acceptable (i.e. higher laser power to increase DC photocurrent, or narrower pass-band to filter off noise at higher frequencies where there is only a small part of the Doppler signal). Additional filtering would block some signal frequencies, so an increase in laser power is preferred. This would move the range of operating currents upwards, which would mean that the buffered front-end would have lower noise over more of the expected signal range.

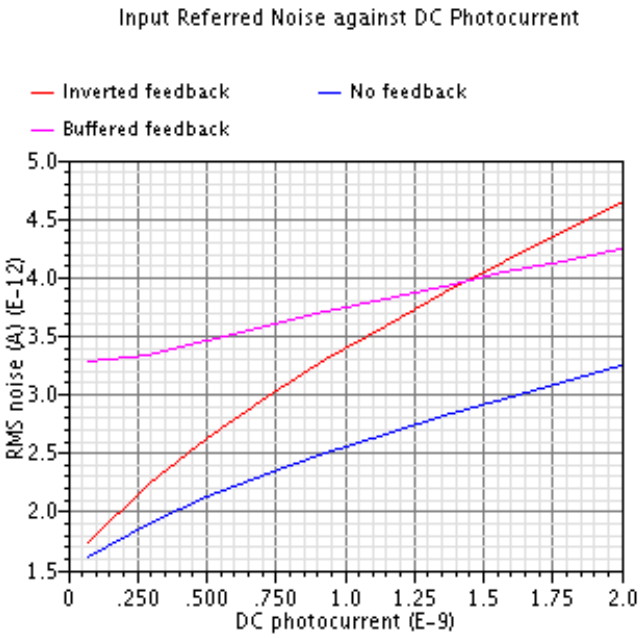

Figure 3.19: Total noise within 100-30 kHz bandwidth for different front-end designs with increasing photocurrent

3.2.5 Linear Front-End Circuit

The photocurrent to voltage converters considered so far have all been logarithmic front-ends, that produce an output voltage inversely proportional to the logarithm of the photocurrent. These have advantages of compact size which is ideal for integrated detectors, and inherent normalisation. However, the non-linear response can make processing more complex, or at least less intuitive, as there is no separate DC channel. A separate DC channel can make some aspects of signal and image processing easier, although the extra channel itself adds some complexity in terms of data acquisition.

This section will investigate an alternative type of front-end, the opamp-based linear front-end. Advantages and disadvantages of this front-end design compared to logarithmic pixels will be considered, along with issues relating to implementation on an IC.

3.2.5.1 Advantages of Linear Front-Ends

One of the biggest issues with logarithmic pixels can be detecting which sections of an image are from the target object, and which are from background (for example the gaps between fingers in an image of a hand). This is because the low light level in the pixels which are

imaging the background causes the AC gain to become very high. This means that the noise in the signal is amplified by a larger gain, which can make thresholding to identify background/foreground more difficult. With a linear pixel, the low light level does not affect gain, so the background pixels give a low AC and DC signal. For flow processing, the AC would then be normalised to give a high AC signal, as occurs naturally in the logarithmic pixel. In the linear pixel case though, the separate reading of the DC channel allows this pixel to be identified as being a low DC pixel, which can be ignored if the light level is below a set threshold.

3.2.5.2 Linear Front-End Schematic

Figure 3.20 shows the schematic of an opamp used as a current to voltage converter. The trans-impedance gain of this circuit is set by the feedback resistor R. C is used to limit the bandwidth of the front-end by setting the RC time constant of the input stage. The opamp on the right of the schematic is a voltage follower used to buffer the input stage from the load on the output.

Circuits such as this are more widely used as trans-impedance stages than the logarithmic circuits already shown. These circuits take advantage of the high common mode rejection ratio and low output noise of opamps specifically developed to be low-noise devices. For example, the Texas Instruments OPA350 has an output noise density of $5\,\mathrm{nV}/\sqrt{\mathrm{Hz}}$. The use of the resistor and capacitor to set gain and bandwidth also make these circuits highly adaptable.

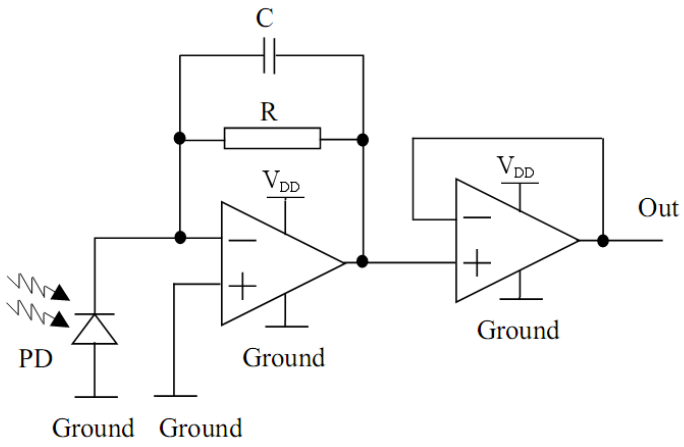

Figure 3.20: Schematic of an opamp based linear front-end (C=0.4 pF, R=20 MΩ for 20 kHz cut-off and 20 MΩ transimpedance) [Kongsavatsak, 2005]

3.2.5.3 Limitations of Linear Front-Ends

The major reason for not using circuits such as those shown in Section 3.2.5.2 is the size requirements of the feedback resistor and capacitor, and to a lesser extent the size of the transistors and compensation capacitor in the opamp. The trans-impedance of the front-end is equal to the value of R. For the signals expected here this resistance must be of the order of 10 MΩ. Using the austriamicrosystems C35 CMOS process, the highest resistance possible is 1 kΩ/□. A 1 μm wide resistor would have to be 10^4 squares (1 cm) long to achieve this resistance. This is possible on an IC, by using snaking resistor paths, but the space consumed is still very large. The resultant resistor also has very high parasitic capacitance. This capacitance is a capacitance to ground rather than between resistor terminals, so does not substitute for the required capacitance. The parasitic capacitance will appear as an additional load capacitance to the first stage opamp, so could affect the frequency response if the total load capacitance was greater than that used in designing the opamp itself. This affect could be used to limit the pixel bandwidth, but this would make performance of the circuit more susceptible to manufacturing variations.

The lack of normalisation is more problematic if on-chip processing is to be used. The division required to perform normalisation based on a DC value is difficult to perform in simple binary operation. If the data is processed off-chip then the processing abilities of a PC or FPGA mean that this is not a significant concern. However, implementing such a circuit on the IC, either using analogue or digital methods, adds to the size and complexity (and hence cost and design time) of the IC.

An additional problem caused by the linear response of the opamp-based front-end is that of saturation with high DC photocurrent. For example, if a trans-impedance of 10 MΩ is used, the output voltage will saturate at a supply voltage of 3.3 V when the photocurrent reaches 330 nA. This is sufficient for this application, however with this stage only, then a photocurrent of 1 nA at 10% modulation depth would give an output voltage of 1 mV. As with the logarithmic front-end circuits, an amplifier would be required to increase the AC signal. If this amplifier had a gain of 50, similar to that of the HDA shown, then the linear front-end based system with 10 MΩ gain would saturate at 6.6 nA. In practice, the operating limit would be below this, as the AC amplification will mean that clipping of the output signal will occur before the input DC photocurrent reaches 6.6 nA. This is above the range of photocurrents calculated for this application, but is considerably lower than the range over

which the logarithmic pixels are capable of operation.

Saturation can be avoided by high-pass filtering the front-end output signal before amplification. This means that the AC can be amplified without also increasing the DC level and saturating the output voltage. However, such filters here would require a very low cut-off frequency (~100 Hz), and hence would require very large capacitors and resistors. The HDA circuits shown could be used in this situation, but the intention of the linear front-end based system is to duplicate conventional photodetector circuits on-chip so conventional op-amp based amplifiers are preferred here.

3.2.5.4 Simulation of On-Chip Linear Pixels

Despite these potential problems, a number of linear front-end based pixels have been implemented on the BVIPS1 IC. The linear array system means that unlike 2D imaging arrays, it is possible to implement pixels requiring components with large footprints without causing unacceptably low fill-factors or limiting pixel count and hence resolution. The implementation of a small number of test pixels allows the advantages of linear pixels to be investigated, and a decision to be taken on whether the advantages justify the increased space. The simulations here for a design that could fit on chip show that on-chip implementation is feasible. The circuit is composed of an opamp based front-end with a trans-impedance of 20 MΩ and an op-amp voltage amplifier gain stage with $A_v = 50$. A schematic of this design is shown in Figure 3.21.

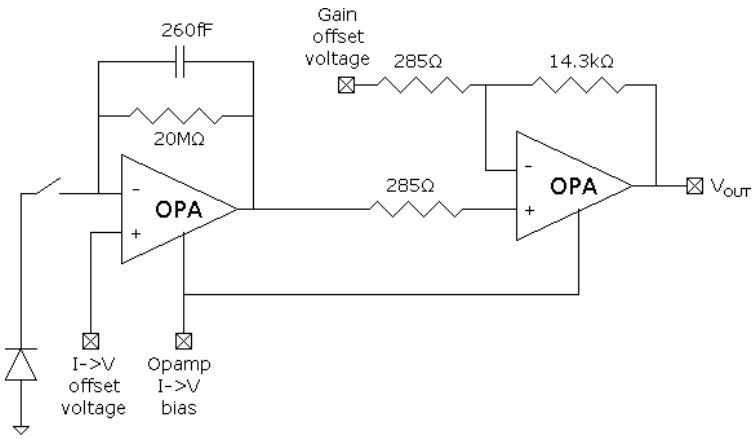

Figure 3.21: Schematic of an opamp based linear front-end suitable for implementation on IC

3.2.5.5 DC Simulation of Linear Pixels

Figure 3.22 shows the DC response from the opamp front-end and gain stage. It can be seen that the output is linearly proportional to photocurrent over the full range of currents expected based on the calculations in Section 1.8.1. The output voltage does not fall to zero because an offset voltage is required to prevent saturation at 0 V, as the IC does not allow dual rail power supplies.

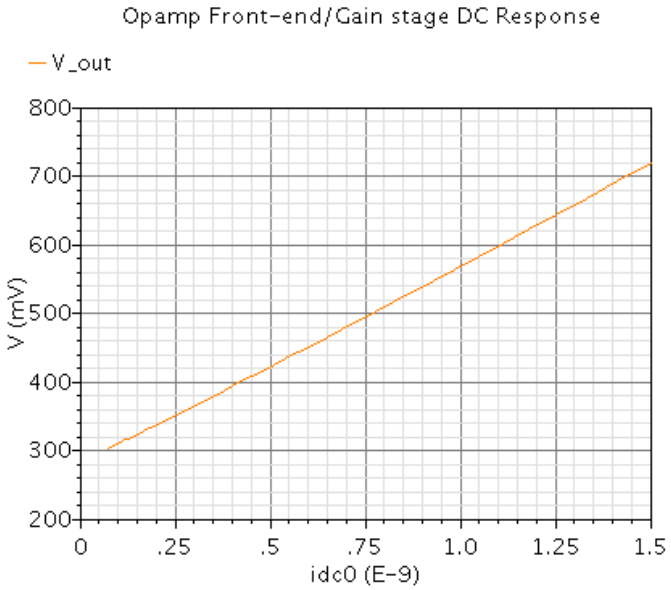

Figure 3.22: DC response of opamp front-end and gain stage

3.2.5.6 AC Simulation of Linear Pixels

Figure 3.23 shows the AC response from the opamp front-end and gain stage. The bandwidth is lower than the design value of 20 kHz. This suggests that the parasitic capacitance of the feedback resistor is reducing the bandwidth below that set by the feedback capacitor. However, the bandwidth is still sufficient to capture most of the LDBF signal, being over 10 kHz. As this circuit does not include any further tuneable low-pass filters, the drop in high cut-off frequency reduces the chances of component variability causing aliasing (by increasing the cut-off frequency above half the sampling frequency), so the circuit design was not changed to raise the cut-off frequency. This does have some impact on the flow calculations, as the higher frequency components that may be blocked would be given increased weighting for

the calculation of flow. This could be compensated for by a scaling factor, but this is only an issue if comparing flow readings from different devices - the proportional change in the calculated flow value for a given change in actual flow should be equal.

If required for future iterations of this design, the cut-off frequency could be increased by reducing the feedback capacitor. Reducing the feedback resistor would also increase bandwidth, but at the expense of reduced gain.

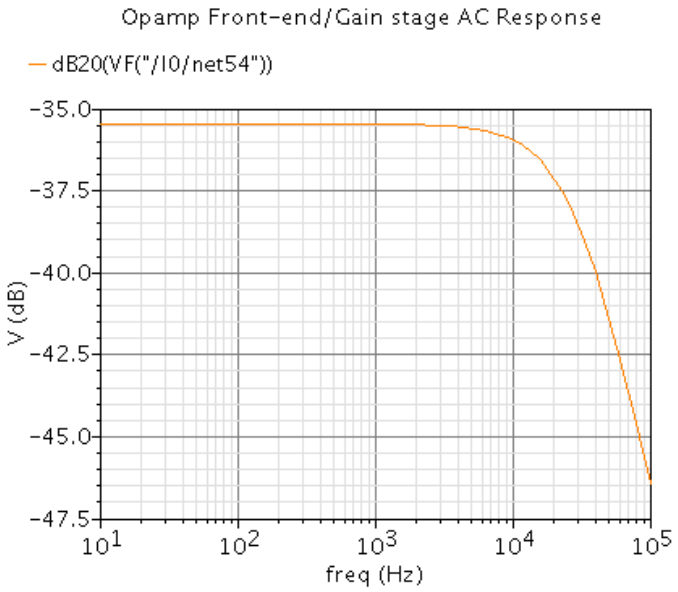

Figure 3.23: AC response of opamp front-end and gain stage

The AC output magnitude below the cut-off frequency is -35.5 dBV, or 16.8 mV (33.6 mV peak-peak). The photocurrent had a DC value of 765 pA, with an AC magnitude of 115 pA peak-peak. This gives a transimpedance of $2.92 \times 10^8\,\Omega$, The expected figure is the product of the 20 MΩ transimpedance of the front-end and the additional gain of 50 from the extra stage, giving $1 \times 10^9\,\Omega$. Both stages were found to have a lower gain than expected, however increasing the front-end transimpedance would require a larger resistor, and increasing the amplification of the gain stage may make saturation more likely. For this reason the circuit shown in Figure 3.21 was not changed before fabrication.

3.2.5.7 Transient Simulation of Linear Pixels

Figure 3.24 shows the transient response of the opamp front-end and gain stage to a typical input photocurrent. The output voltage is sinusoidal and in phase with the input photocurrent as expected, and lacks the slight distortion that is typical in a logarithmic detector. Although the AC magnitude of the output voltage is fairly low at only 25 mV peak-peak, compared to an expected value from a 20 MΩ transimpedance and additional gain of 50 of ~150 mV peak-peak. While low, this is still detectable, and may still give a better signal than the logarithmic pixels if the noise from the opamp based front-end is sufficiently low (see output noise calculations in Section 3.2.5.8). Increasing the gain of the gain stage could make the circuit more susceptible to saturation, so this is to be avoided.

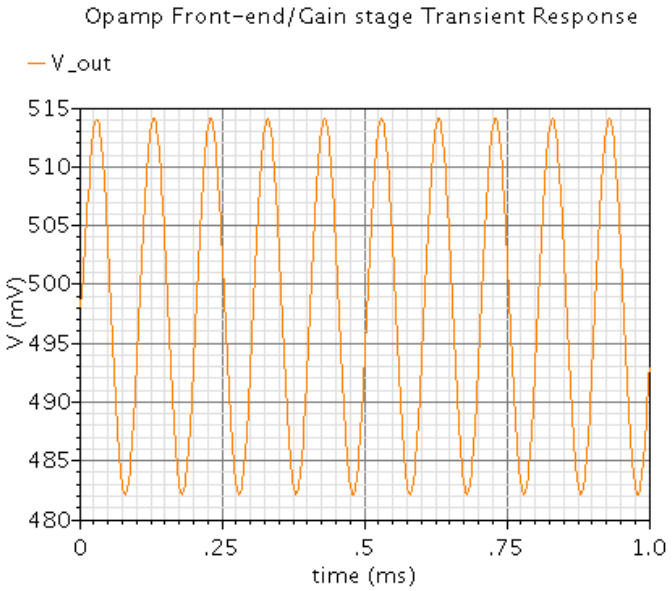

Figure 3.24: Transient response of opamp front-end and gain stage to a 10 kHz, 765 pA DC input signal with 10% modulation depth

3.2.5.8 Noise Simulation of Linear Pixels

Figures 3.25 and 3.26 show the input referred noise spectrum of the linear front-end circuit and gain stage. In the pass band (from around 200 Hz-20 kHz) the noise density is slightly higher than the logarithmic circuits simulated in Section 3.2.4 (30 fA/$\sqrt{\text{Hz}}$ at 1 kHz for the opamp pixel, compared to 15 fA/$\sqrt{\text{Hz}}$ for the buffered logarithmic pixel). However, the response stays flat after the 1/f noise at low frequencies (above ~500 Hz), whereas the noise density of the logarithmic pixels increases as frequency rises - at 10 kHz the opamp and

buffered logarithmic pixels both have an input referred noise density of $\sim 28\,\mathrm{fA}/\sqrt{\mathrm{Hz}}$, above which the logarithmic pixel noise continues to increase. The frequency weighting which will be performed on the output of this circuit for flux calculations may mean this is a significant advantage. The RMS input referred noise over the 100-30 kHz bandwidth is 4.48 pA, giving a peak-peak noise of 29.56 pA. This compares to 5.03 pA RMS / 33.2 pA peak-peak for the buffered logarithmic pixel at 765 pA DC photocurrent, 100-30 kHz noise bandwidth.

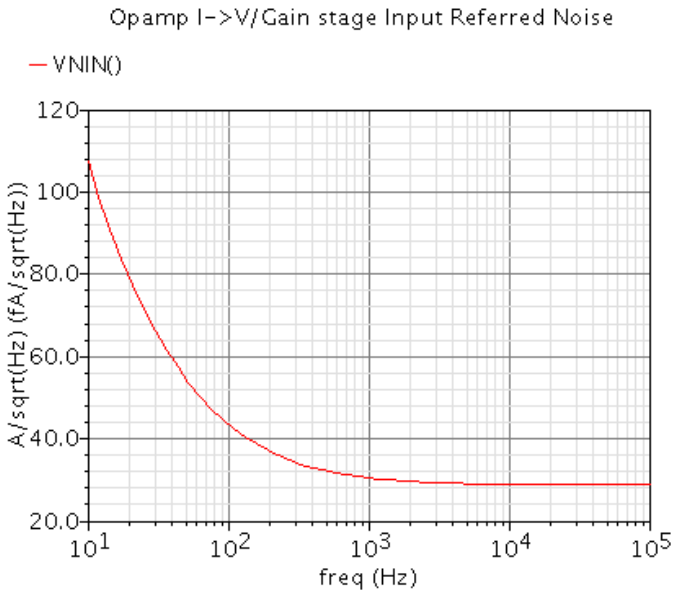

Figure 3.25: Input referred noise spectrum of opamp front-end and gain stage

The opamp front-end may have lower measured noise compared to the logarithmic pixels than predicted by simulation, as the common-mode and power supply noise rejection of the opamp may prevent any additional noise source from appearing on the front-end output. The logarithmic pixels effectively consist of a resistor in series with the photodiode between the power supplies. This means there is little power supply noise rejection. Power supplies will have regulation and bypass/decoupling capacitors, but some noise is still to be expected.

Integrating the output noise voltage from 100 Hz to 30 kHz gives an RMS noise voltage of 1.29 mV, giving a peak-peak noise of 8.514 mV. This is lower than the 25-30 mV peak-peak magnitude seen during AC and transient simulations, suggesting that the SNR from Doppler imaging will be sufficient for successful LDBF use.

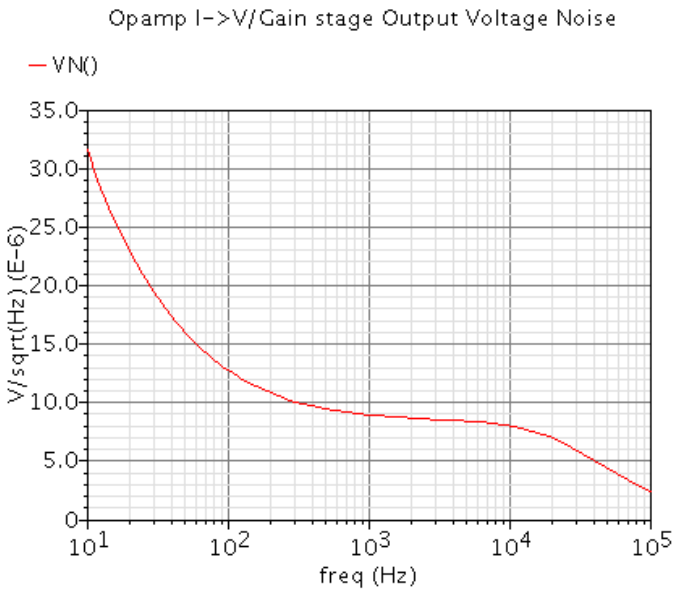

Figure 3.26: Output noise spectrum of opamp front-end and gain stage

3.2.6 Active Pixel Sensors

3.2.6.1 Operating Principle of Active Pixel Sensors

A common architecture used in CMOS cameras, including high frame rate scientific instruments [Modha et al., 2008] and consumer devices is the active pixel sensor, or reset pixel [Holst and Lomheim, 2007, Fossum, 1997].

Figure 3.27 shows the general active pixel sensor design. Rather than converting the photocurrent into an equivalent voltage signal, giving a continuous time output, this pixel samples the light level at set time intervals, giving an output voltage that refers to that sample, rather than the exact moment in time that the output voltage is sampled.

The pixel operates by charging the capacitance of the photodiode to a fixed voltage, usually V_{DD}. The reset signal is then removed, meaning that the photodiode cathode is only connected to the (high impedance) gate of a buffer. The capacitance is then discharged by the photocurrent. After a set time, the voltage on the diode capacitance is measured, generally using an on-chip ADC, through a buffer which isolates the diode capacitance from the read-out stage input capacitance. A high photocurrent leads to rapid discharge, so a larger change in output voltage than for lower light levels.

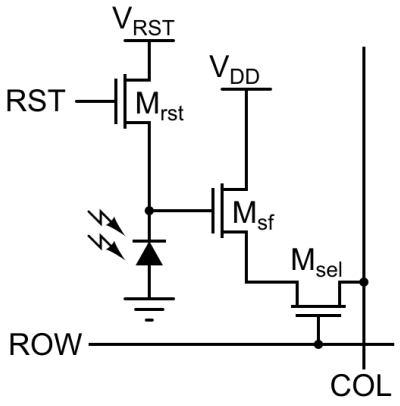

Figure 3.27: General active pixel sensor schematic

3.2.6.2 Advantages of Active Pixel Sensors

This pixel has advantages that its control and output electronics are fairly simple (for example 3 transistors in the example in Figure 3.27) allowing high fill factors. The sampling method also performs averaging, as noise during the discharge time is averaged, with the main sources of noise becoming reset noise (noise on the diode voltage when the reset signal is removed, which affects the voltage from which discharge starts) and readout noise (the noise of the readout sampling circuit). The first of these can be reduced by correlated double sampling, where the diode voltage is read twice, once immediately after the reset is removed, and once at the end of the sampling interval as in the basic version. The difference between this voltage is used, so variations in the original diode voltage have less impact on the output.

The above advantages are generally taken to mean that the active pixel sensor gives superior performance in terms of noise than continuous time pixels such as the logarithmic pixels, mainly due to the averaging inherent to the active pixel [Serov and Lasser, 2005]. However, the time constant of a continuous time pixel such as the logarithmic pixel could be considered to perform a similar type of averaging.

3.2.6.3 Simulation of an Active Pixel

Figure 3.28 shows the photodiode voltage of an active pixel during sampling. The diode voltage is initially driven high by the reset signal. Once the reset switch is made open circuit, the voltage begins to fall, until the next reset pulse sets the voltage back to its high state.

The basic architecture shown here can be modified to add a sample and hold circuit, implementing an electronic shutter. This is based on an additional capacitor in parallel with the photodiode capacitance. During reset and measurement the two are connected, so the voltage on each is equal. After a set sampling interval, a switch breaks the connection between the sample capacitor and the diode capacitance. The diode capacitance continues to discharge due to the photocurrent, whereas the sample capacitor is only connected to the gate of the readout buffer, so only leakage current flows. This means the readout can be measured at different times for different pixels, making data acquisition simpler.

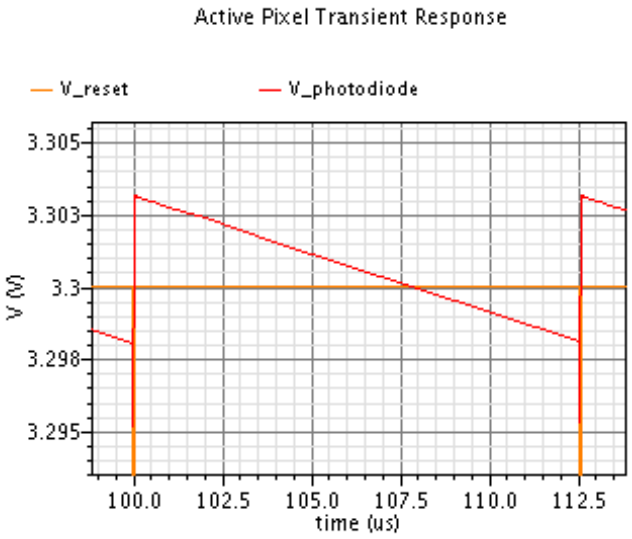

Figure 3.28: Output voltage before during and after discharge interval of active pixel sensor using $1000 \times 50\,\mu$m photodiode with 765 pA DC photocurrent

3.2.6.4 Limitations of Active Pixel Sensors

Figure 3.28 shows a major limitation of active pixels applied to LDBF. The photodiode used here is considerably larger than that common in CMOS cameras, which can have pixel sizes of the order of $1\,\mu\text{m}^2$. Combined with the require sampling rate of 40 kHz ($25\,\mu$s sampling interval) means the change in voltage level over one discharge level is very low. The simulation above has a total discharge of around 5 mV, although this is based on a 40 kHz signal bandwidth, so has half the sampling interval that could be used. However, even doubling this discharge to 10 mV gives a very small signal. The 10 mV change represents the instantaneous DC signal. The AC signal is taken from the change from sample-sample.

Based on a 15% Doppler ratio as used here, this would give a 1.5 mV AC signal. As this is not a sinusoidal signal, amplification as for other circuits is problematic. Instead, an ADC with very high resolution would be required. For this reason, this pixel design will not be used here.

This type of pixel has been used for LDBF, being the type of pixel used in the commercial CMOS sensors used by Serov, Lasser and Draijer for either speckle imaging [Draijer et al., 2009, Serov et al., 2006b] (where higher frame rates but less precision can be achieved with lower sampling rates) or Laser Doppler Flowmetry as used here, but with smaller pixels, higher laser power and lower bandwidth [Serov and Lasser, 2005]. These changes result in more rapid discharge of the diode capacitance, making the use of this detector type possible. However, the device here is required to have large pixels to avoid problems due to mechanical scanning, high bandwidth due to potential high blood velocities in larger vessels, and lower laser power to make its use in a general treatment area safe. One possibility would be to implement a 2D array, such that the full 1000 μm wide light sensitive area is composed of several active pixels, with the pixel with the highest incident light level selected to produce an effective 64x1 array. This would decrease the photodiode capacitance without a proportionate drop in light level (assuming that the line imaged onto the sensor is thinner than 1 mm). The increased complexity of pixel selection, which may need to be changed as the line is scanned, means this approach will not be tested here.

3.2.7 Summary of Current to Voltage Converter Circuits

The simulations shown in this section have demonstrated the basic operating principles of logarithmic and linear current to voltage converters. The weaknesses of the basic logarithmic pixel, mainly insufficient bandwidth at low photocurrents, can be addressed by the designs using buffering in series with the photodiode, or using inverted feedback buffering. These two designs appear to be approximately equivalent in performance, with the main differences being that the inverted feedback design has a lower DC output voltage, lower transimpedance, and input referred noise that is lower at low photocurrents but increases faster than the other buffered design, having higher noise above 1.5 nA. The lower DC voltage is not in itself an issue, but with a large AC signal (after later amplification) this could lead to clipping of the signal at low voltages. The noise performance is partially dependant on the incident light power, and hence the optical setup used, but the increasing noise with rising photocurrent

means that if laser power is increased (e.g. to increase light penetration into the skin, and hence measure flow in deeper tissue), the noise of the feedback buffered circuit deteriorates relative to the series buffered design. For this reason, the preferred logarithmic pixel design is the series buffered pixel.

The linear design has significant drawbacks for use on an integrated circuit, but the 1D array design means that implementation is possible, and so a number of test pixels will be implemented on the first prototype IC to further investigate the performance of these pixels.

Active pixels were considered, given their widespread use in commercially available sensors, including use for LDBF, but the large size of photodiode required results in high capacitance and hence low voltage changes, making them unsuitable for use on the ICs to fabricated during this project.

3.3 Amplifier and High-pass Filter - Hysteretic Differentiator Amplifier

The simulated AC output voltage from the front-end circuits is shown in Figure 3.14 to be less than 10 mV. This voltage may be too small to be accurately measured with commonly available ADCs. For this reason, an amplifier is required to increase the AC voltage. This amplifier must not amplify the DC voltage, as this would saturate the output voltage due to the large input DC voltage. In discrete circuitry this could be done by removing the DC with a high-pass filter, but this is problematic on an IC, where large capacitances or resistances require a large silicon area. Large resistances can be implemented using active circuitry similar to that providing the load in the current-voltage converters, but in this context the high resistance is to be avoided as this causes an increase in thermal noise, where $V_n = \sqrt{4KTBR}$. This section describes two designs of hysteretic differentiator amplifier (HDA), a type of amplifier that amplifies the AC voltage while having a unity gain at DC [Mead, 1989].

Both designs use an OTA (Operational Transconductance Amplifier) as the main circuit element, with either an inverted-inverter and a capacitor [Mead, 1989] or a G_mC low-pass filter (a circuit setting a high-frequency cut-off through the combination of an OTA gain and a load capacitance, described in Section 3.4 [Geiger and SÃ¡nchez-Sinencio, 1985]) as the feedback network [Gu et al., 2008]. The G_mC feedback version has the advantage of greater tunability, as varying the bias of the G_mC in the feedback loop allows the low cut-off to be controlled by an off-chip bias resistor. However, simulation of this design also showed greater variation between different pixels and between ICs due to process variations and IC mismatch. For this reason, the inverted inverter design was used. This section shows the operating principles and simulations results from both circuits to support this decision.

3.3.1 HDA Using Inverted-Inverter Feedback

Figure 3.29 uses an OTA as the main amplifier element, with a feedback path consisting of an inverted inverter and a capacitor. The inverted-inverter (MN0 and MP0) appears at first glance to be a standard CMOS inverter, however the position of the PMOS and NMOS devices is reversed. This means that both transistors are always in the cut-off region, which makes the inverted-inverter act as a very high resistance. The NMOS transistors (MN1 and

MN2) provide a load capacitance of approximately 7 pF on the output that when combined with the large resistance of the inverted-inverter gives a high RC time constant and thus the low cut-off frequency of the HDA is sufficiently low for this application. Signal components below this cut-off frequency are fed back to the OTA inverting input, and hence these signals are not amplified. Above this frequency the RC feedback network blocks the signals, and hence these are amplified by the open loop gain of the OTA, which is set by the OTA transistor dimensions and bias current.

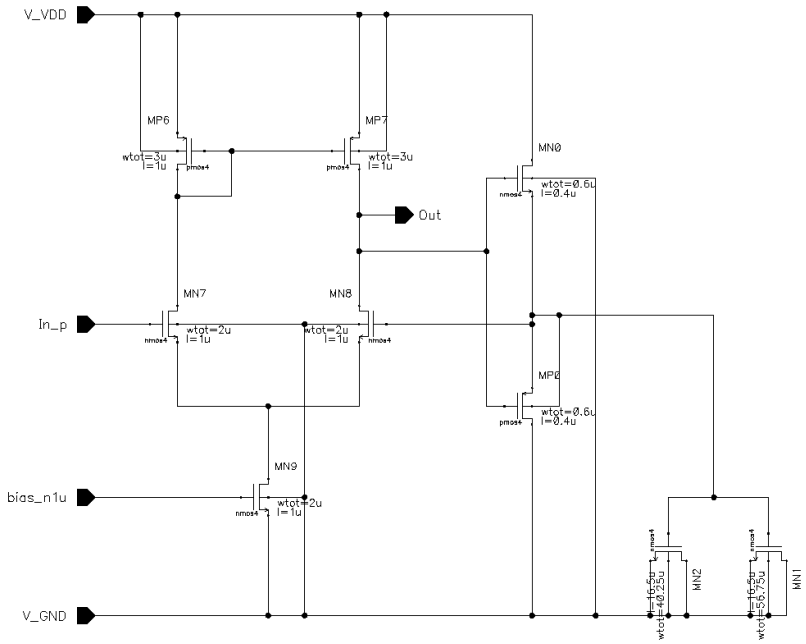

Figure 3.29: HDA using inverted-inverter feedback [Mead, 1989, Kongsavatsak et al., 2008]

3.3.2 HDA Using G_mC Feedback

For this design, the feedback path is provided by a G_mC low-pass filter as shown in Figure 3.30.

The G_mC circuit by itself forms a low-pass filter circuit consisting of an OTA with an additional output capacitor. Filters of this type are used on this IC to provide a compact, tune-able anti-aliasing filter, described in Section 3.4.

The G_mC filter in the feedback path of this type of HDA uses a modified OTA to make its cut-off frequency lower than the version used as an anti-aliasing filter. The cut-off frequency of the G_mC filter is given by:

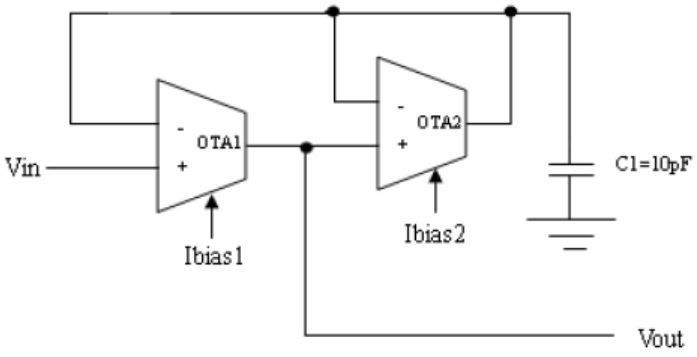

Figure 3.30: HDA using G_mC feedback [Gu, 2007]

$$f_{cut-off} = \frac{g_m}{2\pi C} \tag{3.14}$$

This means that to achieve a low cut-off frequency using such a circuit requires either a large capacitor, making on-chip implementation expensive, or an OTA with a low g_m can be used. A design with low g_m is shown in Figure 3.31 [Gu, 2007, Geiger and SÃ¡nchez-Sinencio, 1985]. The OTA used in the input stage of this design is shown in Figure 3.32.

This design has advantages in that the low-cut off frequency of the HDA can be adjusted by varying the bias current, and hence the cut-off frequency, of the feedback OTA. However the design still requires a 10 pF capacitor which in the process to be used here requires a $90 \times 25\,\mu$m MOS capacitor. Combined with the larger, more complex OTA design this makes the g_mC feedback HDA design the larger of the HDA options.

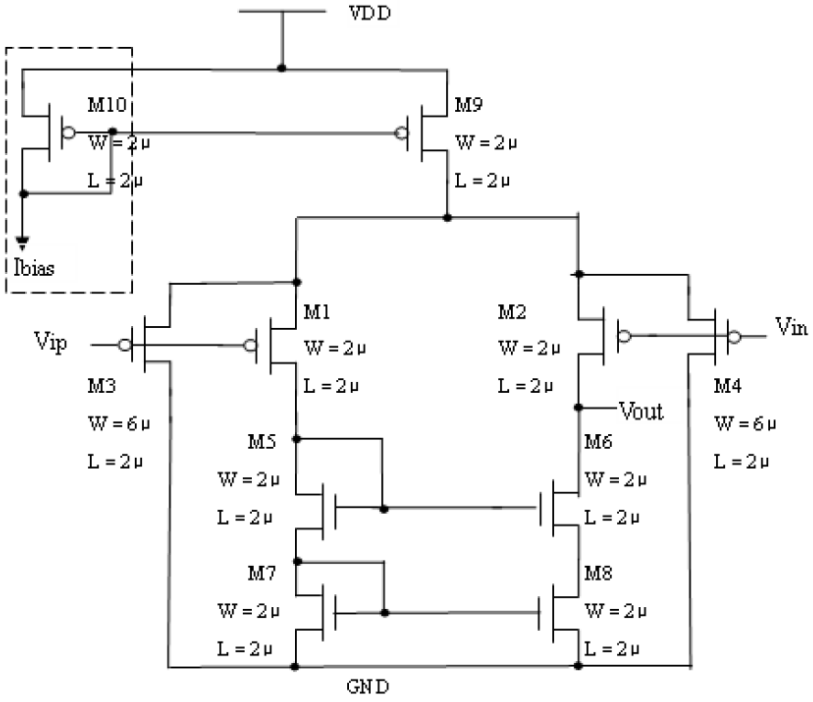

Figure 3.31: OTA2, with reduced $g_m$for HDA-g_mC

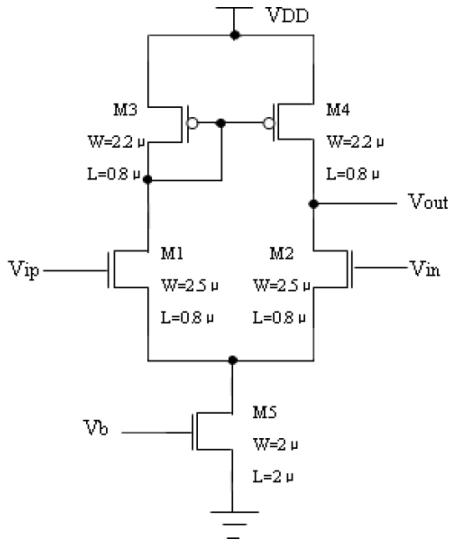

Figure 3.32: OTA1, standard g_m for HDA-g_mC input stage

3.3.3 Simulation of HDA designs

To select a HDA design for use on an IC, simulations were performed in Cadence on the available options. However, the small currents and high resistances involved to compensate

for the small capacitances in these designs can make simulation problematic. Accurate simulation requires setting the minimum possible g_{min} used in the simulation lower than the default value. This default value is generally sufficient to not limit most common circuit designs, while increasing simulation speeds.

This problem is shown by Figures 3.34 and 3.33, which show the significance of g_{min} on the simulation results. The figures show results of four simulations on each device, with the only change between simulations being the value of g_{min} in the simulation settings. For the g_mC feedback design simulations, shown in Figure 3.33, if simulated at the default g_{min} of $1 \times 10^{-12}\,\mathrm{VA}^{-1}$, the circuit attenuates all frequencies heavily. As g_{min} rises from 1×10^{-13} to $1 \times 10^{-15}\,\mathrm{VA}^{-1}$ intended behaviour is observed, but the low cut-off frequency varies from 15 Hz to 200 Hz. For the inverted-inverter circuit simulations, shown in Figure 3.34, as g_{min} falls the low frequency cut-off falls from around 100 Hz at $g_{min} = 1 \times 10^{-13}\,\mathrm{VA}^{-1}$, to 1.5 Hz at $g_{min} = 1 \times 10^{-14}\,\mathrm{VA}^{-1}$, and with no attenuation at lower frequencies when $g_{min} = 1 \times 10^{-15}\,\mathrm{VA}^{-1}$.

Generally g_{min} should be reduced until the simulator reaches convergence with the given settings - i.e. reducing g_{min} further does not affect the simulation results. This suggests that g_{min} should be set to the minimum value of those used here, or lower. However, both circuits simulated here have been fabricated and operate in the expected manner [Gu et al., 2008, Kongsavatsak et al., 2008], showing that using a lower value of g_{min} does not guarantee more accurate simulation of these circuits. One possibility for this is that the large feedback resistances achieved by these circuits to give a low cut-off frequency rely on the very small leakage currents within the fabricated IC. As g_{min} represents a very small conductance added to the simulated circuit, generally to simplify the calculations required by the simulator to achieve convergence. It is possible that the leakage currents and voltages within these circuits represent conductances higher than the lowest g_{min} values used in simulation.

Given the discrepancy between performance of previously fabricated circuits and simulation results at very low g_{min}, for the following simulations g_{min} was set to match known results. For the simulations here g_{min} was set at $1 \times 10^{-13}\,\mathrm{VA}^{-1}$, the closest value to the default at which both circuit designs show similar simulation results to known test results (i.e. amplifying AC signals above a low cut-off frequency of 10-100Hz, while having approximately unity gain at DC).

The characterisation results of the prototype ICs produced following these simulations (shown

later in Chapters 4 and 5) show significant discrepancies between measured and simulated results, and further investigation of possible HDA faults (section 4.6.3) shows that the simulation of the HDA circuit performed in this chapter is potentially flawed. However, this problem only became apparent after the first IC was fabricated. Further simulations on the HDA could not replicate any faults found, regardless of changes of g_{min}. While changes to parameters such as g_{min} may improve simulation accuracy, with no definitive cause of the inaccuracy found, the following simulations remain the most accurate method available for comparing circuit performance.

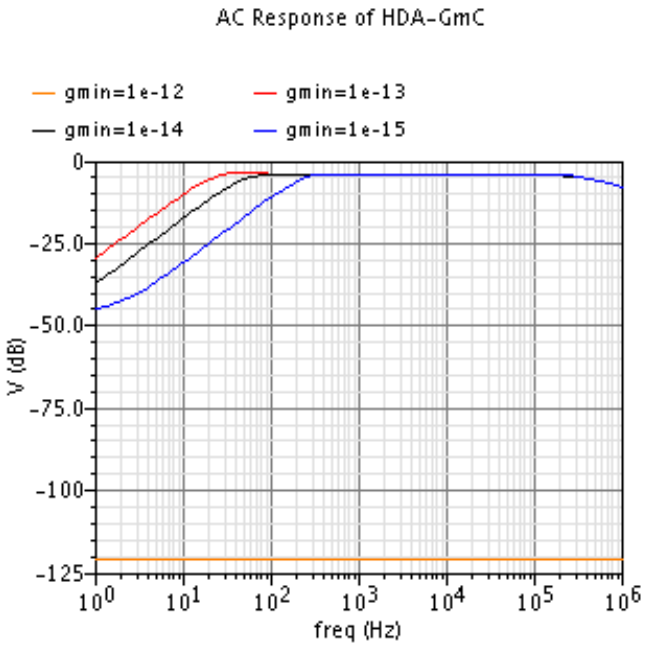

Figure 3.33: AC response of HDA using g_mC feedback, varying g_{min}

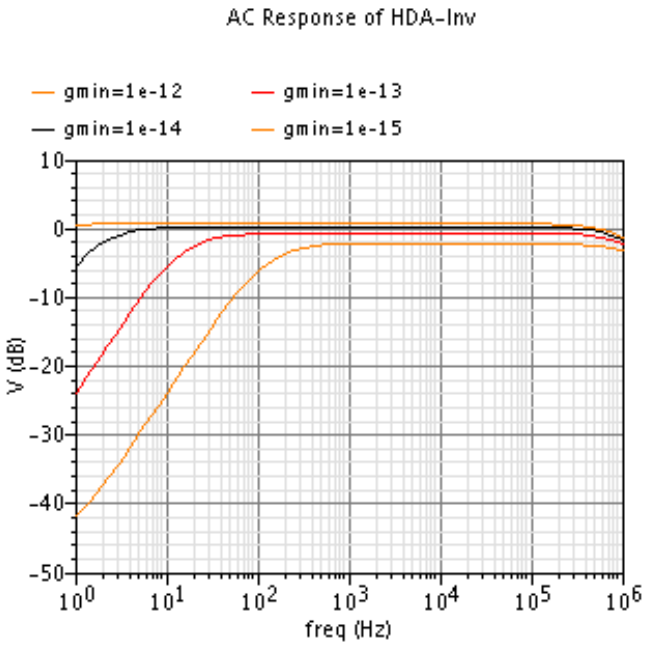

Figure 3.34: AC response of HDA using inverted-inverter feedback, varying g_{min}

3.3.3.1 Transient Response

Figure 3.35 shows the output signals of both HDA designs with an input signal having a 5 mV AC magnitude (approx. signal magnitude expected from front-end) at 10 kHz. Both designs amplify the signal by a similar amount, although the inverted-inverter feedback version has considerably more distortion on the output signal.

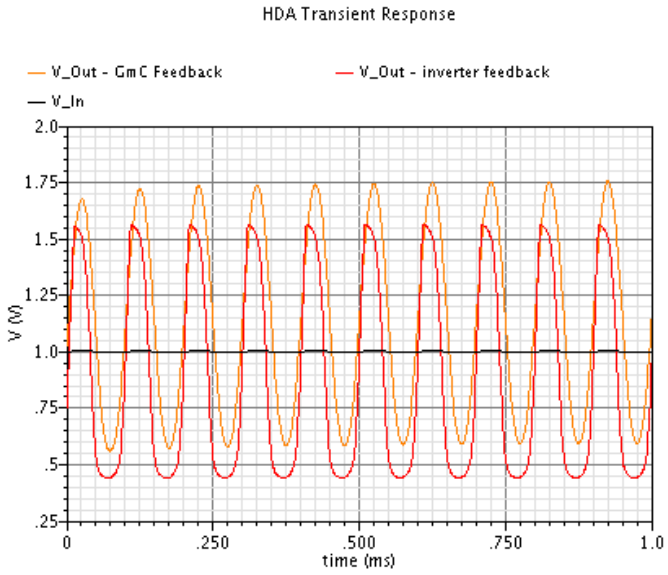

Figure 3.35: Transient response of both HDA designs, with a 1 V DC / 5 mV @ 10 kHz AC input signal

3.3.3.2 DC Response

Figure 3.36 shows the DC response of the HDA designs. The major feature of these plots is the non-linearity shown in the inverted-inverter response at 1.65 V (50% of V_{DD}). This is due to the change of characteristics of the inverted-inverter as the output voltage passes the transistor switching point. This non-linearity is a likely cause of the distortion shown in the transient response.

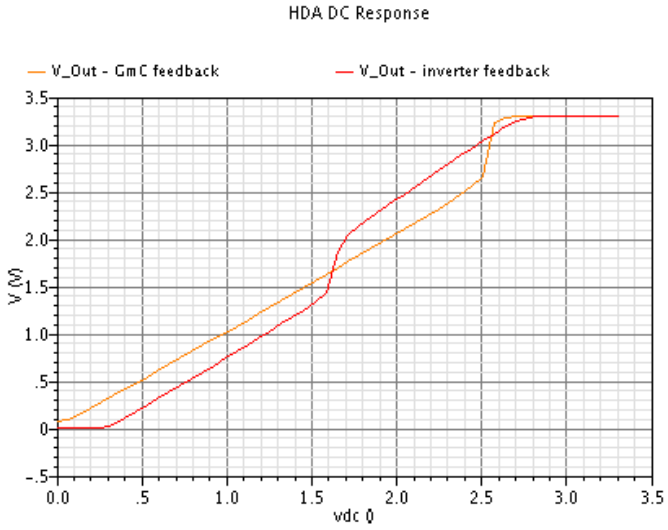

Figure 3.36: DC response of both HDA designs, with input from 0 to 3.3 V

3.3.3.3 Frequency Response

Figures 3.37 and 3.38 show the frequency response of the HDA using g_mC feedback, including the effect of changing the biasing of the feedback OTA. Figures 3.39 and 3.40 show similar results for the HDA using inverted-inverter feedback. (Note that the y-scale units are in dB referred to 1 V, so with a 5 mV input 0 dBV represents a voltage gain of 200, while unity gain is -106 dBV) It can be seen that the basic response is similar, although the inverted feedback design has higher gain and a lower low-frequency cut-off. However, the main advantage of the g_mC feedback design is clearly seen by the relationship between bias current and low-frequency cut-off. This would allow the circuit to be tuned to accommodate chip-chip variation, as well as adjustment to suit the application, such as lowering the cut-off frequency to include more low frequency components of the signal from the front-end in the final output signal. This could potentially increase the signal level at low frequencies in situations where movement artefacts are a less significant problem. This could include monitoring of less mobile subjects (e.g. during sleeping/anaesthesia), or cases where the device itself is less prone to movement - i.e. a device on a stand will be less prone to movement artefacts than a hand-held device.

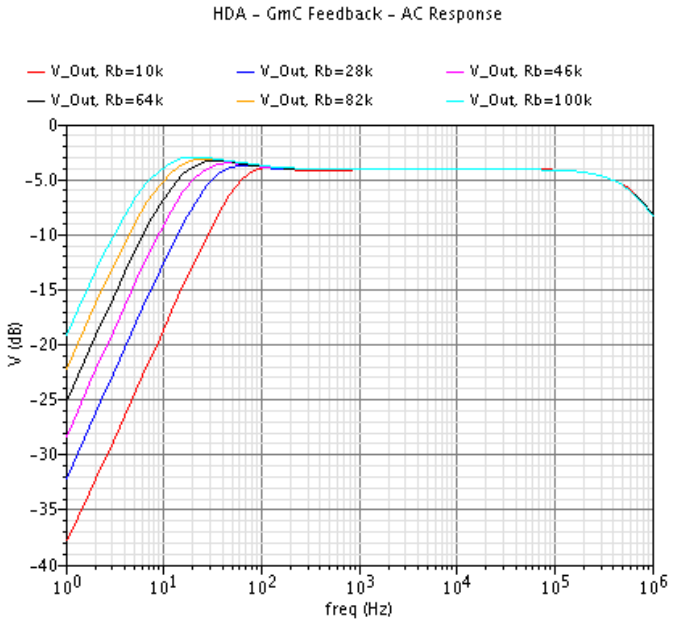

Figure 3.37: AC response of g_mC feedback HDA with varying bias current

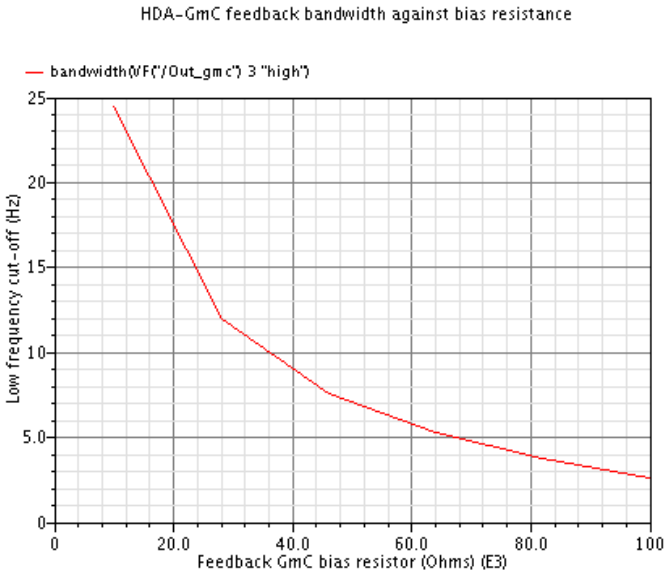

Figure 3.38: Low-frequency cut-off of g_mC feedback HDA with varying bias current

For the inverted-inverter feedback design, changing the OTA bias current has some effect on bandwidth, but this is mainly because the gain in the pass band is slightly altered. The range of adjustment here is much more limited, giving no significant adjustability of the IC.

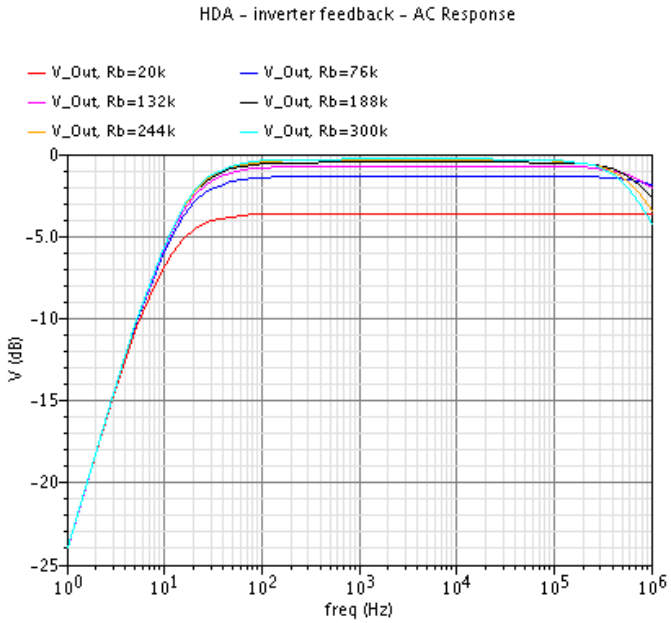

Figure 3.39: AC response of inverter feedback HDA with varying bias current

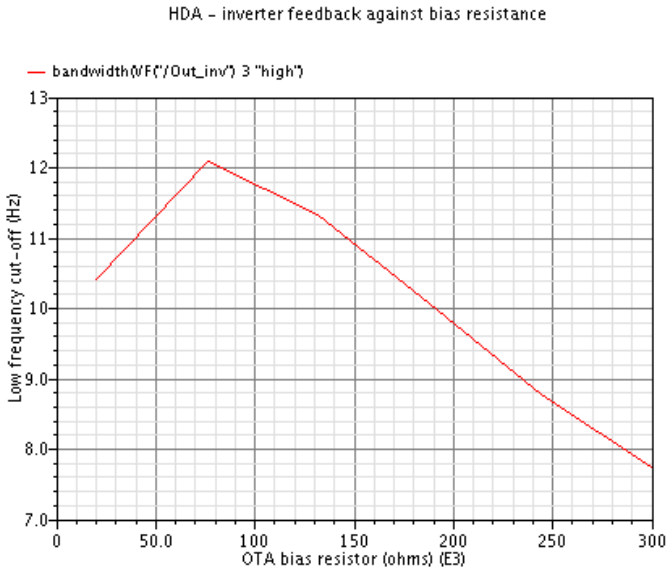

Figure 3.40: Low-frequency cut-off of inverter feedback HDA with varying bias current

3.3.3.4 Noise Response

Figures 3.41 and 3.42 show the input and output and referred noise respectively for both designs. The output noise plot shows higher noise at the output of the inverter-inverter

design. However, this is partly due to the higher gain of the inverted-inverter design, as shown by Figures 3.37 and 3.39. The input referred noise plot, which compensates for gain and is thus a better indicator of signal-noise ratio, shows that the inverter design has lower noise at lower frequencies, and equal noise density at higher frequencies. The higher output noise is therefore due to the higher gain of the inverter design, and the lower low-frequency cut-off. The similar input referred noise densities within the 100-20 kHz range to be used here suggest that neither circuit has a significant noise advantage.

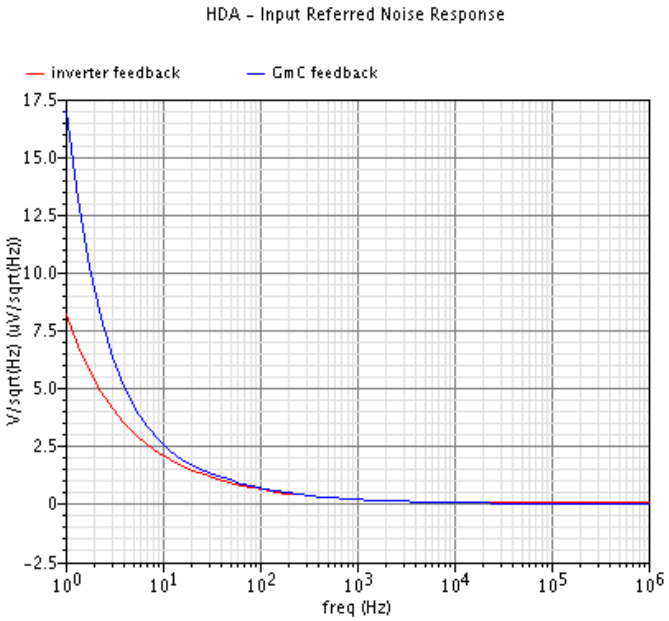

Figure 3.41: Input referred noise spectrum of both HDA designs

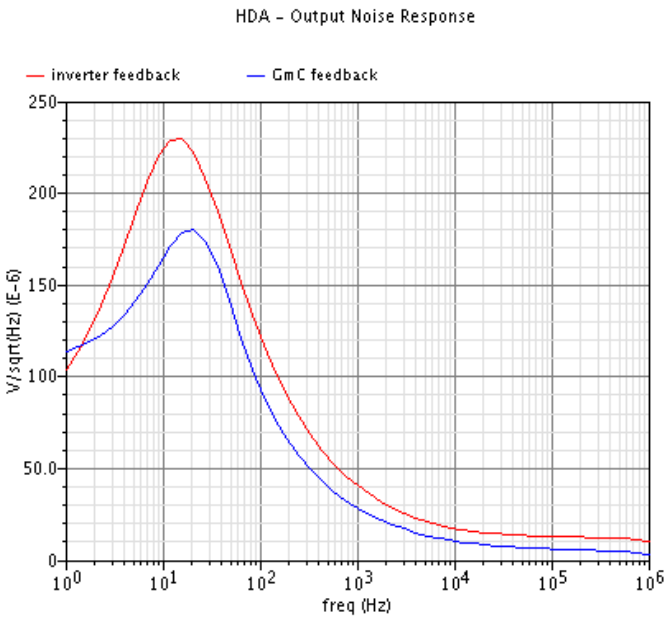

Figure 3.42: Output noise spectrum of both HDA designs

3.3.3.5 Monte-Carlo Analysis of Frequency Response

The simulation results above show the nominal behaviour of each circuit. Given the known problems in simulating these circuits, the statistical variation of these results was also investigated, as the different designs may be affected differently by device variation. This was performed using repeated iterations of the simulations above with random variations in device size and electrical process parameters. This shows the behaviour that can be expected of arrays of these circuits within one IC, and from IC to IC.

Figures 3.43 and 3.44 show the variation in low-frequency cut-off of both designs. It can be seen that the g_mC design has a wider range of cut-off frequencies, which will lead to greater fixed pattern noise. Out of 100 iterations there is also one iteration that does not give a valid cut-off frequency, showing that this circuit does not have the same behaviour within the required range - the gain does not drop with falling frequency to 3 dB below the maximum value, either due to slow roll-off/low cut-off frequency, or due to very low maximum gain such that the expected gain at DC is not 3 dB lower than the maximum. This means either that this circuit will amplify the DC level, causing saturation of the output, or that the AC gain of this circuit is much lower than intended.

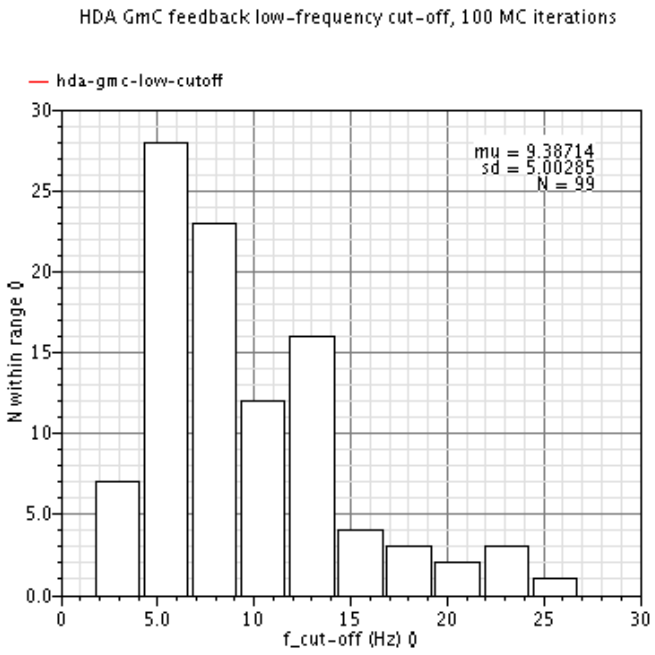

Figure 3.43: Statistical variation of low-frequency cut-off of g_mC feedback HDA

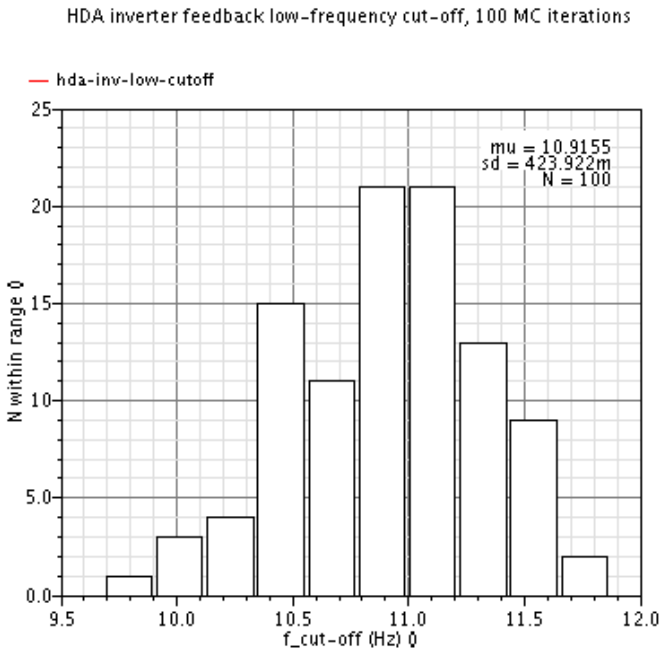

Figure 3.44: Statistical variation of low-frequency cut-off of inverter feedback HDA

3.3.3.6 Corner Simulation of HDA Designs

The Monte-Carlo simulations above show the effect of an expected level of device mismatch. It is also useful to perform corner analysis to investigate circuit behaviour with worst-case device variation/temperature/voltage supply variation.

These simulations are performed in Cadence, using alternative device models provided for the austriamicrosystems C35 process used here. The corners are: cmostm (Typical mean parameters); cmosws (worst case speed, with slower transistors than typical); cmoswp (worst case power, faster transistors but higher power consumption); cmoswo (worst case one, slow PMOS transistors with fast NMOS transistors); cmoswz (worst case zero, fast PMOS transistors, slow NMOS transistors). These simulations generally apply to digital circuits (i.e. inverter behaviour - slowest transmission, worst power consumption per gate, skewed transmission of 1/0 states), but can be used to show the effect of variation on analogue circuits.

Figure 3.45 shows the effect on DC operating point of both designs, by plotting the DC output voltage with a 1 V DC input voltage. The g_mC circuit shows little variation, being approximately 0.05 V from the lowest to highest point. The inverter-inverter circuit shows more variation, with a range from 0.63 - 0.95 V between the worst speed and worst power corners. Given the low modulation depth expected here, this variation of the DC output voltage should not prevent normal pixel operation, although the inverted-inverter circuit already shows high distortion of the output waveform at low voltages (see Figure 3.35), which would be exacerbated by a further drop in DC output voltage. If this level of variation was seen within a chip, the fixed pattern noise seen would be high, although the corner simulation includes process (between chip) variations, so this level of variation across one array is unlikely. It is also possible that the use of an inverter-like structure in this circuit means that these two corners are likely to show more extreme variation than for other analogue circuits, where a corner might not be a genuine worst case scenario.

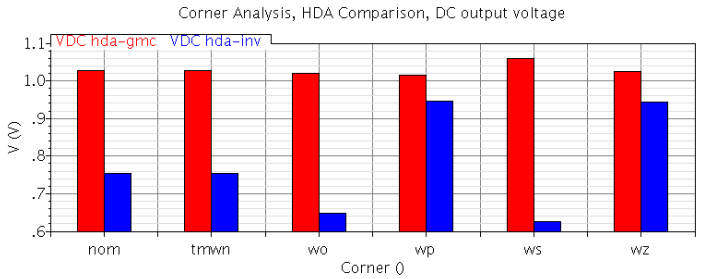

Figure 3.45: Corner Simulation of DC operating point (measured by output voltage) of both HDA designs

Figure 3.46 shows the effect of each corner on AC gain of both HDA designs. Again the g_mC design has lower variation, with a range of output magnitudes from 1.0 - 1.3 V compared to 0.7-1.34 V for the inverted-inverter design. As for the DC operating point simulation, the greatest variation is seen between the worst speed and worst power corners, possibly due to the use of an inverter-like structure in the feedback circuit. While this variation is high, the minimum gain of the inverter-inverter circuit is sufficiently high for detection of the required signals, and the maximum gain is similar for both circuits. The variation would cause high FPN if seen within one IC, but this is unlikely.

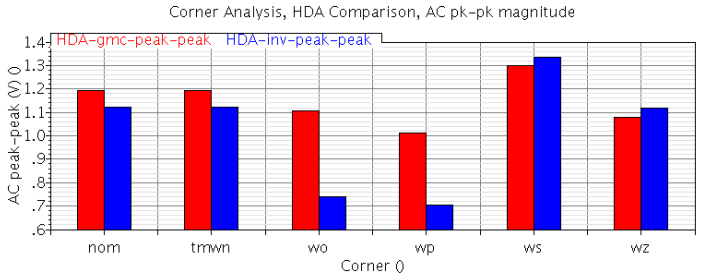

Figure 3.46: Corner Simulation of AC output magnitude of both HDA designs, with a 1 V DC / 5 mV @ 10 kHz AC input signal

Figure 3.47 shows the effect of each corner on the low frequency cut-off of both HDA designs. The Monte-Carlo simulations showed that the variation of this frequency in the g_mC circuit was likely to be higher than for the inverted-inverter circuit, and the corner simulation shows similar but more extreme results. For the inverted-inverter design, the cut-off frequency remains below ~15 Hz in all cases, while for the g_mC the maximum cut-off frequency is 300 Hz. The majority of the Doppler signal may be a frequencies above this, but high variation of the low-frequency cut-off could still have a significant effect on output flow values. However,

if this variation is seen between chips rather than within one array, each individual imager should still have acceptable performance.

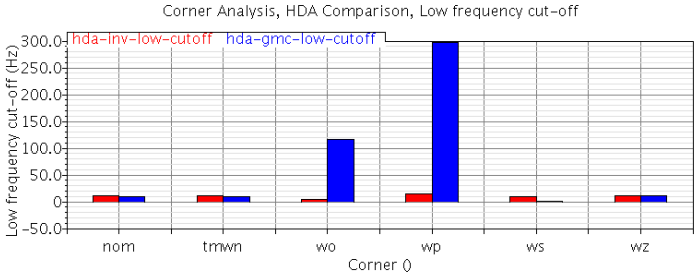

Figure 3.47: Corner Simulation of low frequency cut-off of both HDA designs, with a 1 V DC / 5 mV AC input signal

Figure 3.48 shows the input referred current noise for both HDA circuits. All corners have higher noise than the typical mean case, and in each case the g_mC based circuit has higher noise than the inverter-inverter design, despite the typical mean case showing lower noise for the g_mC based circuit. This suggests that the g_mC based circuit is slightly more susceptible to device variation causing increased noise, although the difference between HDA versions is fairly low.

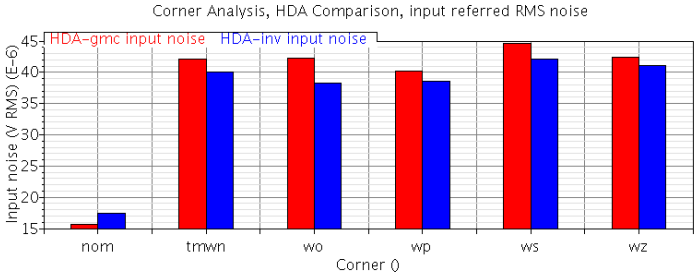

Figure 3.48: Corner simulation of input referred current noise of both HDA designs, with a 1 V DC input signal

Figure 3.49 shows the output voltage noise for both HDA circuits. Most corners show increased noise compared to the nominal case, except for the worst-case-one corner for the inverter-inverter circuit. The g_mC based circuit shows the lowest increase in noise, although as the input-referred noise results are a better indicator of signal-noise ratio, this does not show a major advantage.

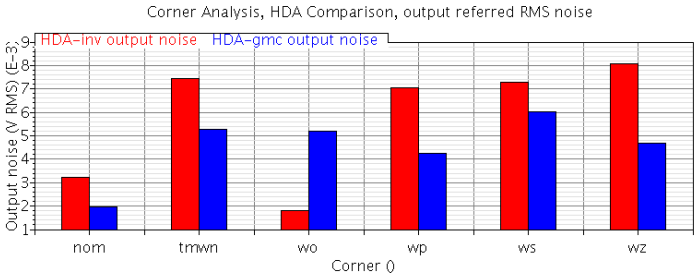

Figure 3.49: Corner simulation of output voltage noise of both HDA designs, with a 1 V DC input signal

3.3.4 Summary of Amplifier Circuits

Two different designs of hysteretic differentiator circuit have been introduced, with both of these circuits offering the voltage gain and low-frequency cut-off required by this application, and both having a relatively compact layout (no large capacitor is required to achieve the low cut-off frequency). The simulation results for the nominal behaviour of each circuit (i.e. not subject to process variation) show considerable advantages for the g_mC based HDA, mainly linearity, reduced distortion and external adjustability. However, the Monte-Carlo and corner simulation results show that this circuit suffers from greater variation between circuits, and hence the inverted-inverter design is chosen to give better reliability (i.e. higher yield).

3.4 Low-pass Filter - G_mC Anti-Aliasing Filter

The HDA circuits shown in Section 3.3 are high-pass filters, with a designed low-frequency cut-off around 100 Hz, and an operating range up to around 1 MHz. At this point a high-frequency cut-off is caused by device parasitic capacitances rather than any aspect of circuit design. This high-frequency cut-off is well above that required by this application, and is not controllable by external biasing. A filter is therefore required to remove AC components above the bandwidth required here (discussed in the Section 1.8.3) to prevent aliasing.

3.4.1 G_mC Filter Schematic

The g_mC filter circuit shown in Figure 3.50 is used for this purpose [Geiger and SÃ¡nchez-Sinencio, 1985]. This circuit provides a compact filter layout ($35 \times 10\,\mu$m) that can be included at the pixel level, and can be controlled by an external bias current to set the bandwidth to the required value. This has also been used successfully on previous Doppler ICs [Gu et al., 2008, Kongsavatsak et al., 2008].

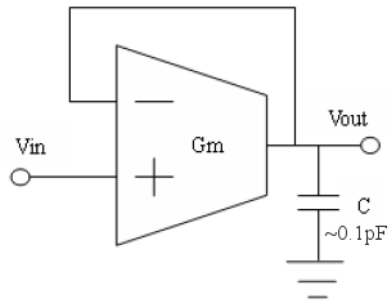

Figure 3.50: General schematic of G_mC anti-aliasing (low-pass) filter

The schematic of the g_mC circuit used on this IC is shown in Figure 3.51, which includes the schematic of the OTA. The OTA used here is a standard circuit, without the changes made to the OTA used in the HDA shown in Section 3.3.1. This is possible as the high frequency cut-off required here of around 20 kHz can be achieved with higher RC values than the low frequency cut-off required of the HDA. The output capacitance in this circuit is implemented using an MOS capacitor with size $5 \times 5.5\,\mu$m, giving a capacitance of 125 fF.

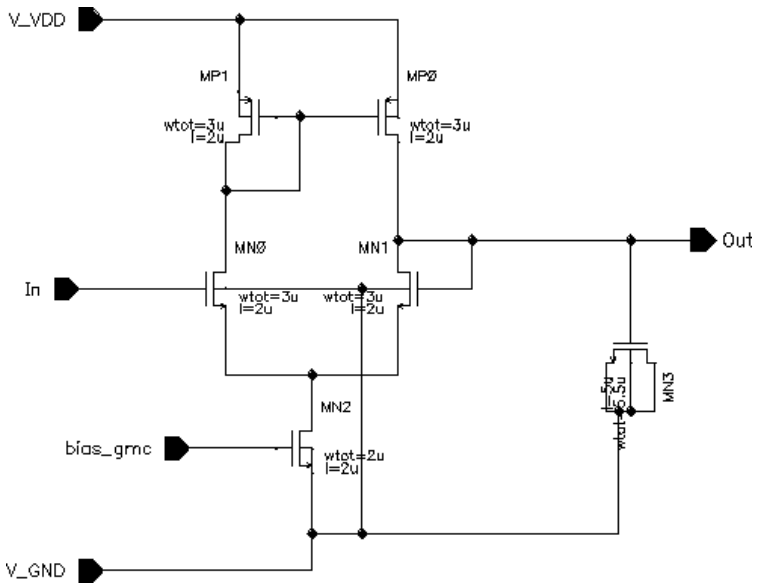

Figure 3.51: Schematic of G_mC filter on BVIPS1

3.4.2 Simulation of G_mC Circuit

The simulations shown below were carried out to verify that the g_mC filter shown in Figure 3.51 is suitable for use as an anti-aliasing filter for an integrated LDBF sensor..

3.4.2.1 Frequency Response

Figures 3.52 and 3.53 show the AC response of the g_mC circuit, with varying bias current caused by varying an external resistance. Figure 3.52 shows that regardless of bias current the gain of the filter remains at approximately unity at low frequencies (simulations were performed with a 600 mV peak-peak AC signal, representing a typical AC magnitude after amplification by the HDA - unity gain therefore corresponds to -10 dBV).

Figure 3.52 shows bandwidth against bias resistor. From this it can be seen that the cut-off frequency can be widely controlled in the range required without using very large bias resistors. This potentially has noise advantages, as the use of large resistors increases thermal noise, which would affect all pixels due to the use of a common bias circuit.

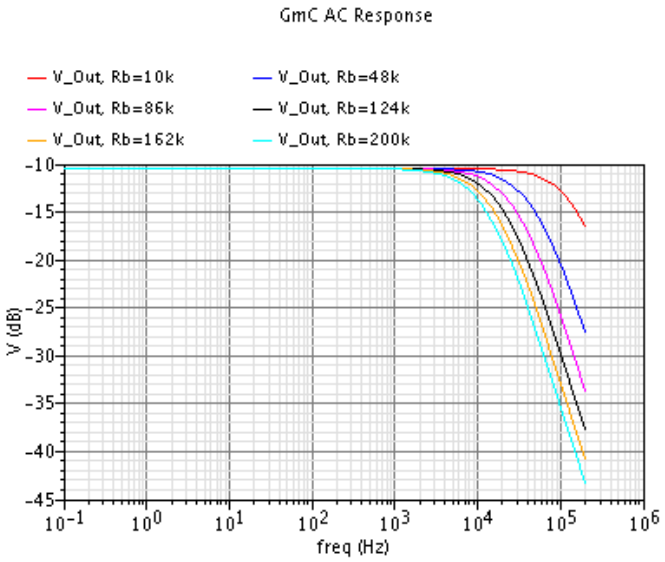

Figure 3.52: AC response of g_mC low-pass filter with varying bias current

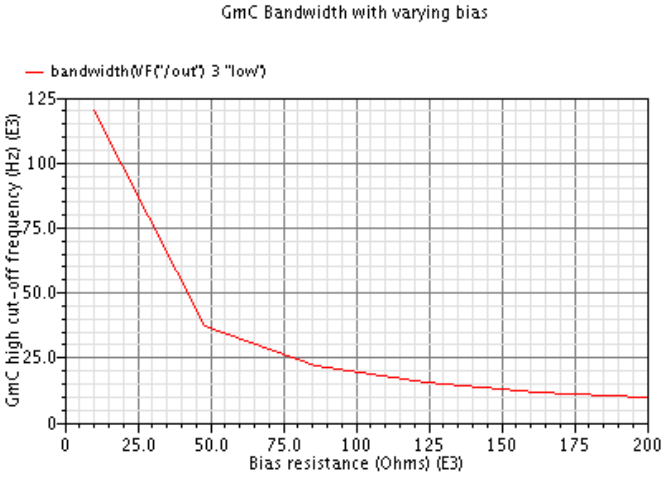

Figure 3.53: High-frequency cut-off of g_mC low-pass filter with varying DC bias current, shown by value of external resistor used to set bias

Figure 3.54 shows the bias current in the OTA over this range of bias points, showing that the current consumption of this circuit is very low, at less than 1 nA to achieve a 20 kHz cut-off frequency.

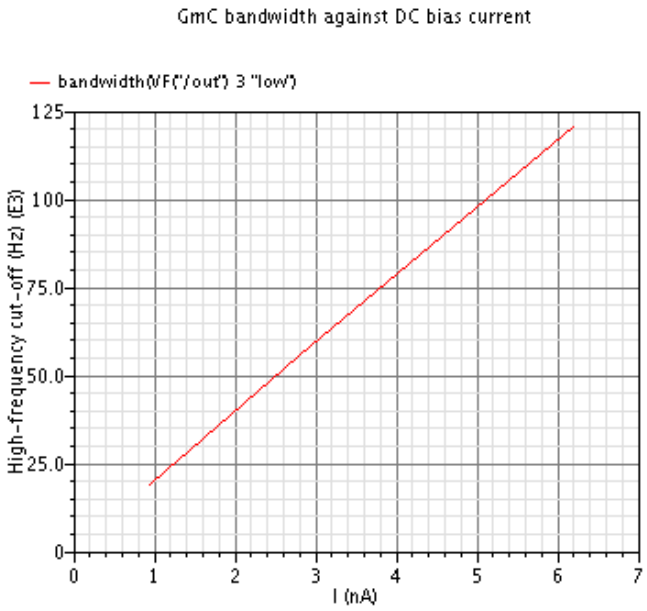

Figure 3.54: High-frequency cut-off of g_mC low-pass filter with varying DC bias current, shown by bias current in OTA

3.4.2.2 Transient Response

Figure 3.55 shows a transient simulation of the g_mC filter. It can be seen that the pass-band signal (5 kHz compared to a cut-off of around 20 kHz) is not significantly attenuated or distorted. Large signals or those at higher frequencies (for example, a signal of the same magnitude as that shown in Figure 3.55, but increased to 20 kHz) tend to be distorted by this filter, as the slew rate of the filter output is limited. This can cause the total harmonic distortion to increase from approximately 2% for the output signal shown in Figure 3.55, to 16% for the same magnitude signal at 15 kHz. As flow processing involves frequency weighting to give higher weighting to high frequency signals (caused by fast moving blood), this could cause incorrect results. However, most of the Doppler signal is significantly below the cut-off frequency so the signal levels at high frequencies should be small, and the harmonics caused by the distortion will be smaller still. Therefore, this is not thought to be a problem, but may need further consideration if the IC response found by testing with increasing flow is not as required.

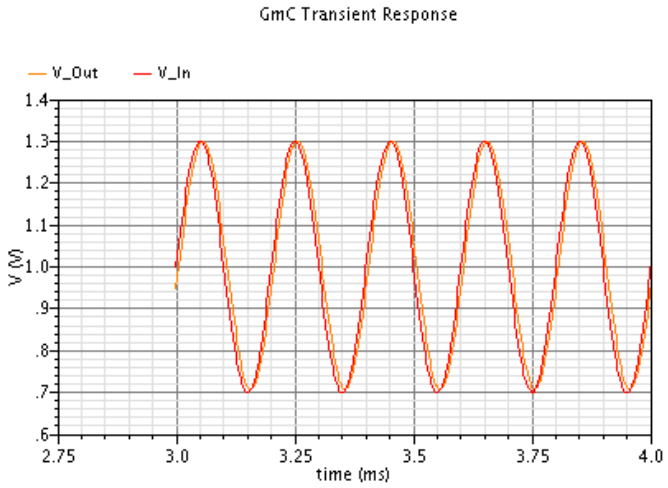

Figure 3.55: Transient response of g_mC low-pass filter to a 5 kHz 600 mV p-p signal

3.4.2.3 DC Response

Figure 3.56 shows the DC output of the g_mC circuit with rail-rail input voltages. It can be seen that the response is linear for all input voltages above 0.15 V, and that the DC gain is unity. The logarithmic response of the front-end means the input to this circuit is very unlikely to go below 0.15 V, even considering the level shifting of the HDA, so this response is suitable.

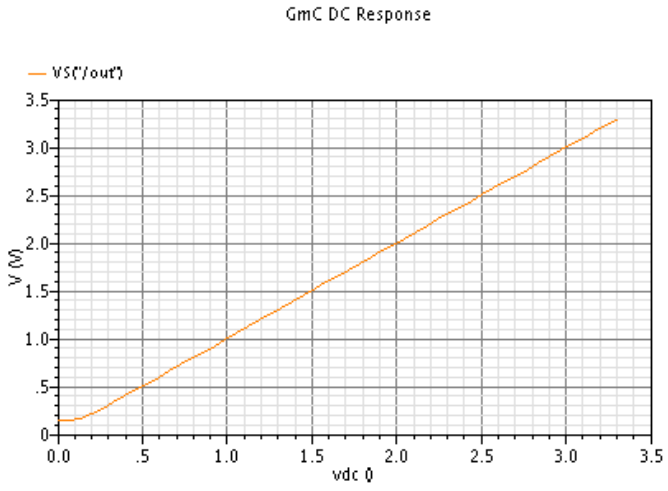

Figure 3.56: DC response of g_mC low-pass filter

3.4.2.4 Noise Spectrum

Figures 3.57 and 3.58 show the noise spectrum of the g_mC filter. It can be seen that the output or input referred noise density in the pass band is below $1.5\,\mu\text{V}/\sqrt{(\text{Hz})}$. Over a noise bandwidth of 30 kHz this equates to an RMS noise of 0.26 mV, compared to an expected signal level of several hundred mV. The unity gain of this filter in the pass-band means that the input and output noise density is expected to be the same. Above the pass band, the input-referred noise rises, but this is an artefact of the falling gain rather than an actual increase in noise. The actual output noise falls, as expected of a low-pass filter. At low frequency the noise density rises due to 1/f noise, and this is the same in the input referred and output noise spectra. The rise in noise here is fairly limited, and this is mostly below the cut-off frequency of the HDA. Signals at these frequencies will also be ignored by post-processing after a Fourier transform has been performed. This means the noise response seen here is acceptable.

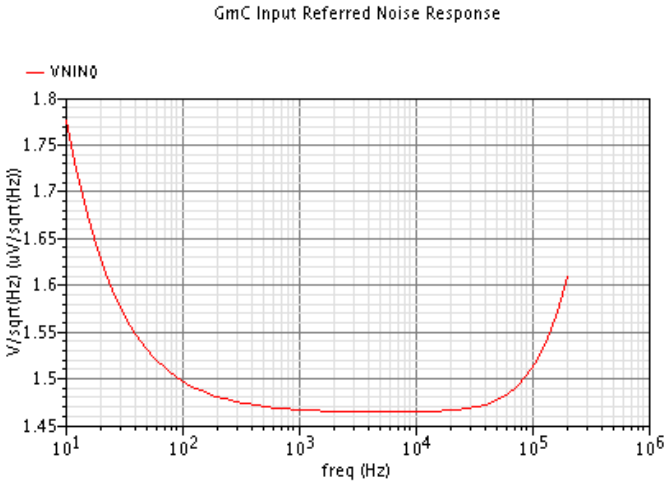

Figure 3.57: Input referred noise spectrum of g_mC low-pass filter

3.4.3 Summary of Low-pass Filter Circuits

The g_mC filter design presented in this section has been introduced, and its suitability for use in this application has been demonstrated through simulations, showing suitable frequency

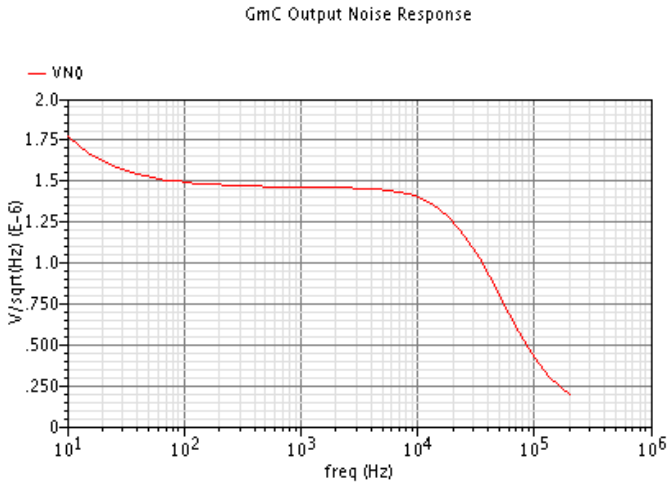

Figure 3.58: Output noise spectrum of g_mC low-pass filter

response which can be adjusted above and below the design frequency of 20 kHz if needed. The circuit also has a suitable DC response (unity gain down to ~0.2 V DC input voltage), low distortion and low noise. Given this adequate performance, and that this circuit is after the amplification stage and hence is not as critical to IC noise performance as other pixel elements, no alternative circuits have been investigated.

3.5 Summary

The basic principle of the logarithmic pixel has been introduced, and a number of variations on the basic pixel design that have improved performance have been described and simulated. The weaknesses of the basic logarithmic pixel design, mainly its limited bandwidth, require the use of a buffered pixel design. Simulations compared the basic pixel with two designs, one using a pair of transistors to buffer the photodiode AC voltage from the load, and one using a load biased by feedback through a CMOS inverter. The former design was shown by simulation to be preferable for use on the BVIPS1 IC.

Several weaknesses of logarithmic pixels were demonstrated, including the limited range of the normalisation principle, which does not occur outside the subthreshold region of operation. There are also non-linearities introduced by the HDA behaviour with larger circuits, and potentially from signals approaching the high cut-off frequency of the g_mC anti-aliasing filter. However, these problems are not sufficiently serious to prevent circuits such as this detecting Doppler signals.

Linear front-ends were discussed, with their advantages and disadvantages considered in terms of behaviour and implementation in a 1D array IC. While there are significant problems in implementing this type of front-end in applications where circuit physical footprint is critical, implementation on chip was shown to be feasible, and simulation results from a possible design were shown.

Active pixel sensors were also discussed, and while these may be suitable for some LDBF sensors, the large size of pixel required here results in a high capacitance relative to the photocurrent expected, and therefore the voltage changes seen are too small to ensure reliable detection of the signal, given that the conventional amplifiers and filters (i.e. continuous time circuits) cannot be used, given the sampling nature of the readout method of these circuits.

Two possible high-pass filters/amplifier circuits were shown, each being an implementation of a hysteretic differentiator circuit. Based on simulations of nominal performance and sensitivity to manufacturing variation, a HDA design using an OTA with a feedback network consisting of a MOS transistor and an 'inverted' CMOS inverter as a high resistance has been selected for use on the BVIPS1 IC.

To select a low-pass filter, a g_mC filter design as used on previous LDBF ICs was simulated. This design showed acceptable performance, and as this area is not critical to pixel performance (being after amplification of the signal, hence sensitivity to noise is reduced), no other designs were considered.

The pixel to be implemented on the BVIPS1 IC will therefore mainly consist of a buffered logarithmic current-voltage converter, a HDA AC-amplifier using an 'inverted-inverter' feedback network to amplify the signal and remove low-frequencies, and a g_mC low-pass filter for anti-aliasing and removing signals above the Doppler bandwidth. An additional array will consist of a small number of opamp-based linear pixels, in order to confirm that this type of pixel can be used as an integrated detector for LDBF.

Chapter 4

Design and Testing of BVIPS1 IC

4.1 Introduction

This chapter describes the first IC produced for the BVIPS project, known as BVIPS1. The pixel designs used on the IC are those described in the previous chapter. This chapter gives specific information on the implementation on this IC of each design, the higher level circuit design making up the overall IC, and how that IC will be used in an imaging system.

Testing to characterise the IC in terms of its gain, DC response, bandwidth and noise level is described and results are compared to simulated behaviour. Finally, the use of the IC in an LDBF imaging setup to measure changes in blood flow will be demonstrated.

4.1.1 Structure of the Chapter

The chapter begins with a top-level description of the IC produced, showing which circuits are included on the IC and showing the overall layout of IC, including the relative position and size of each circuit and the shared photodiode array. The circuits on the IC that are investigated here (i.e. those that are the work of the author) are:

- 64x1 array of buffered logarithmic pixels (main array)
- 32x1 array of buffered logarithmic pixels (duplicate array with additional test points and test analogue filter)

- 4x1 array of opamp-based linear pixels

Each of these circuits are then described in greater detail, including a block diagram of each circuit and a layout plot of one pixel within the circuit. Two other circuits shown on the top level layout plot are not the work of the author (although integration into the overall layout was), hence are not described here. These circuits are an 8-pixel array of linear transimpedance pixels

Characterisation of the IC is described, beginning with a description of the equipment and method used. For each circuit, measurements are made over a range of input photocurrents (i.e. range of illumination intensities) for DC voltage output, AC gain, noise performance and frequency response. Measured response in each case is compared with simulated IC behaviour. As increased noise was seen in the fabricated IC compared to simulations, some efforts made to reduce noise are shown.

A further stage of testing shows the use of the three circuits to measure flow in a series of test targets, both biological and artificial. The equipment setup and method used for this testing is shown, along with results for each of the three circuits for various flow targets. Where applicable, single point measurements are shown (a line plot of flow over time) along with 2D colour plots showing flow at all pixels of an array over time. From these measurements, SNR figures are produced for the processed flux output.

The behaviour shown from the above testing is then discussed for each separate circuit. For the logarithmic pixels, the testing shows a number of problematic aspects of device performance, including signal artefacts, increased noise compared to simulations, and high fixed pattern noise. A series of additional tests are performed to investigate these issues in order to address them in the design of the second prototype IC, and these additional tests and results are shown in section 4.6.

4.2 BVIPS1 IC Design Overview

The IC produced is a two sided device, in that it includes a single 64x1 array of photodiodes in the centre, with two separate sets of front-end and processing circuits - one on each side of the chip. The right hand side of the chip consists of an array of 64 identical pixels, intended for use in a complete line imaging system. The left hand side of the IC contains a number of

prototype circuits, including variations of the right-hand side design and alternative front-end designs by other designers. Four opamp-based linear front-end pixels are included to test a real implementation of the linear design introduced in the previous chapter. These left-hand side designs are included to test circuits elements with a view to inclusion on the second prototype IC.

4.2.1 Layout of IC

Figure 4.1 shows the layout of the entire IC. The large area in the centre is the 64 photodiodes, each of which has a size of $1000 \times 50\,\mu$m (pixel width $\times$pixel pitch) , making the total light sensitive area for 64 sensors $1000 \times 3200\,\mu$m. The high width of each photodiode is required to ensure that light reflected from the target does not wander off the photodiodes during scanning. With a narrower sensor, as the beam is mechanically scanned changes in optical path, reflection angles and imperfections in the mechanical scanning itself can cause the location of the image produced at the sensor to move such that the reflected line misses the photo-sensitive area. annotations

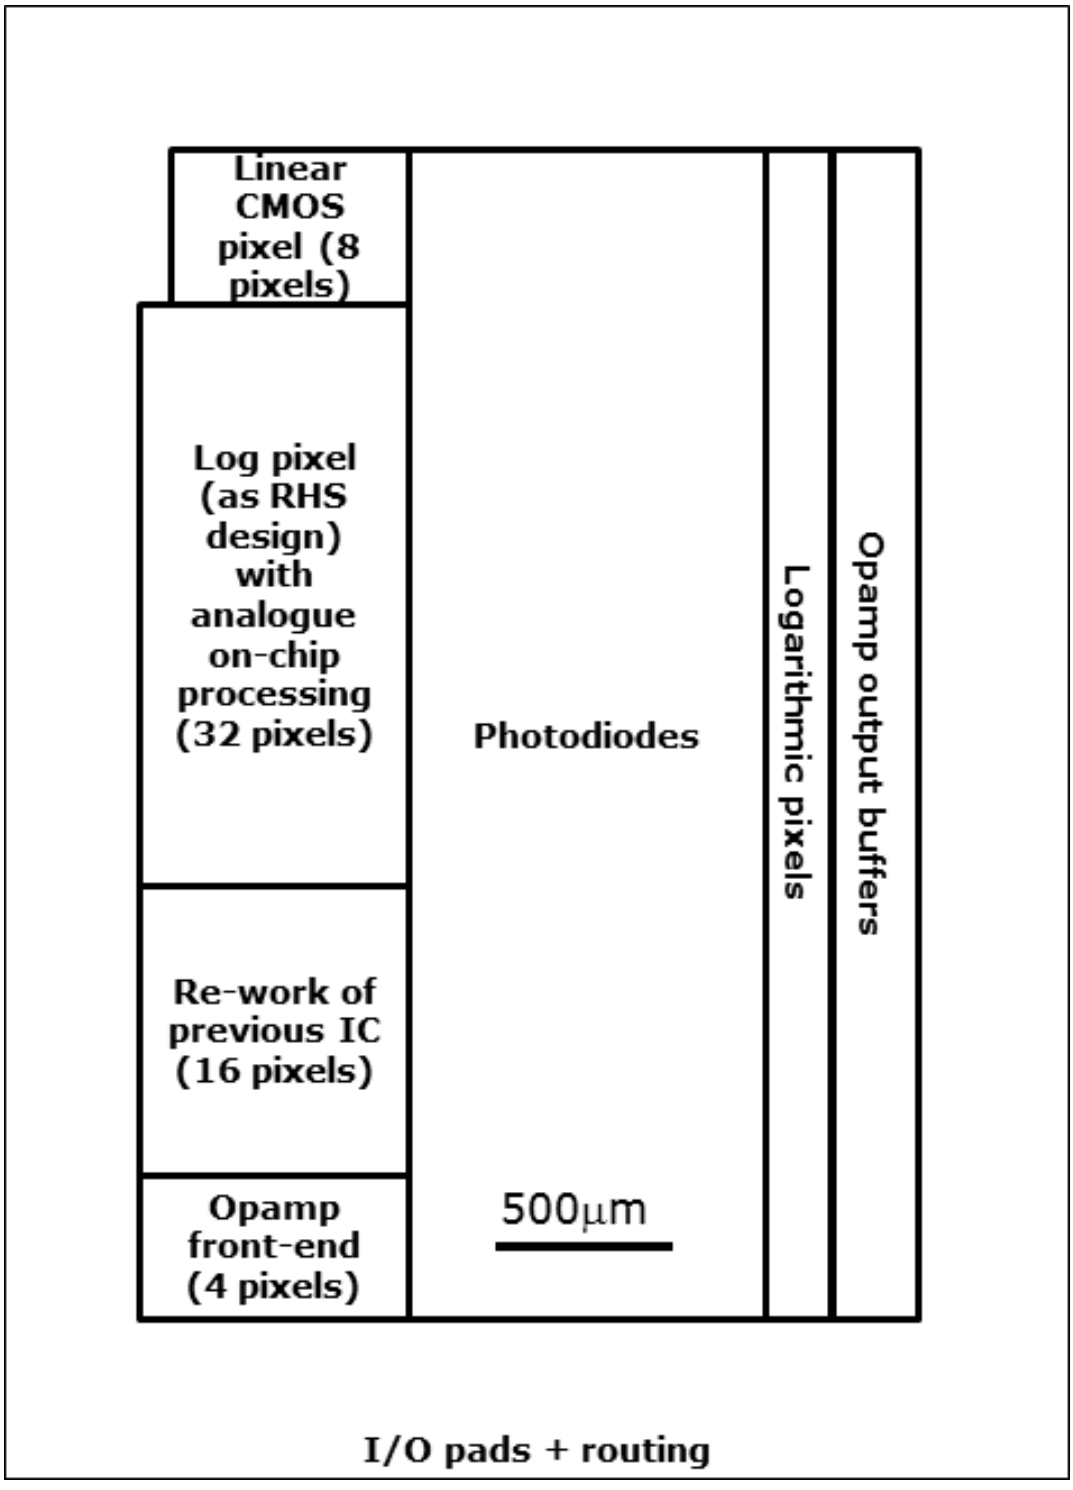
Linear CMOS pixel (8 pixels)
Log pixel (as RHS design) with analogue on-chip processing (32 pixels)
Re-work of previous IC (16 pixels)
Opamp front-end (4 pixels)
Photodiodes
500μm
Logarithmic pixels
Opamp output buffers
I/O pads + routing

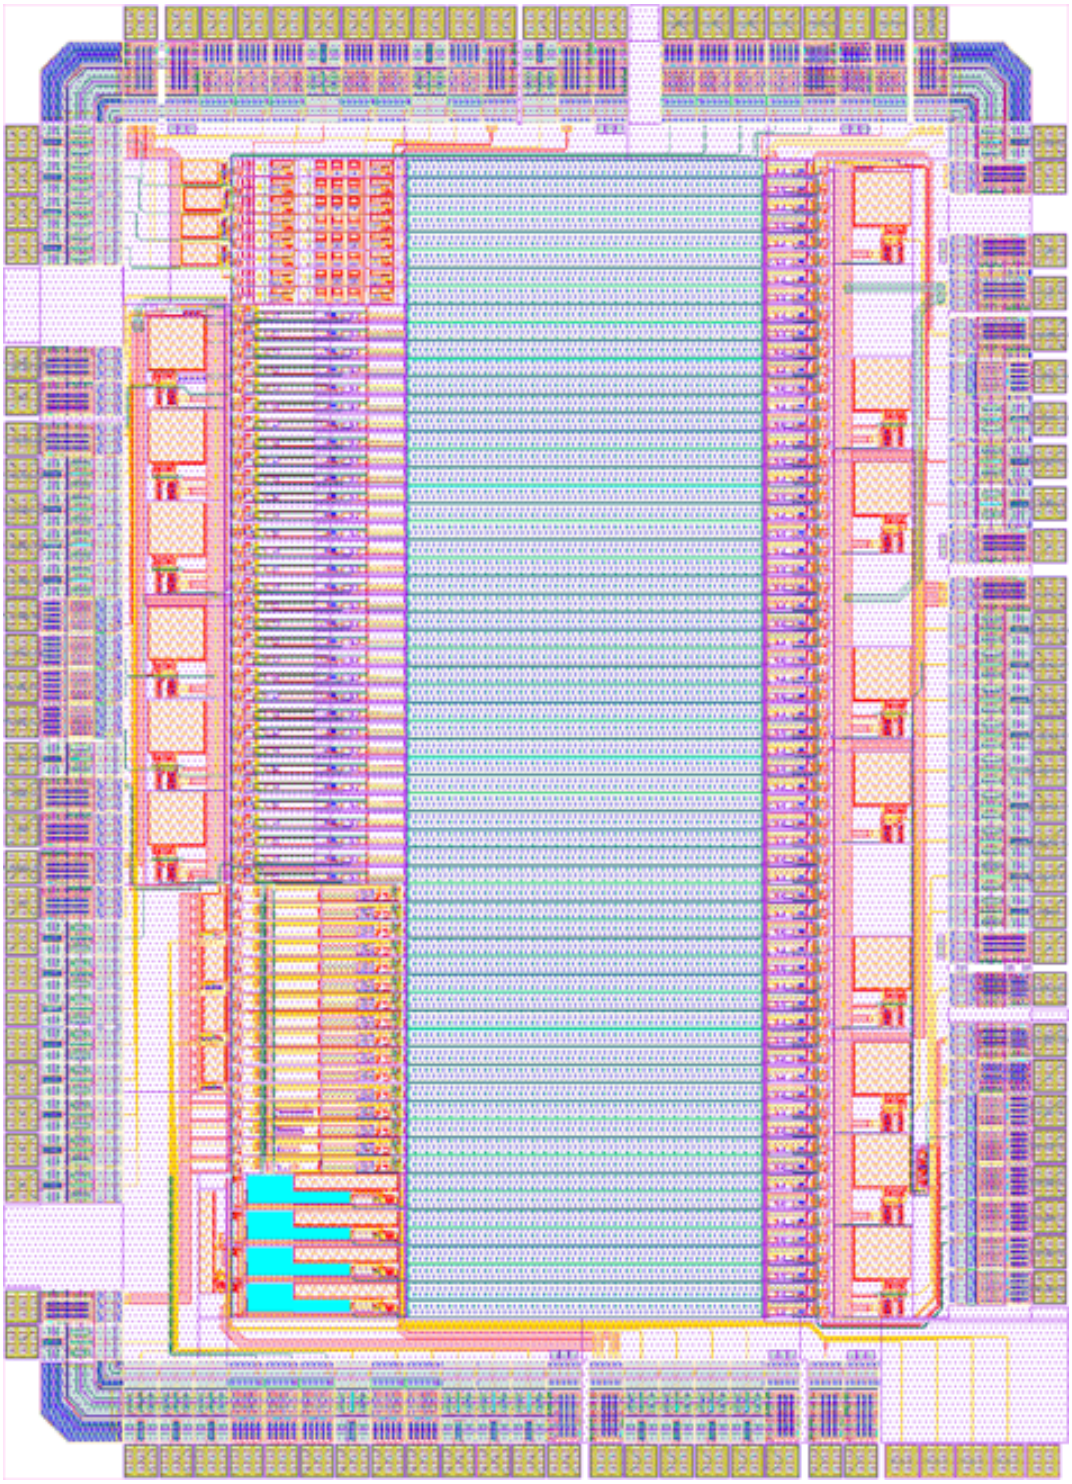

Figure 4.1: Layout of BVIPS1 IC

4.2.1.1 64× 1 Logarithmic Pixel Array (Right-Hand Side)

To the right of the photodiode array are the logarithmic pixels making up the 64×1 array intended as the 'main' array of this IC, which would be used by a line scanning instrument using this IC as a sensor. Each of these logarithmic pixels is $140 \times 50\,\mu$m. Multiplexers and address decoders consume another $55 \times 50\,\mu$m, and there is a further $220 \times 50\,\mu$m for each of the nine output buffers. These components are large due to the requirement for high drive capability for off-chip loads.

The design of these logarithmic pixels is described in more detail in Section 4.2.2.

4.2.1.2 32×1 Logarithmic Pixel Array (Left-Hand Side)

To the left of the photodiode array, towards the centre of the IC, is a second array of logarithmic pixels. This array uses components very similar to the 64×1 array, with some changes to layout to better fit the available space. These pixels have an additional frequency-weighted analogue filter stage, included after the duplicated components. The analogue processing technique is not investigated in this thesis, instead a digital processing method is used throughout. However, this array also provides visibility of intermediate signals that are not output from the main 64×1 array. This allows the behaviour of the front-end itself (i.e. not including the HDA AC amplifier or on-chip low-pass filter) to be measured directly.

The design of these pixels is described in more detail in Section 4.2.3.

4.2.1.3 4×1 Opamp-Based Linear Pixel Array (Left-Hand Side)

Four identical opamp pixels are included to investigate the use of linear front-ends in CMOS LDBF sensors, based on the design introduced in Section 3.2.5.

The opamp pixels are the four pixels in the bottom left of the IC. This section of the IC has a size of $580 \times 400\mu$m, including the front-end, multiplexers and gain stage/output buffer. The pixel pitch used here is $100\,\mu$m rather than the $50\,\mu$m pitch of the logarithmic pixels. This is due to the large size of the feedback resistor required to achieve sufficient gain to detect LDBF signals, which makes designing pixels with a $50\,\mu$m pitch within the area available impossible. The photodiode sharing (between left- and right-side front-ends) means that the photodiodes themselves cannot be made larger, so the opamp pixels connect to photodiodes

1,3,5 and 7. Pixels 0,2,4 and 6 are not connected to a front-end circuit on the left hand side. Instead, they are connected to a bare pad, which if necessary allows the photodiodes to be connected to an off-chip circuit to characterise the photodiode itself.

The design of the linear pixels is described in more detail in Section 4.2.4.

4.2.2 RHS (64x1) Logarithmic Pixel Section

This section describes the circuit design, operating principles and layout of the logarithmic pixels on the right-hand side of the BVIPS1 IC.

4.2.2.1 Circuit Design and Operation

A block diagram of the right hand side (logarithmic pixel) section of the IC is shown in Figure 4.2. The main feature of this design is the 64 logarithmic pixels. Each pixel consists of a buffered logarithmic front-end, a HDA AC amplifier and a g_mC anti-aliasing filter. Outputs from the pixel are taken from before and after the anti-aliasing filter. These signals are then multiplexed in blocks of 16 to produce eight output signals, four from before the filters and four from after the filters. Multiplexing in this manner allows four external ADCs from an existing system to be used to simultaneously sample four pixels, increasing the readout speed or reducing the speed required of the ADCs. Visibility of intermediate points in the pixel (i.e. the outputs from before the anti-aliasing filters) allows more information to be gathered on the operation of the IC, such as allowing characterisation of the filters in isolation of the optical front-end.

In order to allow the use of a single ADC the four post-filter outputs are multiplexed again to produce a single output channel. In a final system it is likely that this will be the main output of the IC.

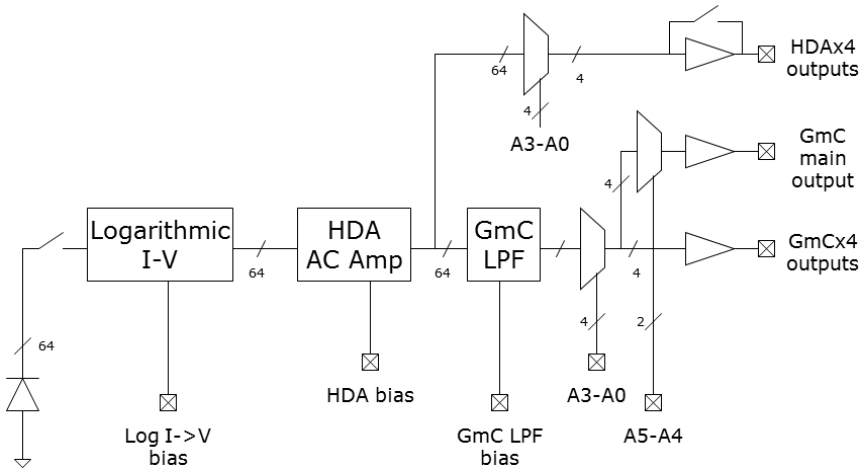

Figure 4.2: Block diagram of 64x1 RHS section of BVIPS1 IC

An external bias pin allows control of the high frequency cut-off of the anti-aliasing filters. Additional bias pins give some limited control of gain in the pass band by adjusting the AC amplifier (HDA) gain and the front-end bias. An op-amp buffer bias pin (not shown) allows the bias current of these buffers to be increased, giving greater bandwidth (i.e. reduced switching times for higher frame rates) at the expense of higher power consumption.

The output pin after the HDA can be used as an input, allowing external signals to be applied to check operation of the g_mC filter, the multiplexers and the output buffers. This bi-directionality is achieved by using transmission-gate switches controlled by an external signal to bypass the HDA output buffer. Additional switches controlled by the same signal (not shown) isolate the outputs of the output buffer and the HDA itself from the input signal.

4.2.2.2 Pixel Layout

Figure 4.3 shows the layout of an individual logarithmic pixel on the right-hand side. The logarithmic front-end is $22 \times 14\,\mu$m including the source-follower output buffer. To the right of that is the OTA used in the HDA AC amplifier, which is $18 \times 19\,\mu$m. The capacitor in the HDA is composed of two large transistors which have gate dimensions of $40 \times 16.5\,\mu$m and $56.75 \times 16.5\,\mu$m. The large capacitor size is used to achieve the low cut-off frequency required.

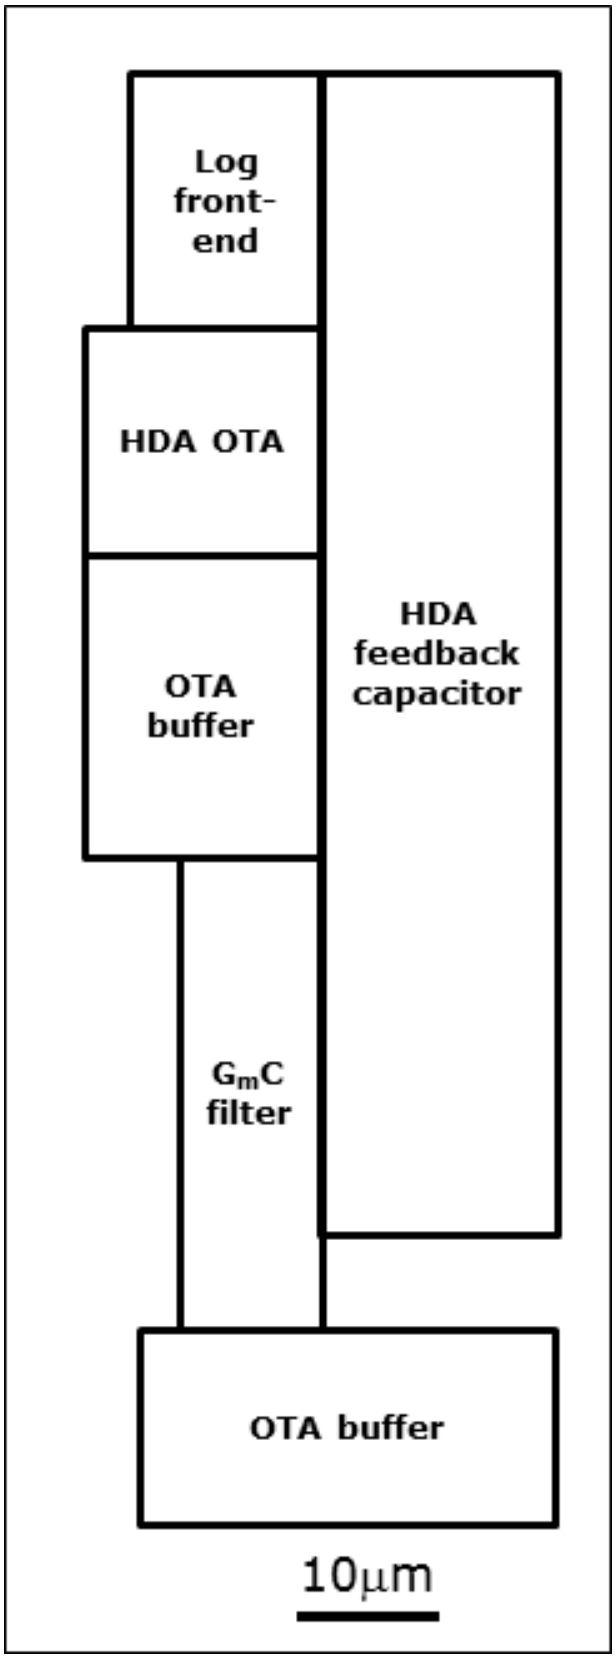
Log front-end
HDA OTA
OTA buffer
HDA feedback capacitor
G_mC filter
OTA buffer
10μm

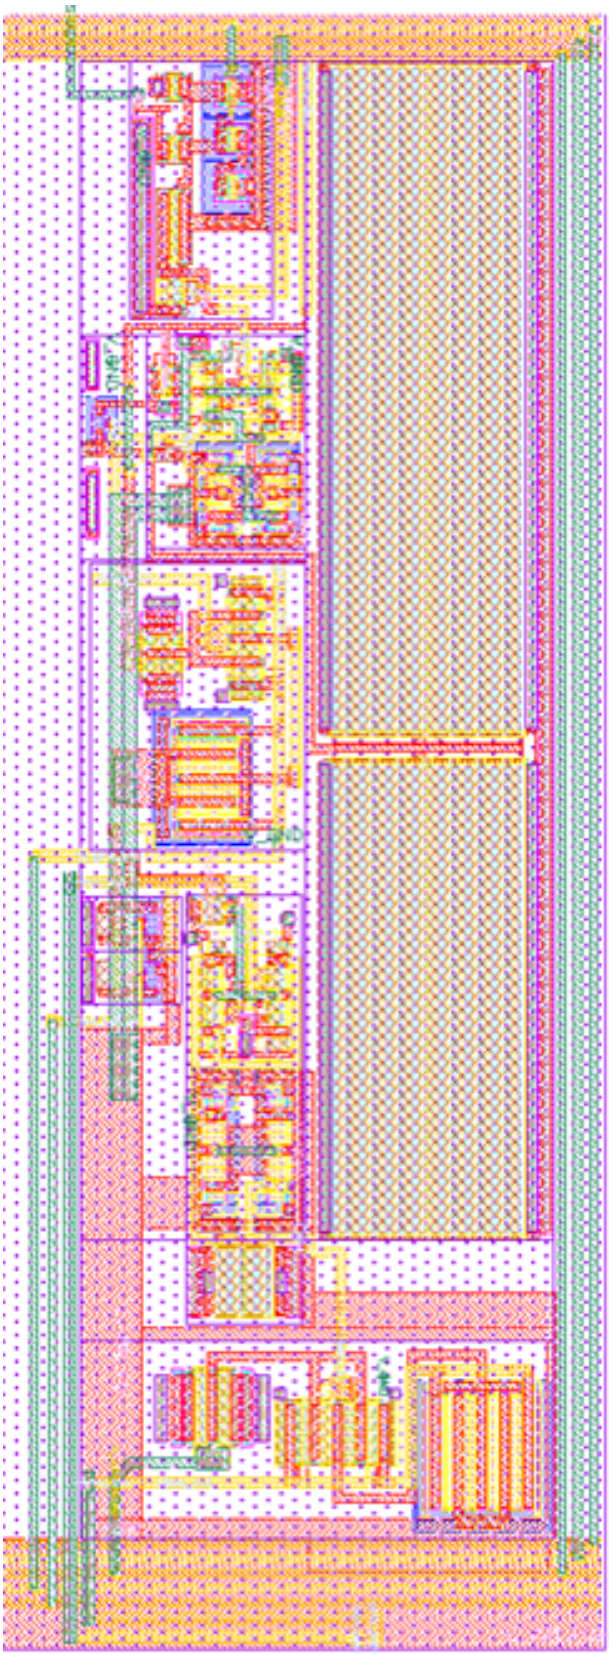

Figure 4.3: Layout of a single RHS logarithmic pixel (not including photodiode) of BVIPS1 IC

The HDA is followed by an OTA buffer used to isolate the HDA from load capacitances, with a size of approximately $24 \times 18\,\mu$m. The g_mC low-pass filter consumes an additional $36 \times 10\,\mu$m.

The final circuit on the right of the pixel is an additional OTA buffer, used to isolate the g_mC from the load capacitance presented by the final opamp output buffers, which have high device size to allow sufficient drive capability for off-chip loads. This buffer is $16 \times 34\,\mu$m.

4.2.3 LHS (32x1) Logarithmic Pixel Section

This section describes the circuit design, operating principles and layout of the logarithmic pixels on the left-hand side of the BVIPS1 IC.

4.2.3.1 Circuit Design and Operation

A block diagram of the left hand side (logarithmic pixel) section of the IC is shown in Figure 4.4. The main difference is the additional processing stage after the G_mC filter, which is a frequency-weighted filter (FWF). This is added to allow investigation of analogue processing methods, as the frequency weighting performs part of the flow calculation method required of a LDBF system. For this work, however, digital methods are used throughout and so the frequency weighted filter is not used. Instead, the output is measured at the three previous steps. This includes an output directly from the front-end circuit, allowing the operation of the individual elements of the logarithmic pixel to be investigated. These extra outputs are another of the differences between the left- and right-hand side, along with the reduced number of pixels.

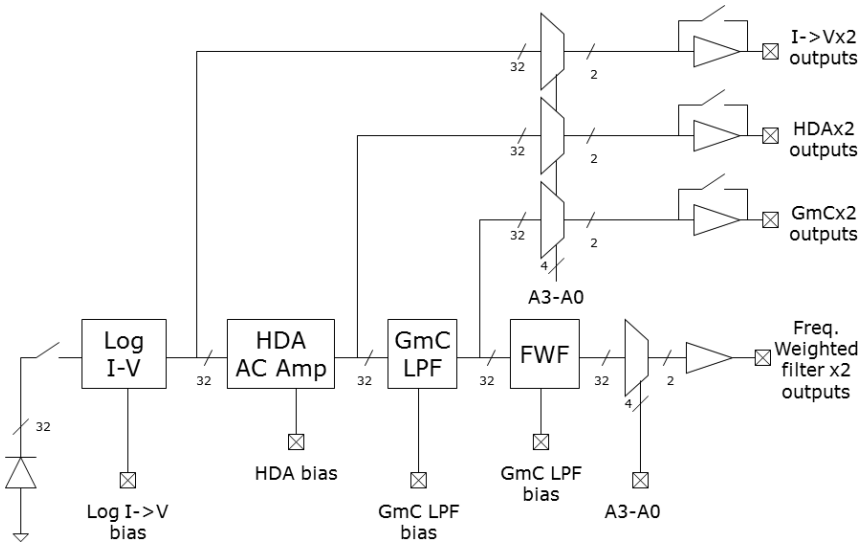

Figure 4.4: Block diagram of 32x1 LHS section of BVIPS1 IC

4.2.3.2 Pixel Layout

Figure 4.5 shows the layout of the left-hand side logarithmic pixels. The main change to the design is that the capacitor in the HDA is moved to be between the pixel circuitry and the photodiode (which is to the right of the layout shown in Figure 4.5). The other circuits used are unchanged, and the active parts of the layout are a mirror image of the design on the RHS. The changes were made to accommodate the RHS design within the space available on the left of the IC, which is subject to additional constraints due to the other circuits present in this part of the IC. Mainly, the multiplexer location being set by the use of shared address lines, and the bias lines being required to run horizontally, as the vertical routing used on the RHS is blocked by the other pixel designs. This results in a gap between the photodiode and the front-end, into which the HDA capacitor is moved to create extra space for routing of bias lines and the front-end power supply.

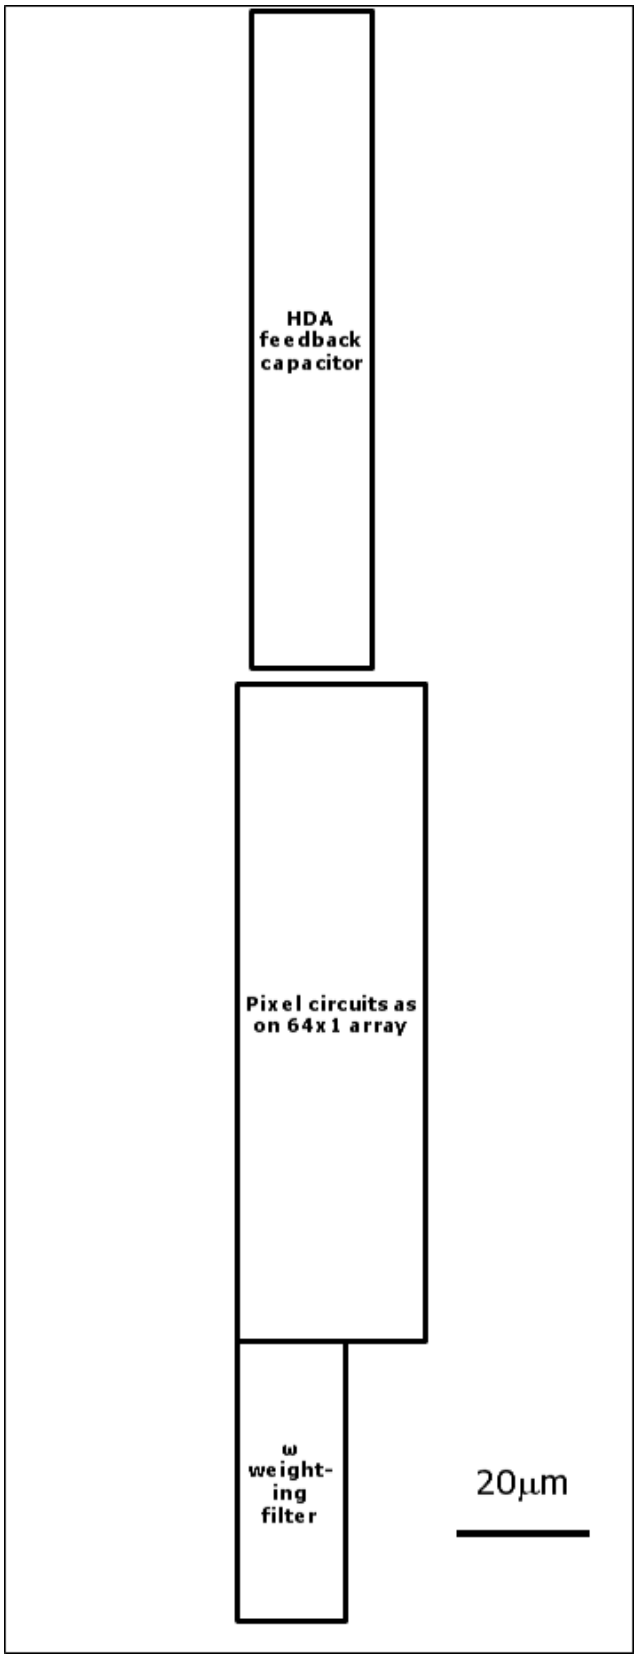
HDA feedback capacitor
Pixel circuits as on 64x1 array
ω weight-ing filter
20μm

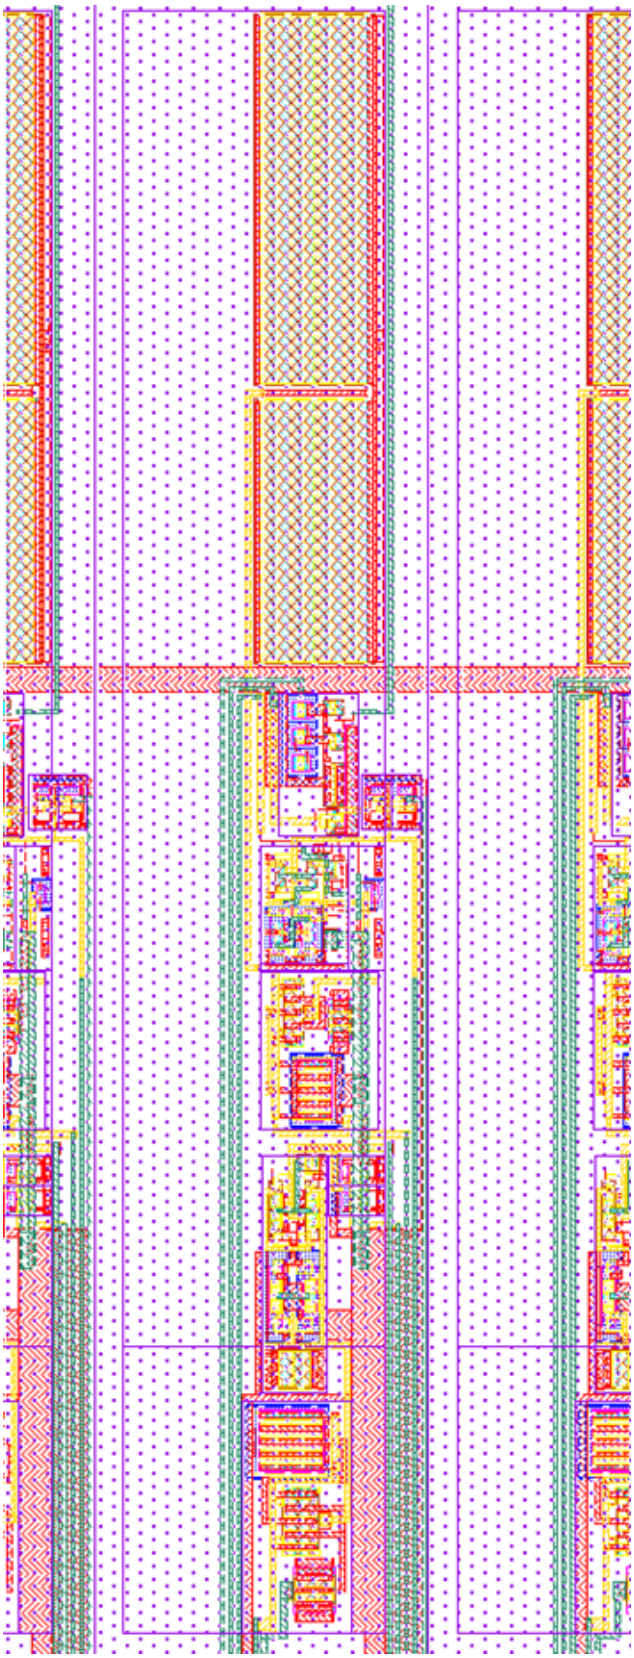

Figure 4.5: Layout of a single LHS (32x1 array) logarithmic pixel (not including photodiode, including edge of adjacent pixels) of BVIPS1 IC

4.2.4 Opamp-Based Linear Front-end Section

This section describes the circuit design, operating principles and layout of the linear pixels on the BVIPS1 IC. The schematic of these pixels is shown in Figure 4.6, in Section 4.2.4.1.

Op-amps are not traditionally used in CMOS imaging applications as op-amps are generally larger than logarithmic or APS pixels, and draw higher current, and hence are not well suited to arrays of sensors on ICs where space is limited. However, the linear array used in the line scanning application means that the op-amp pixels are a more feasible option than they would be for a 2D imaging array, where the front-end circuit size has to be kept to a minimum to allow high pixel fill-factors.

The main advantage of the op-amp front-end is its linear response, producing a larger output voltage for higher photocurrents, and hence larger signal when light power is increased. This can make some aspects of processing easier, for example identifying the background in an image where less light is reflected and hence the photocurrent is lower. With a linear pixel, the low photocurrent results in a low DC and AC voltage, whereas with a logarithmic pixel the DC voltage rises, but the AC gain rises such that noise can give an AC signal level similar to that from blood flow. Conversely, the linearity means that the op-amp pixels do not perform the normalisation inherent to logarithmic pixels.

4.2.4.1 Circuit Design and Operation

The opamp front-end circuit consists of four opamps used as current to voltage converters, a 4:1 multiplexer and a single op-amp based gain stage. The gain stage is required as the gain of the front-ends is set by the size of the feedback resistor used (i.e. transconductance = $R_{feedback}$). A large gain requires a large feedback resistor, which makes the pixel size and hence chip cost high, as well as increasing thermal noise levels. For a 2D array this would also significantly reduce fill-factor, although the linear array layout here means the resistors sit to the side of the photodiodes themselves, so this is not a major consideration for this IC. Using a lower front-end gain followed by an additional gain stage allows high gain without a proportionate increase in resistor size.

Figure 4.6 shows the schematic of the opamp based front-end circuit, with four front-end circuits followed by a single gain stage. One of the key features of this design compared to the opamp front-ends implemented using discrete components in the Moor LDLS imager

is the offset voltages applied to the non-inverting input of the front-end opamp and the inverting input of the gain stage op-amp. These voltages are required as the IC does not have a $-V_{SS}$ rail, and hence signals cannot be centred on, or allowed to drop below, 0 V. The offset voltages are applied to the virtual ground at the input pins of each opamp, such that the differential voltage at the input of each opamp is between signal and V_{offset} rather than between signal and ground. As the offset voltage is set slightly lower than the minimum signal expected, the offset voltage means that the difference between the two is kept to a minimum and a larger increase in photocurrent (and hence voltage) can occur before saturation of the opamp output.

These voltages depend on photocurrent. For the gain stage in particular, the changes in DC input voltage caused by changes of photocurrent will require the offset voltage to be changed to suit the likely range of DC photocurrents.

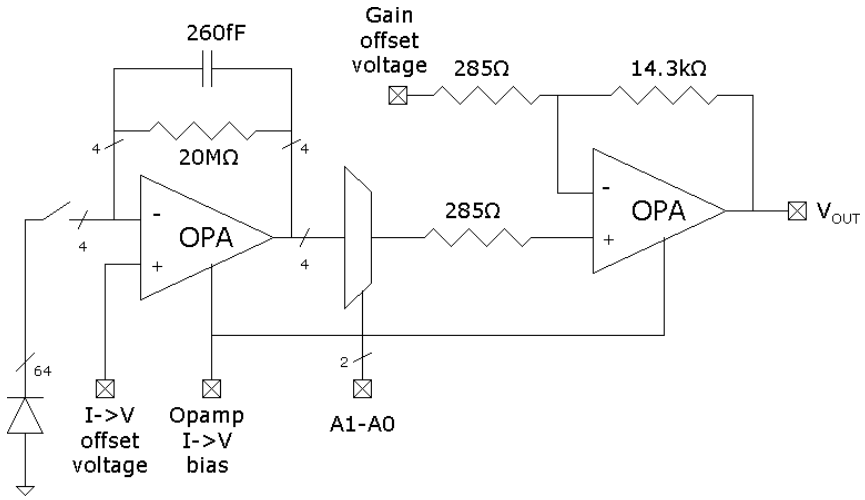

Figure 4.6: Block diagram of 4x1 op-amp front-end section of BVIPS1 IC

4.2.4.2 Pixel Layout

Figure 4.7 shows the layout of the opamp pixel section of BVIPS1. The transistors within each opamp occupy an area of $130 \times 40\,\mu$m, and the compensation capacitor consumes an additional $295 \times 45\,\mu$m. The overall front-end is $100 \times 440\mu$m, with the additional area mostly taken up by the feedback resistor. The feedback capacitor is fairly small compared to the other circuit elements, at $75 \times 9\,\mu$m.

Address decoders and multiplexers are placed between the front-end and the gain stage, adding $40\,\mu$m of width, followed by a single opamp used as a gain stage and output buffer,

which adds an additional 90 μm of width.

The opamp used here is smaller than the design used as an output buffer in the logarithmic pixel section (described in Section 4.2.1.1). The opamp used in the linear front-end circuit is inferior to those used in the logarithmic pixel section in terms of stability (i.e. phase margin) and maximum load capacitance, but the increased size of the larger opamps makes them too large to fit in the space available for the opamp pixels. The smaller design has been used on previous ICs, and has been shown to operate successfully, albeit at lower switching speeds than made possible by the larger opamp buffer design [He et al., 2009].

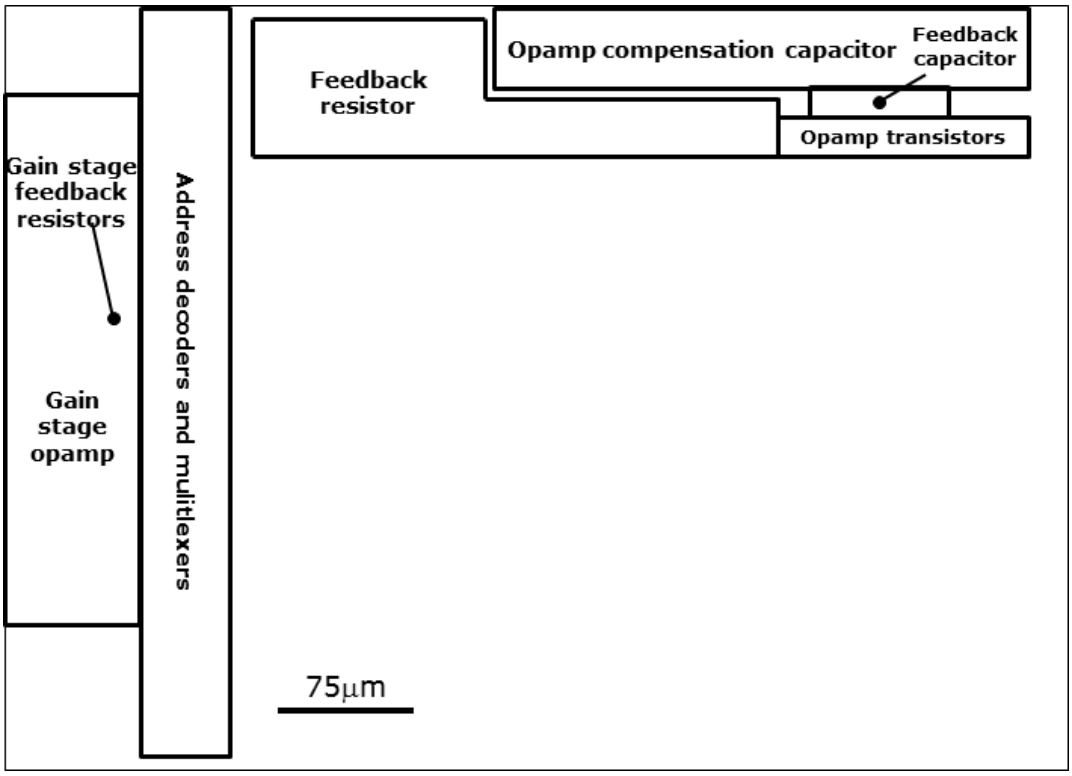
Opamp compensation capacitor
Feedback capacitor
Feedback resistor
Opamp transistors
Gain stage feedback resistors
Address decoders and mulitlexers
Gain stage opamp
75μm

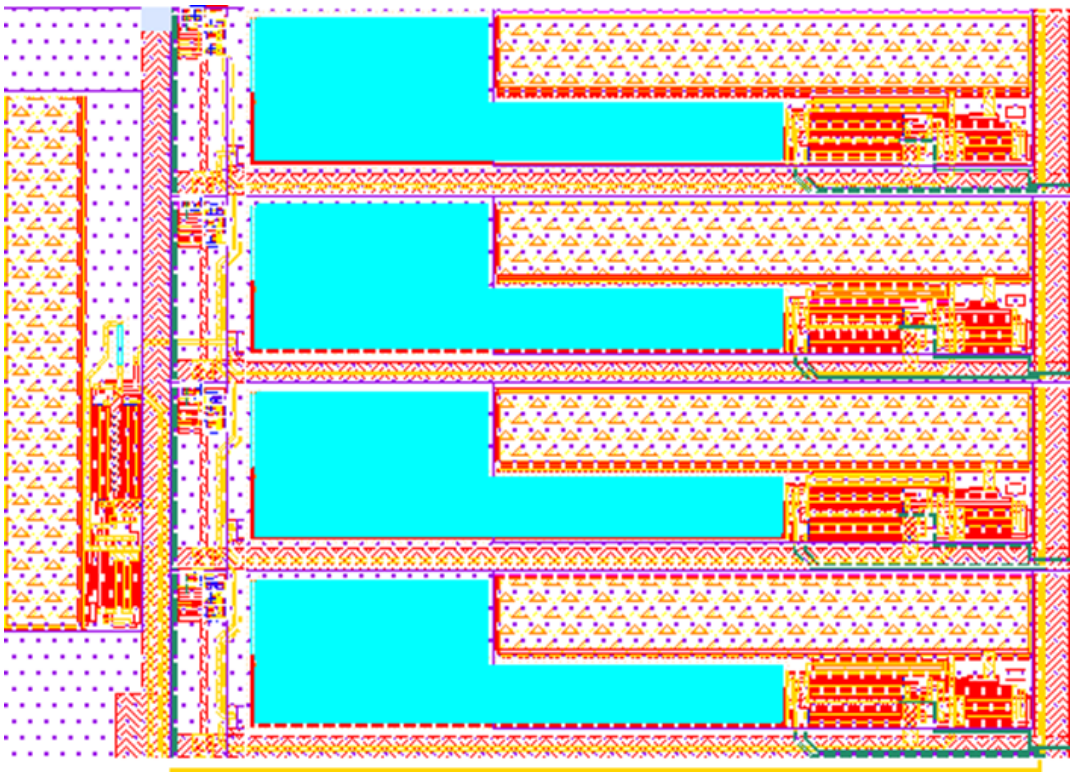

Figure 4.7: Layout of the four opamp pixels (not including photodiode) of BVIPS1 IC

4.3 Characterisation of BVIPS1 IC

This section describes characterisation of the BVIPS1 IC arrays (64x1 log pixel main array, 32x1 log pixel secondary array, 4x1 opamp linear pixel array). This first set of tests used a modulated laser to measure DC response, AC response, bandwidth and noise. Comparisons are made between measured and simulated IC behaviour. Further tests shown in Section 4.5 will show tests performed on the IC using a Doppler imaging setup.

4.3.1 Characterisation Equipment

Figure 4.8 shows the equipment used to perform the testing. For characterisation purposes, rather than illuminating a target with a laser and then imaging the reflected light onto the sensor, the IC was directly illuminated by a low-power laser. This allowed the light incident on the IC to be controlled more accurately and simply than in a reflection setup. The light from the laser was first passed through a beam expander, producing a 5 mm diameter beam allowing the entire sensor area to be illuminated with equal intensity light. After the beam expander a beam splitter was used to direct half of the total light onto a reference photodiode. This was done to give a separate measure of the light power incident on the IC sensor, based on the illumination intensity detected by the reference detector with known gain, responsivity and sensor area.

Modulated illumination was provided by driving the laser with an external signal via a signal generator. This allowed light with a high constant component and 10-20% modulated component to be produced, resulting in a similar modulation depth to that expected of signals from LDBF.

The signals from the IC and reference photodiode were captured using an oscilloscope, allowing up to four signals (i.e. the reference and 3 IC outputs) to be observed simultaneously. One IC output was also captured using an FPGA based data acquisition system, which was also used to drive the IC address lines when capturing data from multiple pixels. The system used for this was based on the Moor LDLS system, consisting of a prototype form of the same FPGA and ADC designs. This system is then linked via USB to a PC running Matlab for data visualisation and storage.

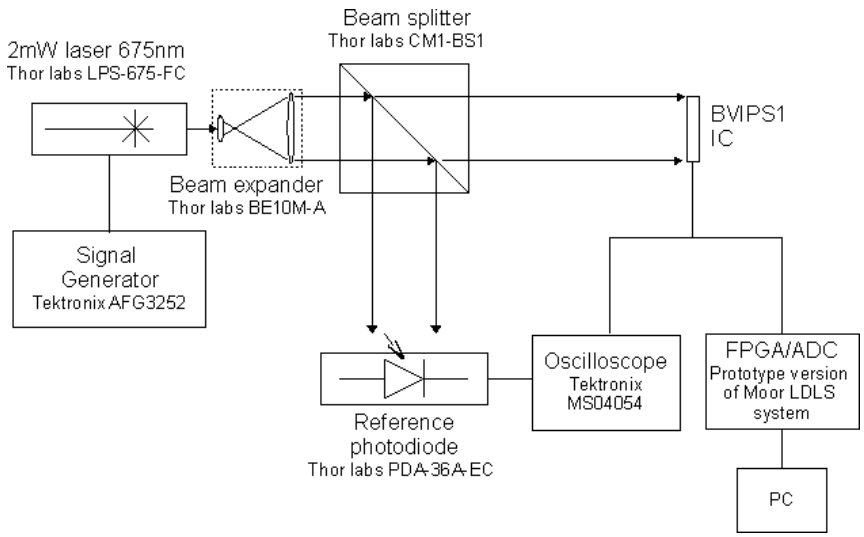

Figure 4.8: Equipment setup for characterisation

4.3.2 Behaviour Tested

The characterisation process looked at five aspects of circuit behaviour:

- DC response
- AC gain
- Noise (including measured voltage noise and input referred current noise)
- Frequency response
- Fixed pattern noise (i.e. variation in circuit behaviour between different pixels within the array)

The results shown here are grouped by these tests, rather than by circuit, allowing comparison of each type of behaviour between the different outputs of the IC. The different outputs shown here are for the logarithmic front-end, the g_mC (after amplification and filtering of the signal) and the opamp linear front-end.

The measured results are an average of the response for all pixels, for three different ICs from the same batch, reducing the susceptibility to chip-chip variation. The error bars on the measured results represent one standard deviation above and below the mean.

Measured results are compared with simulated results, with simulations being performed in Cadence using the same input conditions (photocurrent, bias currents and power supply voltage) as those used during IC characterisation.

4.3.3 Light Levels for Characterisation

In order to determine the illumination power to use for characterisation, measurements were taken of the light level incident on the detector (i.e. illumination intensity) in the Moor Instruments Laser Doppler Line Scanner [MoorInstruments, 2009]. This is the commercial vein imager in which the BVIPS IC is intended to be used, replacing 64 discrete component front-ends. These measurements were carried out to confirm the light levels expected in normal use. While the light level has been previously calculated as corresponding to a DC photocurrent of approximately 1 nA in Section 1.8.1, measurements using the imager (which was not available when the initial estimates were made) are likely to give a more accurate estimate.

The readings are not exact, as the existing discrete photodiodes are larger than those on the IC. However it is assumed that given appropriate modification of the imaging optics in the LDLS (such as a shorter focal length lens to focus the same length line onto a smaller detector) that a similar amount of light will be captured by both systems.

IC photocurrent is calculated by measuring the power of the laser line incident on the Moor LDLS sensor. This was done by removing the original sensor and front-end circuits and using a power meter (Thor labs PM120) in its place. The power meter sensor is circular and has a 9 mm diameter. The LDLS detector is 23 mm long so this power is scaled up to give total power in the line. This is then scaled by the ratio of the IC array length (3200 μm) : LDLS array length (to give the power incident on the IC array). This is divided by 64 to give the incident power per pixel, which is then multiplied by the IC responsivity of 0.3 A/W [Kongsavatsak, 2005] to give pixel photocurrent.

Target	P_{meas} (μW)	P_{line} (μW)	$P_{IC\,Array}$ (μW)	$P_{IC\,Pixel}$ (nW)	$I_{DC-IC\,Pixel}$ (nA)
Black Card	2.4	6.2	0.88	13.8	4.13
Hand (i.e. entire line on skin)	18.0	46.0	6.56	103.0	30.8
Finger (i.e. only small part of line on skin)	6.0	15.3	2.19	34.2	10.3
White Card	21.2	54.2	7.73	121.0	36.2
Calculation method	P_{meas}	$P_{Meas} \times \frac{23}{9}$	$P_{line} \times \frac{3.2}{23}$	$\frac{P_{IC\,Array}}{64}$	$P_{IC\,Pixel} \times 0.3$

Table 4.1: Light power measurements from Moor LDLS and corresponding predicted IC photocurrent for $50 \times 1000\,\mu$m pixel (all taken using constant intensity illumination - no modulated source)

The photocurrents obtained using this method, shown in Table 4.1, are significantly larger than those estimated from previous measurements in Section 1.8.1. The estimated values were around 1 nA for a typical situation, from which a range of 75 pA - 1.5 nA was chosen for initial simulations. For the measurements taken using the Moor LDLS shown in Table 4.1, the typical situation (imaging a hand) results in a photocurrent of approximately 30 nA. The range of incident light power estimated from these measurements is approximately 4 nA to 35 nA, representing the change from a dark, matte target (black card) to light reflected from a more reflective, light coloured surface (white paper). This represents an increase by a factor of 30 in the typical photocurrent. The range of photocurrents is larger, at 30 nA compared to 1.5 nA from the initial estimates, but this is a smaller fraction of the lower limit.

This increase in expected photocurrent may be due to an increase in laser power, or an error or poor assumption in the scaling used in the original calculations. The estimate was based on a previously measured reflected optical power density for a single point system, which was then scaled up by the increased laser power (7 mW to 45 mW), then divided by 64 to account for the change from a single point system to an array detector. The most likely flaw in this method is that it is assumed that the optical power density is equal at all angles of reflection, which may not be the case. While the reflection of light from skin appears to be isotropic/lambertian rather than specular, the reality is in-between these cases. It is therefore possible that the original method (the exact setup of which is not recorded) measured the power density away from the direct reflection path, i.e.. only the isotropic part

of the reflected light, whereas the design of the Moor imager, using a similar optical path for the incident and reflected light through the same scanning mirror, means that more of the specular part of the reflected light reaches the detector, hence increasing the measured illumination intensity.

If this is the cause of the discrepancy, there could be further implications in terms of modulation depth, as the direct reflections are more likely to be from light reflected from surface tissue (with no blood vessels), whereas light that penetrates further into the skin and undergoes more scattering events is more likely to cause isotropic reflection. An increase in light level due to more specular reflection would therefore suggest a reduction in modulation depth and hence a reduction in measured Doppler signal.

The increase in photocurrents could lead to severe problems, as this means the IC was designed and tested to work with a different photocurrent range to that for which it is required to operate. However, one advantage of the logarithmic pixel response of the main array design is high dynamic range, and hence the design should be able to operate at higher photocurrents than intended. Additionally, if excess light is an issue the light level could be reduced through the use of absorptive filters, although this approach is counter-intuitive to the most accurate signal acquisition.

Because of the uncertainty caused by this discrepancy, the currents used during characterisation will range from the lower end of the calculated expected photocurrent to higher than the measurements taken from the LDLS, as the changes to the optical system or a further increase in laser power could increase the photocurrent beyond these. The DC photocurrents used were therefore set from 200 pA−80 nA.

4.3.4 DC Response

This section gives results for each element of the IC in response to varying DC photocurrent.

4.3.4.1 Logarithmic front-end output

Figure 4.9 shows the experimental and simulated variation of DC voltage as DC photocurrent rises for the logarithmic front-end output. The LHS circuit is used for this measurement as this circuit allows direct observation of the front-end output. The shape is the logarithmic response expected, giving a straight line on a log scale. The simulated response is shown for

comparison. The slope of the measured results match that of the simulated results fairly well, showing that the DC operating point of the pixel is following that predicted from theory. However the offset between the two shows that the exact DC level is difficult to accurately predict. This is partly due to the fairly small range of the DC output voltage. The large size of the error bars relative to the total response demonstrates the need to calibrate each pixel.

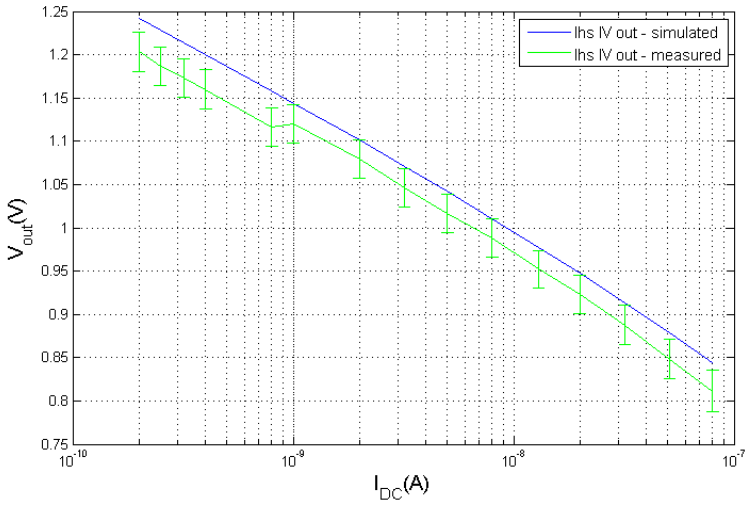

Figure 4.9: Measured and simulated V_{DC} against photocurrent, logarithmic front-end output of BVIPS1 IC.

4.3.4.2 g_mC output

Figure 4.10 shows the variation of DC output voltage as DC input photocurrent rises for the g_mC output of the logarithmic front-end circuit. At this point the signal has been amplified by the HDA amplifier and low-pass filtered by the g_mC. Again, the simulated response of the design is shown for comparison.

This part of the circuit is duplicated in 32 pixels of the left hand side (LHS) half of the IC, so the response of both circuits is shown. Some variation is to be expected due to pixel-pixel variation, as with fixed pattern noise within any CMOS array. The variation between LHS and RHS arrays could be larger than that between pixels due to the greater physical distance between the circuits. The layout of the circuits is also slightly different, with some circuit elements being moved around to better fit the space available on the LHS, as shown in Figure 4.5. The LHS circuit also has an additional stage (a frequency weighted filter) after the g_mC. The OTA buffers between the g_mC filter and the frequency-weighted filter should mean that this extra stage has no effect on the previous stages' behaviour, but

problems such as insufficient buffering (buffers that cannot drive the input capacitance of the frequency weighted filter and the large output buffer) may cause a difference in circuit behaviour.

As for the IV output behaviour, the slope of the graphs for simulated and measured behaviour is fairly similar, but there is an offset between the two. By comparison of Figures 4.9 and 4.10 it can be seen that the simulations predict a level shift occurring through the HDA and g_mC, which is not observed in the experiment as the measured range of the g_mC DC output voltage is very similar to that of the IV output.

The simulation of the HDA is potentially inaccurate due to the problems of simulating the very high resistance caused by the inverter-inverter arrangement - the low currents in the transistors making up the feedback resistor are similar in magnitude to the sub-threshold leakage current. The process used here gives sub-threshold leakage of $0.5 - 2\,\mathrm{pA}/\mu\mathrm{m}$ of transistor length [austriamicrosystems, 2007]. Both transistors in the HDA used in the BVIPS1 logarithmic pixels are $0.6\,\mu$m wide, potentially resulting in 1 pA leakage current. The simulated source-drain current in these transistors is between 0 and 2 pA, depending on the DC voltage at the output of the HDA, and therefore the leakage current can significantly affect the DC operating point and hence behaviour of these transistors.

This possible inaccuracy means the offset shown in simulation may not be an accurate result. Alternatively, the change in HDA input voltage caused by the variation between measured and simulated front-end output voltage may mean the level-shift caused by the HDA is reduced.

However, the absolute level of the DC output is not important, instead it is the change in level between light and dark that is useful for LDBF, as it allows changes in illumination intensity to be detected. For example, this change can be used to detect whether a part of an image is foreground (the target tissue) or background (e.g. gaps between fingers), based on the lower light level from the background which is not directly illuminated by the laser. In addition, the inherent normalisation in log pixels (see Section 3.2.1.2) mean that the DC level is not as crucial for correct operation as in a linear pixel. Therefore, this behaviour is acceptable for correct operation of the IC.

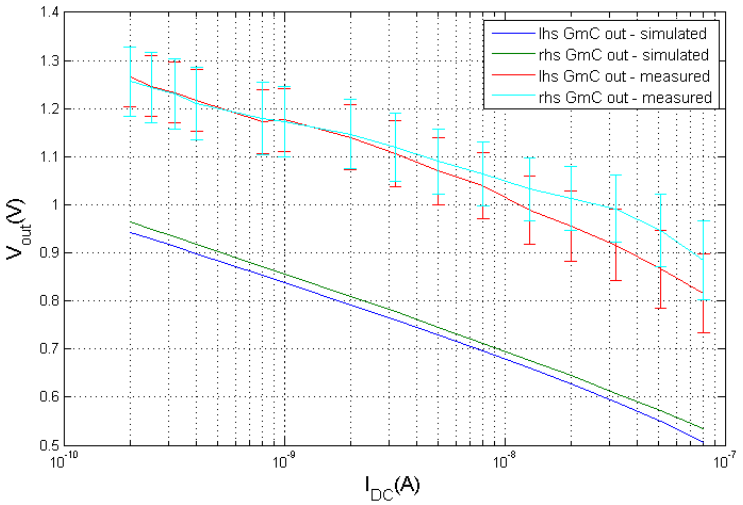

Figure 4.10: Measured and simulated V_{DC} against photocurrent, g_mC output of LHS and RHS logarithmic pixel arrays on BVIPS IC.

4.3.4.3 Opamp front-end output

Figure 4.11 shows the variation of DC voltage as DC photocurrent rises for the opamp front-end output. This is a linear front-end design, so the response is very different to that seen for the logarithmic pixels.

While Figure 4.11 does not show a linear response over the full photocurrent range, this is the response expected. The gain of the front-end alone is not sufficient for the output voltage to be reliably measured. A non-inverting opamp gain stage is used after the front-end to increase this signal. This is problematic, as unlike the HDA the amplifier here amplifies the DC as well as the AC voltage, so is prone to saturation as the DC voltage from the front-end varies. To account for this, an offset voltage can be applied to the amplifier, which is set close to the output from the front-end. This means that the relative DC voltage is low, so the output of the gain stage does not saturate. The output of the opamp circuit is linear over the range of DC photocurrents where the offset prevents saturation. Outside of this range the output voltage will saturate at near 0 V or 3.3 V.

The offset voltages used here are 0.1 V for the front-end offset, and 1 V for the gain-stage offset. The front-end offset is low as the AC voltages expected here are small, so a 0.1 V offset is sufficient to prevent saturation, while keeping the front-end DC output voltage low, meaning that higher magnitude signals can be amplified by the gain stage before saturation occurs. The gain-stage offset is set to be just below the lowest voltage expected from the

front-end for the expected range of input photocurrents. If this value is too low, the large difference in DC voltages at the gain stage input means that saturation of the gain stage output at V_{DD} will occur. If the offset is too high, then the gain-stage output will be fixed at 0 V (or 0 V + V_T if the opamps used do not have rail-rail operation) until the photocurrent rises sufficiently to cause the front-end output voltage to rise above the offset voltage.

Figure 4.11 shows that the output voltage is linear for a range of currents around 10 nA (region B of Figure 4.11), which is the typical photocurrent expected for the line scanner system. However, the linear range can be adjusted to account for changes in light power by changing the offset voltage.

There is some difference between the simulated and measured results. The measured results show the circuit saturating at a slightly higher DC current than the simulations predict (saturation occurs between regions B and C in Figure 4.11) . This could be due to a variation in the gain of the gain stage, as higher gain would result in the linear operation being over a narrower range of photocurrents, as a relatively smaller change in DC voltage would be required to cause saturation. Alternatively, the DC response of the front-end stage could be different to that predicted from simulation due to device variation, such that the DC output voltage has a higher range and gradient in the actual circuit. This is difficult to ascertain, as the opamp circuit does not have an intermediate output from the front-end, in order to keep the silicon area and number of pins of the IC down.

The measured response, unlike the simulated response, does not drop to 0 V below 3 nA (region A in Figure 4.11). This could be because the opamp used does not have rail-rail operation, and hence the output saturates above 0 V. However, this would have been expected to have been seen from simulation, and the difference between the two responses is more than that explained by manufacturing variations. The fact that the measured response does not become flat at lower currents suggests that the gain of the system is lower in the measured case. This would cause the slope of the linear region to be less steep, and hence this region will cover a wider range of DC photocurrents. This, in combination with the lack of rail-rail behaviour making saturation occur gradually rather than at a definite point, explains the behaviour seen in Figure 4.11.

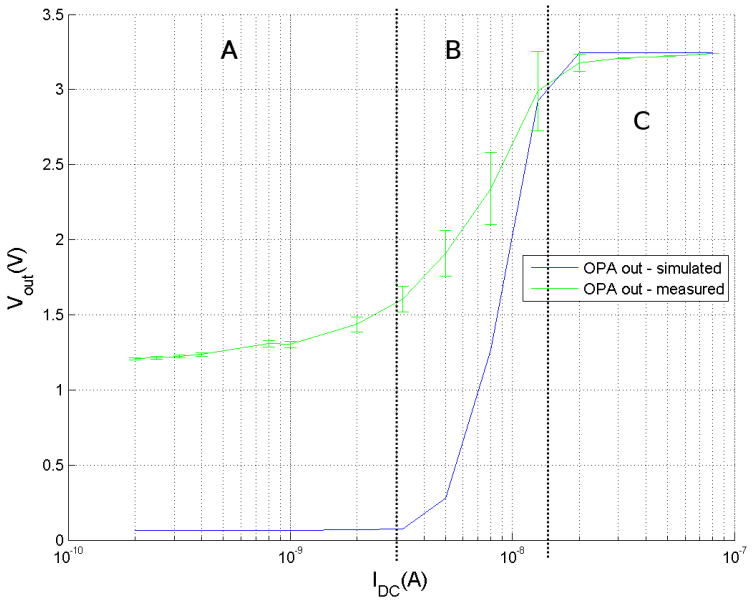

Figure 4.11: Measured and simulated V_{DC} against photocurrent, Opamp front-end output of BVIPS1 IC. Regions shown are: below linear region, saturated at 0 V (A); in linear operating region (B); above linear region, saturated at V_{DD} (C).

4.3.4.4 HDA behaviour

Figure 4.12 shows the DC behaviour of the HDA used to amplify the signal from the logarithmic front-end. This is found by dividing the DC voltage measured at the g_mC output of the LHS (32x1) and RHS (64x1) logarithmic pixel by the DC voltage at front-end output of the LHS logarithmic pixel. This gives an accurate measure for the LHS HDA, but for the RHS HDA this assumes that the front-ends on both sides have equal performance. This means the values for the RHS HDA are approximate. The similarity in circuit design and layout means that this is a reasonable approximation.

It can be seen that the measured gain is higher than the simulated results, and is more than unity such that the DC voltage after the HDA (and therefore the voltage at the main IC output) will be around 5% larger than that from the front-end itself. Given the inaccuracy in simulating the HDA described in Section 4.3.4.2, some non-ideal behaviour is to be expected. However, the gain is close enough to unity to not cause saturation of the output. The gain also stays relatively constant with DC photocurrent, suggesting reliable behaviour at a range of photocurrents.

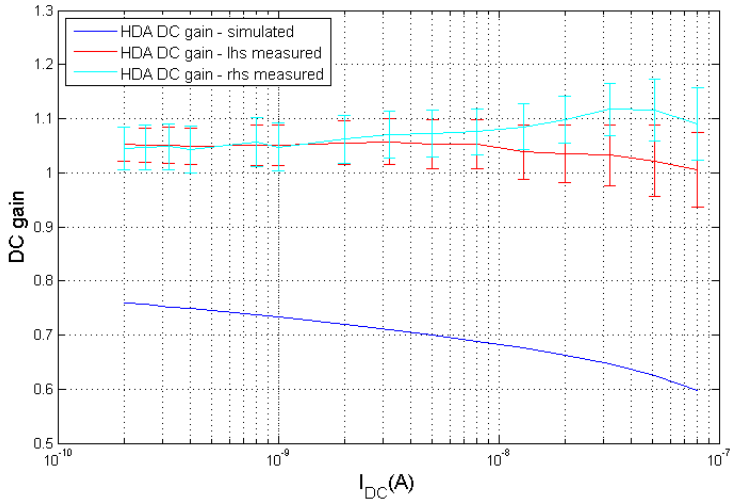

Figure 4.12: Measured and simulated DC gain against DC photocurrent, HDA in logarithmic pixel on BVIPS1 IC

4.3.5 AC Response

4.3.5.1 Logarithmic front-end output

Figure 4.13 shows the AC response for the log front-end output, again showing comparison of real and simulated results and also variation across the IC. The AC photocurrent was set to give a constant modulation depth, with the AC peak-peak photocurrent being set to 15% of the DC photocurrent. The frequency used was 5 kHz to ensure that the signal is within the pass-band of all circuits.

It can be seen that gain is higher than expected, although there is also considerable variation in gain across the array. The normalisation predicted from theoretical operation of log pixels does not appear to work as well as expected at first sight, as shown by the slope on the graph, showing an increase in AC output voltage as the light level, and hence DC photocurrent, rises. This is because perfect normalisation assumes that the model of MOSFET behaviour in the subthreshold region is exact. In reality, the model is an approximation of more complex behaviour, and the accuracy of the model varies within the subthreshold region, which does not have definite and sharp upper and lower limits [Allen and Holberg, 2002]. However, it should be noted that the increase in peak-peak voltage is from 7 mV to 10.5 mV, a factor of 1.5, compared to a change in DC photocurrent of 200 pA to 80 nA, a factor of 400. This suggests that the smaller variations in power due to laser fluctuations or skin remittance will

not cause a major change in signal levels. The slope on the graphs may be caused by the inaccuracies of the transistor sub-threshold simulation model, such that the AC resistance of the active load is not exactly inversely proportional to the DC photocurrent. The higher AC gain compared to simulation results is not predicted from the similar gradients of the DC responses. This is because the diode connected transistors used as a load in the pixels mean that the AC behaviour cannot be exactly predicted from the gradient of the DC response. The AC parameters are controlled by, but not the same as, the DC parameters. This is discussed in Section 3.2.1.

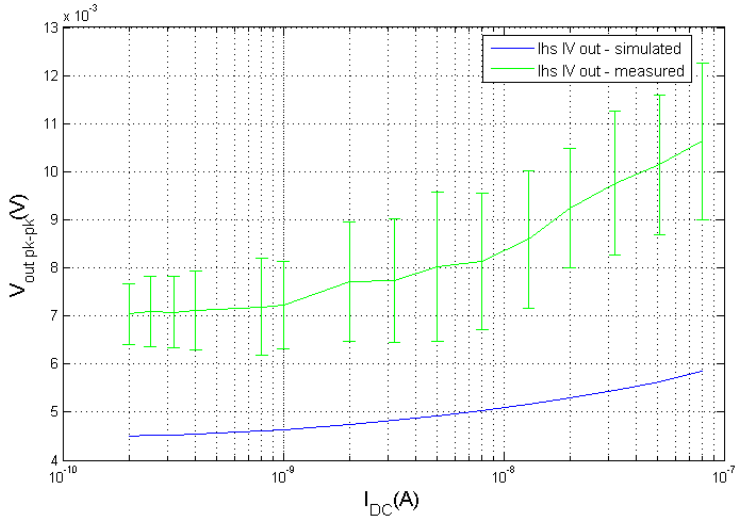

Figure 4.13: Measured and simulated V_{AC} against DC photocurrent, fixed modulation depth, logarithmic front-end output of BVIPS1 IC

Figure 4.14 shows the AC transimpedance (gain) in ohms for the logarithmic front-end. This is calculated by measuring the AC output voltage (as shown in Figure 4.13) and dividing by the corresponding AC photocurrent for each DC photocurrent value. It can be seen that while the IC has higher transimpedance than expected, the relationship between transimpedance and I_{DC} is of the form predicted from theory and simulation. There is an offset between the measured and simulated transimpedance, with the measured transimpedance being roughly 60% higher than that found from simulation. This is similar in some ways to the results shown for the DC response - the measured DC voltage is higher than the simulated response, which is similar to the effect that would be caused by a reduction in photocurrent. The reduction in photocurrent would also cause an increase in transimpedance. This suggests that either the conversions between light level and photocurrent are inaccurate, which could be caused by effective pixel area being lower than expected due to non-ideal behaviour at the

edges of the photodiodes, or a lower than expected responsivity, which could be caused by coatings on the IC or the glass used in the packaging. Alternatively the differences could be electrical, for example process variations causing a reduction in gain of the front-end circuit.

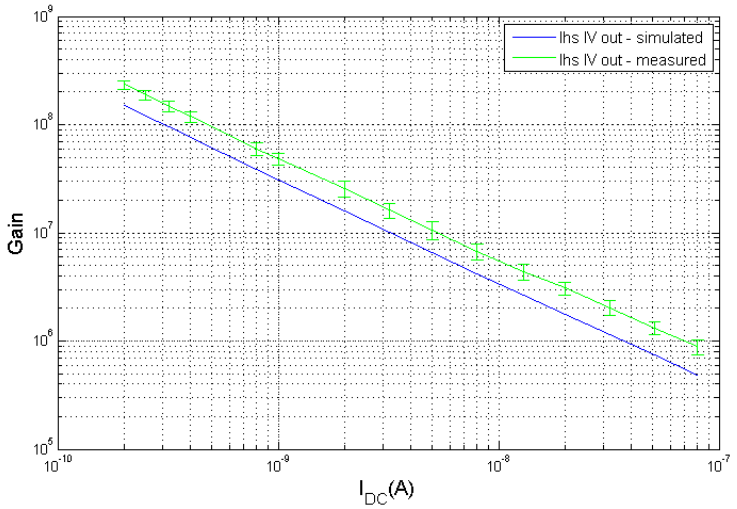

Figure 4.14: Measured and simulated transimpedance against DC photocurrent, logarithmic front-end output of BVIPS1 IC

4.3.5.2 GmC output

Figure 4.15 shows the AC output voltage against DC photocurrent for the g_mC output of the logarithmic pixel array, showing comparison of real and simulated results and then variation across the IC. In contrast to the front-end output results, the AC gain measured here is lower than that predicted by simulation. This suggests that the gain of the HDA circuit on the IC is lower than that expected from simulation. The design of the HDA, namely the 'inverted-inverter' feedback arrangement of two fully turned off MOSFETs, means that this circuit is difficult to simulate accurately due to the very low currents involved, so some difference in behaviour here is to be expected.

The AC voltage here shows an opposite trend to that of the front-end output, with the AC voltage falling as DC photocurrent increases. This is caused by the HDA gain varying with input DC voltage - as the photocurrent rises, the front-end output voltage falls, which causes a drop in the gain of the HDA and therefore a small drop in overall AC output voltage. Again, the change in AC voltage is small compared to the corresponding change in light level (a factor of 3 compared to 400). A potentially more significant issue shown here is the large pixel-pixel transimpedance variation. While calibration can compensate for some of this,

pixels with particularly low gain will have poor signal to noise ratios and will therefore be less sensitive to changes in blood flow.

One notable feature on the graph is the difference in response between the simulated output on the right and left hand sides of the IC, which are intended to be duplicate designs for all the stages used here. However, after fabrication it was found that the LHS array lacked a dedicated output buffer at the g_mC stage (as this was not intended as the final output of this array), such that the large off-chip outputs are only driven by the small buffer at the pixel level. It is possible that the increased load resistance presented to this buffer results in an additional AC gain that causes the upward curve at higher photocurrents.

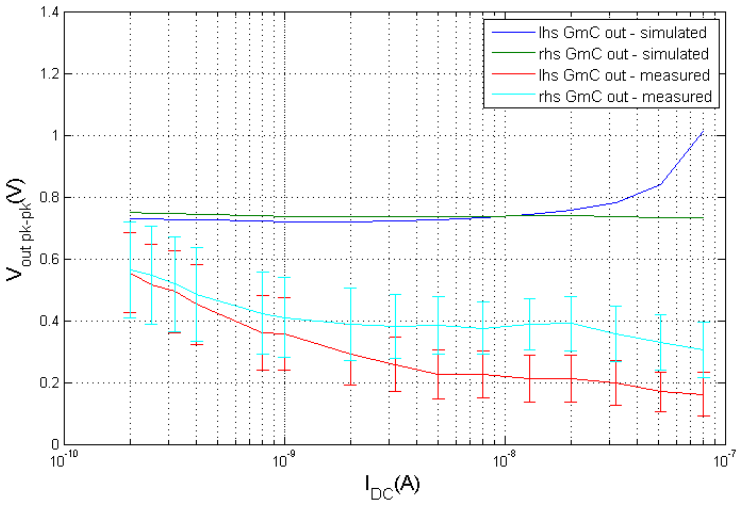

Figure 4.15: Measured and simulated V_{AC} against DC photocurrent, fixed modulation depth, g_mC output of BVIPS1 IC logarithmic pixels

Figure 4.16 shows the AC transimpedance for the g_mC output of the logarithmic pixel array. This is calculated by measuring the AC output voltage (as shown in Figure 4.15) and dividing by the corresponding AC photocurrent for each DC photocurrent value. This shows similar behaviour to the AC voltage graph, with the measured gain lower than the simulated gain, and also falling faster with photocurrent such that the normalisation is not ideal. However, it does show that the behaviour of the IC follows the same trends as the simulations predict, further demonstrating the normalising behaviour of logarithmic pixels.

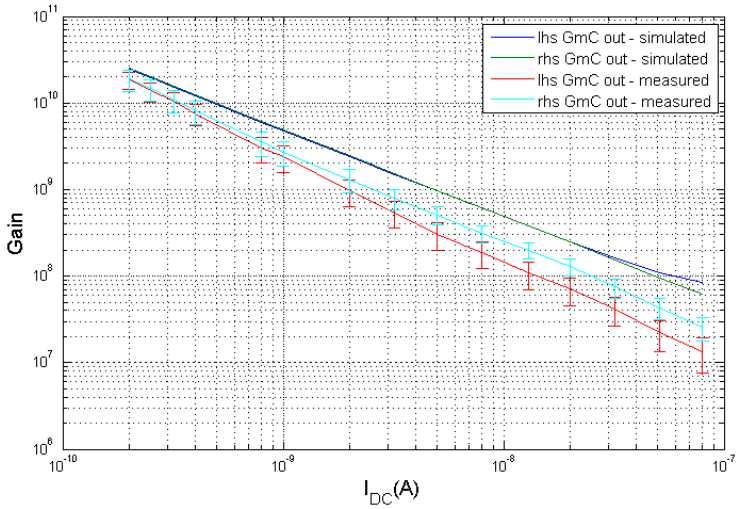

Figure 4.16: Measured and simulated AC transimpedance against DC photocurrent, g_mC output of BVIPS1 IC logarithmic pixels

4.3.5.3 Opamp front-end output

Figure 4.17 shows the AC response for the opamp front-end section of the IC. This clearly shows the high gain of this circuit while the photocurrent is in the range which matches the offset voltages applied to the front-end and gain stage (region B, 10 nA photocurrent for 100 mV front-end offset and 1 V gain-stage offset). The transimpedance in this range is lower than that found by simulation, which makes the opamps less likely to be sensitive enough to detect Doppler blood flow signals, although this obviously also depends on noise levels. The decrease in transimpedance also explains the behaviour seen in the DC response, where saturation at low DC photocurrents (region A) was more gradual in the measured case than the simulated case. Reduced transimpedance means smaller AC voltages, and therefore the DC voltage has to drop further before saturation occurs. Above the operating range (region B), saturation occurs as expected and the AC voltage becomes very low.

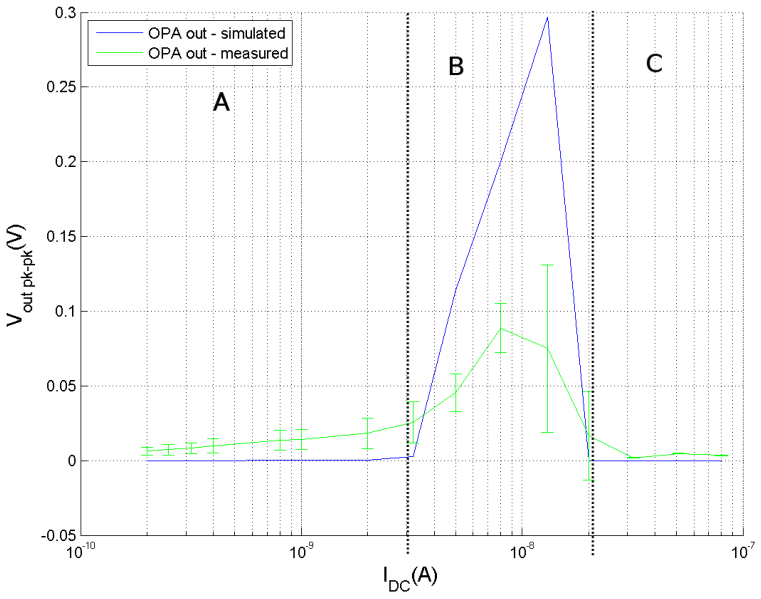

Figure 4.17: Measured and simulated V_{AC} against DC photocurrent, fixed modulation depth, opamp front-end output of BVIPS1 IC. Regions shown are: below linear region, saturated at 0 V (A); in linear operating region (B); above linear region, saturated at V_{DD} (C).

Figure 4.18 shows gain against I_{DC} for the opamp front-end output of the IC. The shape of this graph follows that of the AC output voltage graph shown in Figure 4.17 in that the transimpedance is higher in the operating region (region B) before dropping off to nearly zero above this (region C), as this design does not normalise the AC voltage. However, at lower DC photocurrents (region A) the transimpedance is higher than that expected, which again corresponds to the more gradual saturation caused by reduced transimpedance. The apparent increase in transimpedance at very low photocurrents is due to noise voltages giving a non-zero AC output voltage, which with the low AC photocurrent appears as a high transimpedance. This would not be expected for the linear pixel, as the constant transimpedance should mean that the AC output voltage is very low with very low AC input photocurrents. The increasing transimpedance values may be caused by a high noise level being divided by a very low AC photocurrent. Another unexpected aspect of the transimpedance graph is that the measured response is not flat in region B, where the circuit should not be saturated and gain should be constant. It is possible that within the operating region but approaching saturation the non rail-rail operation of the opamps causes some compression of the output signals, and therefore a small reduction in gain.

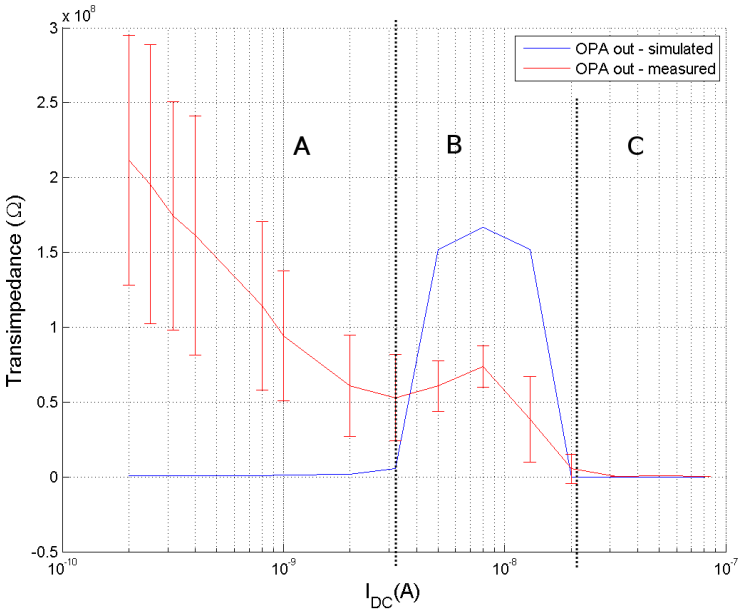

Figure 4.18: Measured and simulated transimpedance against DC photocurrent, opamp front-end output of BVIPS1 IC. Regions shown are below linear region, saturated at 0 V (A), in linear operating region (B), and above linear region, saturated at V_{DD} (C).

4.3.5.4 HDA behaviour

Figure 4.19 shows the AC gain of the HDA used to amplify the signal from the logarithmic front-end. As for the DC behaviour, this is found by dividing the AC voltage measured at the g_mC output of the LHS (32x1) and RHS (64x1) logarithmic pixel by the AC voltage at front-end output of the LHS logarithmic pixel. This gives an accurate measure for the LHS HDA, and an approximation for the RHS HDA based on the assumption that the LHS and RHS front-ends have the same characteristics.

It can be seen that the measured AC gain is lower than the simulated results which may cause problems with the signal level being too low for accurate detection. This can be compensated for to some extent by using an ADC with a higher dynamic range, such that the smaller signals can still be accurately measured, assuming that the SNR of the HDA output signal is high enough. However, the increase in dynamic range itself places additional constraints on the system - the purpose of the HDA is to increase the signal level to a magnitude that can be measured with a commonly available ADC - otherwise a very accurate ADC could be used to measure the signal from the front-end directly. However, increasing the accuracy of an ADC generally results in a longer acquisition time, and hence lower bandwidth.

The gain here of between 20 and 80 should be sufficient to not place unreasonable constraints on ADC selection. However, the reduced gain may need to be addressed for future iterations of the design as the additional accuracy required of the ADC may increase overall system cost or make the IC unsuitable for integration into existing systems.

The increase in simulated AC gain at higher DC photocurrents may be due to a change in the DC operating point due to a change in voltage at the HDA feedback network. The DC response measured at the g_mC output shown in Figure 4.10 shows lower simulated DC voltages than those measured. While this is not a direct measurement of the DC operating point at the HDA, it shows that the simulation inaccuracies mean the DC operating point used in simulation may be different to that occurring on the IC. By contrast, DC response of the logarithmic front-end output (see Figure 4.9), which is the HDA input voltage, shows a good match between simulation and measurement, suggesting that the DC behaviour of the HDA is not as predicted. This makes direct comparison of AC behaviour problematic, as it cannot be assumed that the DC operating points for corresponding light levels are matched.

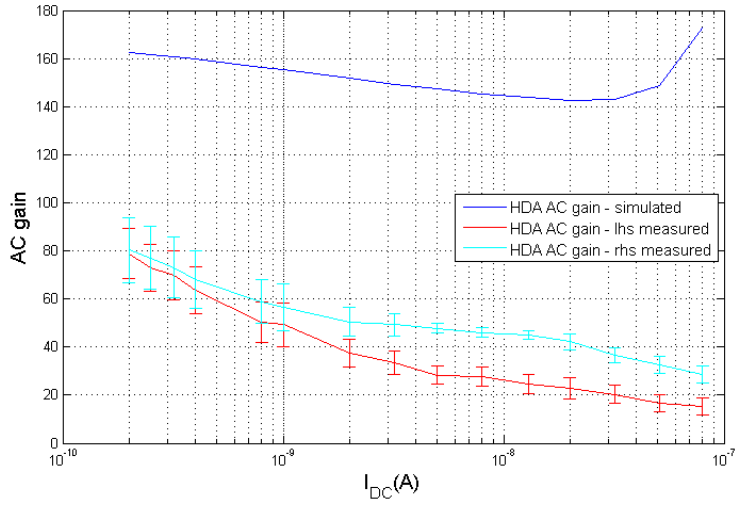

Figure 4.19: Measured and simulated AC gain against DC photocurrent, HDA in logarithmic pixel on BVIPS1 IC

4.3.6 Noise Response

Noise measurements were taken by measuring the IC RMS output voltage, while illuminating with constant intensity light. This allows a measurement of total noise against varying DC photocurrent. Further investigation of noise shown in Section 4.4 also uses Fourier transforms

to investigate the spectral content of the noise at various points). From the voltage noise measured at the output, the input referred current noise can be found, by dividing output voltage noise by total measured transimpedance. This allows a value for input current noise to be found that can be compared to AC (i.e. signal) photocurrent.

As well as measured and simulated noise, the theoretical minimum noise level is shown on all plots. This minimum is the shot noise density as a function of DC photocurrent. For input referred current noise plots, this is plotted directly. For output noise voltage plots, the shot noise current density is multiplied by the simulated gain for the corresponding photocurrent.

In all of the plots shown in this section, the measured noise level is significantly higher than that expected from simulation. Efforts to reduce or identify the source of this noise prior to characterisation are described in Section 4.4. As these noise reduction efforts were largely unsuccessful, the additional noise in the system is likely to be due to wider faults in the circuit which are investigated in Section 4.6.3.

4.3.6.1 Logarithmic front-end output

Figure 4.20 shows measured output voltage noise from the logarithmic front-end. Figure 4.21 shows input referred current noise.

The difference between the simulated and measured noise is very large, with the difference being over two orders of magnitude at high photocurrents (Figure 4.21). The voltage noise results (Figure 4.20) show twice the noise level at the lowest photocurrent (0.1 nA), rising to approximately two orders of magnitude worse at the higher photocurrents (80 nA).

The difference in noise response may be due to noise on the power supply of the IC PCB, however, various steps were taken to reduce this such as additional decoupling, inductive smoothing of the power supply and using a battery power supply to isolate the IC from mains interference (see Section 4.4). Some noise may be added by the buffers on the IC PCB, and some noise may be picked up on the power supply and signal lines, although these are shielded where possible. Given that the measured output voltage noise level rises with photocurrent, whereas the simulated/theoretical noise shows a reduction in noise with rising photocurrent (due to falling transimpedance), the noise level may be linked to the incident light level. As the effect of shot noise at the output reduces with rising photocurrent, this may suggest that stray light may cause additional noise through generating electron-hole

pairs in the substrate. While the photodiode has a guard ring, and most of the IC is covered by top layer metal, some stray light is inevitable, and not all components have guard rings.

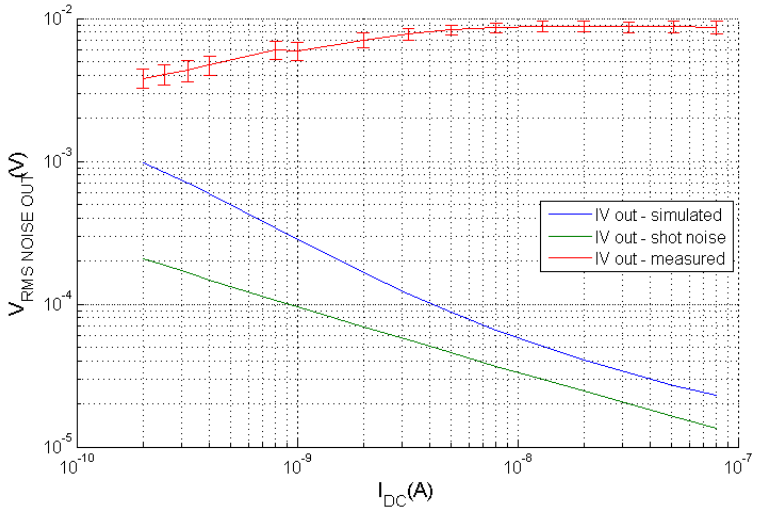

Figure 4.20: Measured and simulated output voltage noise and theoretical shot noise limit against DC photocurrent, logarithmic front-end output of BVIPS1 IC

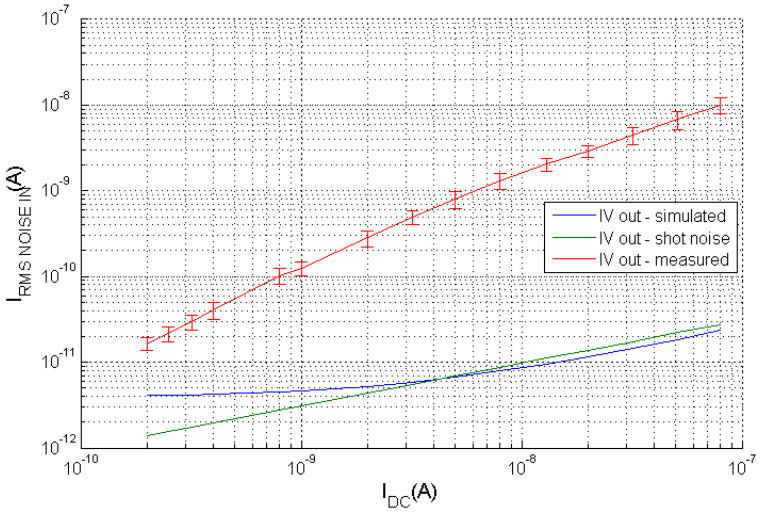

Figure 4.21: Measured and simulated input referred current noise and theoretical shot noise limit against DC photocurrent, logarithmic front-end output of BVIPS1 IC

4.3.6.2 GmC output

Figures 4.22 and 4.23 show the output voltage and input referred current noise responses of the IC as measured at the g_mC output.

The measured output voltage noise is similar to the simulated noise level at low photocur-

rents, partly due to reduced gain compared to simulated results, but does not fall as sharply as the simulated noise level. This suggests that the causes of the front-end output noise rising rather than falling with photocurrent still apply here, although these noise sources do not appear to be as significant, given the smaller discrepancy between the measured and simulated noise here. Noise reduction efforts are shown in Section 4.4, while additional possible faults are investigated in Section 4.6.3.

The input referred noise graphs are very similar to those seen for the front-end output in Figure 4.21, with input referred current noise being the same or lower after the HDA and low-pass filter. A small improvement is to be expected, as any noise sources affecting signal output/sampling now only effect the amplified signal, so are proportionally smaller. This demonstrates that while the HDA has lower gain than expected, it still leads to an increase in SNR.

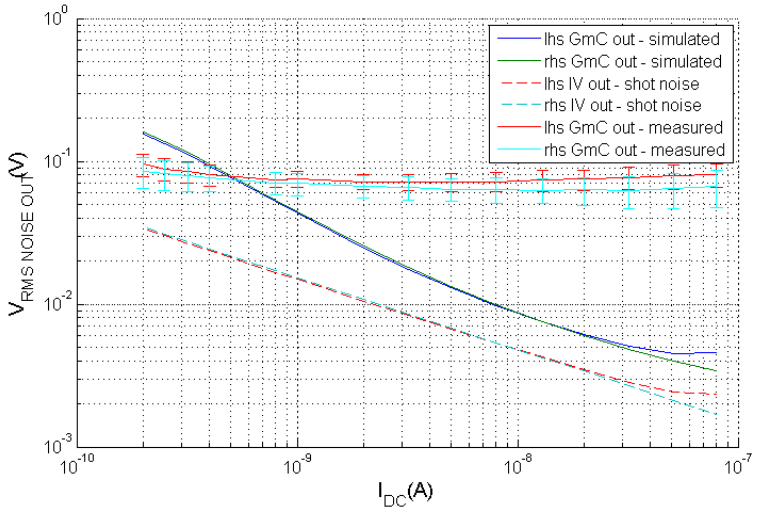

Figure 4.22: Measured and simulated output voltage noise and theoretical shot noise limit against DC photocurrent, g_mC output of BVIPS1 IC logarithmic pixels

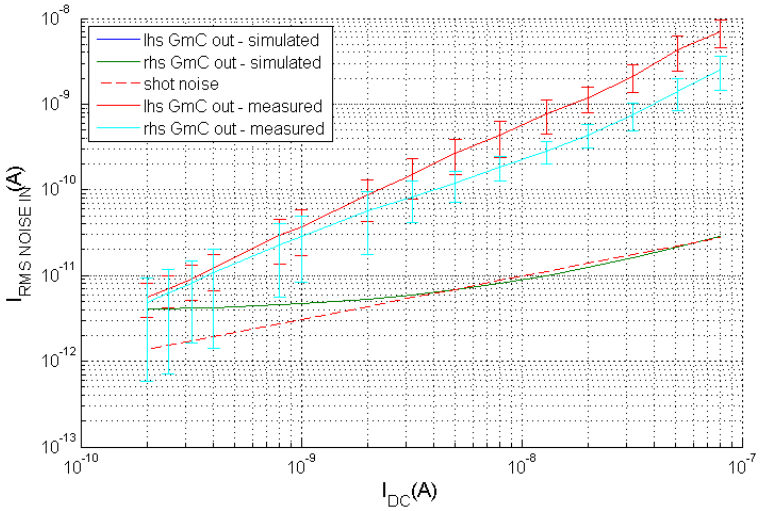

Figure 4.23: Measured and simulated input referred current noise and theoretical shot noise limit against DC photocurrent, g_mC output of BVIPS1 IC logarithmic pixels

4.3.6.3 Opamp front-end output

Figures 4.24 and 4.25 show the output voltage and input referred current noise responses for the opamp front-end output of the BVIPS1 IC. The voltage noise graph shows similar characteristics to the AC output voltage graph in Section 4.3.5.3, Figure 4.17, with higher noise in the operating region (region B on the graph) and lower noise when the output is saturated (regions A and C on the graph). This is to be expected, as the saturation of the output voltage reduces the noise voltage as well as the signal voltage. Again, the measured noise is approximately an order of magnitude higher than the simulated noise due to additional noise such as that from the power supply and the data acquisition system (see Section 4.4).

The input referred noise graph (Figure 4.25) does not follow the common trends, such as noise being proportional to $\sqrt{I_{DC}}$ (as expected if dominated by shot noise) or following the AC gain against I_{DC} relationship. The measured response here shows a linear trend on the response, which suggests that this may be caused by laser noise (which is proportional to I_{DC} rather than $\sqrt{I_{DC}}$, although this is expected to be several orders of magnitude below that measured here. However, the measured noise response is similar to that seen for the logarithmic front-end and g_mC outputs, suggesting that the dominant noise sources are shared by all circuits. This could be laser noise, power supply noise, or noise added by the on-chip buffers, PCB buffers and data acquisition hardware.

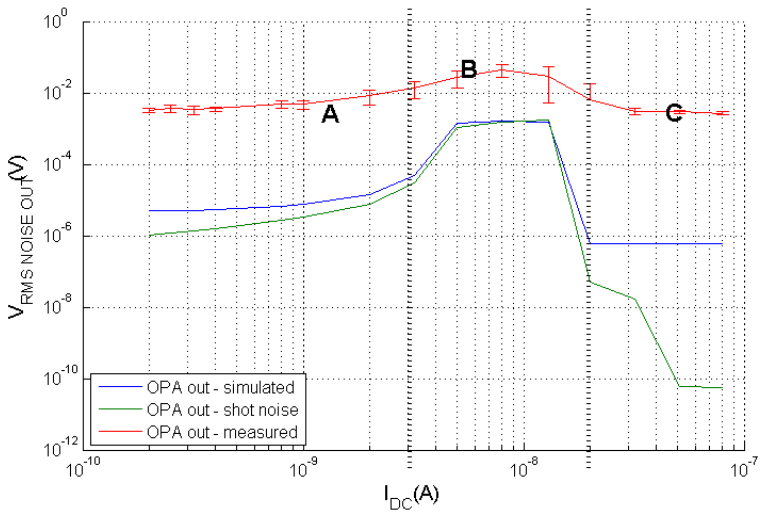

Figure 4.24: Measured and simulated output voltage noise and theoretical shot noise limit against DC photocurrent, Opamp front-end output of BVIPS1 IC

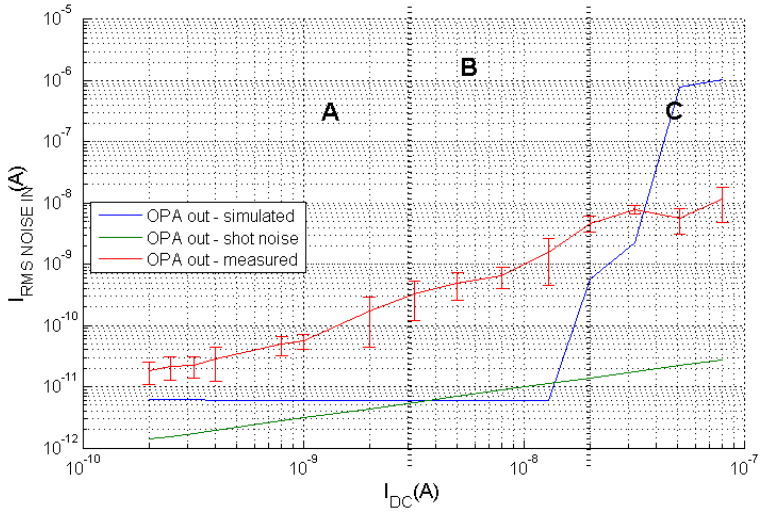

Figure 4.25: Measured and simulated input referred current noise and theoretical shot noise limit against DC photocurrent, Opamp front-end output of BVIPS1 IC

4.3.7 Frequency Response

This section shows the frequency response of the IC at a typical DC photocurrent of 13 nA from three different IC outputs. As well as the full frequency plot shown for typical DC photocurrent, the 3 dB and 6 dB cut-off frequencies are given at high (1 nA) and low (51 nA) DC photocurrents, showing the behaviour expected during normal operation as well as how this behaviour will vary with changes in laser power, skin reflectivity etc. The AC

photocurrent used for these measurements was set to give a modulation depth of 15%.

4.3.7.1 Logarithmic front-end output

Figure 4.26 shows the frequency response as measured at the LHS logarithmic front-end output at typical DC photocurrent. The response is relatively flat, and is dominated by the pixel-pixel variation. This is to be expected, as the cut-off frequency of the front-end itself is intended to be higher than the system bandwidth, such that the bandwidth is independent of photocurrent, and can be set by the following g_mC low pass filters.

The error bars on the graph also show the high fixed pattern noise across the IC, with a standard deviation in transimpedance of around 2 dBΩ - compared to an average transimpedance of 155.6 dBΩ. Given concerns about fixed pattern noise and 'dead' pixels (shown more clearly in Section 4.5.3), it is possible that the high variation is caused by the performance of the dead pixels, rather than variation between 'normal' pixels. To investigate this, a third line is added to the plot showing the measured performance of array after discarding the outlying pixels. This is done by finding the pass band gain of each pixel (at 3 kHz) and discarding the top and bottom 10th percentiles - i.e. those pixels with a gain sufficiently far from the mean to be considered 'dead' pixels. However, the plot of measured performance from these pixels is not significantly different from the plot of all pixels (except for a small drop in spread as would be expected). This shows that the variation seen is down to pixel-pixel variation of normal pixels, rather than being specifically caused by a small number of malfunctioning pixels.

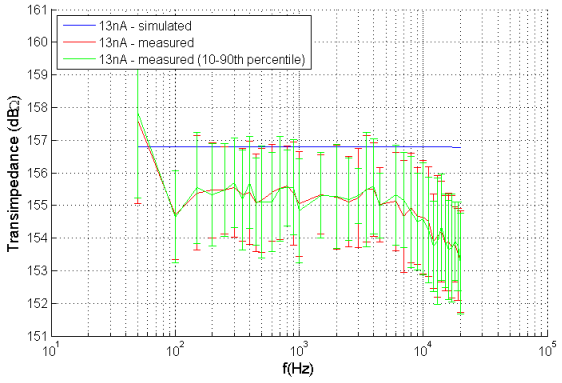

Figure 4.26: Measured and simulated frequency response at normal (13 nA) photocurrent, 15% modulation depth, logarithmic front-end output of BVIPS1 IC

4.2 shows the variation of the frequency response over a range of photocurrents. The flat response means that the high and low bandwidth limits are not met within the range tested. This is to be expected, given that this output is before the high or low-pass filters, although the results confirm that the front-end used has sufficient bandwidth at all photocurrents in the range needed (although this range is higher than the original design range - see Section 4.3.3).

	$I_{DC} = 1\,nA$	$I_{DC} = 13\,nA$	$I_{DC} = 51\,nA$
6 dB-low (Hz)	not reached	not reached	not reached
3 dB-low (Hz)	not reached	not reached	not reached
3 dB-high (Hz)	not reached	not reached	not reached
6 dB-high (Hz)	not reached	not reached	not reached

Table 4.2: 3 dB and 6 dB cut-off frequencies at high (1 nA), typical (13 nA) and low (51 nA) DC photocurrents, front-end output of BVIPS1 IC logarithmic array

4.3.7.2 GmC output

Figure 4.27 shows the frequency response of the logarithmic pixel, measured at the g_mC output of the RHS array, for a typical DC photocurrent, while Table 4.3 shows the variation of the frequency response over a range of photocurrents.

As for the front-end output, the frequency response shows a high level of fixed pattern noise, with standard deviation of around 5dBΩ (the same standard deviation as found on the front-end output, shown in Figure 4.26), compared to a total transimpedance of approximately 180 dBΩ. The response is also shown after pixels with pass-band gain in the top and bottom 10th percentile pixels are removed. As for the front-end output, the filtered results show no change other than the slightly lower spread expected, showing that the variation is not due entirely to a small number of dead pixels.

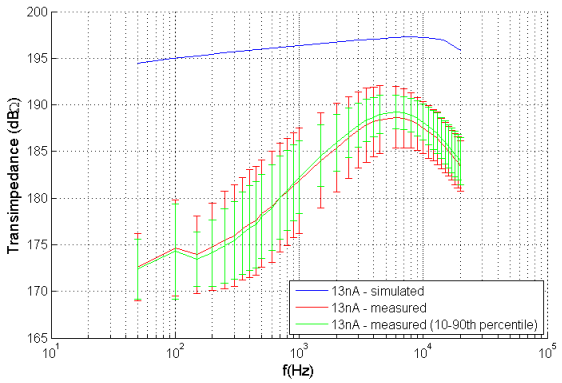

Figure 4.27: Measured and simulated frequency response at normal (13 nA) photocurrent, 15% modulation depth, g_mC output of RHS logarithmic pixel array on BVIPS1 IC

It can be seen from Table 4.3 that the low frequency cut-off (set by the HDA) in all cases is higher than that expected from simulation, being around 1.5-2.5 kHz. This is a concern given that a significant proportion of the Doppler signal is found at frequencies lower than this. However, the gain falls fairly gradually (10dBΩ/decade) below this frequency, so these lower frequency components will not be fully filtered out.

The high frequency cut-off is slightly lower than 20 kHz in all cases, being between 15 and 17 kHz. The cut-off frequency was set by externally adjusting the bias of the g_mC low-pass filter, with this result showing that the cut-off was set slightly lower than intended, or that the bandwidth has drifted since initial setting. Given that Doppler signal is mostly at low frequencies, the cut-off frequencies found are acceptable, with minor adjustments possible if required. Additionally, the results show that this high cut-off is not significantly affected by DC photocurrent, and that the upper cut-off frequency is set by the low-pass filter rather than being restricted by insufficient bandwidth of other circuits.

	$I_{DC} = 1$ nA	$I_{DC} = 13$ nA	$I_{DC} = 51$ nA
6 dB-low (Hz)	1000	1000	1500
3 dB-low (Hz)	1500	2000	2500
3 dB-high (Hz)	15000	15000	17000
6 dB-high (Hz)	not reached	not reached	not reached

Table 4.3: 3 dB and 6 dB cut-off frequencies at high (1 nA), typical (13 nA) and low (51 nA) DC photocurrents, g_mC output of RHS logarithmic pixel array on BVIPS1 IC

4.3.7.3 Opamp front-end output

Figure 4.28 shows the frequency response of the opamp pixel at a typical DC photocurrent. It can be seen that at the DC operating point expected the measured results show a lower (but acceptable) gain of 170dBΩ compared to a simulated gain of 190dBΩ, and the high frequency cut-off is sufficient for this application if lower than that set for the logarithmic pixels. The response does not show any low frequency roll-off, as there is no high-pass filter, such as a HDA, in this design. This means that the system amplifies all photocurrents down to DC, and hence motion artifacts caused by movement of the tissue relative to the detector will not be filtered out as well as in the other systems, which may cause problems for flow imaging.

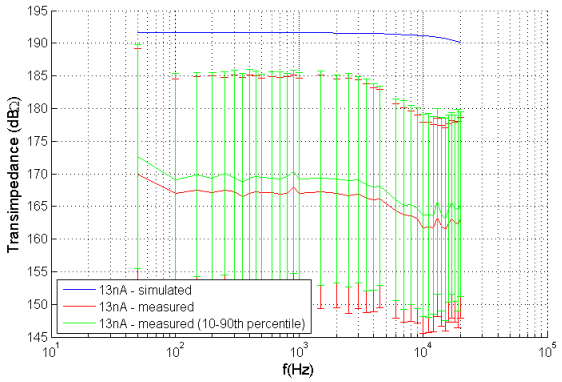

Figure 4.28: Frequency response of the IC at normal (13 nA) photocurrent, 15% modulation depth, opamp front-end output of BVIPS1 IC

Table 4.4 shows the variation of cut-off frequencies with photocurrent for the opamp detector. At the low photocurrent, the low-cut off frequency rises into the signal bandwidth and the high cut-off also falls, while at high photocurrents the 6 dB cut-off point is not reached showing that the bandwidth is increasing with photocurrent so that there is no longer a sharp cut-off. This circuit does not include any specific filters, with the performance of the opamp being sufficient to limit bandwidth in the typical photocurrent range. The opamp design also does not have the same inherent dynamic range advantages of the logarithmic pixel, and hence more variation with DC photocurrent is to be expected.

	$I_{DC} = 1\,nA$	$I_{DC} = 13\,nA$	$I_{DC} = 51\,nA$
6 dB-low (Hz)	2000	not reached	not reached
3 dB-low (Hz)	4000	not reached	not reached
3 dB-high (Hz)	11000	6000	12000
6 dB-high (Hz)	15000	15000	not reached

Table 4.4: 3 dB and 6 dB cut-off frequencies at high (1 nA), typical (13 nA) and low (51 nA) DC photocurrents, opamp front-end output of BVIPS1 IC

4.4 Noise Reduction

As shown in Section 4.3.6, the noise levels found from all circuits were significantly higher than expected from simulations. This had been noticed in preliminary measurements, and various techniques were applied to reduce the noise levels. This section discusses this work, with results showing overall noise reductions as well as showing noise spectra to identify any specific sources of noise. The changes made included adding additional bypass/decoupling capacitors, using shielded power supply cables, and using a battery power supply.

4.4.1 Original Configuration

The initial configuration uses the IC on a PCB as originally designed, with bypass/decoupling capacitors at each IC power supply but without decoupling capacitors at the input/output of the voltage regulators themselves. The power supply was provided by a bench power supply, using non-shielded cables. The IC was illuminated with a uniform laser illumination giving a DC photocurrent of 30 nA. Figure 4.29 shows the noise signals recorded on the power supply to the logarithmic front-end (64 x 1 array), and the noise on the output of one of the log pixels (both signals AC coupled to remove DC).

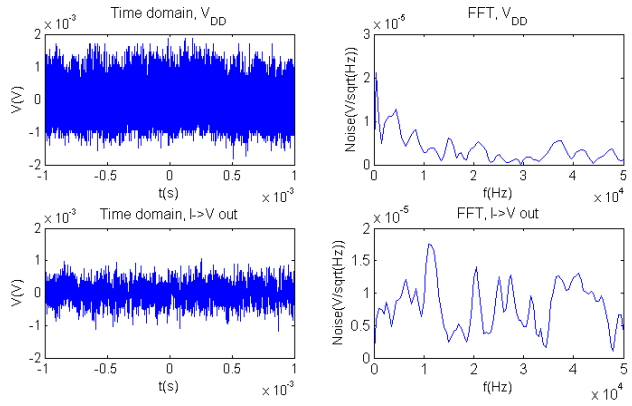

Figure 4.29: Noise on power supply and I->V output, time and frequency domain, basic configuration

Figure 4.30 shows the same power supply and front-end voltage noise plots after adding additional capacitors to the voltage regulators on the PCB, between $V_{IN} - V_{GND}$ and $V_{OUT} - V_{GND}$. There is some improvement in noise (see table 4.5), but the effect is limited and is unlikely to have a significant effect on IC performance.

While no clear reduction in noise is shown between the basic configuration and the modified board, the spectra for the power supply and the output for both configurations do not show an obvious pattern that could be linked to a noise source affecting the signal in all cases. There is significant 1/f noise visible on the power supply in both cases, but this is not seen on the front-end output. While this appears to show that there is no specific noise source or fault causing the increased noise levels, it also means that there is no obvious way of reducing the noise level without design changes to the IC.

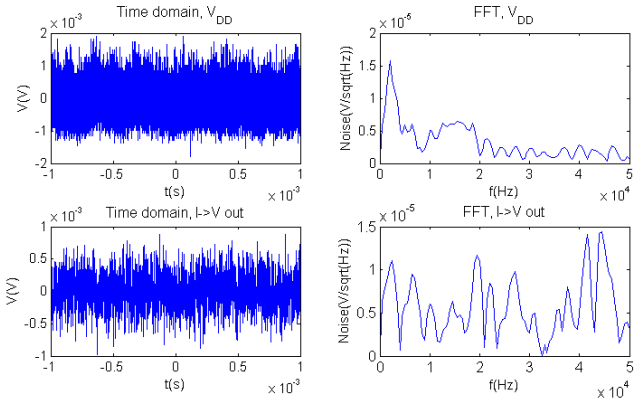

Figure 4.30: Noise on power supply and I->V output, time and frequency domain, shielded power supply

Table 4.5 shows the RMS noise voltage on each point, for the original configuration and after applying each change to reduce the noise. Each change shows some reduction in either output or power supply noise, but the reductions are very small and could be down to measurement variation rather than a significant change. For the final characterisation shown in Section 4.3, the shielded power supply and additional capacitors were retained, but the battery power supply was not used.

Noise (V_{RMS})	V_{DD}	V_{I-V}
Original configuration	0.404	0.281
Shielded power supply	0.399	0.270
Additional capacitors	0.398	0.239
Battery power supply	0.352	0.301

Table 4.5: Voltage noise from logarithmic front-end power supply and output, before and after applying noise reduction modifications

4.5 Doppler Experiments

4.5.1 Doppler Imaging Setup

The optical setup used for the following tests is a bench top system using a 5 mW HeNe laser and a line generator, producing a 30 mm line in a fixed position at the imaging target. The line on the target was then imaged onto the IC detector using a single 25.4 mm focal length lens, with a diameter of 25.4 mm. Data acquisition was performed using the ADC and FPGA back-end of the Moor Laser Doppler Line Scanner (LDLS) system, in order to investigate the feasibility of integrating the IC into an imaging system.

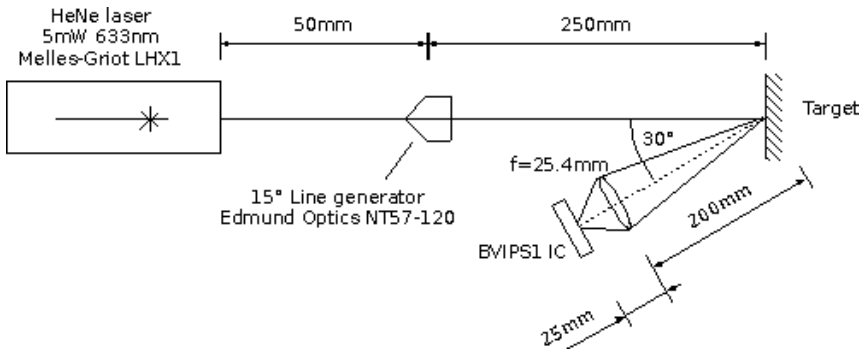

Figure 4.31: Equipment setup for imaging using the BVIPS1 IC

Four test targets were used in this setup. Two non-biological targets were used - a motility

phantom, which consists of a suspension of polystyrene spheres that mimic blood flow through Brownian motion, and a static reflector, which is a piece of plastic which has similar colour and reflectance to skin. This is used to represent the no-flow condition. Two corresponding biological targets were used, these being an occluded and non-occluded finger. This is less repeatable than the phantom tests, as the flow for a finger will be affected by physiology, temperature, etc. However, this is also fairly representative of the changes required to be measured by the final system. For instance, from biological samples the flux measurement would not be expected to go to zero with no flow, due to some randomness in the scattering of light through static tissue. This non-zero flux reading is referred to as the biological zero [Zhong et al., 1998, Liebert and Maniewski, 1996]. Also, it is possible to occlude flow to the finger during measurement to observe transitions between high- and low-flow in real-time.

4.5.2 Raw Signals from Doppler Imaging

Figure 4.32 shows the output from the BVIPS1 IC measured using the FPGA and ADC in the Moor LDLS system. The changes in the AC component of the signal from a motility target, a static skin phantom, a finger and an occluded finger can be seen. The lower flux signal expected from the occluded finger and particularly from the static phantom can be seen in the lower peak-peak of the signals, at approximately 80 mV for the phantom and 120 mV for the occluded finger . Conversely, the motility phantom and the non-occluded finger give signals with a peak-peak of $\sim 200 - 300$ mV.

However, the falling spike in the signal from the motility target at 3 ms shows a fault that occurs in this circuit under certain input conditions. This issue is discussed in Section 4.6. However, this does not appear to mask the change in Doppler signals from the different targets.

4.5.3 Processed Flux output from Doppler Imaging

The images in Figure 4.33 are produced using the test system described in Section 4.5.1, Figure 4.31, with the flux results calculated on the FPGA and sent to the PC. Further detail of the flux values from individual pixels is shown in Figure 4.34. The system uses fixed line illumination, such that the figures are built up by repeated sampling of the linear array, giving a map of pixel response over time. Figure 4.33a shows the plot from the static

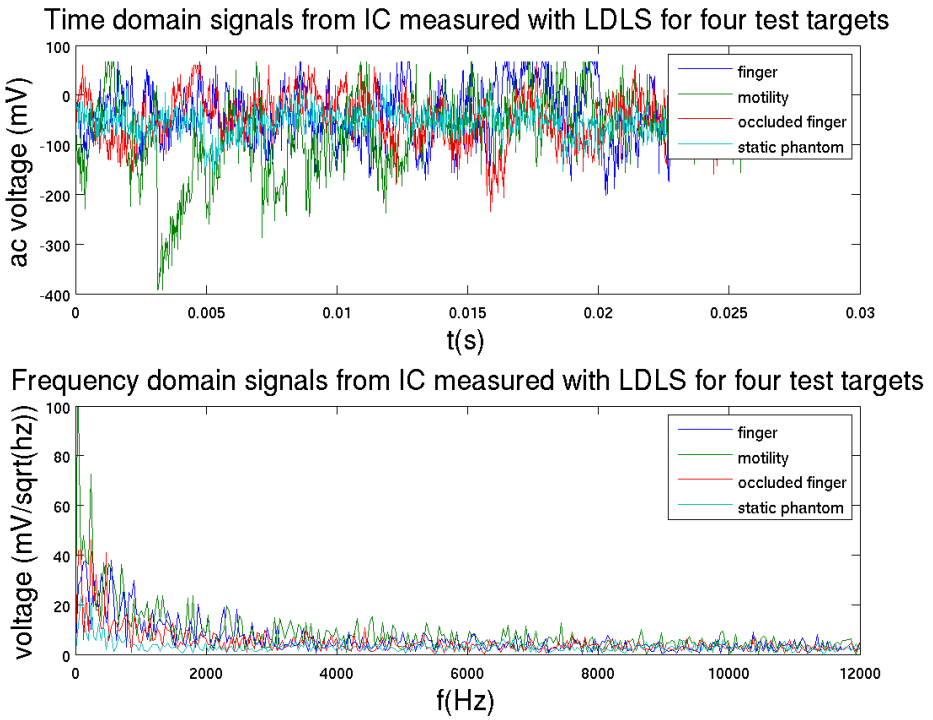

Figure 4.32: Doppler signals recorded from logarithmic pixels on BVIPS1 IC, with data captured using the Moor LDLS FPGA system

skin phantom, while Figure 4.33b shows the plot produced from the motility target. The darker blue of the skin phantom plot shows the lower flux measurement from this target. Figure 4.33c shows a plot taken from fingertips, with blood flow occluded (OCC) and released (REL) twice during the imaging period. In this plot the changes during occlusion and release can be clearly seen, even though no spatial information can be seen - the line was projected across part of three fingers, and no spatial pattern corresponding to the target can be seen in the plots produced. The technique was repeated with the line projected on a part of the back of the hand including veins, and again no spatial relationship could be seen. This shows the IC is capable of measuring changes in flux, but cannot be used to detect veins in this configuration. This may require a change to the optics such as the use of a higher powered or longer wavelength laser to increase light penetration into the skin, or changes to the design of the IC itself.

A reduction in spatial resolution is expected from a full-field system compared to a single point system [Steenbergen et al., 2010]. The Doppler shifts which allow the measurement of flow are caused by light being scattered by moving blood cells near the skin surface. This

scattering means that light incident on a single point is reflected from a wider area of the skin after scattering through tissue (see Section 1.6, Figure 1.2). In a full field system, the reflected light from any single point can be seen as the sum of the scattered light from all surrounding tissue within the area over which this scattering occurs. Additionally, this area will be larger for higher penetration depths caused by higher power and longer wavelength - features which are desirable for this application, as higher penetration depth is required to image veins rather than just surface flow.

One issue shown by these plots is that calibration of fixed pattern noise is required to prevent variation in the signal due to actual variations in the target being masked by pixel-pixel variation of the IC.

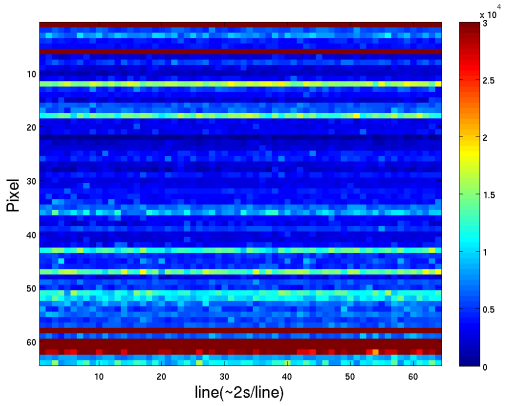

(a) Static reflector

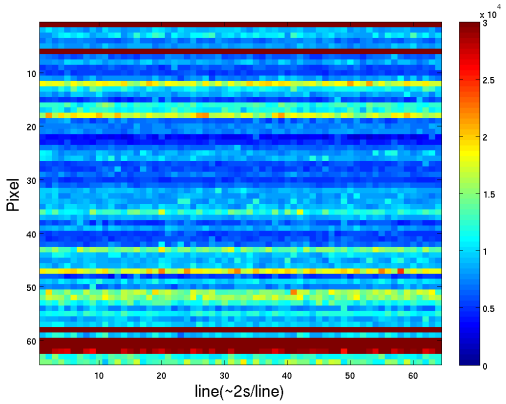

(b) Motility phantom

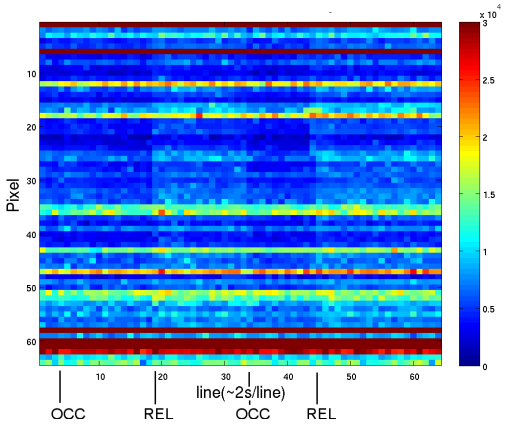

(c) Finger, free-occluded-released-occluded-released

Figure 4.33: Processed flux readings across array and over time from logarithmic pixels on BVIPS1 IC, with data captured using the Moor LDLS FPGA system

4.5.4 Quantified Flux Measurements and SNR

Figure 4.34 shows the flux reading from an individual pixel over time (i.e. one row of the previous plots), along with the mean flux from all pixels of the array at each time point. Plots are shown for the static reflector (i.e. zero flow, giving a measure of noise on the processed flux value), motility reflector (giving a reasonable maximum flow value), and occlusion/release of flow in a hand, showing the changes to be seen between high and low biological flow.

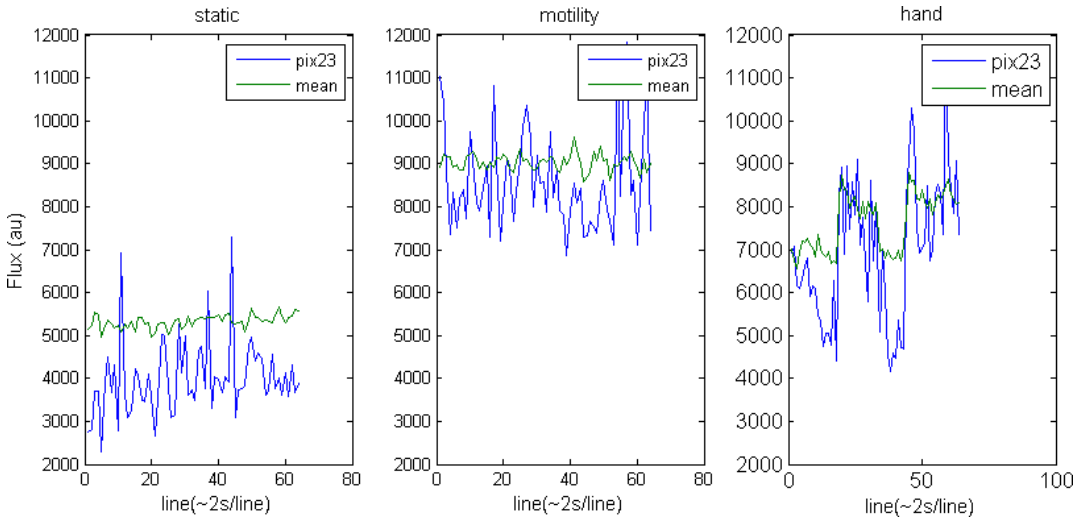

Figure 4.34: Plots of flux from a single pixel of BVIPS1 IC logarithmic array

Table 4.6 gives quantified flux readings for each plot above, giving a mean and standard deviation for each flow level, effectively giving signal and noise levels for each.

For the single pixel plots, the standard deviation of the flux readings is approximately 1000, compared to a minimum flux value of 4000 and a maximum of approximately 8500. This gives a SNR of between 4 and 8.5, showing that the design is capable of detecting flux changes, but with poor flux resolution and poor repeatability.

The averaged results show significant improvement, with standard deviation of approximately 135 compared to flux readings between 5300 and 9000, hence SNR between 40 and 67. However, averaging would be expected to improve the SNR performance, and the level described here (averaging all pixels) effectively turns the IC into a single point detector. Some level of imaging could be applied, such as averaging adjacent pixels to improve SNR at the expense of spatial resolution, or averaging over time to improve SNR at the expense of frame rate. However, these options are available to all systems so do not represent an advantage of the integrated array in particular.

The commercial systems described in Section 2.7 have SNRs between 17 (Perimed PIM3, 250 ± 15 PU) and 10 (Moor LDI2, 0-5000 PU ± 10%), although the specification should represent a worst case accuracy measurement. The IC produced here is therefore significantly worse than the commercial systems, but the SNRs have a similar order of magnitude.

Target	Mean flux (all)	Mean flux (1pixel)	Flux std.dev. (all)	Flux std.dev. (1pixel)
Static	5326.9	3996.8	152.2	879.6
Motility	9035.0	8598.3	188.0	1177.6
Hand (Free)	8136.1	7731.3	139.6	920.6
Hand (Occ)	6900.3	5063.9	55.2	769.4

Table 4.6: Mean and standard deviation (over time) for one pixel of BVIPS1 IC logarithmic array

4.5.5 Flux Measurements using Opamp Pixels

Changes in flux were also measured using the opamp section of the IC described in Section 4.2.4. Figure 4.35 shows time and frequency domain output from this section of the IC with high and low flow and with high and low power illumination.

The time domain signals show that the opamp circuit, due to its lack of a high-pass filter element, is more susceptible to changes in DC light level - the voltage change between lowest detected power (10 mW laser illuminating a static phantom) and highest detected power (20 mW laser illuminating a motility target, which being white gives a higher reflected light level) is 1.7 V. This can make the changes in AC signal harder to identify by eye, which also shows the higher dynamic range requirement for the ADC to be used with this circuit.

The DC voltage output as seen in Figure 4.35 does not appear to be linear - a doubling of illumination power does not lead to a doubling of output voltage. This is expected as the range of output voltages seen is large compared to the output voltage range over which linear behaviour was observed during characterisation (shown in Figure 4.11), which ranged from approximately 1.4 to 3 V. The minimum output voltage shown in Figure 4.35 is lower than the minimum seen in the characterisation plot, as the offset voltages used required adjustment for the different range of light levels encountered.

The frequency response plot shows a clear change between high and low flux with either power, although the higher power does make the change larger. This shows that the opamp front-end is a viable alternative to the log pixel as long as the dynamic range of the ADC is satisfactory.

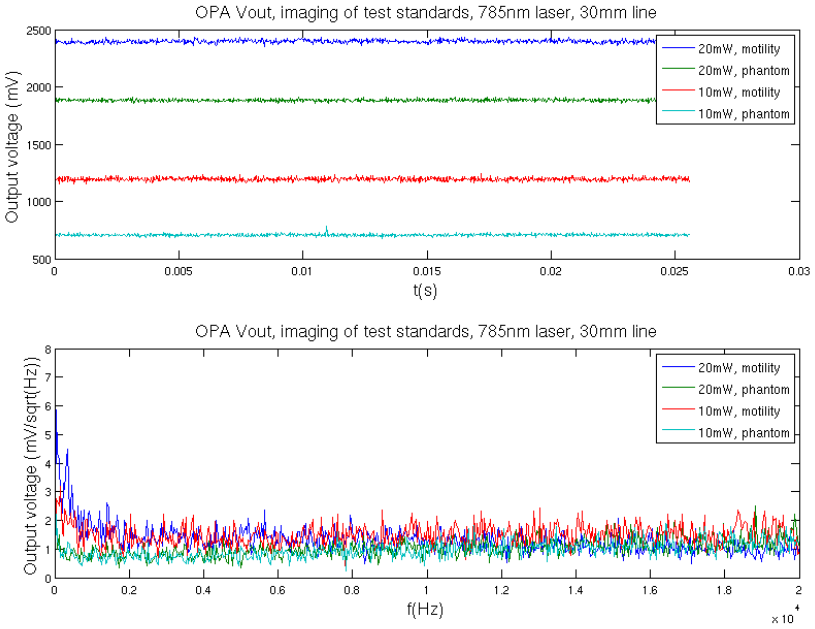

Figure 4.35: Doppler signals recorded from opamp pixels on BVIPS1 IC, with data captured using the Moor LDLS FPGA system

Figure 4.36 shows the processed flux measurement from the opamp output over occlusion (OCC) and release (REL) of flow in a finger, using the opamp front-end. Again, a clear change in flux can be seen. However, it can also be seen that there is a downward drift on the signal which does not appear related to changes in flow. This may be due to changes in DC level (e.g. a gradual movement of the finger such that the reflected light level falls), showing that as this pixel does not perform the normalisation inherent to log pixels, normalisation either on- or off-chip will be required.

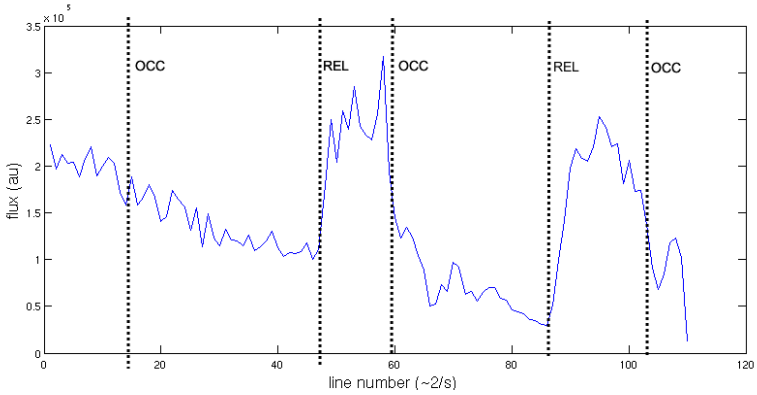

Figure 4.36: Processed flux readings over several cycles of occlusion and release of blood flow in a finger, from an opamp pixel on BVIPS1 IC, with data captured using the Moor LDLS FPGA system, 20 mW 785 nm laser, 30 mm line

4.6 Discussion and Fault Investigation

This section summarises the behaviour of the IC observed during characterisation and testing, considering the behaviour compared to that expected from simulations, and the effect this could have on an imager using this IC. In addition to this, a number of issues not directly shown by the testing already given are discussed, and testing performed to investigate this behaviour is described.

4.6.1 Discussion of Logarithmic Pixels

The logarithmic pixel design (as used in arrays on the right- and left-hand sides of the IC) has been shown to be capable of detecting changes in blood flow when using line illumination with a power density similar to that expected in a line scanning vein location system. Characterisation results show that the behaviour mostly follows that expected from simulation - the DC response from all outputs is generally similar in gradient and range to that expected, although the measured values are frequently offset from those seen in simulation. As sensors such as this are mainly concerned with detecting modulated light due to the Doppler shift, offsets such as those seen are a minor issue as long as the AC behaviour is as required.

The AC behaviour found from characterisation of the front-end itself (i.e. not including amplifier stages) showed higher gain than expected from simulation, and with a frequency response similar to that seen in simulations. This confirms the suitability of this design for use in LDBF integrated sensors. While the pixel noise level was higher than that expected, this did not prevent the detection of changes in blood flow.

The AC behaviour measured from the g_mC output (the main output after amplification and filtering of the signal) showed lower overall gain than expected from simulation, despite the increased gain from the front-end. The frequency response also showed a higher low-frequency cut-off, above 1 kHz rather than the 200 Hz design value. However, these differences did not prevent successful detection of blood-flow changes, so are not considered major faults. The decrease in gain is accompanied by a slight reduction in input referred current noise between the front-end output and the g_mC output, suggesting that the decrease in gain does not increase SNR. The reduction is also not sufficient to require a specialist ADC (i.e. higher resolution), as would be required if the output voltage AC magnitude was not significantly larger than that from the front-end output.

However, the difference observed between measured and simulated AC amplifier behaviour are sufficient to require further testing. This testing also investigated behaviour seen during measurements of the IC that did not appear to be consistent with the characterisation results. Details and results of this testing are shown in Section 4.6.3.

4.6.2 Discussion of Opamp Pixels

Characterisation of the opamp based linear pixels shows good agreement with the simulated behaviour. The most significant difference is that the saturation occurs more gradually when the input photocurrent rises above or falls below the range in which linear behaviour is predicted. This is thought to be due to reduced gain in the front-end stage, leading to a reduced voltage swing at the input of the gain stage opamp, and hence a wider range of photocurrents in which the gain stage opamp output voltage is not saturated.

The successful detection of changes in blood-flow using the opamp demonstrates that this type of circuit can be implemented on an integrated LDBF sensor. This testing verifies that the combination of gain, noise level and bandwidth achieved are sufficient for LDBF applications, and also that the more limited range of operation against DC input photocurrent compared to logarithmic pixels was not overly problematic.

4.6.3 Spikes on Logarithmic Pixel AC amplifier Output

This section discusses an aspect of the logarithmic pixel behaviour that was not demonstrated during characterisation, and that only affected the output voltage signal in certain situations and configurations.

4.6.3.1 Faults Observed During Flux Imaging

During testing of the logarithmic pixel section of the BVPIS1 IC with the Moor LDLS back-end, a problem was observed during switching of the IC multiplexer. This problem did not always occur in the same way, having a dependence on light level and modulation depth. Figure 4.37 shows output voltage from the RHS G_mC output during multiplexing, with the IC directly illuminated by an LED driven by a DC current. This provides un-modulated illumination at a slightly higher intensity than that expected from an imaging setup where

light is reflected from a target. The IC array was sampled repeatedly, where in one set of sampling all pixels are sampled 1024 times at a sampling frequency of 40 kHz each (a single ADC sampling at 2.56 MHz), with sampling of all pixels interleaved (the entire array is sampled once before the multiplexer switches back to pixel 0 for the second sample). This gives a time to sample all pixels 1024 times of approximately 25 ms. Between blocks of sampling there is a delay while the signals are processed, and where the mirror in the scanning imager would be moved.

As the illumination intensity is fixed, the only changes expected in the output voltage during such a test would be due to the fixed pattern noise of the array, such that a slightly different DC voltage is measured from each pixel. In addition, transients from the switching can propagate to the output voltage causing some oscillations on switching, but these oscillations should be small, and should fade away in less than 200 ns in order to not affect the sampled output voltage. Ideally, fixed pattern noise would be negligible, and no switching transients should be present such that no change in output voltage would be seen during multiplexing.

The bottom plot of Figure 4.37 shows the voltage before, during and after one block of multiplexing. The top plot shows the voltage over a short period (0.4 ms) at the start of multiplexing, showing repeated spikes on the output voltage (note that t=0 on the plot refers to the time at which the oscilloscope was triggered, not the start of the multiplexer switching). This effect became more severe (larger spikes,causing the output voltage to drop by a larger amount during multiplexing) when the DC photocurrent was increased.

Figure 4.38 shows the output voltage seen during a similar test, using a modulated visible red LED to provide illumination, giving a 1 kHz modulated photocurrent at approximately 20% modulated depth, with a similar power to that used for the test shown in Figure 4.37. The multiplexer switching speed was also reduced - the sampling frequency was unchanged, but the multiplexer was switched after 256 samples rather than after every sample. The points at which the multiplexer switches are shown. It can be seen that on switching the output voltage drops by ~ 0.15 V. It then begins to return to the original voltage, but this takes ~ 5 ms. This long settling time (compared to a multiplexer switching period of around 400 ns to sample 64 pixels at 40 kHz per pixel) means that the output voltage will sampled near the bottom of the voltage spike, giving an inaccurate sample.

The effect of this problem can be seen in Figure 4.39. This shows the output from one pixel with the same modulated LED illumination (1 kHz visible red LED, 20% modulation depth),

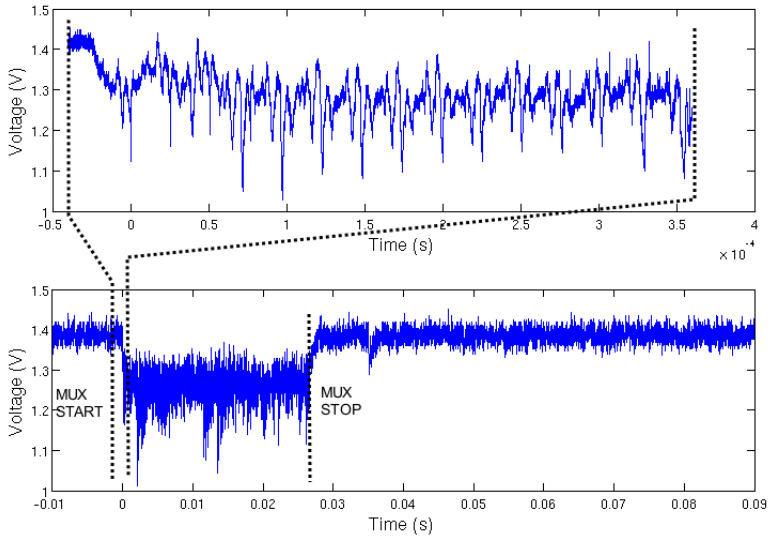

Figure 4.37: Voltage measured before, during and after multiplexing, RHS logarithmic pixel g_mC output of BVIPS1 IC

with and without the multiplexer in use. Without multiplexing a large AC component is observed at 1 kHz. When the pixels are multiplexed and the signal from the same pixel is shown, the AC signal is mostly suppressed, and there is also a drop in DC level.

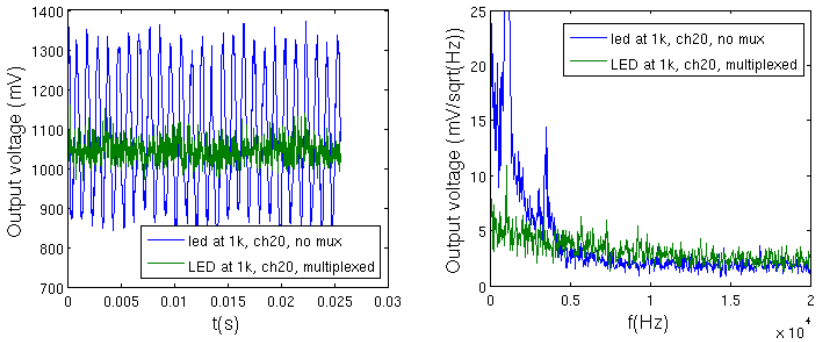

Figure 4.39: Voltage measured with and without multiplexer switching, 1 kHz modulated LED illumination, RHS logarithmic pixel g_mC output of BVIPS1 IC

4.6.3.2 Fault Location

To investigate the configurations under which spikes occur, a series of tests were performed on other sections of the IC with various illumination conditions. The points at which signals were recorded were:

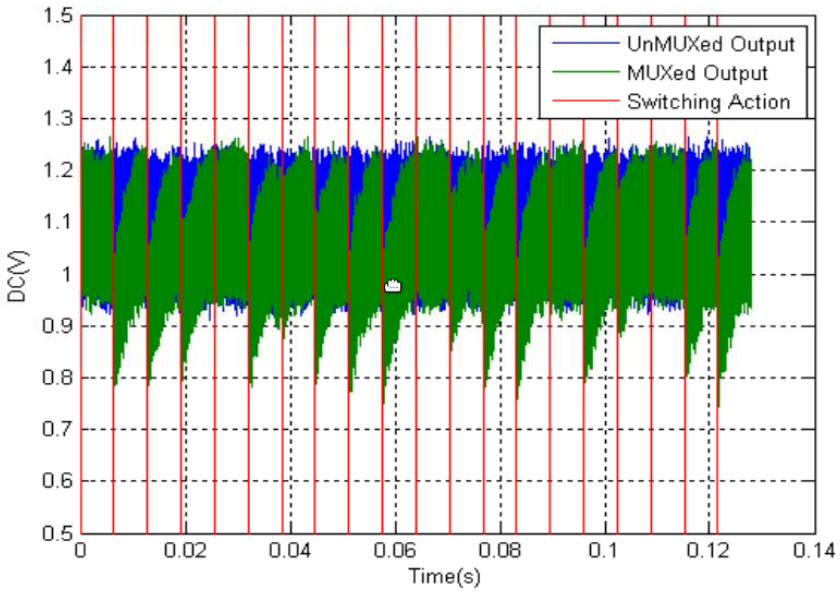

Figure 4.38: Voltage measured during multiplexing at reduced speed (~6 ms/256 samples between MUX switching), RHS logarithmic pixel g_mC output of BVIPS1 IC RHS logarithmic pixels

- Right hand side, main output (2 multiplexer stages) - this is the output used in the tests shown in Figures 4.37,4.38 and 4.39.
- Right hand side, measured at one of the four outputs from the first stage multiplier (see block diagram in Figure 4.2) - this is used to see if the number of multiplexer stages used affects the spikes.
- Left hand side (1 multiplexer stage) - this should give the same results as the above two configurations, but will show if the fault is due to a difference in the LHS/RHS layout, as these are intended to be duplicate designs but have some layout changes to fit in with other LHS circuits.
- Opamp - this may show if the fault applies to an individual circuit, or if similar behaviour occurs for other circuits on the IC.

The test configurations were:

- LED, DC drive - this is used to provide a constant and low power source of illumination, giving a DC photocurrent of approximately 8 nA
- LED, DC with AC modulation - this gives an optical modulated input at 2 kHz, giving 8 nA DC photocurrent with a 20% modulation depth.

- Laser, constant - this gives an illumination intensity (and therefore DC photocurrent) similar to that expected in normal Doppler imaging ($\sim 15\,\mathrm{nA}$), with no AC component
- Dark - this eliminates any effects of stray light, and also produces the highest AC gain from the logarithmic photodiodes.
- Electrical input - this is used to investigate the later stages of the pixels (HDA and g_mC) in isolation of the front-end behaviour. This test was performed with either a 1 V DC only input voltage, and with an additional AC voltage, aiming to mimic the voltages produced by the front-end for both constant intensity and modulated optical sources. The AC amplitude used was set higher when input to the g_mC (500 mV, rather than 20 mV when input to the HDA), simulating the gain of the HDA. The frequency used was 2 kHz for all modulated signals.

Table 4.7 shows a summary of the results of this testing. A tick shows that the configuration behaved as expected (no spikes on the output). A cross shows a large number of significant spikes. Intermediate results (some small spikes, but of a similar order of magnitude to expected noise) are marked '~'.

IC Output Point	IC Sections Tested	Electrical Input	LED, DC	LED, modulated	Laser, Constant (DC)	Dark
g_mC RHS (64x1) 2 MUX stages	RHS front-end, HDA, g_mC , all MUX stages	-	X	√	X	√
g_mC RHS (64x1) 1 MUX stage	RHS front-end, HDA, g_mC , first MUX stage	-	X	√	X	√
g_mC LHS (32x1) 1 MUX stage	LHS front-end, HDA, g_mC , MUX	-	X	√	X	√
Opamp 1 MUX stage	Opamp front-end, gain stage and MUX	-	√	√	√	√
logarithmic front-end LHS (32x1) 1 MUX stage	LHS logarithmic front-end, 1 MUX stage	-	√	√	√	√
g_mC LHS (32x1) 1 MUX stage	LHS HDA, g_mC, 1 MUX stage	1 V DC into HDA	X	X	X	X
g_mC LHS (32x1) 1 MUX stage	LHS HDA, g_mC, 1 MUX stage	1 V DC + 20 mV p-p @ 2 kHz into HDA	~	~	~	~
g_mC LHS (32x1) 1 MUX stage	LHS g_mC, 1 MUX stage	1 V DC into g_mC	√	√	√	√
g_mC LHS (32x1) 1 MUX stage	LHS g_mC, 1 MUX stage	1 V DC + 500 mV p-p @ 2 kHz into g_mC	√	√	√	√

Table 4.7: Testing of BVIPS1 logarithmic pixel circuits in various conditions to investigate MUX spike occurrence - '√' shows correct (artifact free) operation, 'X' shows artifacts occurred

From Table 4.7 it can be seen that the spikes occur mainly when there is no AC component in the input signal. Also, spikes do not occur on the opamp output, so the fault appears to be in the logarithmic pixel circuitry. There is no difference between LHS and RHS log pixels, or between one and two stages of multiplexing.

From the electrical testing, it can be seen that the fault only occurs when a signal is applied to the HDA input, and not when applied to the g_mC input. This suggests that the problem may be related to the HDA. Figure 4.40 shows the output signals from testing with a 1 V DC electrical input to the HDA and the g_mC. For input to the g_mC, small changes can be seen when switching between pixels due to pixel-pixel variation, but there are no spikes. However, when the signal is applied to the HDA there are very large spikes on all transitions. This confirms that the fault is in the HDA, and mainly affects DC operation.

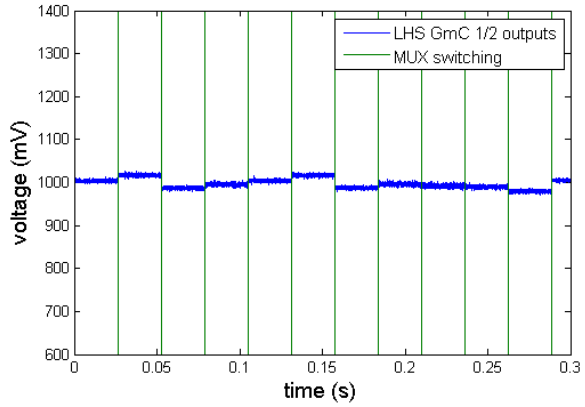

(a) 1 V DC Electrical input to g_mC, dark

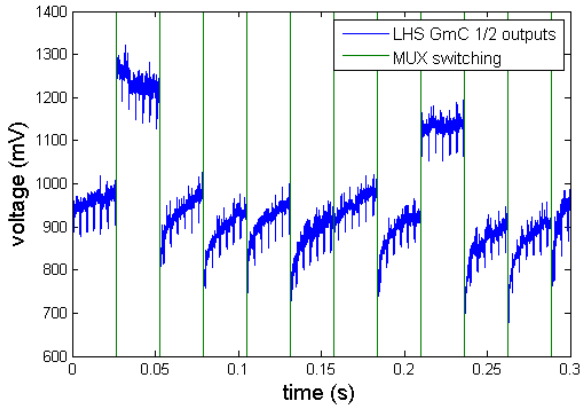

(b) 1 V DC Electrical input to HDA, dark

Figure 4.40: Output voltage during multiplexer switching, measured at LHS logarithmic pixel g_mC output of BVIPS1 IC, with a 1 V DC Electrical input applied to the g_mC input or the HDA input.

Figure 4.41 shows that there is a link between the spikes occurring and the AC component of the signal. This shows the output signal when a 20 mV, 20 kHz peak-peak AC component is added to the 1 V DC electrical signal input to the HDA. The output signal would be

expected to consist of the output voltage from the DC photocurrent as shown in Figure 4.40b, with an additional AC output voltage. The envelope of the resultant signal can be seen in Figure 4.41a. The output signal has a similar shift in DC levels to that seen when applying a DC electrical input to the g_mC, where the HDA is not included in the signal path (Figure 4.40a). This is caused by fixed pattern noise giving each pixel a different DC output voltage. However, the envelope of the signal between switching is flat, showing that the output voltage spikes seen when applying a DC input voltage to the HDA (shown in Figure 4.40b) are not present. A close up of one of the switching operations seen in Figure 4.41a (at 0.05 s) is shown in Figure 4.41b. This shows that there is still a small perturbation of the signal on switching, but it is reduced in size and the signal returns to the original level in less than 0.2 ms, compared to approximately 5 ms for the DC input case.

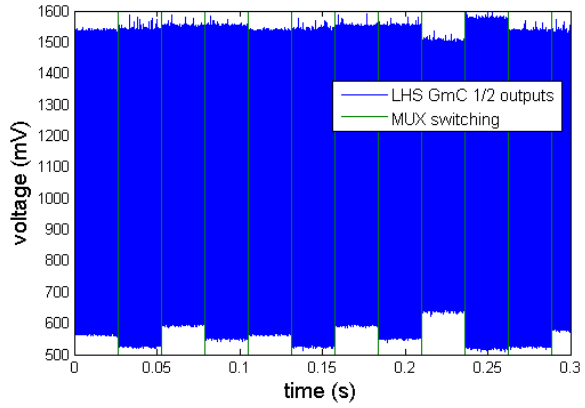

(a) Modulated electrical input to HDA, dark

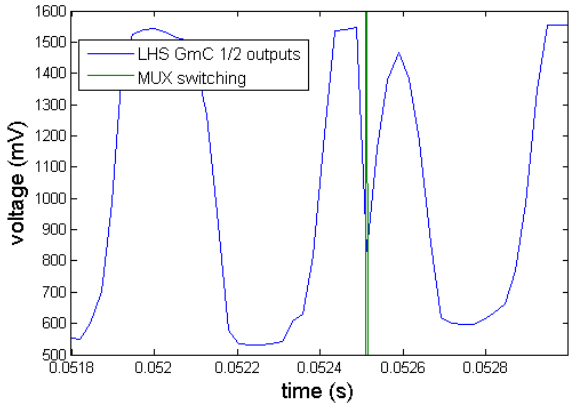

(b) Modulated electrical input to HDA, dark, closeup of one address transition

Figure 4.41: Output voltage during multiplexer switching, measured at LHS logarithmic pixel g_mC output (after first stage multiplexer only) of BVIPS1 IC, when a 1 V DC & 2 kHz/20 mV AC pk-pk electrical signal is applied to the HDA input.

To confirm that there is a link between AC signals and the spikes, Figure 4.42 shows the output signal when illuminated with an LED resulting in a DC photocurrent of approximately 8 nA, with and without a 2 kHz modulated drive signal giving a 20% modulation depth. The optical signal gave a DC photocurrent of approximately Again, the test with modulated input (Figure 4.42b) shows reduced spikes - the troughs of the output waveform form a horizontal line, as opposed to the DC situation (Figure 4.42a) where there are downward spikes of 0.3 V. The peaks of the output signal with a modulated optical input shown in Figure 4.42b do not form a horizontal line, but there is no trend in this behaviour that corresponds to the spikes seen during this testing. The varying height of the output signal is likely to additional noise caused by the optical rather than electrical signal (leading to additional noise sources in the front-end), combined with non-linear behaviour of the amplifier causing the signal at lower voltages (the troughs of the signal) to be compressed, resulting in a flatter envelope than that seen from the peaks.

One other aspect of the fault is shown by Figure 4.42 - the spikes do not only occur on switching, but also occur between multiplexer switching. Also, electrical testing has shown that the fault occurs when a signal with small or no AC component is applied to the HDA, and does not occur when the signal is applied to the g_mC despite the same signal path through the multiplexers. This suggests that the fault is in the HDA circuit itself, rather than being a multiplexing problem. It is possible that switching noise on multiplexing triggers the fault, hence the fault initially appearing to be a multiplexer problem.

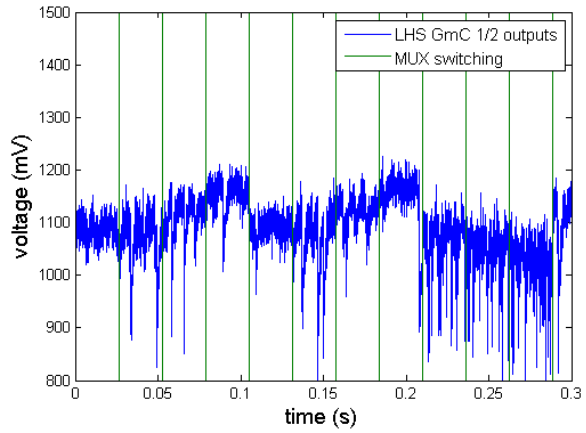

(a) LED constant illumination, g_mC output

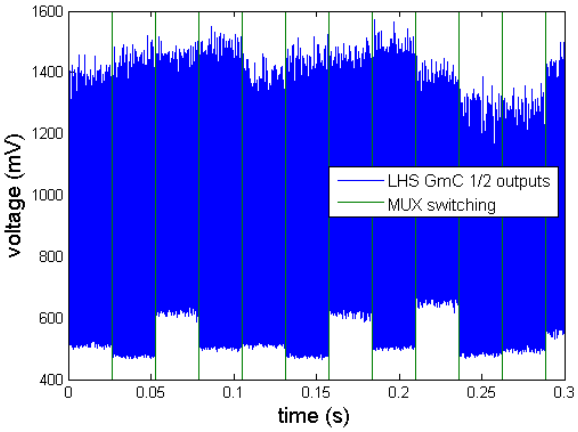

(b) LED modulated illumination, g_mC output

Figure 4.42: Output voltage during multiplexer switching, measured at LHS logarithmic pixel g_mC output (after first stage multiplexer only) of BVIPS1 IC, with 8 nA DC photocurrent and with/without 2 kHz AC photocurrent at 20% modulation depth

4.6.3.3 Cause of AC Amplifier Artifacts

The switching artefacts seen during measurement have not been seen during simulation, making determining their cause problematic. A similar design to the HDA circuit has been used on previous ICs, where this behaviour has not occurred [He et al., 2009].

The main change to the HDA for this IC was an increase in the capacitor size to reduce the low-frequency cut-off. It is possible that this increase in size causes unexpected behaviour, due to a further increase in an already large RC time constant from the capacitor and the inverted-inverter circuit. One possible cause, given the link shown between the

fault and modulation of the input signal, is that the inverter-inverter-capacitor arrangement within the HDA is in some respects similar to a switched capacitor (SC) circuit [Allen and Holberg, 2002]. The basic schematic of a switched capacitor circuit is shown in Figure 4.43a, with its equivalent circuit shown in Figure 4.43b. A block diagram of the inverted-inverter HDA is shown in Figure 4.43c. If C in the HDA circuit is considered to be equivalent to C_1 in the SC circuit, and the two switches in the SC circuit (ϕ_1 and ϕ_2) are equivalent to the two transistors of the inverted-inverter, then a similarity between the circuits can be seen, although the output of the filter connected to the inverting input of the HDA input OTA is taken from between C_1 and C_2 rather than V_{out} in the SC filter schematic.

SC circuits require drive signals consisting of two non-overlapping clocks to control for the switches with higher frequency than the signal bandwidth, [Gu, 2007, Allen and Holberg, 2002], which may mean that when the HDA is presented with a DC only signal, there is no AC drive signal and hence the circuit malfunctions. One of the major problems with SC circuits is charge injection on switching [Gu, 2007], which may explain why the larger capacitor causes this problem - the larger capacitor may increase the amount of charge injection, hence the severity of the spikes. However, the similarities between these circuits do not present an obvious means of simulating this effect. In addition, the simulations performed would have shown the effects of charge injection, suggesting that either this effect is not having any significant impact, or that it only causes the spikes seen when combined with the non-ideal behaviour of the high resistance of the inverted-inverter feedback network.

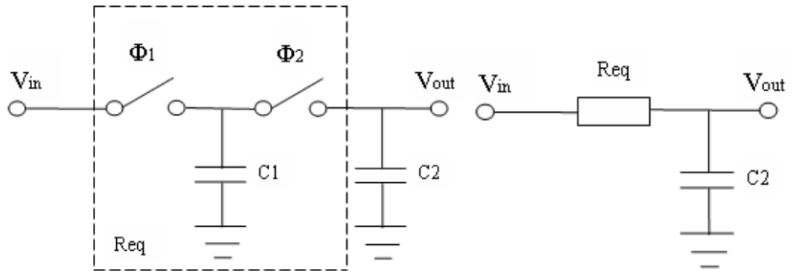

(a) Switched capacitor circuit

(b) Equivalent circuit of SC circuit

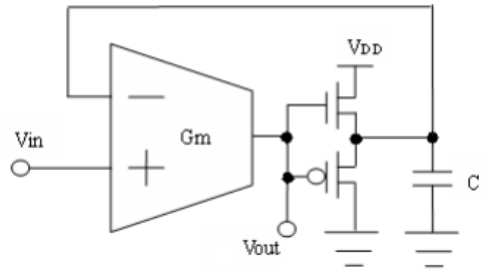

(c) HDA showing inverted-inverter-capacitor feedback network

Figure 4.43: Switched capacitor filter circuit compared to HDA feedback network

While this change to the design was made to reduce the low frequency cut-off, the characterisation results shown in Section 4.3.7.2 show that this has been ineffective, due to the known problems in accurately simulating the HDA, which relies on non-linearities to function [Mead, 1989]. Because of this, and because of uncertainties regarding the cause of the fault, further iterations of this circuit revert to the previously tested HDA design [He et al., 2009].

4.6.4 Fixed Pattern Noise

It can be seen from the logarithmic pixel flux images (Figure 4.33) that some pixels do not show a clear change in flux between occlusion and release, and there are some pixels that show a red line across the image, showing a flux measurement that is very high at all times. Some of this variation can be removed by calibration, however if some pixels have either insufficient gain or excessive noise then this may prevent any signal from being detected regardless of calibration.

To investigate this further a modulated LED was added to the imaging test setup (shown in Section 4.5.1, Figure 4.31). This results in a DC photocurrent, and hence DC bias point, similar to that expected during blood flow imaging (10 nA I_{DC}) by using the same 5 mW,

633 nm laser to illuminate the static skin phantom used in Doppler measurements, and with the same imaging optics. An additional visible red LED is then used to directly illuminate the IC, producing a modulated photocurrent. The LED increases the DC photocurrent, but as its power is lower than that from the laser and it is not focused with a lens, this change is relatively small (approximately 3 nA, giving a total photocurrent below the higher estimated photocurrents given in Table 4.1). The LED was driven from a signal generator with a DC voltage of 2.5 V and an AC voltage of 100 mV p − p or 1000 mV p − p. This setup is used to give an approximation of the photocurrent expected from blood flow imaging, but with more direct control of the AC modulation depth. The lower AC drive voltage gives a modulation depth of approximately 10%, rising to approximately 50% for the higher AC drive signal.

Figure 4.44a shows the signals from pixels 1 and 2 for the smaller modulation depth. In the absence of fixed pattern noise, both output signals would be identical. Pixel 1 is one of the pixels that does not show a change in flux between occlusion and release. It can be seen from the time and frequency domain plots below that pixel 1 has insufficient SNR to detect the modulated LED signal, suggesting that it is not sensitive enough to detect Doppler signals from blood flow, with noise rather than a spike at 1 kHz being the dominant feature of the output spectrum.

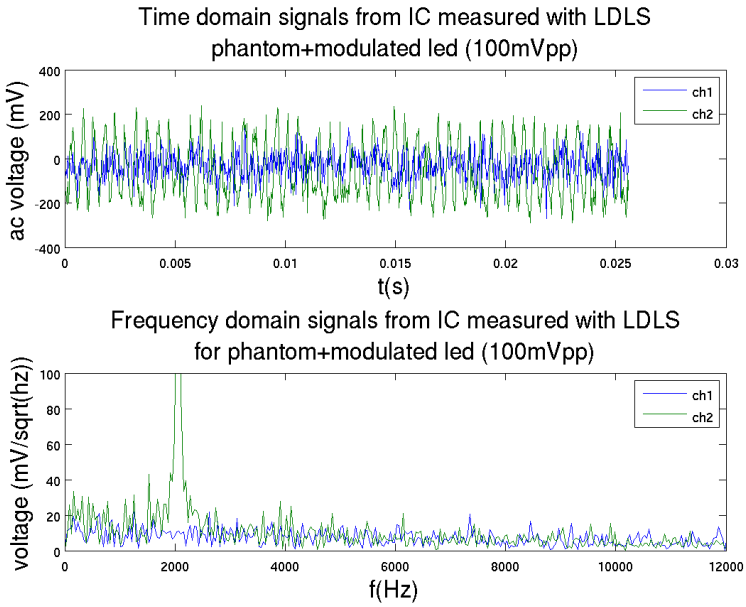

(a) Illumination from a reflected laser and a modulated LED driven by 100 mV p − p signal

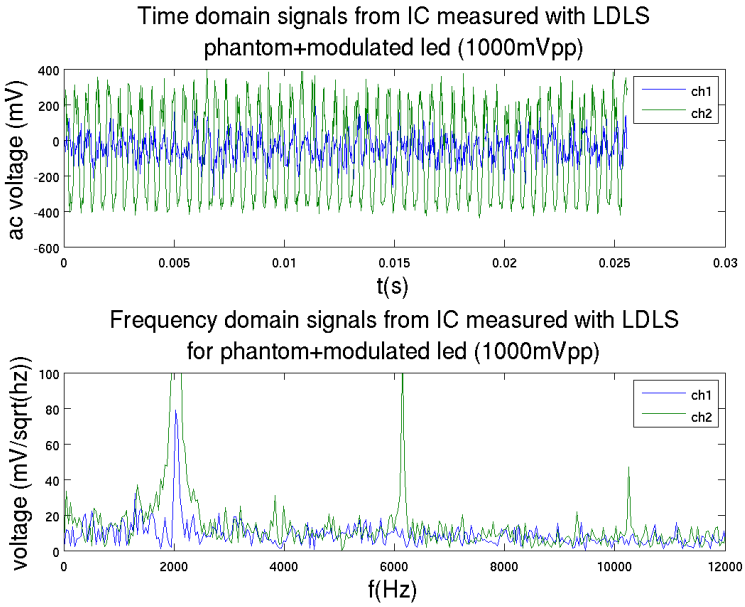

(b) Illumination from a reflected laser and a modulated LED driven by 1000 mV p − p signal

Figure 4.44: Time and frequency domain signals from pixels 1 and 2 with illumination from a laser reflected off a static target and an additional modulated LED giving different modulation depths, main g_mC output of RHS logarithmic pixel, BVIPS1 ICs

Figure 4.44b shows the same signals with the higher modulation depth (40%), which is

considerably larger than the modulation depth expected from blood flow imaging (10-15%). The modulated signal can be seen on the pixel 1 output, but the AC output signal is still very low. In comparison, the AC signal from pixel 2 is now high enough to cause distortion, although some of this may be from over-driving the LED, which would result in the odd harmonics that can be seen in the frequency domain output. The successful detection of modulated light shows that pixel 1 is functioning partly as intended, albeit with severely impaired performance, as it is capable of detecting modulated light. This suggests that there is no major fault in the pixel which would be expected to cause behaviour totally different to that seen from the other pixels. Therefore, this behaviour does appear to be severe fixed pattern noise rather than a design or manufacturing fault.

To check that the fixed pattern noise is due to manufacturing variations rather than to a fault on the design, the flux measurements shown in Figure 4.33 were repeated with two different ICs (30 ICs were made using this design, allowing IC-IC variation to be investigated, or for later use of the IC in other experiments).

The resulting plots in Figure 4.45 show pixel response maps for measured flow over two cycles of occlusion and release of blood flow in a finger. As for the blood flow measurement results shown in Section 4.5.3, Figure 4.33c, darker coloured bands representing occluded blood flow can be seen, showing successful detection of blood flow changes by all ICs. However, all plots also show horizontal lines caused by the variation in flux reading from pixel-pixel being greater than that seen between occlusion and release. While a pattern of this type occurs in all cases, different patterns of low and high sensitivity pixels are seen on each plot. The differences between the images show that the effects are not related to the position of the pixels within the IC. If that had been the case, for instance if the same pixels on each IC had low sensitivity, or if low-sensitivity pixels tended to be in one area of the array, this would have shown a problem with the overall layout of the IC. Instead, the fixed pattern noise is due to random variations within the pixel circuitry itself.

Some variation of this sort is to be expected in any CMOS process, however the variation shown here is too much to be acceptable for a production system, with 10% of pixels showing poor sensitivity (no visible change in flux reading between blood flow occlusion and release), so changes to the pixel design to reduce fixed pattern noise will be required.

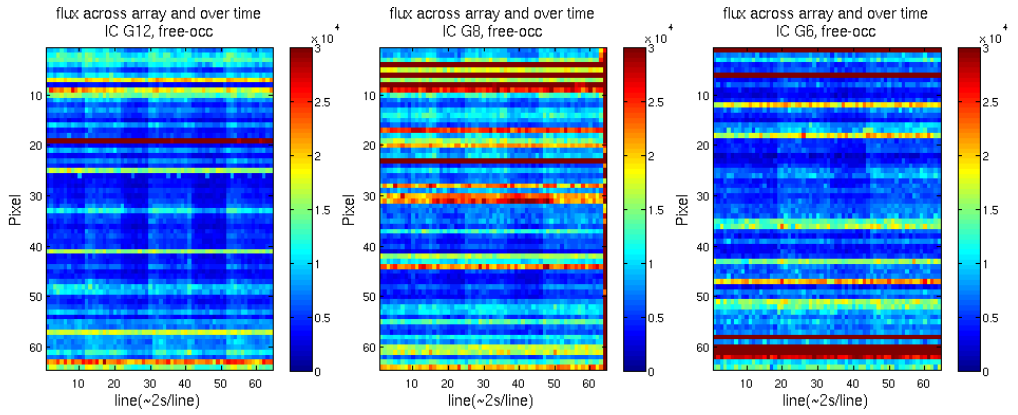

Figure 4.45: Plots of flux measurements from all pixels over time during occlusion-release of flux in a finger, main $g_m C$ output of RHS logarithmic pixel, three different BVIPS1 ICs

4.7 Summary

The first prototype IC designed for this thesis, including a 64 x 1 array of logarithmic pixels and four prototype opamp based pixels, has been described. Characterisation of the various elements of this IC were shown, with post-layout simulation results compared with corresponding test results, measured using a modulated laser to illuminate the IC.

The characterisation results show that actual IC behaviour, while similar to that expected in some respects, has significant differences to simulated behaviour, mainly in terms of higher noise levels than those expected. Possible sources for this include PCB power supply noise, noise in the data acquisition system, and distortion caused by the AC amplifier. Work was done to identify and reduce these noise sources, with relatively limited effect. Further investigations focused on unexpected behaviour seen from the IC, which showed a significant flaw in the AC amplifier used on the IC. This flaw may be significant in introducing additional noise and causing other discrepancies with simulated behaviour.

Other discrepancies between measured and simulated behaviour include offset DC output voltage responses at various points, although as the slope of the characteristics is similar, this is not a problem. The normalising function of the logarithmic detectors was shown to function correctly, although not perfectly over the full photocurrent range. The AC voltage output for a signal with fixed modulation depth falls slightly as the photocurrent rises, rather than staying constant as required for perfect normalisation. However, the drop is from 0.6 to 0.3 V over a photocurrent range of 200 pA to 80 nA, so over the smaller range of currents

expected within a Doppler blood flow measurement, the normalisation should be sufficient.

The frequency response of the logarithmic pixel showed a higher than expected low-frequency cut-off, with the increase due to leakage currents in the feedback network of the hysteretic differentiator amplifier used as high-pass filter and AC amplifier. This feedback network has a very high resistance resulting in operating currents of a similar magnitude ($<$3 pA) to the leakage currents in the transistors used. As the roll-off below this low-frequency cut-off is gradual, it did not prevent the IC from performing its intended function of blood flow measurement, although it may have restricted the signal amplitude and hence the SNR. Other issues with the HDA mean that a change of design will be made for other reasons, which will address this problem. The logarithmic pixel high-frequency cut-off was set in the range required, and this bandwidth was maintained at all photocurrents, showing that the high-frequency cut-off does not become limited by front-end bandwidth at low photocurrent.

For the opamp pixel, the response goes to DC as no low-pass filter was included in the circuit. This is acceptable for testing, but a high-pass filter will be added for the next prototype of this design to prevent movement artefacts and remove 1/f noise. The high-frequency cut-off, while lower than for the logarithmic pixels at approximately 6 kHz for typical I_{DC}, was sufficient for Doppler imaging.

Both circuits were used for Doppler Blood Flow measurements. The 64x1 array was capable of detecting changes in blood flow from occlusion and release of blood flow in a hand when using line illumination. However, the sensitivity and spatial resolution was not sufficient for vein location. The logarithmic pixels had an SNR between 4 and 8.5, compared to 10-17 for commercial systems (although the commercial devices have better spatial resolution, and the SNR figures given are likely to be conservative). The four pixels of the opamp design could detect changes in blood flow in a hand over time, but the low number of pixels in the array on this IC means that a larger array (higher pixel number) is required to determine whether vein location is achievable with this pixel design.

A number of problems were shown with the logarithmic pixel design, including high fixed pattern noise, a number of pixels with sensitivity too low to detect modulated light at the levels expected and spikes on the output voltage from the HDA. These will be addressed on the next IC by a change of HDA design, and increased redundancy to prevent 'dead' pixels.

Chapter 5

Design and Characterisation of BVIPS2 IC

5.1 Introduction

This chapter describes the design, simulation and testing of the second prototype imaging IC designed for this project. This IC uses a 64x1 array of photodiodes, with two arrays of photodetectors and analogue processing circuits. One of these is based on the logarithmic pixels used in BVIPS1, the second is a linear detector array using opamps. This design is a development of the opamp pixels tested on BVIPS1. The IC also includes on-chip analogue to digital conversion, and digital flux processing.

5.1.1 Structure of the Chapter

The chapter begins with a description of the overall IC design, given in section 5.2.1. Descriptions and block diagrams for the logarithmic and opamp pixel structures are given in Sections 5.3.1 and 5.4.1 respectively. Following the description of the IC, the individual changes made to the logarithmic detector section are discussed in more detail in Sections 5.5 (regarding changes to the log pixel part, including a new HDA, higher order filters and multiple front-ends), while the design changes for the modified opamp front-end detector are shown in Section 5.6.

As the full IC uses integrated processing, the processing itself and how it is used on the full IC are described in Section 5.7. Note that while inclusion of the processing block into the full IC design was done by the author for this project, the design of the processing block itself is not the work of the author.

Characterisation results for the IC using a known illumination source is shown in Section 5.8, and results from Doppler imaging of flow targets is shown in 5.9. However, time available for testing of this IC was limited, so the results shown here are less thorough than for BVIPS1.

Finally, Section 5.10 summarises the design changes made and the results obtained.

5.2 IC Design

This section will discuss the overall design of the BVIPS2 IC, in particular the designs of the logarithmic and opamp pixels incorporating the changes described previously in this chapter.

5.2.1 IC Top Level Block Diagram

Figure 5.1 shows the overall circuit design of the BVIPS2 IC. Bias and control lines are not shown to make the main design elements clearer.

The main omissions are the control lines of the multiplexers. These can be controlled by the digital processing section, or by external inputs to the analogue section of the IC. Similarly, the ADC can be controlled by the digital section, or by external control lines.

The IC consists of two main photodetector arrays. The top left of Figure 5.1 shows the logarithmic pixel array, described in Section 5.3.1, while the bottom left shows the linear opamp pixel array described in Section 5.4.1. Both arrays use a common set of photodiodes, with the unused set of photodetector circuits isolated from the photodiodes by a transmission gate switch (as used on the BVIPS1 IC).

Each array has 64 AC outputs and 64 DC outputs. These are then multiplexed down to two AC and two DC outputs from either array (using five of six total address lines, and a log/opamp select line). An additional multiplexer selects either both AC or both DC signals, which are each passed to an ADC. At this point an additional multiplexer also selects one of the two signals (using the 6th address line) for passing to an analogue output pin.

The digital outputs of both ADCs are multiplexed to select one ADC at a time to connect to the digital processing section. As the bottleneck in this system is the ADC sampling speed, the digital block is sufficiently fast to read data from both ADCs sequentially. The signals input to the digital section are also passed to off-chip buffers, allowing the sampled signal from the ADC to be read and processed by off-chip devices. Alternatively, the processed flux and DC readings output from the digital section can be read directly, avoiding the requirement for off-chip processing. Note that the DC values output from the on-chip digital processing are extracted from digital filtering of the same input signal used to calculate flux values, as the algorithms were designed for logarithmic pixels without a separate DC channel. This will restrict the use of the on-chip processing to analogue channels that maintain a DC value. While this is not ideal, it was not practical to re-design the digital processing section to address this issue in the time available for design.

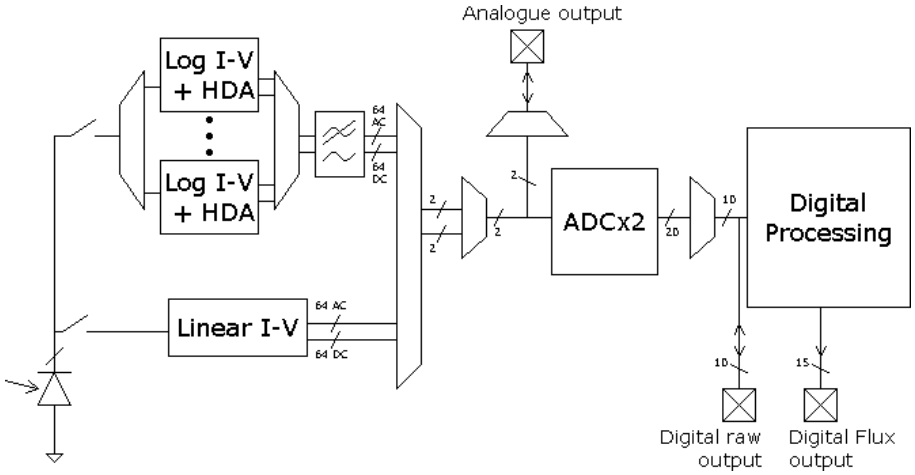

Figure 5.1: Block diagram of full BVIPS2 IC

Figure 5.2 shows the layout of the BVIPS2 IC. The analogue section of photodiodes, both front-ends and multiplexers has a size of $2100 \times 3250\,\mu$m. The front-end selection circuitry is $70 \times 3250\,\mu$m, which is located directly right of the analogue section such that the control lines connect directly. The main digital processing circuit, in the lower right of the layout, is $950 \times 2000\,\mu$m. Both ADCs are above the digital section on the top right of the design, with an area of $640 \times 1150\,\mu$m. Output buffers, and extra multiplexers for the analogue section fit between the main analogue section, the digital section and the ADCs. The layout results in various voids around the edge of the IC (e.g. to the right of the ADCs and around the edges of the IC in the gap between core and pads used for routing), and these spaces are used to add capacitors for decoupling of power supplies and analogue bias connections.

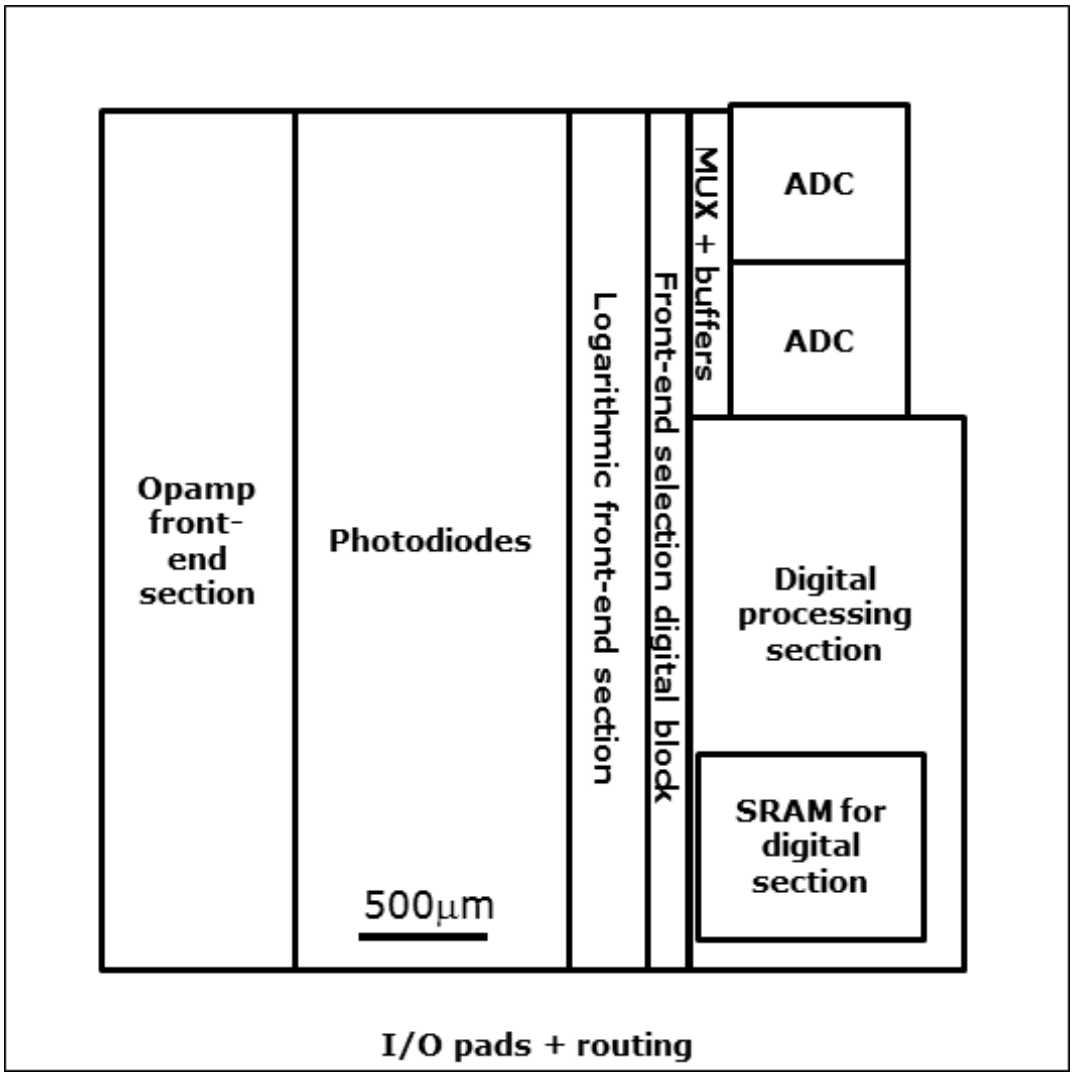
Opamp front-end section
Photodiodes
500μm
Logarithmic front-end section
Front-end selection digital block
MUX + buffers
ADC
ADC
Digital processing section
SRAM for digital section
I/O pads + routing

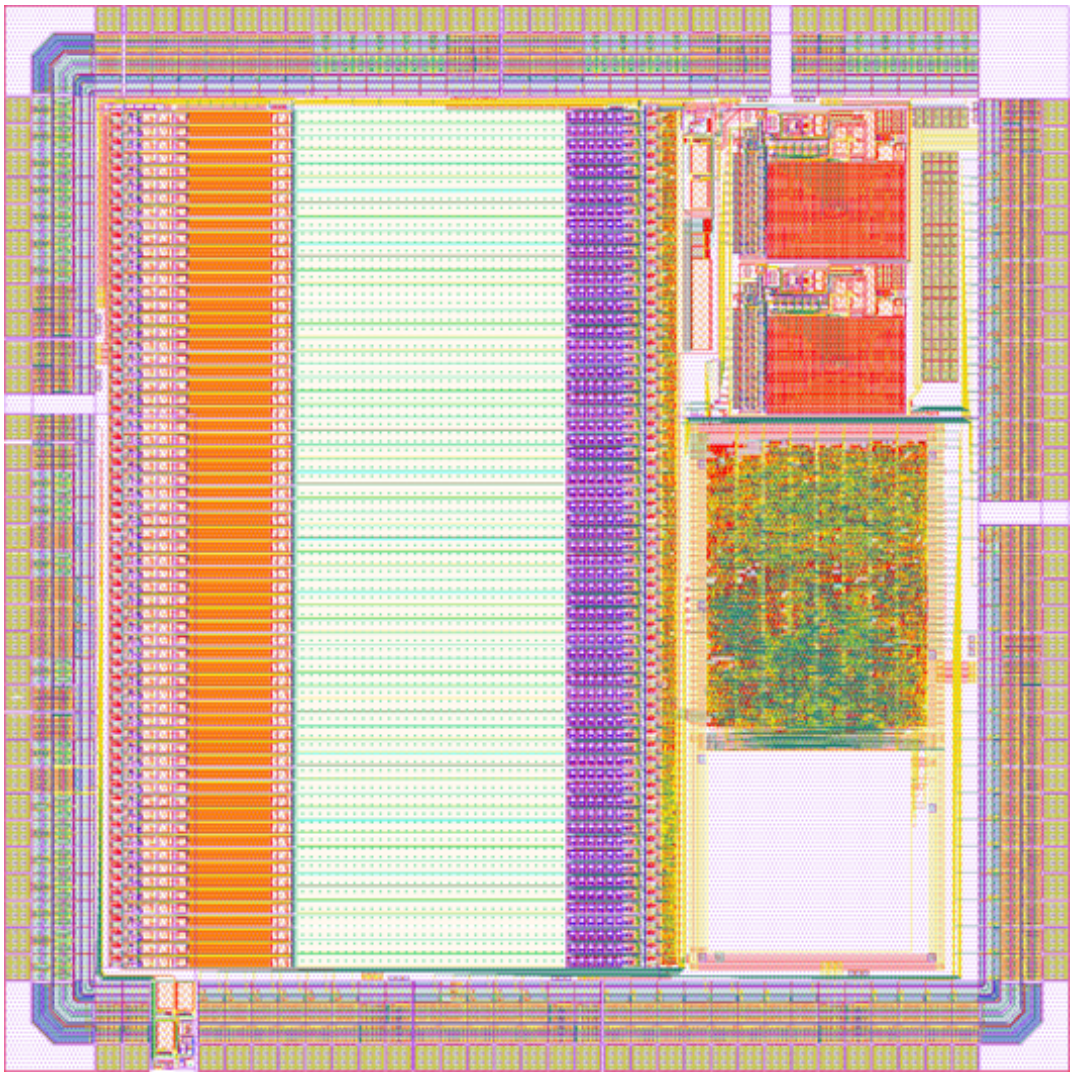

Figure 5.2: Layout of BVIPS2 IC

5.3 Logarithmic Pixel Design

This section describes the structure of the logarithmic pixels implemented on BVIPS2, including a brief description of the changes made to address the problems identified during testing of BVIPS1. These changes are discussed in more detail in Section 5.5.

5.3.1 Logarithmic Pixel Structure

A simplified block diagram showing the topology of the logarithmic pixel circuit is shown in Figure 5.3. This shows the design of a single pixel, with buffers after the HDA and low-pass filters omitted to clarify the main circuit elements. The buffers used are OTA based voltage follower buffers, as used on the BVIPS1 IC.

Each photodiode can be connected through a multiplexer to one of six front-end/HDA blocks, the outputs of which are connected to common filters through an additional multiplexer. Each block has separate AC and DC channels, where the DC channel bypasses the HDA. This feature of the pixel is discussed in Section 5.5.3, while the modifications made to the HDA are discussed in Section 5.5.1. Low-pass filters ($f_c = 20\,\text{kHz}$) are used to prevent aliasing and also reduce noise. These filters are described in Section 5.5.2.

The multiplexers used for front-end/HDA selection can be driven by either a digital section on-chip which allows each pixel to have a block set individually, or the block can be set globally using a six bit input bus.

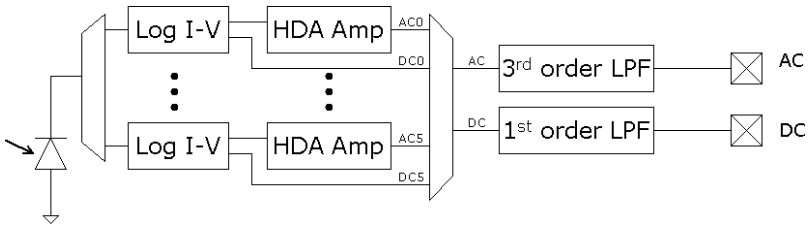

Figure 5.3: Block diagram of logarithmic front-end section of BVIPS2

Figure 5.4 shows the layout of the multiple front-end logarithmic pixel design. It has an overall size of $300 \times 50\,\mu\text{m}$.

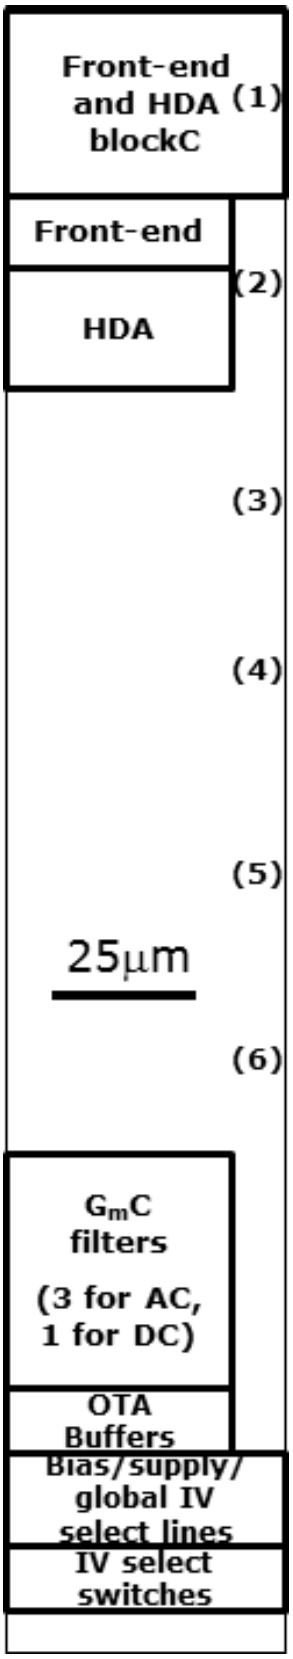
Front-end and HDA blockC
(1)
Front-end
HDA
(2)
(3)
(4)
(5)
25μm
(6)
GmC filters
(3 for AC, 1 for DC)
OTA Buffers
Bias/supply/ global IV select lines
IV select switches

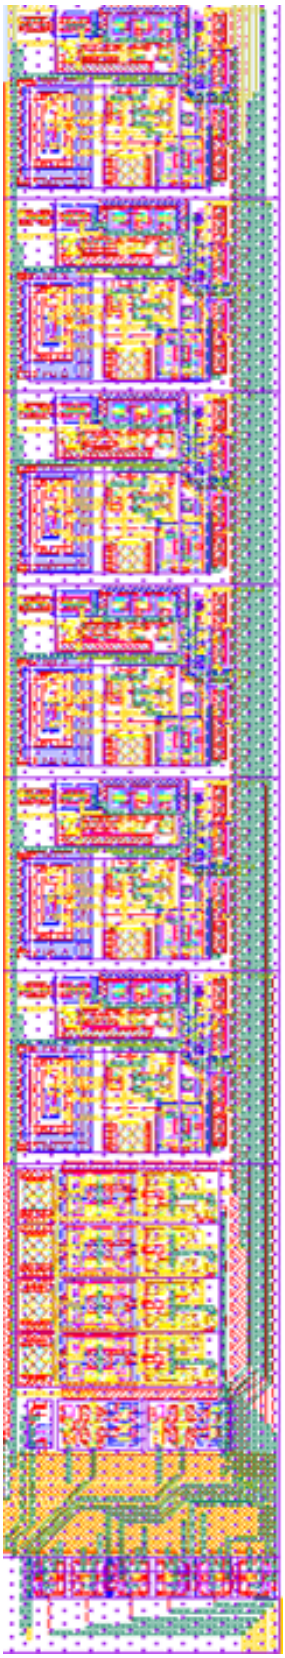

Figure 5.4: Layout of logarithmic front-end section of BVIPS2

5.4 Opamp Pixel Design

This section describes the structure of the linear opamp pixels implemented on BVIPS2. This includes the design of a compact opamp for use in the front-end stage, and a high-pass filter for use in the AC channel.

5.4.1 Opamp Pixel Structure

Figure 5.5 shows the topology of the linear detector circuit. Again, OTA voltage follower buffers are included after all amplifiers and filters, but these are not shown to clarify the main circuit elements.

The pixel consists of a front-end stage consisting of a compact opamp (the design of which is shown in Section 5.6.2, Figure 5.21), used in the same circuit as the opamp pixel on BVIPS1 (shown in Figure 4.6). The output of this circuit is fed to an AC and DC channel. The DC channel uses a non-inverting opamp amplifier configuration using the same compact opamp as the front-end, with feedback resistors giving a gain of 10. This gain is chosen to give sufficiently large signal range at the output to detect changes in the DC level over the range expected within an LDBF image, while not saturating over that same range.

The AC channel uses either the high-pass filter (described in Section 5.6.4, Figure 5.27) to remove the DC level, followed by non-inverting opamp amplifier using the same compact opamp design. This channel has a higher gain than the DC channel, as the fixed input DC level means an offset voltage can be set at the amplifier input to prevent saturation due to a large DC voltage at the opamp input.

An alternative AC channel is implemented using the HDA AC-only amplifier, as used in the logarithmic pixel. This is included due to the possibility of manufacturing variability causing unexpected behaviour of the high pass filter in the opamp AC channel - for example, if the DC voltage output from this filter varies widely across all pixels in the array, the DC offset voltage applied to the amplifier to prevent saturation will not be suitable for all pixels, causing saturation of some pixels regardless of the voltage applied.

G_mC filters are included on the DC and AC channels, with a higher order filter used on the AC channel as it is this channel used for flux measurements. A lower order filter is used here compared to the logarithmic pixel (which uses a 3rd order LPF), as the bandwidth of the front-end itself is set to 30 kHz by the RC feedback network, effectively implementing an additional order of filtering.

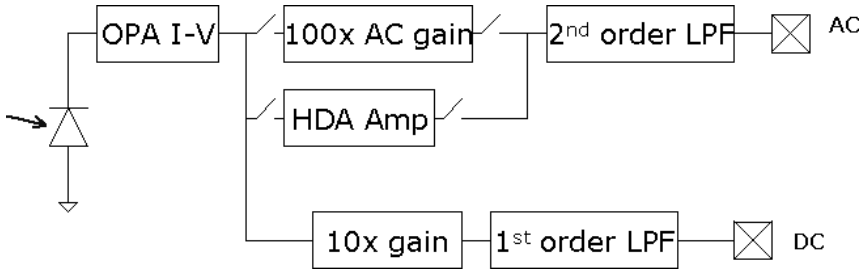

Figure 5.5: Block diagram of linear front-end section of BVIPS2

Figure 5.6 shows the layout of the opamp based linear pixel design. It has an overall size of $635 \times 50\,\mu$m. $390\,\mu$m of this is taken up by the front-end, including a $300\,\mu$m wide resistor, $20\,\mu$m wide capacitor and $55\,\mu$m wide opamp. The high-pass filter in the opamp based AC channel is $45\,\mu$m, and the amplifiers for the AC and DC channels are each $62\,\mu$m wide, including the opamp and feedback elements. The HDA for the AC channel adds $15\,\mu$m of width, and the g_mC filters and buffers add $50\,\mu$m.

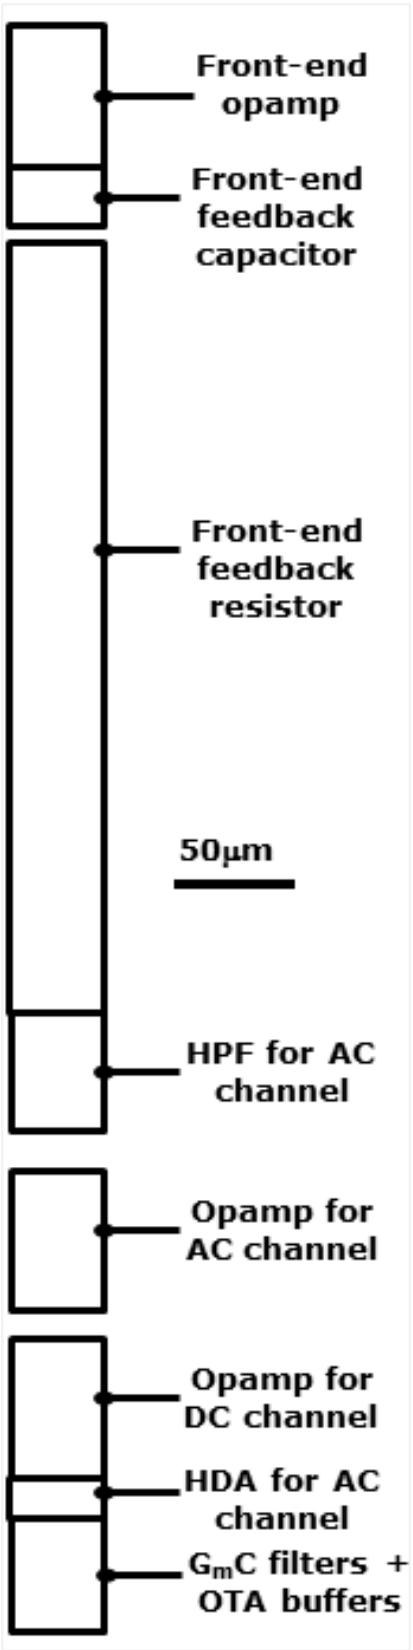
Front-end opamp
Front-end feedback capacitor
Front-end feedback resistor
50μm
HPF for AC channel
Opamp for AC channel
Opamp for DC channel
HDA for AC channel
G_mC filters + OTA buffers

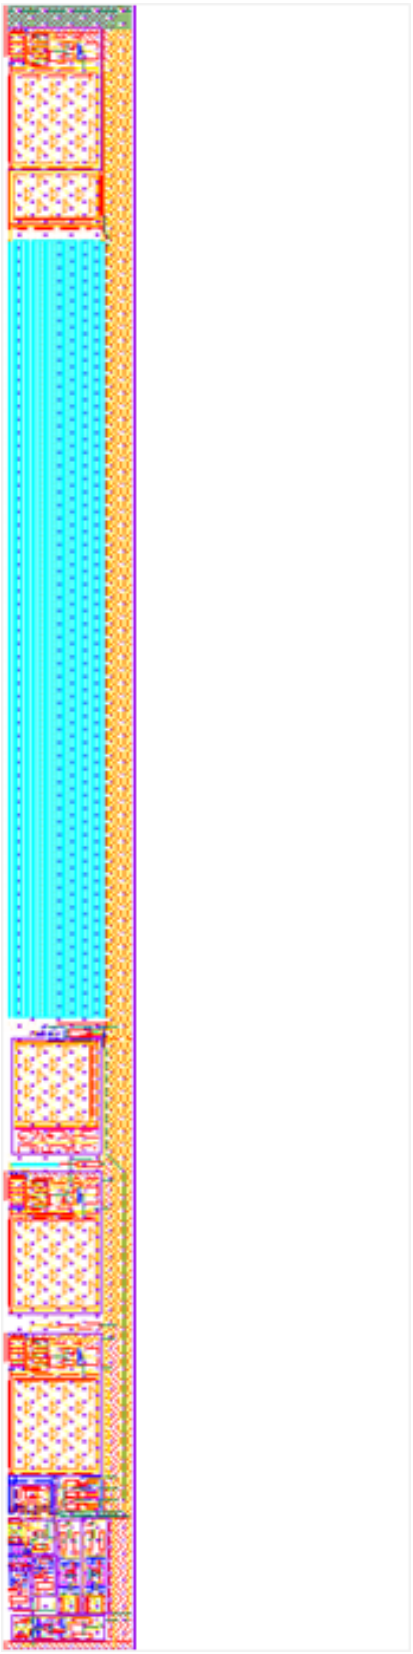

Figure 5.6: Layout of linear front-end section of BVIPS2

5.5 Changes to Logarithmic Detector

As shown in chapter 4, the logarithmic pixels used on BVIPS1 are capable of detecting blood flow, but with insufficient accuracy or repeatability for the IC to function as required in an array system. A number of changes to that pixel design are required to solve the issues identified.

5.5.1 HDA Alternative Design

A major cause of additional noise and distortion in BVIPS1 was shown (in section 4.6.3) to be the HDA AC amplifiers. This erroneous behaviour was not shown in simulations, including additional simulations performed after the fault was shown during testing. This makes addressing this problem through design changes problematic, as simulations cannot be used to reliably show the effect of any changes made.

However, similar circuits to the HDA on BVIPS1 have been used on other ICs manufactured at the University of Nottingham (64×64 DOP3 IC, designed on DTI NEAT project [Hoang, 2009, He et al., 2009]). These use the same OTA-inverted inverter-capacitor arrangement, but with different capacitances and with changes to the layout.

The main change to the HDA design for the BVIPS1 IC, shown in Figure 5.7, was an increase in the capacitor value, intending to reduce the low-frequency cut-off which had been shown through simulation with the original capacitance (of 120 fF) to be 300 Hz. The capacitance was increased to 5 pF, bringing the low-frequency cut-off down to ~100 Hz.

The characterisation of BVIPS1 has shown that as well as causing additional noise and distortion, the desired fall in cut-off frequency has not been achieved. Because of this, for the second prototype IC (BVIPS2) this change was abandoned in favour of reverting to a HDA design that had already been successfully tested.

The design used for the BVIPS2 IC uses the same OTA as that used in the BVIPS1 log pixels (shown previously in Figure 3.29), with the MOSFET used as a capacitor reduced in sized to $5.5 \times 6\,\mu$m. Apart from the change in capacitance altering the performance of the circuit, the reduction in physical size allows a guard ring to be placed around the inverted-inverter

section and the capacitor. This may help to isolate these parts of the circuit from noise sources in the rest of the IC, including substrate currents caused by stray light generating electron-hole pairs in regions of the IC where there is no field such as that in the photodiodes to contain the charge carriers. Figure 5.9 shows the layout of the HDA designs on BVIPS1 and BVIPS2, showing the changes to the feedback capacitor and the additional guard ring. The schematic of the old and new HDA circuits are shown in Figure 5.7 (original on BVIPS1) and Figure 5.8 (new HDA with smaller feedback capacitor).

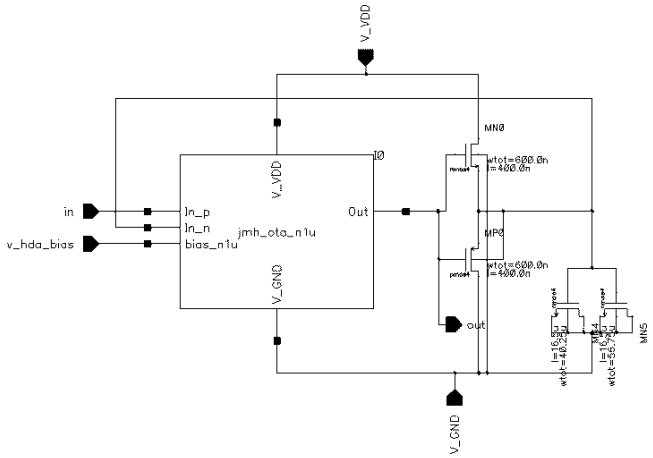

Figure 5.7: schematic of original HDA used on BVIPS1

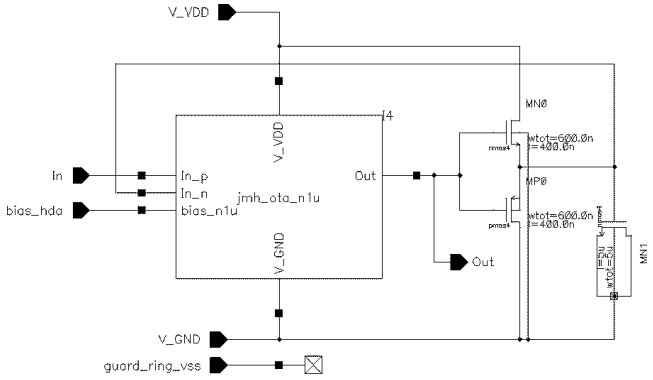

Figure 5.8: schematic of new HDA used on BVIPS2

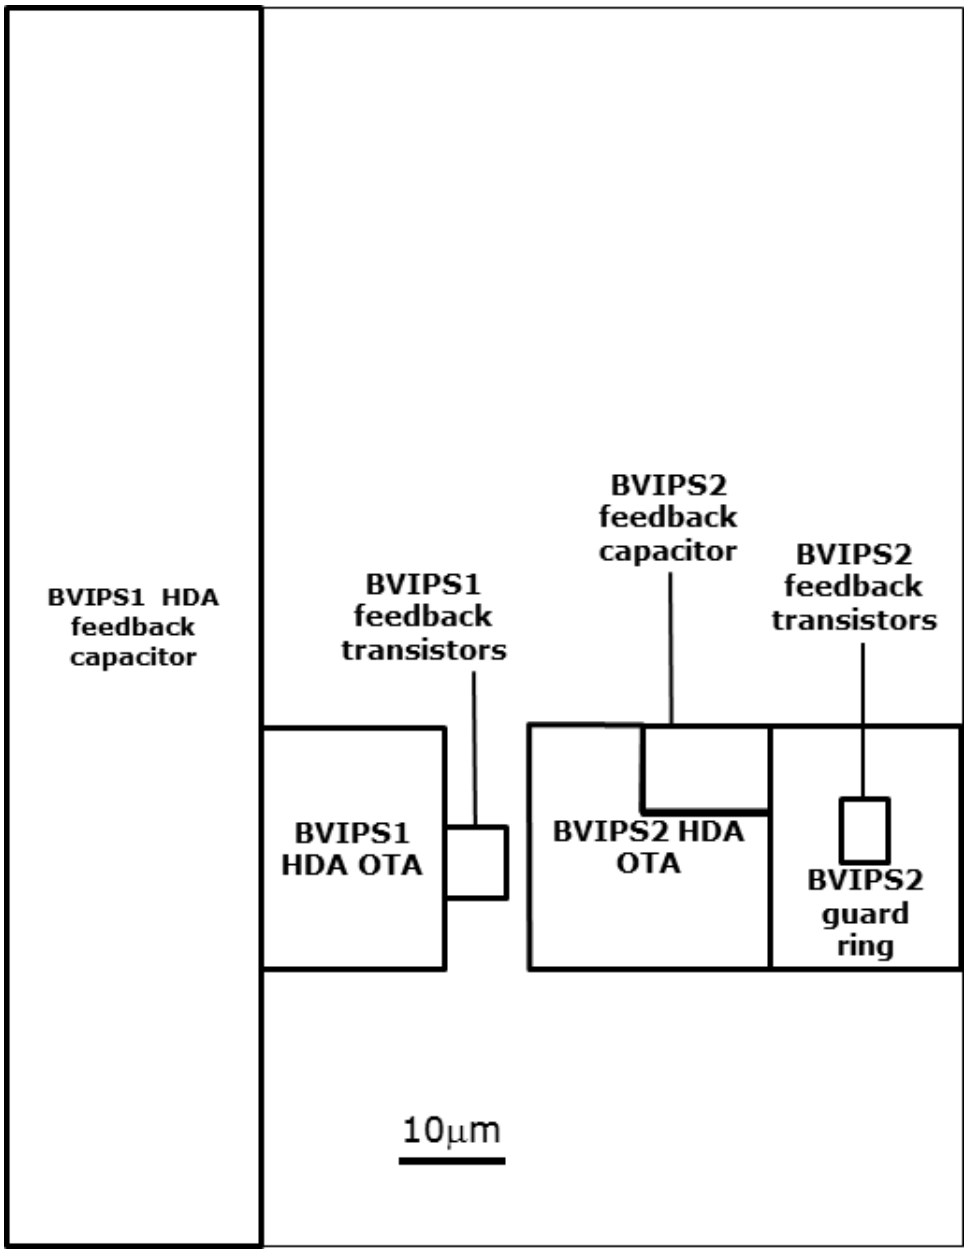
BVIPS1 HDA feedback capacitor
BVIPS1 feedback transistors
BVIPS2 feedback capacitor
BVIPS2 feedback transistors
BVIPS1 HDA OTA
BVIPS2 HDA OTA
BVIPS2 guard ring
10μm

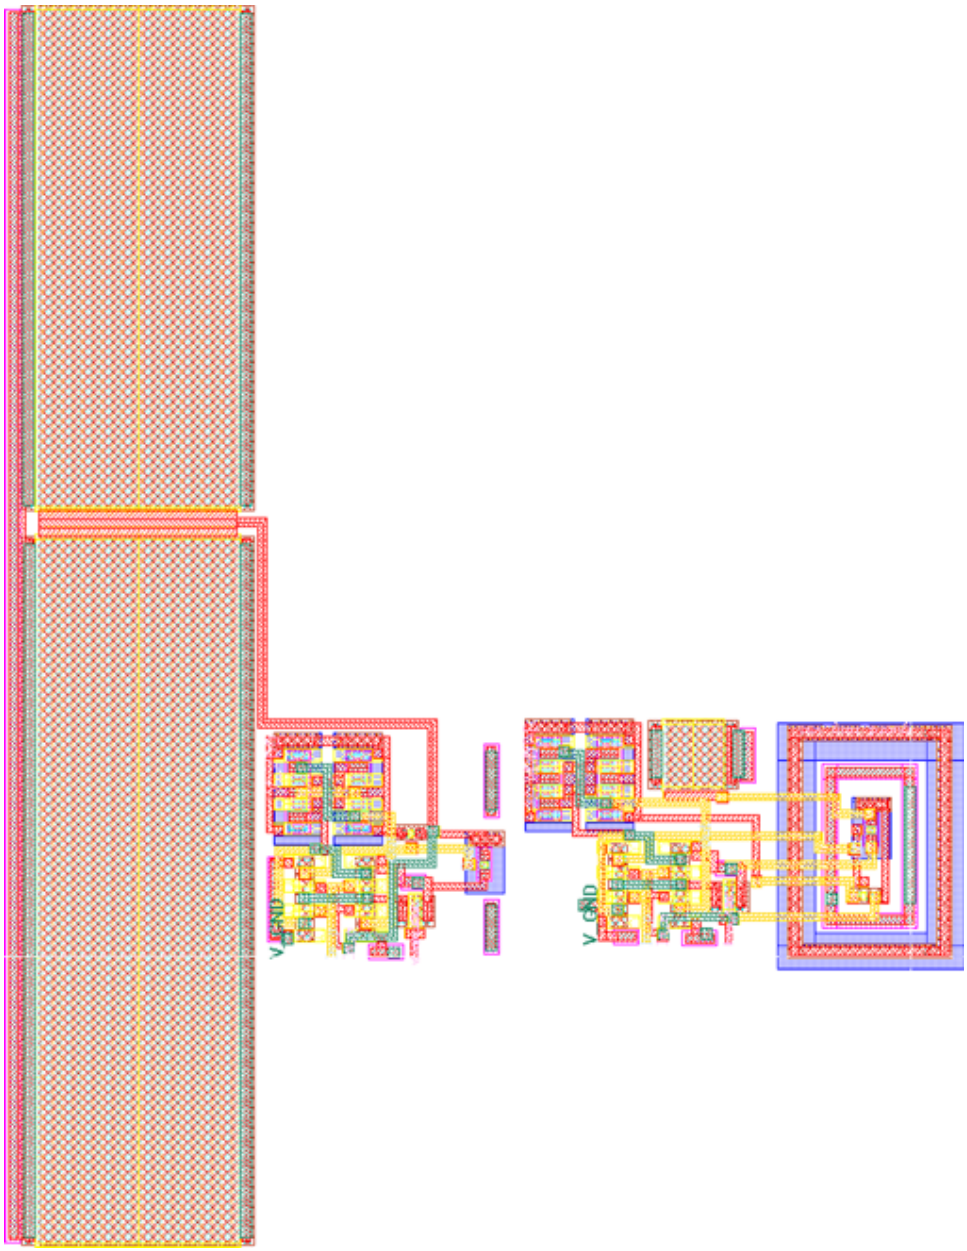

Figure 5.9: Layout of HDA circuits on BVIPS1 (left) and BVIPS2 (right)

Simulations were performed to show the difference in performance between the original and new HDA designs, although the accuracy of these simulations is known to be limited given the discrepancies seen with BVIPS1.

Figure 5.10 shows the AC response of both HDA designs, with a 1.2 V DC / 3 mV AC input signal. It can be seen that the original design has higher gain in the pass band than the new design, and a lower low-frequency cut-off (100 Hz instead of 2.5 kHz). This reduction in cut-off frequency was the purpose for the original design change, however given the limited success off this change (with a low cut-off frequency of approximately 1.5-2.5 kHz), this change has clearly been unsuccessful.

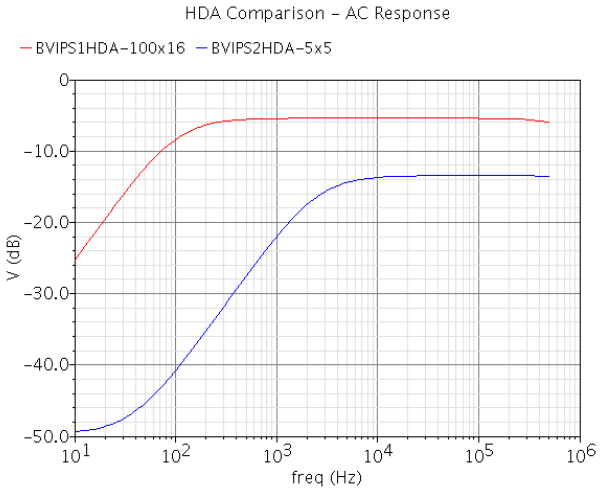

Figure 5.10: AC response of original (BVIPS1-100x16) and new (BVIPS2-5x5) HDA designs

Figure 5.11 shows the DC response of both HDA designs, showing no difference between the responses. This is to be expected given the changes apply to the AC properties of the feedback path, but the simulation confirms that the change will not have a negative impact on the DC operating point of the surrounding circuits.

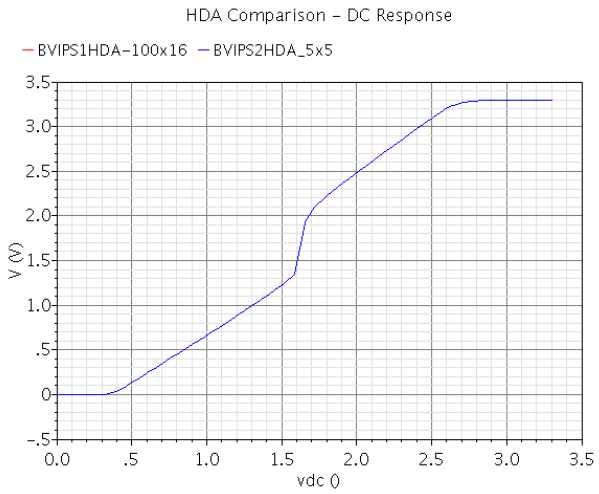

Figure 5.11: DC response of original (BVIPS1-100x16) and new (BVIPS2-5x5) HDA designs

Figure 5.12 shows the input referred noise spectra of both HDA designs, showing lower noise from the original circuit. The increased noise at low frequencies is likely to be linked to the lower gain below the new circuits simulated cut-off, as this gain reduction means noise at the output is proportionally greater when input referred.

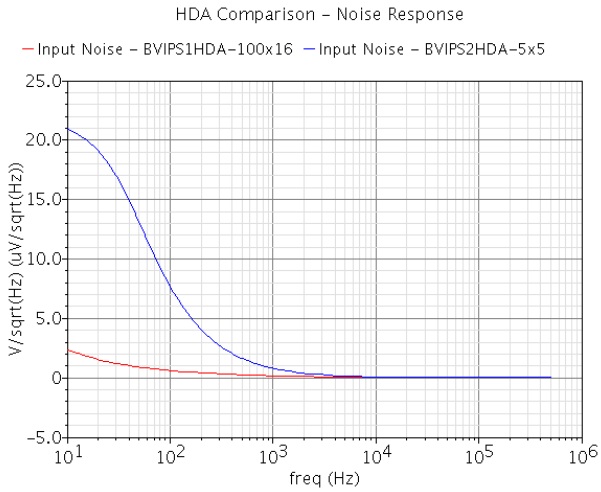

Figure 5.12: Input referred noise response of original (BVIPS1-100x16) and new (BVIPS2-5x5) HDA designs

The AC and noise simulations shown here suggest that the original design is superior to the new design presented here. However the successful use of this design on other projects, the problems seen with the first IC, and the improvements to the layout to reduce susceptibility to external noise sources mean that the new design was chosen for use on BVIPS2.

5.5.2 Higher Order Filters

A major difference between the linear detector arrays designed here and the 2D arrays that are more common in LDBF is that the area available for each pixel is considerably larger. The pitch of each pixel is subject to similar limits, being the size of IC (as dictated by budget), divided by the chosen number of pixels. However, here each pixel can be long and thin. This can be used to give a wide array, simplifying beam alignment, but the additional area can also be used to implement additional (light insensitive) circuitry without reducing the fill factor of the photosensitive area itself. This means that additional circuitry such as extra filters can be included at each pixel with relatively little impact on IC size.

For this reason, the single order g_mC circuits used in BVIPS1 have been cascaded to give higher order filters. This results in sharper cut-off, which more effectively filters off noise just above the signal bandwidth, and limits the effect of aliasing. This should result in a lower noise floor, and hence better SNR without any additional gain.

Rather than design a new filter for this purpose, g_mC filters of the same design and layout as those already used on BVIPS1 were cascaded to produce filters of various orders. To verify that cascading filters in this way does not cause problems in turns of level shifting or attenuation, simulations were performed on cascaded filters.

Figures 5.13 and 5.14 show the transient output after each stage of the cascaded filter, including and excluding the DC component. The input signal is a 3 kHz sine wave, well below the cut-off frequency of a single filter of 20 kHz. The peak-peak magnitude is 100 mV. It can be seen that the AC part of the output signal is the same after all stages, except for a small phase change. This phase change is to expected, and given a roughly linear phase response shown by simulation of a single filter, a similarly linear phase response should be expected from a cascaded filter.

The signal including the DC component shows the major effect of cascading filters, that of a DC level shift through each filter of roughly 4 mV. Given the DC level expected from typical

LDBF imaging of roughly 1 V, this shift is not enough to significantly affect the output AC signal.

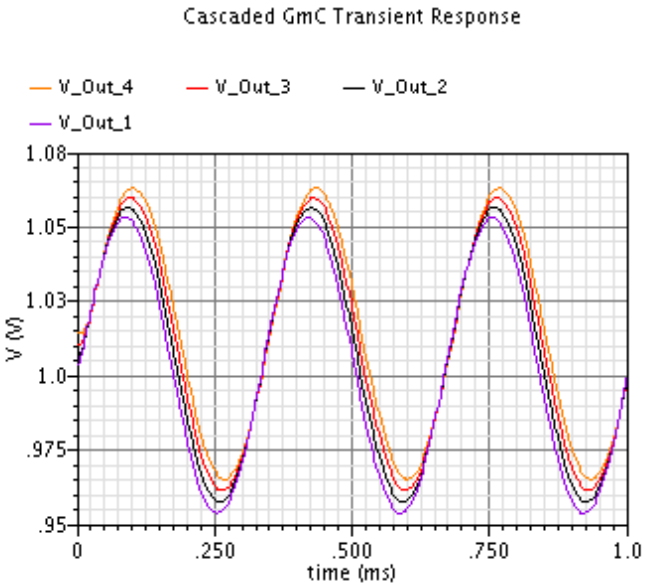

Figure 5.13: Transient response of cascaded g_mC filters (1 V DC / 100 mV AC @ 3 kHz input)

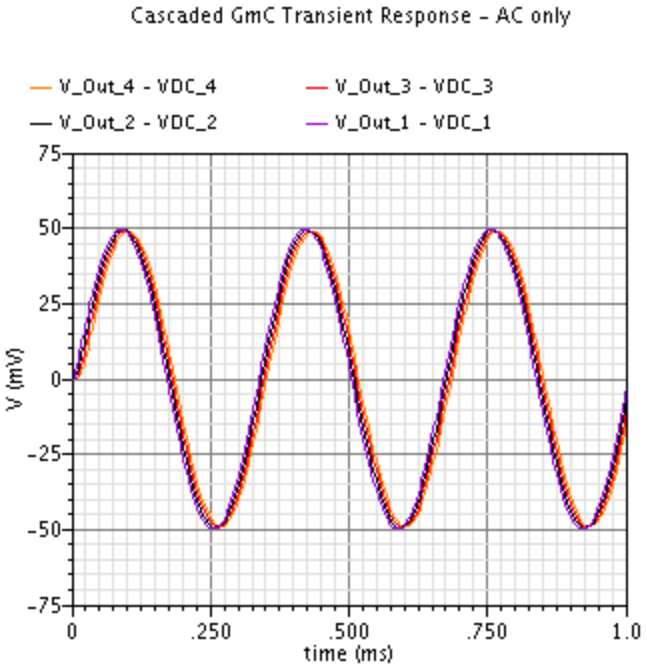

Figure 5.14: Transient response of cascaded g_mC filters, with DC component removed (1 V DC / 100 mV AC @ 3 kHz input)

Figure 5.15 shows the DC response of different numbers of cascaded filters. It can be seen

that in the middle of the response the only difference is the 4 mV offset already shown. There is a more significant difference when the input voltage approaches either power supply rail, particularly V_{SS}. However, as the single filter is non-linear at these points this does not represent a major change in filter behaviour. Also, the signals expected here have a DC of around 1 V, with a smaller AC component, such that operation at these limits is not critical for this application.

The linear pixel used in one array on this IC means that a wider variety of DC voltages may be expected, but the range is unlikely to be wide enough for the non-linear filter behaviour to be an issue. The offset voltages used in the opamp pixel design, and the non-linear behaviour of the opamps themselves at input voltages near ground or V_{DD}, mean that filter input voltages in this range are unlikely to occur.

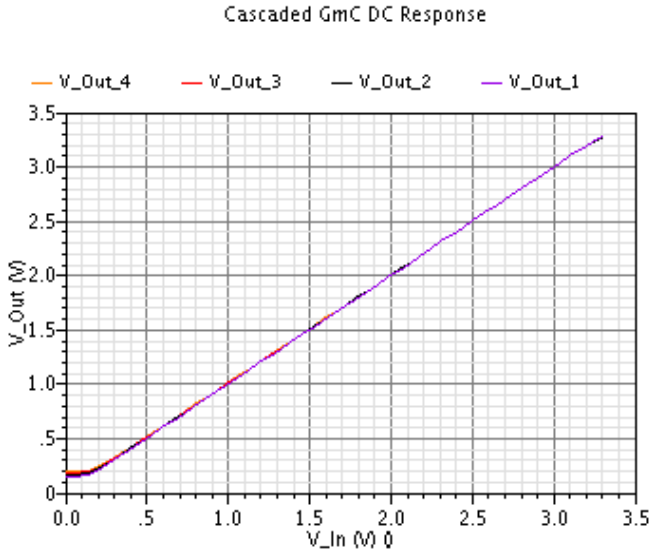

Figure 5.15: DC response of cascaded g_mC filters

Figure 5.16 shows the frequency response of different numbers of cascaded filters. As expected, it can be seen that the pass band gain is equal in all cases, but that the cut-off frequency falls as the number of filters rise, and that the roll-off is much steeper with a higher order filter. This does show that to set a cut-off of 20 kHz will require the cut-off of each filter to be set higher than 20 kHz, but as this is adjustable by an off-chip bias current, this is already possible with this design.

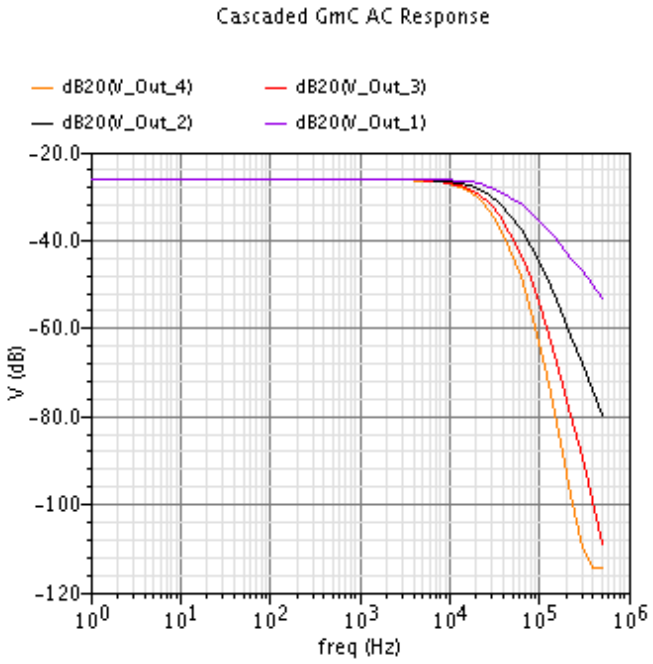

Figure 5.16: Frequency response of cascaded g_mC filters

5.5.3 Multiple Front-ends and Amplifiers for FPN Reduction

5.5.3.1 Motivation for Multiple Front-ends

Fixed pattern noise on the BVIPS1 IC can be considered as two different problems. The most severe of these was 'dead' pixels, which exhibited some gain when tested with high modulation depth illumination, showing some degree of correct operation, but with gain too low to detect LDBF signals. The second of these, gain variation from pixel-pixel, is to be expected in CMOS sensors, and can be compensated for using calibration techniques. This is possible as the issue here is mismatch between pixels affecting a flux map, rather than the quality of output signal from any particular pixel. Calibration is not possible for the very low gain pixels, as a calibration routine applied to signals with such low SNR will results in a corrected signal with very high noise levels.

The dead pixels are likely to be linked to the HDA issues identified in chapter four, and as such should be resolved by the change of HDA design used in the BVIPS2 IC. However, as the cause of this issue was not fully determined, it is possible that there will still be low-gain pixels with the altered design.

The problems of FPN are further complicated by on-chip processing. Calibration generally

involves measuring output signals from all pixels with equal illumination of all pixels using a modulated source. All pixels can then be characterised in terms of DC output and frequency response. This means that when an unknown illumination pattern is applied, correction algorithms can be applied to obtain an image with little or no fixed pattern noise.

This technique requires access to raw signals, which can be processed before flux processing. If on-chip processing is to be performed, the raw-signals cannot be altered before processing, as the processing electronics are directly connected to the ADC outputs. It should be possible to include some calibration electronics in the digital processing section of the IC, but this adds to the computational demands on this section, and also requires a high degree of configurability, as the calibration coefficients will have to be calculated for each IC and then stored in registers in the digital section. This adds to the complexity and hence size and design time of the digital section, making it unsuitable for a prototype IC at this stage.

5.5.3.2 Principle of Operation of Front-end Selection

An alternative approach made possible by the small size of the logarithmic detectors, combined with the linear array topology, is to use duplicate front-ends. This is a similar technique to the shared photodiode design used in BVIPS1, where a single photodiode is connected to one of several front-end circuits as required. That technique is already used in BVIPS2 to switch between logarithmic and linear pixels. The multiple front-end design expands on this, having a number of logarithmic front-end circuits at each pixel. A digital control section will then be added allowing the front-end to be used by each pixel to be individually set.

By doing this, the front-end used can be selected to be that which most closely matches the average behaviour for all front-ends on the IC. This cannot fully remove fixed pattern noise, but should reduce it considerably, which may be enough to provide a significant improvement. One aspect of this is that if 'dead' or low-gain front-ends are found, an alternative front-end can be selected, bypassing the poorly performing circuit.

5.5.3.3 Application of Pixel Selection to an Existing 2D Array IC

To investigate the FPN reductions possible using such a circuit, pixel selection techniques were tested using characterisation data from an IC fabricated at the University of Nottingham which uses similar detectors in a 32×32 array (one of four arrays on 64×64 DOP3 IC, designed

on DTI NEAT project [Hoang, 2009, He et al., 2009]) . The characterisation was performed using illumination with a modulated source with an intensity similar to that expected from LDBF applications. In order to test the pixel selection technique on this data, each row of 32 pixels is considered as 32 alternative front-ends, with N of these available for selection as the 'in-use' detector for each pixel.

N is varied from 1 upwards in order to determine the optimum number of alternative front-ends. Higher N means reducing the FPN after circuit selection. However, the increase in N means more silicon area consumed by un-used front-ends, as well as increasing complexity and size of the front-end selection circuitry.

Figure 5.17 shows the effect on FPN from the front-end of increasing the number of pixels available for selection. Each column of the images shows the DC voltage or gain at each pixel of a linear array made up of all selected pixels. The leftmost column where N=1 shows the case of no circuit selection. N=2 represents a choice between two different pixels, continuing up to N=10.

Two images are shown, as FPN can be measured in different ways. The two most obvious used here are DC output voltage and gain. Gain would appear to be the most useful, as flux processing uses the AC component of the signal. However, if the DC variation means that the signal from a pixel is saturated against ground or V_{DD} then this is an important measure of FPN. In Figure 5.17 the maximum and minimum DC voltages are all sufficiently close that this does not appear to be a problem. At this stage, both cases can be considered, as the IC design is not affected by how the decision on which circuit to be used is made.

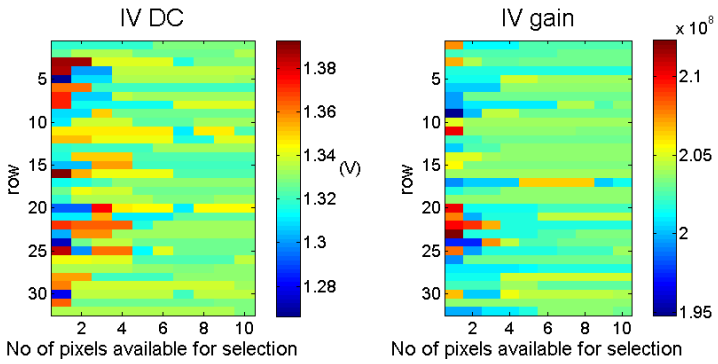

Figure 5.17: FPN from front-end after selection of 1 from N circuits

The decreasing variation in the colours within each column as N increases show the effectiveness of this technique. It can also be seen that the decrease in FPN is most rapid at lower

values of N. The change between N=1 and N=2 is very obvious, whereas the change from N=6 to N=7 is much less clear, particularly in the gain case. The improvement in FPN is approximately proportional to $1/\sqrt{N}$ - this is shown further in Figure 5.19.

5.5.3.4 Multiple HDA Selection

So far the circuit selection has only been considered in terms of the front-end. However, the HDA behaviour seen in BVIPS1 suggests this part of the circuit is likely to be a cause of significant FPN, and is also likely to cause low-gain pixels. A reduction in FPN from the front-end could therefore be made irrelevant if variation between HDAs in different pixels causes significant FPN. As the HDA circuit itself is fairly compact (roughly $20\,\mu\text{m} \times 35\,\mu\text{m}$), it is possible to include multiple HDAs at each pixel. This could be done with no extra circuit selection circuitry, instead combining each front-end with a HDA, and then selecting 1 of N signal paths, each consisting of a front-end and a HDA.

Figure 5.18 shows the effect on FPN from the HDA and front-end combination in terms of DC output voltage, and from the HDA along in terms of gain, in a similar manner to that already shown for the front-end. The HDA can be seen to be as significant as the front-end in terms of FPN, with the DC output from both circuits varying by 0.1 V in the N=1 case here. When measured by gain, the HDA is a more significant source of FPN, with the variation from mean to maximum (when N=1) being around 10% for the HDA but only 5% for the front-end (shown in Figure 5.17).

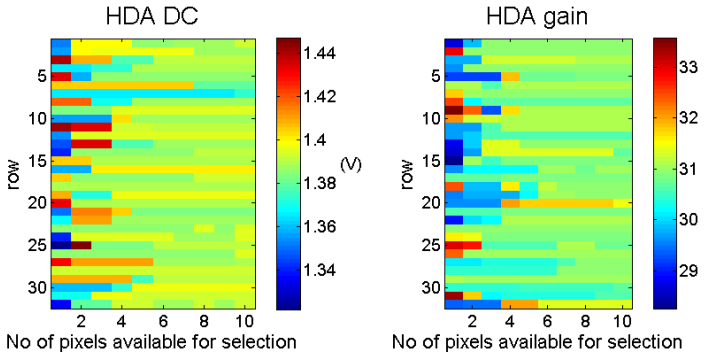

Figure 5.18: FPN from HDA after selection of 1 from N circuits

Again, it can be seen that as N increases the rate of FPN reduction falls, with the change after an increase of N by 1 difficult to see by eye after N=6.

Rather than rely on observation of the images produced by pixel selection, the effect of increasing N is measured by calculating the standard deviation of the chosen measure of FPN across all in-use pixels of the array after pixel selection.

The reduction in standard deviation of DC and gain from the front-end and the HDA is shown in Figure 5.19. It can be seen that in all cases the fall in standard deviation roughly follows a $1/\sqrt{\mathrm{N}}$ relationship. This is a similar relationship to the standard error of a population mean:

$$SE_{\overline{x}} = \frac{s}{\sqrt{n}} \tag{5.1}$$

where $SE_{\overline{x}}$ is the standard error of the mean, s is the standard deviation of the population, and n is the number of observations in the population. Increasing the number of alternative front-end circuits increases the population, hence reducing the error between the population mean (that of the circuits available for selection) and the actual mean (i.e. the nominal front-end/HDA response).

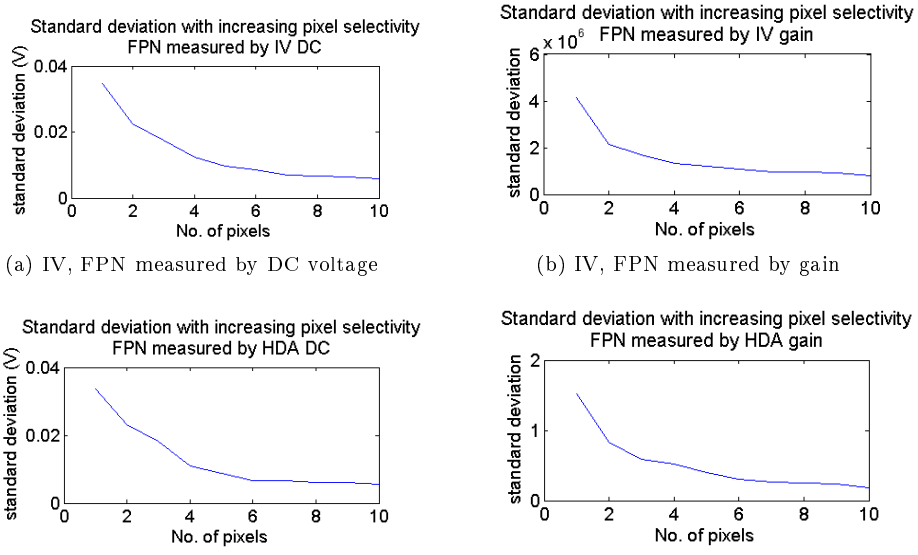

(a) IV, FPN measured by DC voltage

(b) IV, FPN measured by gain

(c) HDA, FPN measured by DC voltage

(d) HDA, FPN measured by gain

Figure 5.19: FPN measured by standard deviation of DC voltage or gain, from front-end and HDA, after selection of 1 from N circuits

5.5.3.5 Selection of N for BVIPS2

From Figure 5.19, amongst other design considerations, the number of front-end and HDA pairs to be included at each pixel is chosen as 6. This gives a high degree of FPN reduction,

while being low enough to not consume large amounts of silicon area. Using N=6 also keeps the pixel selection circuitry simple enough to not make design and testing excessively time consuming, and does not require a prohibitive number of additional I/O pins.

5.5.3.6 Digital Control of Circuit Selection

The logarithmic front-end and HDA circuit to be used will be selected using a set of buffers with a serial interface to an off-chip programmer, such as an FPGA. The analogue section will have 6 inputs per pixel, with 1 of 6 being set high to select a signal path. This could be done using three pins and an address decoder, but requiring an address decoder at each pixel would require extra space, without a major reduction in complexity of the digital section.

Also, this technique potentially allows more than one set of HDA and front-end to be made active at once. This could cause problems due to connecting multiple outputs, but as all circuits should have the same input this might not cause major issues, and may provide some degree of averaging. This requires additional investigation of the potential problems and benefits, and has not been done as part of the work presented here.

The digital section has five external pins - clock, reset and enable signals, and two pins for a serial interface. There are 384 outputs, 6 for each of the 64 pixels. These are entirely on-chip, so the large number is not a problem. The digital section is designed such that the outputs for circuit selection correspond to the locations of the input lines to each pixel, making routing of all select lines simple.

The pixel selection circuit was designed as part of the BVIPS project at the University of Nottingham by Dr. Kevin Xu. It uses a serial input from an external device (such as an FPGA or PC) or from a serial connection to the main on-chip processing block. This allows a series of latches to be set such that after a brief initialisation procedure the select lines to each pixel are fixed, and the selection circuit clock can then be stopped to avoid switching noise being introduced to the analogue signal.

5.5.3.7 Fault Tolerance

While the circuit selection technique provides considerable improvements to FPN, the extra complexity increases the chance of a fault in design or manufacture. This could potentially make all circuits unusable, so some degree of fault tolerance is required in case this occurs.

The most likely cause of problems comes from the interfacing of analog and digital sections of the IC. Mixed signal simulations will be performed to check the operation of this manner, but there is more chance of errors in the simulation settings or models here than for purely analogue or digital sections. Similarly, the digital design has not been previously implemented, so this should be considered a first prototype of this technique.

One source of redundancy here is to provide a purely analogue means for pixel selection. This is similar to a 2D array IC having row and column addresses. Each pixel has two sets of circuit selection inputs - one connects to the digital section for individual selection, the other to a set of global address lines set by an off-chip address bus. The two sets of select lines are isolated by a set of transmission gate switches, with an external pin allowing the circuit selection to be set to either global or individual selection.

While this in itself provides additional complication, the global circuit selection method can be simulated entirely within the analogue simulation environment, providing a high degree of confidence in this configuration.

5.6 Changes to Opamp Pixel Design

This section describes changes made to the opamp pixel design to make the design more suitable for inclusion on an IC, and to improve the performance of the pixel design.

5.6.1 Limitations of Opamp Pixel on BVIPS1

The test results from BVIPS1 show that the opamp based pixel can be used on an IC for detecting LDBF signals. However, it also showed the two major limitations. The size of the opamps used makes designing pixels with a 50 μm pitch difficult - BVIPS1 used a 100 μm pitch, with those pixels interleaved with photodiodes that were only connected on the opposite side of the IC. Secondly, the lack of DC rejection of the gain stage can cause problems with saturation due to changing DC light levels. This in turn leads to limited gain, in order to ensure that the range of DC photocurrent over which saturation does not occur is large enough to not require constant adjustment of the system.

5.6.2 Compact Opamp Design

Approximately 50% of the pixel area is due to the components in the feedback network, mainly the feedback resistor. As it is this resistor that sets the gain, this cannot be reduced in size without redesigning the other parts of the pixel to compensate for reduced gain. However, the opamp used in the BVIPS1 pixels is itself a large component, being around $100 \times 200\,\mu$m.

The opamp used for the front-end on BVIPS1 was not originally designed as a component for a front-end, but as a buffer for driving off-chip outputs. This means it was designed to be capable of driving high capacitance loads. That is not required for the opamps in the front-end, so the use of a more compact opamp designed for this application allows a more compact pixel to be designed. Reduction in the maximum rated drive capacitance allows for a reduction in the compensation capacitance (and hence physical size) as well as reducing the physical size of the transistors used in the opamp due to the lower currents involved.

5.6.2.1 Design Parameters

The opamp design was based on the following parameters: $V_{DD} = 3.3\,V$, $V_{SS} = 0\,V$ (i.e. single supply), $V_{Out(Max)}$=3 V, $V_{Out(Min)}$=0.5 V, Slew Rate = 10×10^6 V/s, Gain-Bandwidth product GB = 10 MHz, $V_{ICM(Max)}$= 2.5 V, $V_{ICM(Min)}$=1 V, phase margin = 60^o, $C_L = 3\,pF$. The load capacitance is smaller than that required to be driven by an opamp driving off-chip loads, but is considerably larger than that likely to be found on-chip - for example, if the next stage is a $10 \times 10\,\mu$m MOSFET transistor, with a gate capacitance of $4.5\,fF/\mu m^2$, the input capacitance would be less than 0.5 pF. This is a larger input transistor size than used in all the OTA buffers on this IC, so can be considered a worst case scenario for this design. The maximum and minimum output voltages do not allow rail-rail operation, but as this is not generally required by LDBF systems (where the light level can be controlled to be within a suitable range) this is not a major limitation, while simplifying the design procedure. Similarly, the input common mode minimum and maximum voltages specify a relatively narrow range, but the feedback circuit used for a photodetector circuit and the light level controls mean this is not expected to be a significant drawback. The slew rate and gain-bandwidth product is larger than that required here, but reducing these values results in larger components (with higher RC constants), rather than a more compact design.

The main process parameters for the austriamicrosystems C35 process are: Gate oxide capacitance $C_{ox} = 4.5\,\text{fF}/\mu\text{m}^2$, hole mobility $\mu_p = 126\,\text{cm}^2/\text{Vs}$, electron mobility $\mu_n = 370\,\text{cm}^2/\text{Vs}$. These can be simplified to $k'_p = 56.7 \times 10^{-6}\,\text{A/V}^2$ and $k'_n = 166.5 \times 10^{-6}\,\text{A/V}^2$, where $k' = C_{ox}.\mu$. Additional parameters include the channel length modulation factors, $\lambda_n = 0.04\,\text{V}^{-1}$ and $\lambda_p = 0.05\,\text{V}^{-1}$, and threshold voltages, $V_{Tn} = 0.5 \pm 0.1\,\text{V}$, $V_{Tp} = 0.65 \pm 0.1\,\text{V}$ [austriamicrosystems, 2007].

The gate length chosen for the transistors here is $1\,\mu$m. This could be reduced to the minimum allowed by the process of $0.35\,\mu$m, but increasing the length used makes the devices less susceptible to process variations. It also allows W/L ratios slightly less than 1 if required.

5.6.2.2 Design Process

The design process followed here is taken from 'CMOS Analog Circuit Design' by Allen and Holberg [Allen and Holberg, 2002]. The circuit is an un-buffered two-stage opamp, using the circuit design shown in Figure 5.20.

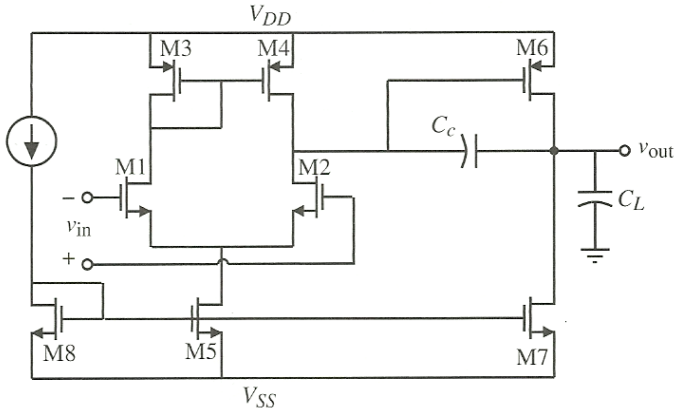

Figure 5.20: Circuit template of compact opamp design [Allen and Holberg, 2002]

The process starts by finding the minimum compensation capacitor required for stability with the given load:

$$C_{C\,(min)} = 0.22.C_L = 660\,\text{fF} \tag{5.2}$$

To provide an additional stability margin, a compensation capacitor of 1 pF will be used. If implemented using a polysilicon-polysilicon capacitor where $C \sim 1\,\text{fF}/\mu\text{m}^2$, this capacitor

would be around $32 \times 32\,\mu$m. This is still a large component, but is feasible for inclusion on a $50\,\mu$m pitch pixel.

The design process to set transistor parameters then follows the standard procedure from Allen and Holberg, giving $(W/L)_1 = (W/L)_2 = 2.37$, $(W/L)_3 = (W/L)_4 = 0.87$, $(W/L)_5 = 2.07$, $(W/L)_6 = 24.6$, $(W/L)_7 = 29.3$. The schematic for this design is shown in Figure 5.21.

To check that the design meets the specification at this point, the actual minimum V_{out} can be found analytically, and is lower than the 0.5 V required:

$$V_{out(actualmin)} = \sqrt{\frac{2.I_7}{K'_n.(W/L)_7}} = 0.24\,V \tag{5.3}$$

It is also possible at this point to check the power consumption, from the supply voltage and DC current:

$$P_{diss} = (V_{DD} - V_{SS}).(I_5 + I_7) = 0.5\,\text{mW} \tag{5.4}$$

This is higher than ideal, and higher than that in the logarithmic pixels which have very low power consumption, due to bias currents below $1\,\mu$A. This is to be expected for an opamp pixel, and is one of the reasons they are not more commonly used on array ICs, particularly 2D array ICs where the pixel count is generally higher. However, while the power consumption is high, the total figure of 96 mW if three opamps are used for each of the 64 channels is not unreasonable. This lower total power due to a lower pixel count is an advantage of linear array systems.

Finally, the gain with this configuration can be checked. The open loop gain found here is sufficiently large for normal opamp behaviour to be expected.

$$g_{m1} = g_{m2} \tag{5.5}$$

$$A_v = \frac{2.g_{m2}.g_{m6}}{I_5(\lambda_n + \lambda_p).I_6(\lambda_n + \lambda_p)} = 6895 \tag{5.6}$$

5.6.2.3 Compact Opamp Schematic

The schematic of the designed opamp is shown in Figure 5.21.

From the above calculations, the correspondence between the transistor numbering in the design process and the Cadence schematic are: M1=MN4, M2=MN5, M3=MP6, M4=MP5, M5=MN3, M6=MP4, M7=MN6.

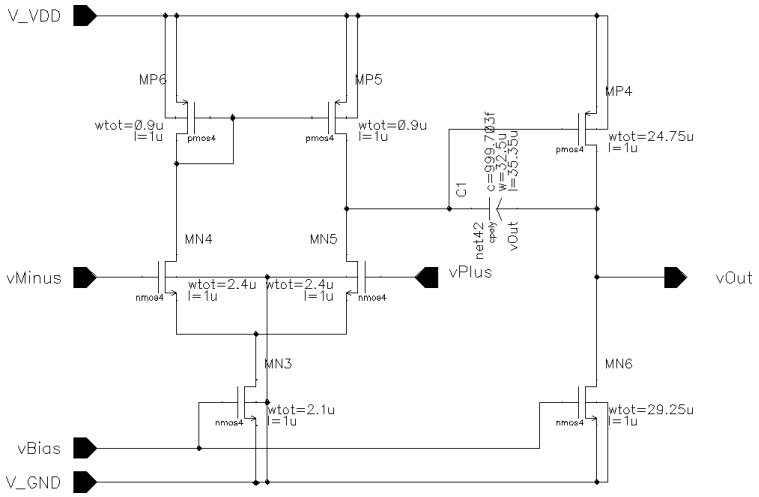

Figure 5.21: Schematic of compact opamp design

5.6.2.4 Compact Opamp Layout

The layout produced for this opamp is shown in Figure 5.22. It is 37 μm tall and 55 μm wide. This leaves space for routing of signals and power between pixels, and the width means that using several per pixel is possible.

The large component on the right of the layout is the compensation capacitor. The smaller transistors at the top left are the input stage, and the larger transistors in the bottom left are the output stage. The input pair and load pair of transistors use a common centroid layout.

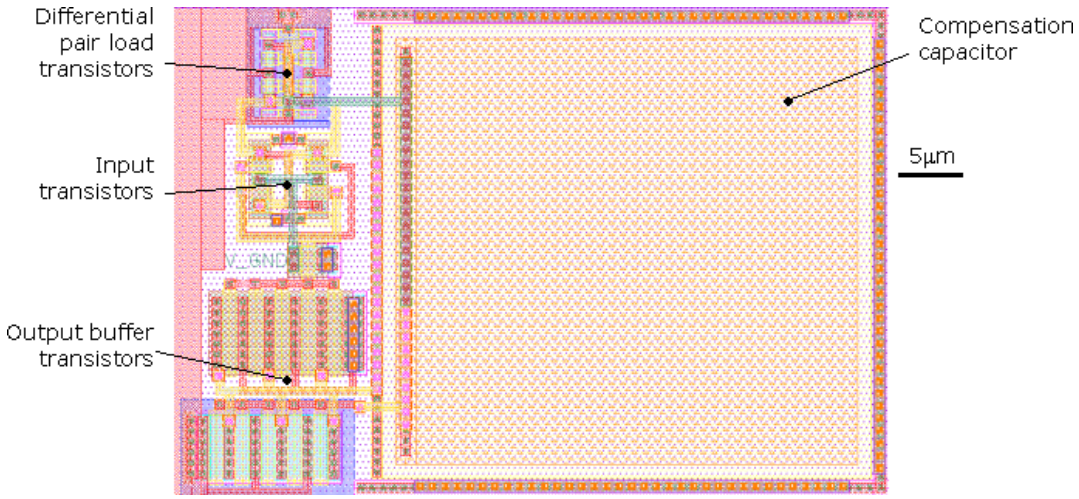

Figure 5.22: Layout of compact opamp design

5.6.3 Simulation of Compact Opamps

To verify the design of the compact opamps, simulations were performed on the opamp separately to the front-end circuit, comparing the compact design to that used in the opamp front-end on BVIPS1.

5.6.3.1 Common Mode Range

Figure 5.23 shows the DC response of both opamps when connected as unity gain buffers, showing the input common mode range of both designs. For the compact opamp, the common mode range (CMR) is from 0.1 V to 3.0 V, and is slightly non linear from 0.2 V to 2.8 V. For the original opamp design, the CMR is 0.1 V to 3.25 V. The range of the new opamps is slightly lower than the original design, but this range is sufficient for use as a front-end circuit, as rail-rail operation is not required.

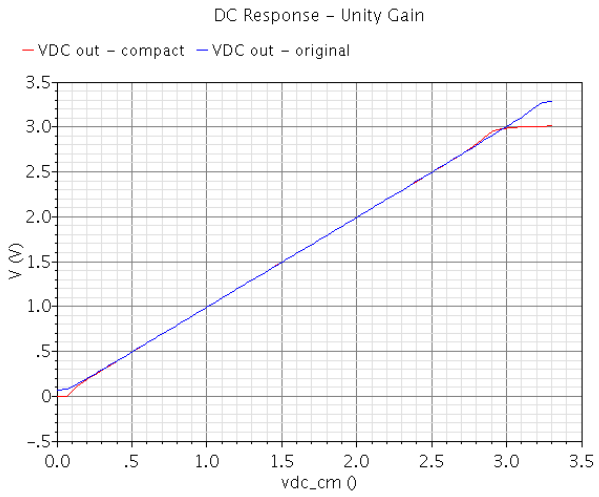

Figure 5.23: DC response of compact and original opamps with unity gain feedback, showing input common mode range

5.6.3.2 Open Loop Gain

Figure 5.23 shows the DC response of both opamps when connected with no feedback, showing the transition from high to low output voltage. A virtual ground was set at 1.65 V, i.e. $V_{DD}/2$. The open loop gain, as measured by the slope of this transition, is 1290 for the original opamp design and 7960 for the compact opamp, showing that the compact opamp is closer to the ideal case in this respect, although the open loop gain is significantly lower than the calculations during the design process suggested.

The input offset voltage of both circuits can also be found for this by the input voltage at which the output voltage is equal to the virtual ground voltage of 1.65 V. For the compact opamp this is 1.6489 V, giving V_{OS} of 1.122 mV, and for the original opamp this is 1.6495 V, giving V_{OS} of 0.509 mV. The compact opamp is therefore less-ideal than the original, but the offset voltage is small enough to not be a problem in an application such as this where DC voltage is not critical.

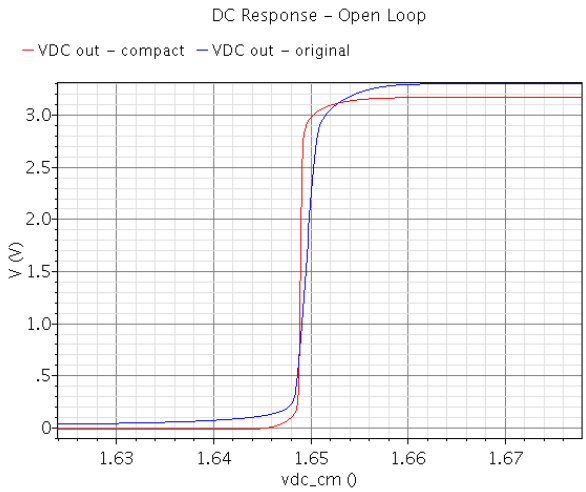

Figure 5.24: DC response of compact and original opamps with no feedback (open loop), showing input common mode range

5.6.3.3 Phase Margin

Adjusting the DC operating point to account for the offset voltages allow an AC simulation to be performed on the opamps in the open loop condition, giving the phase margins of both opamps. Figure 5.25 shows the frequency response of both circuits. For the compact opamp the gain falls to 0 dB at 8.5 MHz where the phase is -115º, giving a phase margin of 65º. For the original opamp the gain falls to 0 dB at 40.7 MHz, where the phase is -100º, giving a phase margin of 80º. This shows that both designs are stable. Additionally, the open loop gains from these simuations are 77.7 dB (7636) for the compact opamp and 62.2 (1287) for the original opamp, showing good agreement with the DC simulations. The discrepancies are due to the offset voltages not being perfectly set, resulting in a DC operating point slightly different to the optimum (where V_{OUT} is at exactly 1.65 V, i.e. the virtual ground voltage).

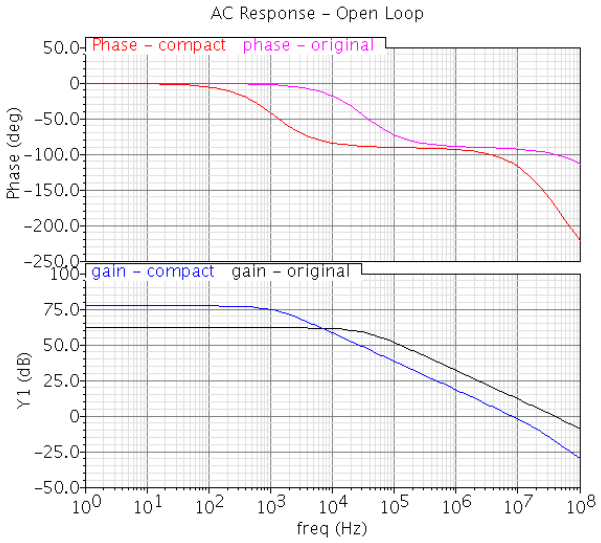

Figure 5.25: Frequency response of compact and original opamps with no feedback (open loop), showing phase margins

5.6.3.4 Noise Response

Finally, the noise performance of the new opamp should be considered. Figure 5.26 shows the input referred noise spectra of both designs in the unity gain configuration with an input voltage of 1.65 V. The original design shows lower noise, particularly lower 1/f noise. However, for the new design the noise density drops to ~$1\,\mu V/\sqrt{Hz}$ at 100 Hz, with an average of $43\,nV/\sqrt{Hz}$ over the 100 Hz - 30 kHz LDBF bandwidth. This gives an RMS noise of $7.5\,\mu V$, which is insignificant compared to the noise levels seen during characterisation of BVIPS1.

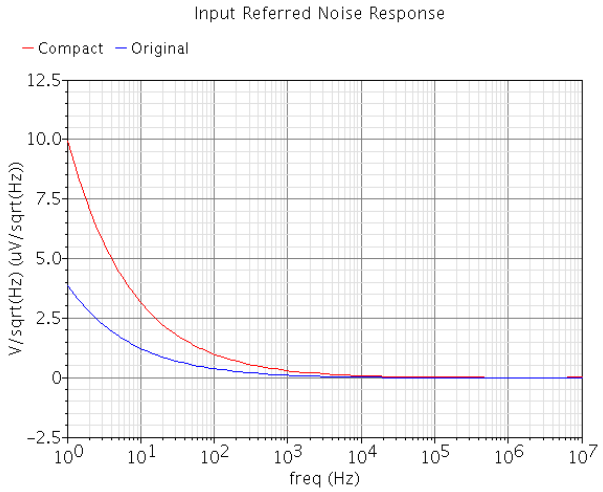

Figure 5.26: Input referred noise response in unity gain configuration, 1.65 V DC input, of both opamp designs

5.6.4 High-Pass Filter for AC Channel

In discrete designs, the DC component of the input signal to the gain stage can be removed with a simple CR high-pass filter. This replaces the varying DC component with a known, fixed DC value. However, this application requires a low-cut off of around 100-200 Hz, requiring large RC values. This makes implementation on-chip problematic, as the large R and C values require large silicon area.

Implementing a HPF on chip would allow the DC value of the input signal to the gain stage to be known. This allows an offset very close to this known value to be used, with only the differential voltage between the two values amplified. This means that a higher gain can be used before saturation occurs. It also means the circuit is less susceptible to changes in DC photocurrent, which would otherwise change the DC input voltage to the gain stage.

One option here is to use diode connected transistors to provide a high resistance, in a similar manner to the logarithmic front-end. The schematic of this circuit is shown in Figure 5.27. This circuit can be considered as a basic CR high-pass filter, with transistors used to implement the resistors. Sets of three diode connected transistors are used in series to increase the total resistance. Using one resistor network from the output node to ground, and one from the output node to V_{DD} allows a DC bias current to flow in the transistors, setting up their DC operating point. This DC bias also introduces a DC output voltage (of half the supply voltage), replacing the original DC voltage. The capacitor used in Figure 5.27 is a polysilicon

capacitor, as MOS capacitors (which have a higher capacity per unit area) have to have one connection to either ground or V_{DD}, and are therefore not suitable for this circuit.

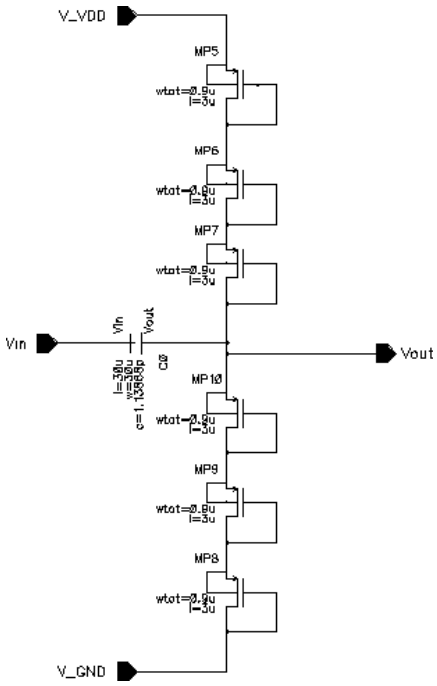

Figure 5.27: High-pass filter circuit using diode connected transistors to form a CR circuit

Figure 5.28 shows the AC response of the high pass filter designed, with a signal having a 0.5 V DC level and 1 V AC magnitude. This shows that the cut-off is sufficiently low for this application. The voltage falls to zero at DC, showing that the DC input voltage is completely blocked.

It can also be seen that the cut-off is not very sharp. This means that the signal is attenuated slightly at frequencies around 1 kHz, which are within the pass-band. While the attenuation is fairly small, the signal at this point (i.e. before the AC amplification stage) can be $<$10 mV, so this may cause a reduction of the SNR.

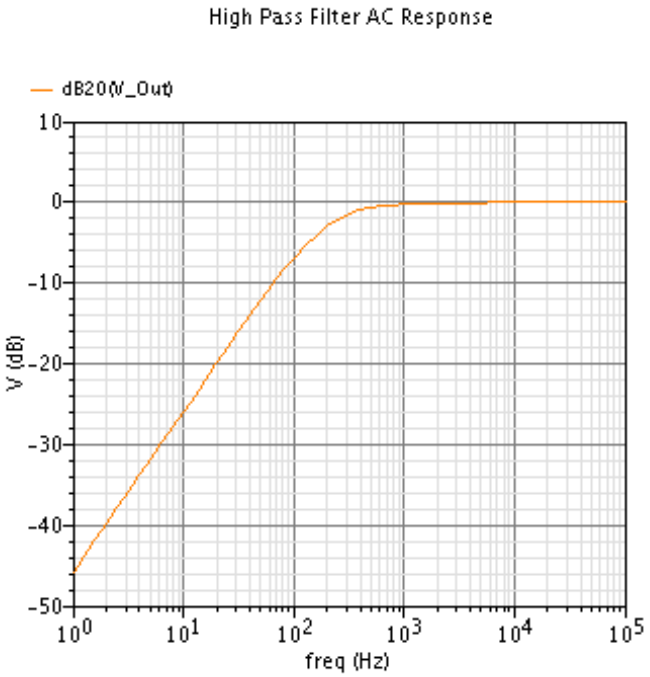

Figure 5.28: AC response of high-pass filter circuit

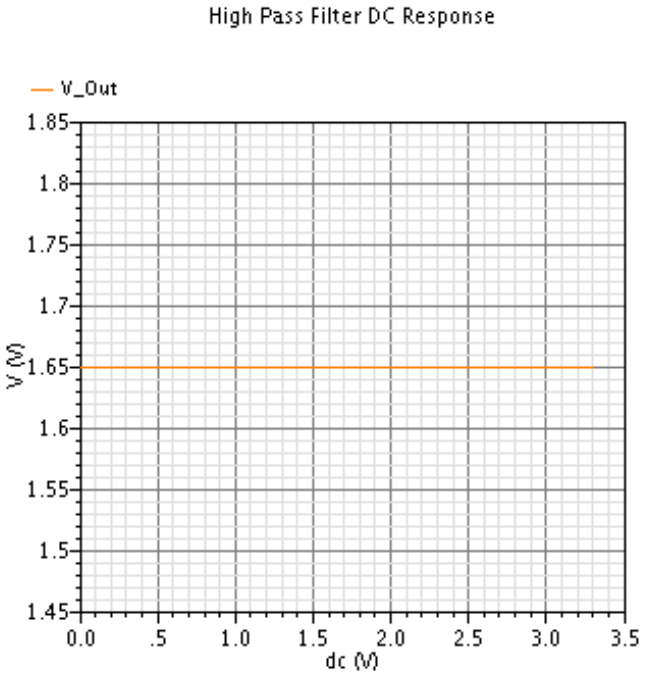

Figure 5.29: DC response of high-pass filter circuit, showing constant DC output voltage

Figure 5.29 shows the DC response of the high pass filter. The output voltage is 1.65 V at all input voltages, showing that the output voltage will not depend on photocurrent,

and is halfway between the supply voltages. This allows the highest gain to be used before saturation, as an AC peak-peak signal of around 3 V would be required to cause saturation.

Figures 5.30, 5.31 and 5.32 show the transient input and output voltages for the filter with three different input frequencies. The output voltage AC component can be seen to rise with frequency, while the DC component remains constant. This further demonstrates the attenuation caused to the input signal at signals within the pass band. As this range is where a large part of the power in the Doppler signal is to be found, this may be problematic.

There is also some distortion evident in the output signal, particularly at 1 kHz and 3 kHz, where the input signal is closer to the cut-off frequency of the filter. This distortion could affect the flow measurement of a signal, as the distortion will introduce higher frequency harmonics of the fundamental. The effect of signals at these frequencies will be increased by the frequency weighting of the flow calculation algorithms. However, the harmonics are themselves related to the magnitude and frequency of the original signal, and therefore the flow measurement is still proportional to the actual flow level. Because of this, the distortion would not be expected to prevent the IC from operating as intended.

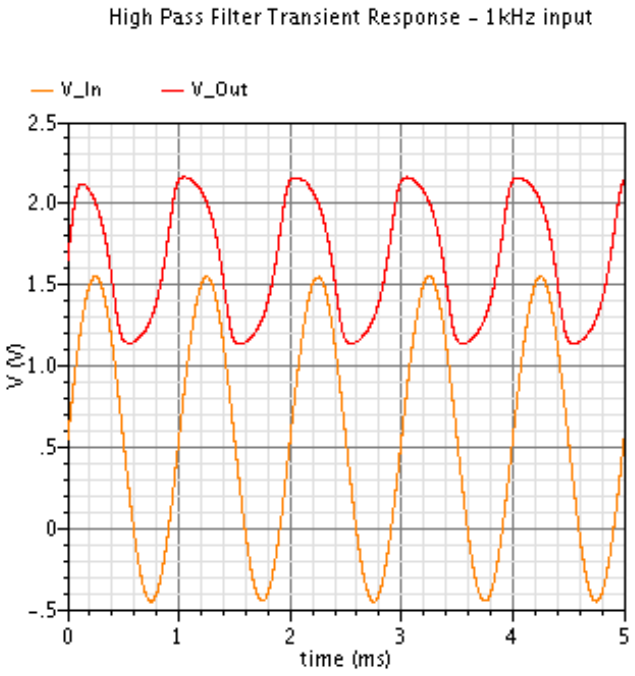

Figure 5.30: Transient response of high-pass filter circuit with 1 kHz input

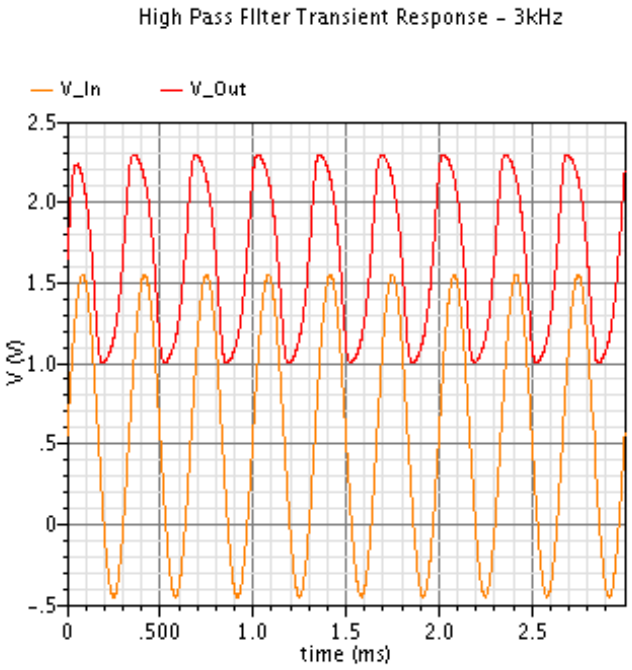

Figure 5.31: Transient response of high-pass filter circuit with 3 kHz input

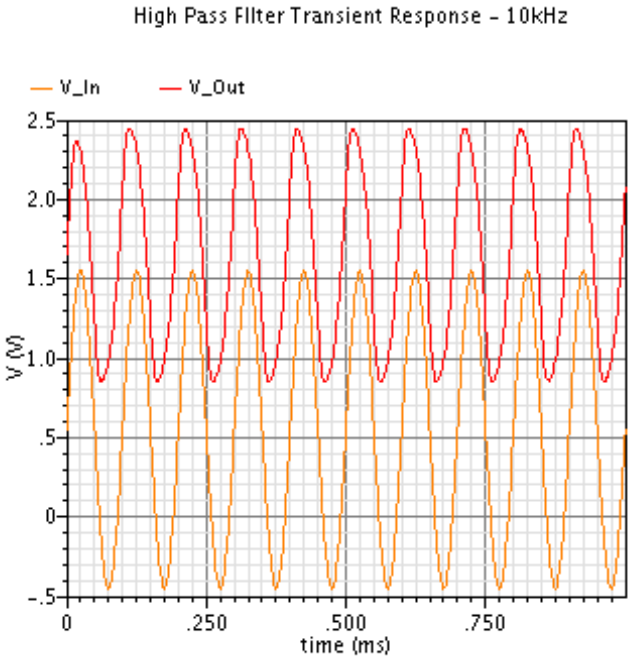

Figure 5.32: Transient response of high-pass filter circuit with 10 kHz input

A possible problem with this circuit is that the small currents involved due to the high resistances required make the design susceptible to process variations. In this situation a

variation in cut-off frequency would cause high fixed pattern noise, as pixels with closely matched front-end circuits could have different amounts of the Doppler signal blocked by filters with different cut-off frequencies.

Figure 5.33 shows the results from Monte-Carlo analysis of 100 iterations of this design. This shows that the low cut-off frequency does not vary greatly from circuit-circuit, with a mean of 199 Hz and a standard deviation of 13.4 Hz.

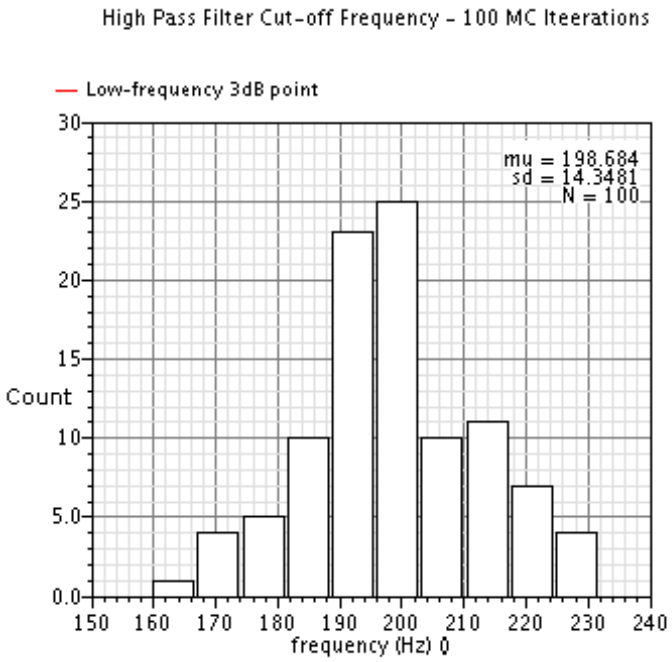

Figure 5.33: Variation of cut-off frequency of HPF, 100 Monte Carlo Iterations

5.6.5 HDA Amplifier AC Channel

As the HPF designed above is an untested design, which relies on very low currents to provide the low cut-off frequency, it is possible that the behaviour of the fabricated IC will not match that expected from simulation. For this reason, the linear pixel was designed with an alternative AC channel. This channel uses the same HDA design as the logarithmic pixel, providing a relatively large AC gain without amplifying the DC.

While this does not follow the approach of implementing a common linear photodetector design on-chip, this approach may combine the good power supply and common mode noise rejection of an opamp photodetector, with the small size and good dynamic range of the hysteretic differentiator circuit. As the smaller HDA design used here is very compact com-

pared to the other components of the opamp pixel, its inclusion has little impact on the size of the pixel. Additionally, using this channel means the linear pixel matches the logarithmic pixel in having an output channel that combines AC and DC. While separate channels may prove to be more useful, a single channel means that a single ADC with no channel switching can sample AC and DC simultaneously. The two components can then be separated by the digital processing section.

5.7 On-Chip Processing

A major change between the 1st and 2nd prototype ICs is the use of integrated processing. This is one of the advantages presented by linear array detectors - as less space is required for the analogue signal detection circuits, additional circuitry can be included on the IC to perform flow processing without making the IC too expensive for inclusion in a commercial imager.

This part of the IC is not the work of the author, but as part of a parallel project at the University of Nottingham developing a 2D blood flow imager [Hoang, 2009, Hoang et al., 2010]. It is briefly discussed here to give a description of the type of full system that can be implemented on chip when using a linear detector array of the type designed by the author.

5.7.1 Motivation for On-Chip Processing

Combining the processing onto one IC has a number of advantages. These advantages are partly in terms of circuit design, allowing a technically better system to be built. In addition to this, the integration can result in smaller component counts, allowing the cost of a full system to be reduced and making a commercial device more feasible. The advantages of using on-chip processing are:

- Data bottlenecks - for this system, if raw data is sent off chip for every pixel of a 64x1 system, sampling at 40 kHz with a 10-bit ADC, the data rate required is 25.6 Mbps. While this is possible (USB 2.0 theoretical data rate is 380Mbps), it adds additional constraints to the design. Using integrated processing means that only one value for flux needs to be sent off-chip per pixel per frame. If the 1024-point FFT processing from earlier results is used, the transmitted data is reduced by 1024 times (or 512 if a

DC value is also sent). This advantage becomes more significant as pixel count rises, although the lower pixel count of the 1D array means this aspect of integrated systems is more advantageous for 2D detector arrays than 1D arrays.

- Reduced component count - As the output from an integrated system is a flow value that does not require any further data acquisition or processing, the demands on an external system are greatly reduced. While an external system will still be required to handle data transmission/display, the lower computation power required potentially allows the use of a low-cost, low-power microcontroller to fulfill this role. Additionally, the lower component count can make the device physically smaller. This could lead to more applications for the sensors developed here, such as developing battery powered hand-held devices. However, the continual development of low power but powerful microprocessors may reduce this advantage, and an external processing system may not add significant physical size compared to components such as the laser and power supply.

5.7.2 Implementation of Processing

Processing of data from BVIPS1 was performed on a PC, using a 1024-point FFT. This process is computationally intensive and can require proprietary algorithms, so a simpler method is used for implementation on chip, using time-domain filters [Hoang et al., 2010].

Integrated processing is performed using 3rd order IIR time domain filter processing ($\omega^{\frac{1}{2}}$ + square + average filter). The RMS error in flow values from this algorithm compared to the 1024 point FFT processing is approximately 2% (RMS error compared to mean normalized flow)[Hoang, 2009], showing that on chip processing can be used in place of off-chip processing on a PC, and hence the system-on-chip approach is suitable for this application.

5.7.3 Flexibility of the Integrated System

While the BVIPS2 IC is designed to integrate all functions onto one IC, the IC topology is sufficiently flexible to be used in other modes - bypassing the flow processing block, or the ADCs and processing block.

To test the analogue detector components only, the inclusion of an analogue output means that all data acquisition, digitisation and processing can be performed off-chip, with the

multiplexers also controlled by an off-chip processor. The same output is used for both linear and logarithmic front-ends, and AC and DC channels. These can be selected by electrical input signals, allowing a single data acquisition system to operate the IC in all modes.

Alternatively, the on-chip ADCs can be used to convert the raw analogue signal into digital data which is then output from the IC. This still requires off-chip processing, but removing the requirement for analogue-digital conversion simplifies the off-chip parts of the full system.

The digital section itself can be controlled externally, and test signals can be applied through the ADC outputs. This is intended to be used to test the digital section, but could also be used to process signals from an external analogue front-end using the on-chip processing system. This may be useful for prototyping future sensor designs.

Finally, the digital section can perform both flux processing and multiplexer control functions. This is the highest level of integration.

5.8 Laboratory Characterisation

5.8.1 Characterisation Equipment and Procedure

Characterisation of the BVIPS2 IC was performed using the same optical setup as that for BVIPS1 (described in Section 4.3.1). The IC was directly illuminated with a red laser, with the power and modulation depth varied to provide illumination over the range of photocurrents expected from LDBF. However, the FPGA data acquisition system used for BVIPS1 was replaced with a National Instruments ADC card (PCI-6259), which allowed sampling of the data from the ICs analogue output (see Figure 5.1) as well as control of the various digital control signals. The full characterisation setup used for the BVIPS2 IC is shown in Figure 5.34.

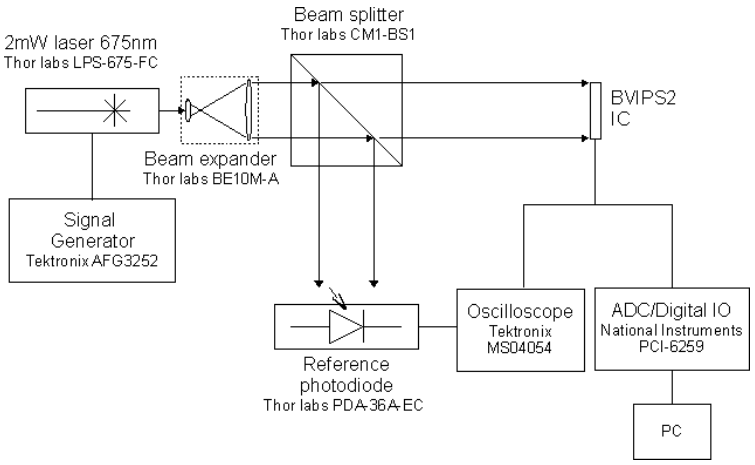

Figure 5.34: Equipment setup for characterisation of the BVIPS2 IC

The range of photocurrents used was from 250 pA - 80 nA. For AC characterisation a 15% modulation depth signal at 3 kHz was applied. These values are based on measurements of light power using the Moor LDLS, described in Section 4.3.3, with a range extending above the highest current found from these measurements, and a low current limit below that measured as typical from earlier (and much lower) estimates shown in Section 1.8.1. The modulation depth is that accepted as typical of LDBF taken from Belcaro [Belcaro et al., 1994], using a frequency which is in the pass-band of the filters used by a sufficient margin for variation of filter roll-off steepness and cut-off frequency to not have a disproportionate effect on the characterisation, while still being in the lower part of the frequency range where the majority of the Doppler signal is found.

5.8.2 DC Response

5.8.2.1 Logarithmic Pixel

Figure 5.35 shows the DC behaviour of the logarithmic pixel design from pre-layout simulations. The slope is as seen in the log pixels on the previous IC (Figure 4.9), with a straight line when plotted on a logarithmic scale. The DC and AC channels have the same slope, as while the AC channel has an approximately unity gain at DC, there is a level shifting effect.

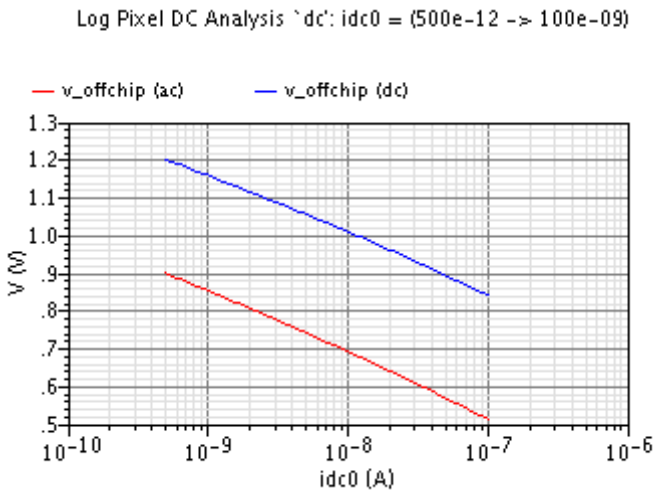

Figure 5.35: DC Response from pre-layout simulation of the logarithmic pixels on the BVIPS2 IC

Figure 5.36 shows the DC measured behaviour of the logarithmic pixel DC outputs for each of the six channels, compared to post-layout simulations. The post-layout simulation results are very similar to the pre-layout results. The measured results show a higher DC voltage at all DC photocurrents than the simulated results, although the slope is similar. The six front-end circuits give a similar response, although the offset between different blocks represents a significant proportion of the photocurrent range, demonstrating the need for accurate calibration (or front-end selection to reduce FPN).

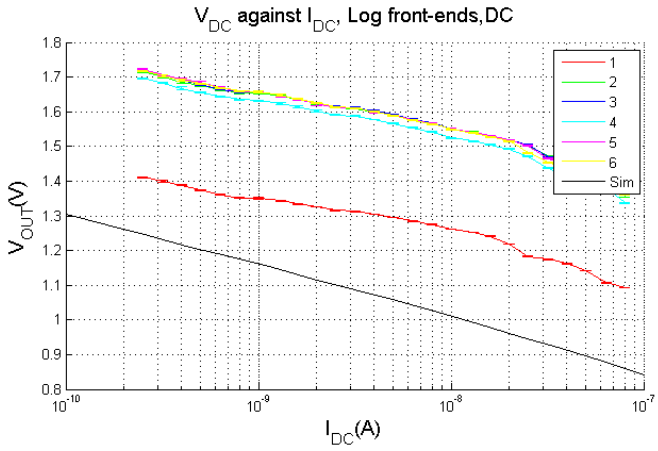

Figure 5.36: Measured and Simulated DC output voltage against DC photocurrent of the logarithmic pixel DC output of the BVIPS2 IC

Figure 5.37 shows the DC measured behaviour of the logarithmic pixel AC outputs for each of the six channels, compared to post-layout simulations. Again, the post-layout simulation

results are very similar to the pre-layout results. The measured results show a higher DC voltage at all DC photocurrents than the simulated results, with the response of all front-end/HDA circuits having a similar response shape. There is also little or no slope on the measured DC response below 7 nA, while above this the DC output voltage drops from around 2.25 V at 7 nA to 1.75 V at 30 nA. As no corresponding change in the slope of the DC output (i.e. the logarithmic front-end output), this appears to be a change in the DC offset caused by the HDA. A step in the DC response of the HDA was shown in the initial simulations of this circuit (see Figure 3.36 in Chapter 3), although this step did not appear to affect the response of the BVIPS1 IC. As this output is not intended for DC measurements this aspect of IC behaviour is not a significant limitation on IC use, but it should be considered when setting bias voltages for the on-chip ADCs or any off-chip processing circuitry.

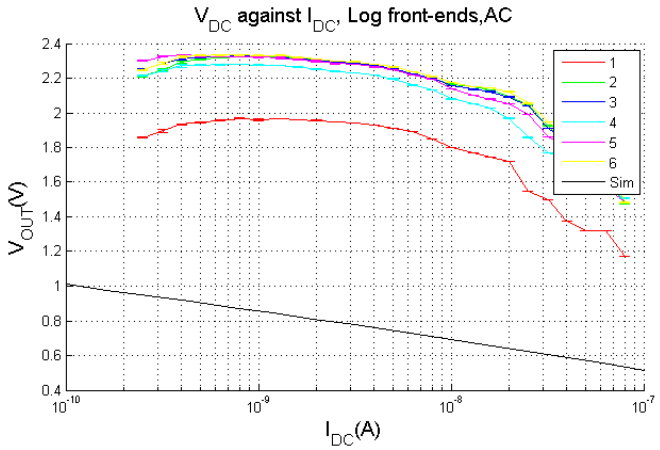

Figure 5.37: Measured and Simulated DC output voltage against DC photocurrent of the logarithmic pixel AC output of the BVIPS2 IC

5.8.2.2 Linear Pixel

Figure 5.38 shows the DC behaviour of the linear opamp based pixel design from pre-layout simulations. The flat response of the AC (opamp gain stage) channel is due to the DC blocking of the high-pass filter. The DC channel has a linear response in the region where the DC gain stage is not causing saturation of the output. The HDA based AC channel has a response that is similar to the DC channel in that it gives its intended operating response only over a fairly narrow range. It is less linear in this range than the opamp gain stage, as the HDA gain varies with the changes in input voltage. This reduced linearity is contrary to the linear design of this pixel, but the HDA channel still offers more linear performance than

the logarithmic pixels, while having lower size and power consumption than the opamp AC gain channel, which may be useful for applications where the pixel size has to be lower than that used here.

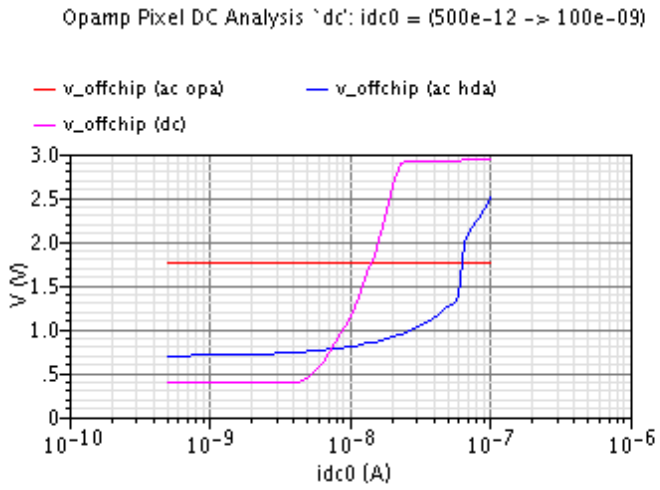

Figure 5.38: DC Response from pre-layout simulation of the opamp pixel outputs on the BVIPS2 IC

Figure 5.39 shows measured DC behaviour compared to post-layout simulation results. The post-layout simulation results are similar to the pre-layout results shown in Figure 5.38. The measured results for the AC (opamp gain stage) output (Figure 5.39a) show the same flat response, although at a slightly lower DC level of 1.66 V compared to 1.75 V. The response is also not totally flat, dropping to 1.64 V at 40 nA I_{DC}, but this variation is relatively low. The AC (HDA gain stage) response shown in Figure 5.39b shows a similar shape, but the DC output voltage is higher at low DC photocurrents (1.2 V measured compared to 0.7 V simulated) and does not rise as rapidly when the DC photocurrent approaches the highest intensity tested (>50 nA). This lower slope is probably due to the non rail-rail operation of the opamp being more pronounced in the actual circuit than predicted from simulation.

For the DC output (Figure 5.39c), the measured DC output voltage starts higher than the simulated value and rises as expected from the simulation results, albeit over a narrower voltage range, and a broader DC photocurrent range.

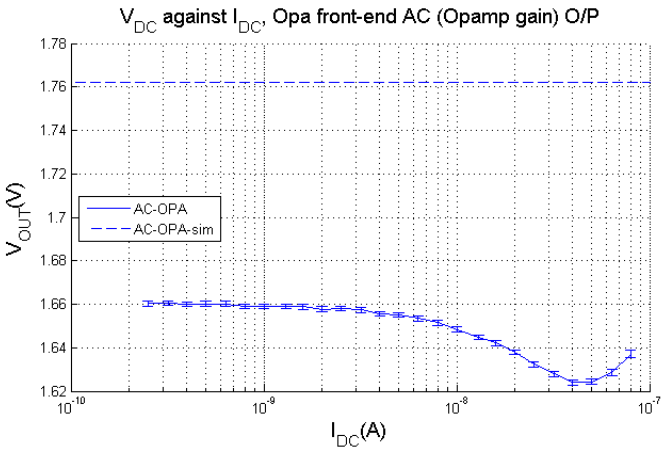

(a) AC (opamp gain) output

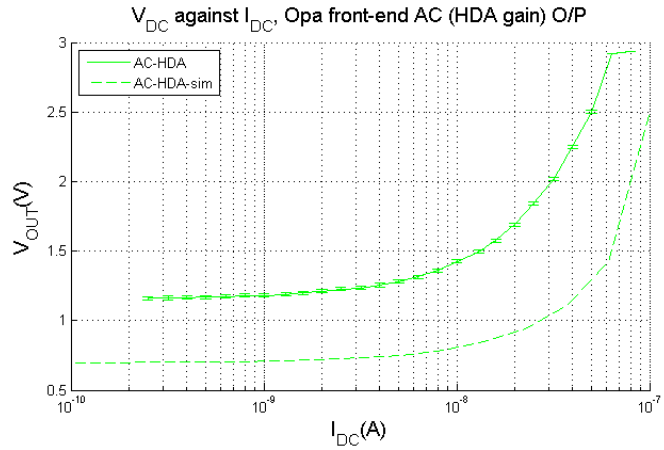

(b) AC (HDA gain) output

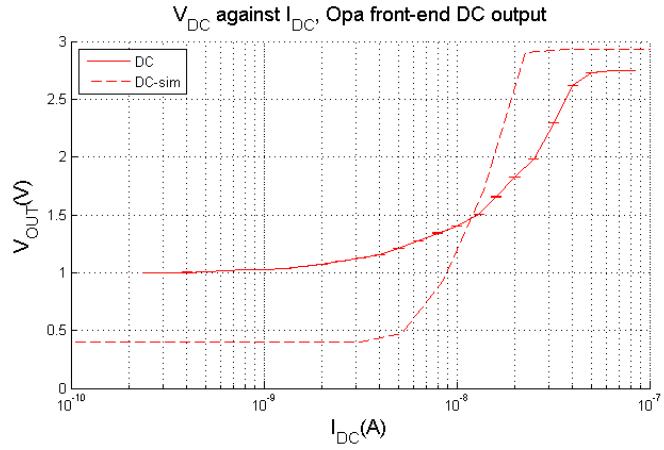

(c) DC output

Figure 5.39: Measured and Simulated DC output voltage against DC photocurrent of the opamp pixel outputs of the BVIPS2 IC

5.8.3 AC Response

5.8.3.1 Logarithmic Pixel

Figure 5.40 shows the frequency response of the logarithmic pixel design from pre-layout simulations. The plot shows a frequency sweep with a DC photocurrent of 10 nA and a 3 kHz AC component with a modulation depth of 15% (i.e. 1 nA peak-peak). The increased gain of the AC channel compared to the DC channel can be clearly seen from the simulation results. The high and low cut-off frequencies can also be seen. The low frequency cut-off shown by simulation of approximately 2 kHz is higher than the 100-200 Hz required. However given the problems with simulation accuracy shown by the HDA faults in the design used on BVIPS1, and successful testing of this design on another IC (discussed in Section 5.5.1), the HDA design was not changed on the basis of these results.

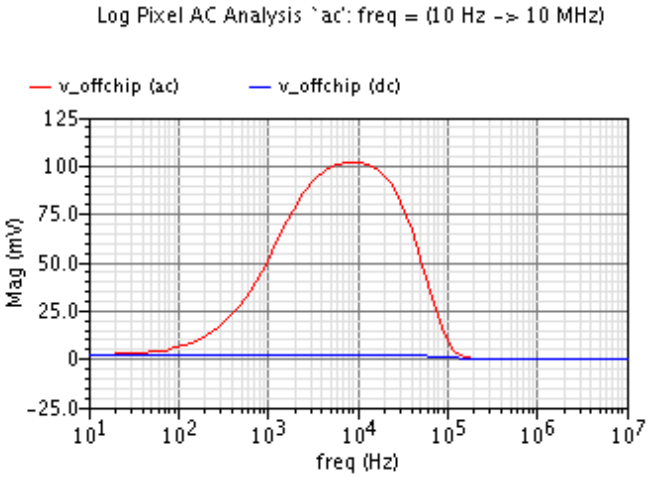

Figure 5.40: AC Response from pre-layout simulation of the logarithmic pixels on the BVIPS2 IC, AC and DC output channels

Figure 5.41 shows the same results as Figure 5.40, re-scaled to display the DC channel AC response. It can be seen that the bandwidth of this output includes DC (as the magnitude does not fall at lower frequencies), and with a much lower gain than the AC channel such that the DC value is easier to determine from a signal with AC and DC components. The cut-off frequency is also similar to that of the AC channel, reducing aliasing during sampling of the output signal. A lower cut-off frequency could be used for the DC channel, however this would require modifying the g_mC filter to use a larger output capacitor, or reducing the bias current (see Figure 3.52 in Chapter 3) by using a separate bias input. Because of

this, the cut-off was not changed. This has an additional advantage in that the raw signal from the front-end can be sampled for flux calculations if required, for example if there was a problem with the HDA or if a higher resolution ADC was available.

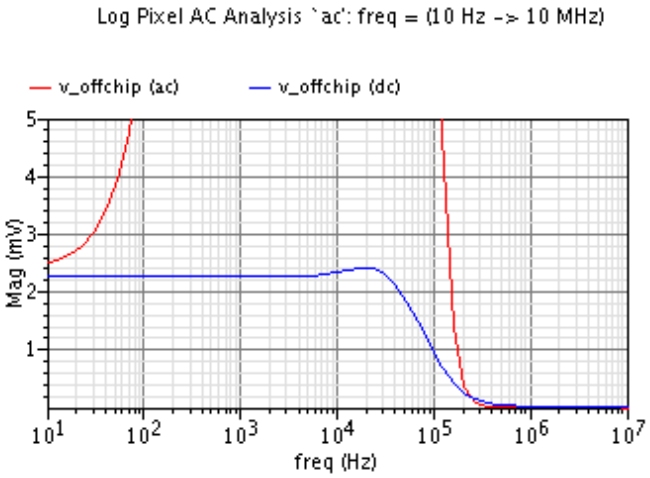

Figure 5.41: AC Response from pre-layout simulation of the logarithmic pixels on the BVIPS2 IC, AC and DC output channels

Figure 5.42 shows the AC peak-peak output voltage against varying DC photocurrent for the BVIPS2 IC logarithmic front-end DC output, comparing measured response with that from post-layout simulations. The modulation depth used was 15% at a frequency of 3 kHz (see Section 5.8.1). Figure 5.43 shows the corresponding response for the logarithmic front-end AC output. The normalising behaviour of the logarithmic pixel means the AC output voltage should remain constant with changing DC photocurrent. This is the case at photocurrents below 1 nA, as shown by the low gradient of the responses in this DC photocurrent range. Above this range the AC voltage drops, suggesting that the normalising behaviour is not working ideally, possibly as the increase in DC photocurrent means the load transistors are operating at the edge of the sub-threshold region where the normalising behaviour breaks down. However, the DC photocurrent in these transistors should be unchanged from that in the BVIPS1 IC (the corresponding response is shown in Figure 4.13), where the slope on the AC output voltage with DC photocurrent was less steep. It should be noted that while the AC output magnitude falls from 40 to 15 mV (falling by a factor of 2.7), the corresponding change in DC photocurrent is from 1 nA to 10 nA (an increase by a factor of 10), suggesting that some level of normalisation is occurring even if not as fully as intended.

The post-layout simulation responses for both DC and AC outputs (Figures 5.42 and 5.43

respectively) show a similar magnitude to the peaks of the response shown from pre-layout simulations in Figure 5.40.

The measured results for the log pixel DC output (Figure 5.42) show higher measured AC magnitude than that expected from simulation, while the log pixel AC output (Figure 5.43) has a similar AC output magnitude between the measured and simulated results. The shape of the measured AC response is different to that expected, with AC output magnitude falling with increasing DC photocurrent. This suggests that the normalisation principle is not operating as intended, although the fall in output magnitude is proportionally much small than the increase in DC photocurrent, the range of which covers several orders of magnitude.

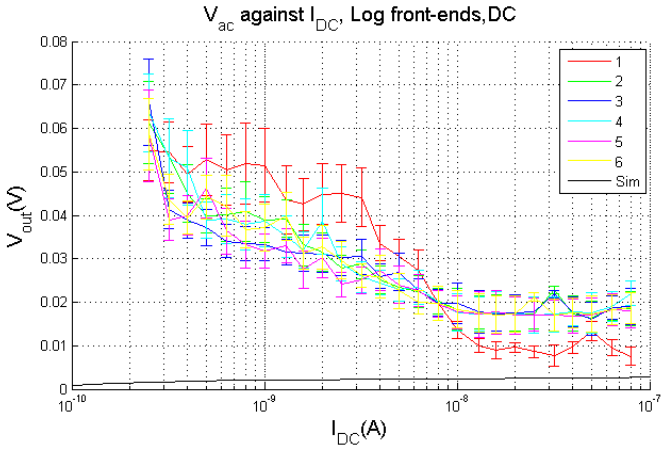

Figure 5.42: Measured and Simulated AC output voltage against DC photocurrent of the logarithmic pixel DC output on the BVIPS2 IC

The gain of the HDA circuit is lower than expected from simulation, as shown by the greater difference between the output AC magnitude for the DC and AC channels in the simulated case compared to the measured case - for the DC output (pre-HDA, Figure 5.42), the measured gain is higher than the simulated gain, whereas at the AC output (post-HDA, Figure 5.43), the simulated output has risen to be equal to or greater than the measured response over the full DC photocurrent range measured. This suggests that the front-end itself has a higher transimpedance than expected from simulation, while the HDA has a lower AC voltage gain.

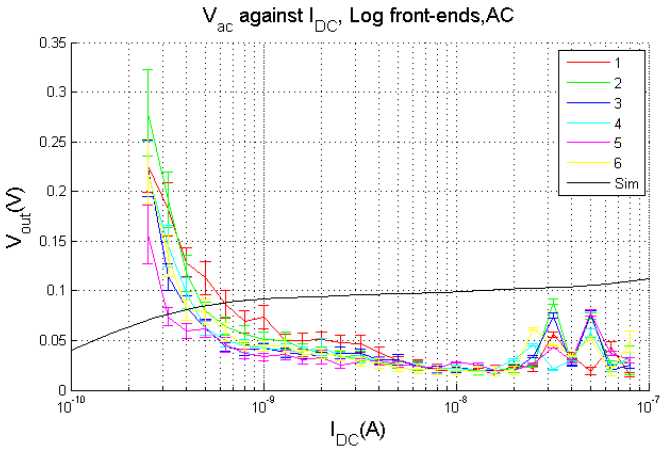

Figure 5.43: Measured and Simulated AC output voltage against DC photocurrent of the logarithmic pixel AC output on the BVIPS2 IC

In addition to considering the AC magnitude of the output signal, the transimpedance of the circuit should be investigated. Figures 5.44 and 5.45 show the transimpedance of the logarithmic front-end DC and AC outputs respectively. The measured transimpedance plots show the expected inverse relationship between DC photocurrent and transimpedance. For the DC output (Figure 5.44) the measured transimpedance is significantly higher than that measured, which corresponds to the higher measured AC magnitude compared to simulation results shown in Figure 5.42.

The increased slope at all photocurrents further demonstrates that the normalisation is not occurring to the same level as expected, as it is this increase slope that means the AC output magnitude drops more rapidly than expected as DC photocurrent rises.

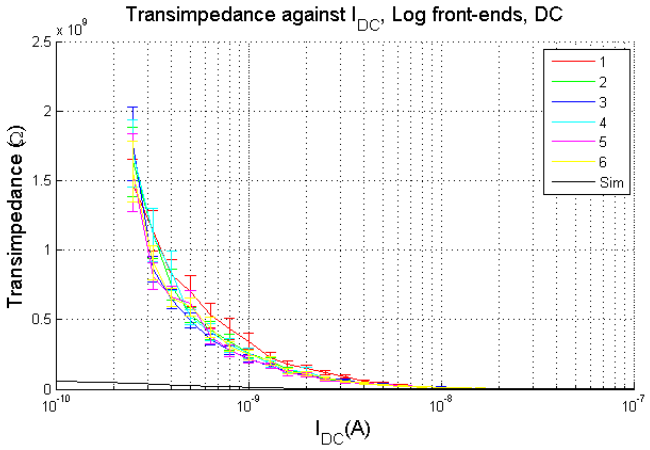

Figure 5.44: Measured and simulated transimpedance against DC photocurrent of the logarithmic pixel DC output on the BVIPS2 IC

For the AC output (Figure 5.45) there is a close match between simulated and measured transimpedance, although the increase in transimpedance at low DC photocurrents ($< 300\,\text{pA}$) is more pronounced in the measured case. Again, this corresponds to the results for simulated/measured AC magnitude shown in Figure 5.43.

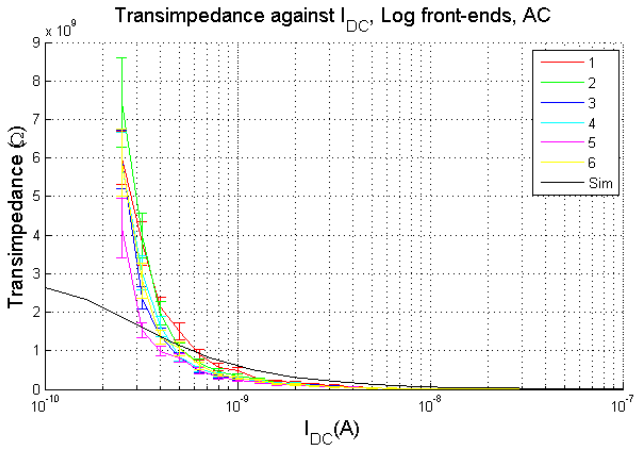

Figure 5.45: Measured and simulated transimpedance against DC photocurrent of the logarithmic pixel AC output on the BVIPS2 IC

5.8.3.2 Linear Pixel

Figure 5.46 shows the frequency response from pre-layout simulations of the linear opamp based pixel design. The plot shows a frequency sweep with a DC photocurrent of 10 nA and a modulation depth of 15%. The high (35 kHz) and low (200 Hz) cut-off frequencies of the opamp gain stage output are similar to the intended values (20 kHz and 200 Hz respectively). The simulated high frequency cut-off is over the required 20 kHz, but this can be adjusted by changing the bias current of the low-pass filters, which will be done based on experimental results. The HDA gain stage response has a similar shape to the opamp gain channel, but with a higher low-frequency cut-off. This is similar to the effect shown in Figure 5.40, so the design was not changed on the basis of these results.

The DC channel has a similar frequency response to that shown for the logarithmic pixel DC channel in Figure 5.40, although the AC gain is higher due to the high gain used by the front-end stage to prevent the gain stage causing saturation of the output voltage.

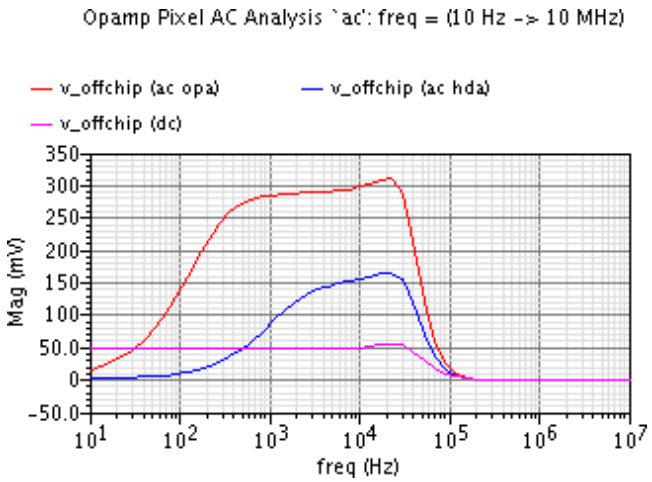

Figure 5.46: AC Response from pre-layout simulations of opamp pixels on the BVIPS2 IC

Figure 5.47 shows AC output voltage magnitude against DC photocurrent for the opamp pixels on BVIPS2, comparing measured response with that seen from post-layout simulations. The AC component of the photocurrent used for this measurement had a modulation depth of 15% and a frequency of 3 kHz.

For the AC (opamp gain stage) output (Figure 5.47a), the measured and simulated responses have a similar shape, although the measured magnitude is approximately 30-50% of the corresponding simulated magnitude. This may be due to manufacturing variations giving a lower feedback resistance in the front-end, or variations in the gain stage resistors. The measured response is still linear - for example, from 10 nA to 50 nA (an I_{DC} increase by a factor of 5), the output voltage rises from approximately 0.15 V to 0.7 V, a factor of 4.67. For the AC (HDA gain stage) output (Figure 5.47b), the increase in measured output magnitude with rising DC photocurrent is less linear than the simulated response. The measured response has a higher output AC magnitude at lower DC photocurrents, and begins to follow the rise of the simulated response, but the simulated response rises faster at higher photocurrents (>10 nA). This follows the DC response shown in Figure 5.39, which shows the DC output voltage of this circuit also rising more rapidly in simulation than measured behaviour.

For the opamp front-end DC output (Figure 5.47c) the measured and simulated behaviour show a similar magnitude at lower photocurrents. Above 10 nA the simulated AC magnitude falls rapidly back towards zero (as the opamp saturates), while the measured magnitude begins to fall but at a lower rate - this follows the results seen from DC characterisation in Figure 5.39, where the measured results did not show saturation occurring at the same point

as predicted from simulation.

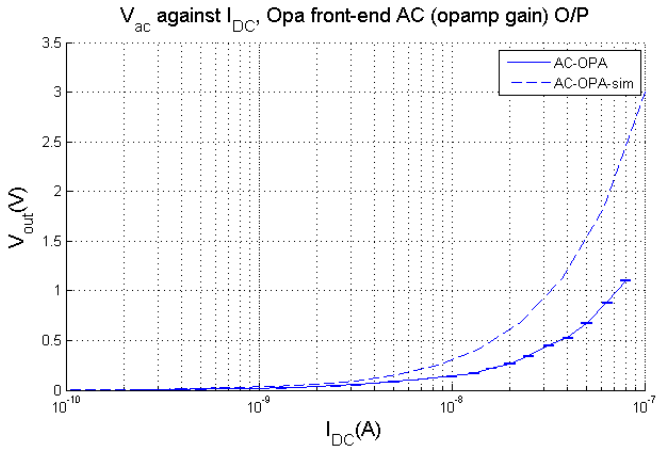

(a) AC (Opamp gain) output

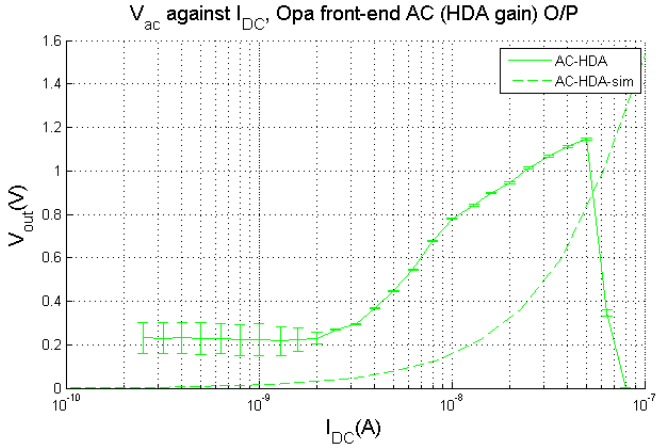

(b) AC (HDA gain) output

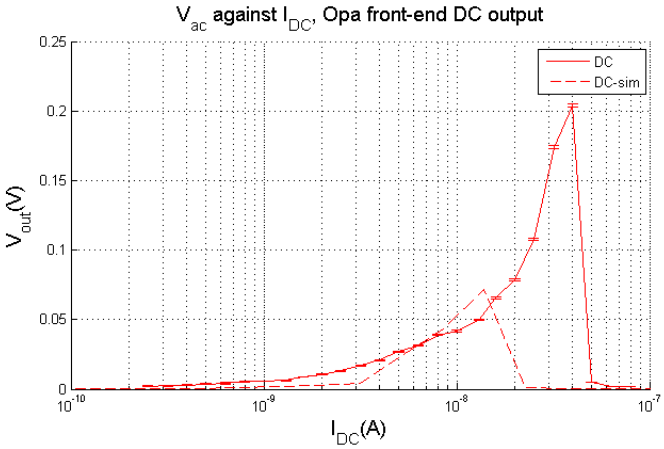

(c) DC output

Figure 5.47: Measured and simulated AC output voltage against DC photocurrent of the opamp pixels on the BVIPS2 IC

Figure 5.48 shows AC transimpedance against DC photocurrent for the opamp pixels on BVIPS2, comparing measured response with that seen from post-layout simulations. The AC component of the photocurrent used for this measurement had a modulation depth of 15% and a frequency of 3 kHz.

The flat response of the simulated results for the AC outputs (seen most clearly in Figure 5.48b) show the expected constant transimpedance over varying DC photocurrent. The simulated DC output is mostly flat, but with an increase in transimpedance from the DC output from approximately 4-14 nA. Outside this range the DC output saturates (or, at the edges of this DC photocurrent range, is clipped), reducing the transimpedance.

The measured results for the DC and AC (opamp gain) channels, shown in Figure 5.48b, show a relatively flat response, although the AC (opamp) channel has lower transimpedance than expected from simulation. This channel also shows higher transimpedance at low photocurrents, although the combination here of very low photocurrents and output voltage noise leads to erroneously high transimpedance values. For the DC channel, the transimpedance remains at the level predicted from simulation in the 4-14 nA DC photocurrent range over a wider photocurrent range, suggesting that saturation beyond this range does not occur as abruptly as expected from simulation.

For the AC (HDA gain) channel, the measured transimpedance falls continuously with increasing photocurrent, from a higher level than predicted from simulation ($6 \times 10^9\Omega$ measured / $1 \times 10^8\Omega$ simulated at 250 pA), to almost zero at 80 nA. This suggests that the changing DC voltage at the HDA input affects the AC gain of this circuit. There is a flatter section of the response between 3 and 10 nA, where the opamp front-end is intended to operate, suggesting that the HDA gain channel is still operating as required, albeit over a narrow DC photocurrent range.

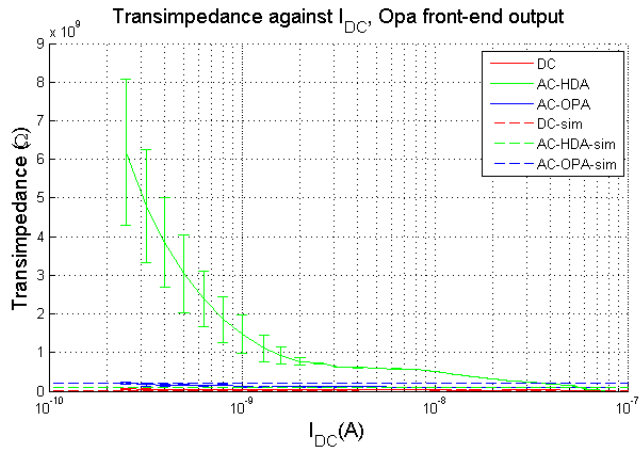

(a) Full transimpedance range

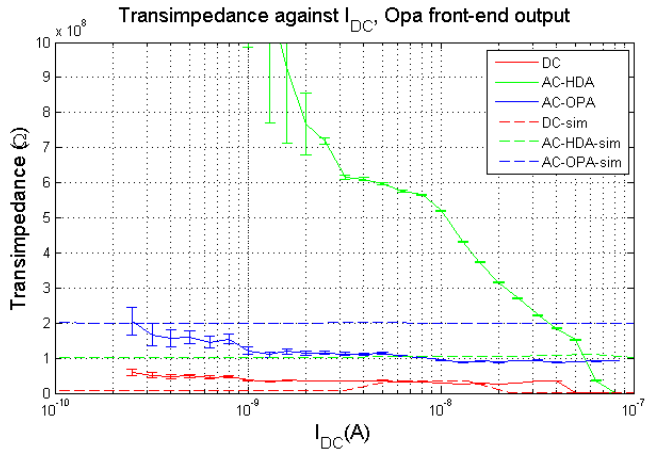

(b) Detail of lower transimpedance values

Figure 5.48: Measured and simulated transimpedance against DC photocurrent of the opamp pixels on the BVIPS2 IC

5.8.4 Noise Response

5.8.4.1 Logarithmic Pixel Noise Spectra

Figures 5.49 and 5.50 show the input referred current noise spectrum and output voltage noise spectrum of the log pixel at 10 nA DC photocurrent, as found from pre-layout simulation. The spectra are a similar shape to those seen for the previous IC, with an input noise density of approximately $70\,\mathrm{fA}/\sqrt{\mathrm{Hz}}$ in the pass band for the AC channel. This is higher than that seen previously (shown in Chapter 3, Figure 3.17) as the extra elements added to the circuit (additional front-ends, switches for circuit selection, higher order filters) will increase the theoretical noise. However, the characterisation of the BVIPS1 IC showed that

the actual noise levels were significantly higher than the simulated noise levels (shown in Chapter 4, Figure 4.23). The changes should result in the measured noise being closer to the theoretical noise (for example, due to better rejection of external noise sources and reduced noise bandwidth), and therefore resulting in an overall fall in measured noise levels.

The DC channel has a higher input referred pass band noise density than the AC channel due to its reduced gain. The noise density of this channel does not rise as sharply at low frequencies due to the lack of a high-pass filter on this channel, meaning that near DC this channel has higher gain, which results in lower input referred current noise.

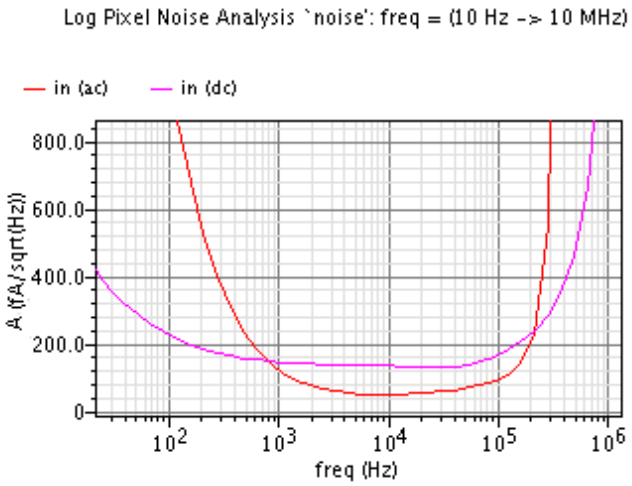

Figure 5.49: Input referred current noise response from pre-layout simulation of the logarithmic pixels on the BVIPS2 IC

The output voltage noise spectrum (Figure 5.50) shows a lower noise level on the DC channel, as noise at the input to this channel is not amplified by as much. The noise spectra for both channels drops off above the designed cut-off frequency, as expected.

The AC channel has a drop in the noise density spectrum between 1 and 4 kHz. This suggests that some element of this circuit has a reduction in gain above this frequency. The cause of this is not clear, although this frequency does correspond to the peak at a higher frequency than expected of the AC gain shown for the AC channel in Figure 5.40. This was considered to be caused by inaccuracies in the simulation of the HDA, which may also cause the irregularity shown here.

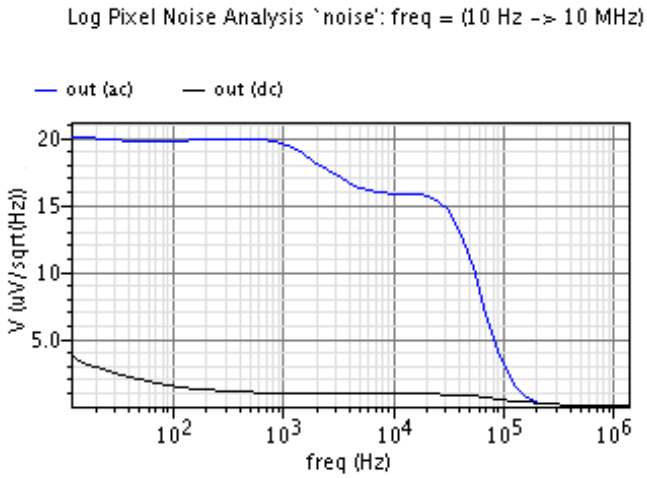

Figure 5.50: Output voltage noise response from pre-layout simulation of the logarithmic pixels on the BVIPS2 IC

5.8.4.2 Logarithmic Pixel Noise against DC Photocurrent

Figure 5.51 shows input referred RMS current noise over 100 Hz-30 kHz for the logarithmic pixel DC outputs of the BVIPS2 IC, comparing measured results with post-layout simulations. The simulated RMS noise increases proportionally to the square root of DC photocurrent, as expected for logarithmic pixels. The measured results also follow this shape, but the noise current is approximately five times greater than predicted from simulation. While this is a significant increase, the discrepancy is lower than that seen for BVIPS1, where the measured noise values were an order of magnitude greater than simulated values (see Section 4.3.6).

All alternative front-end circuits show approximately equal measured noise levels, although as for the previous results the variation with increasing photocurrent does not follow a smooth curve, suggesting problems with the measurement process (such as variation in illumination intensity during measurements).

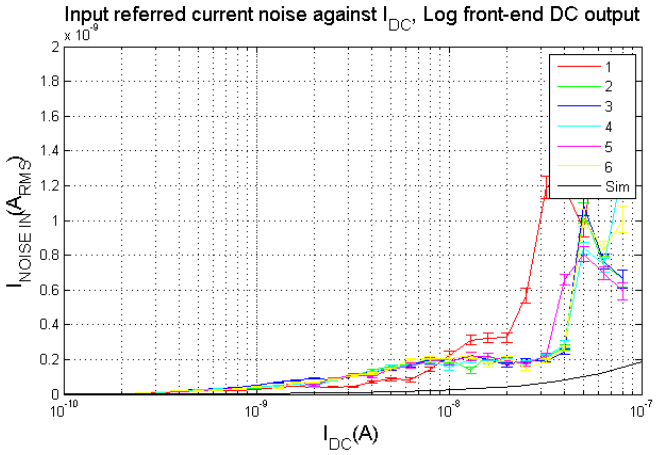

Figure 5.51: Measured and simulated input referred RMS current noise against DC photocurrent of the logarithmic pixel DC outputs on the BVIPS2 IC

Figure 5.52 shows input referred RMS current noise over 100 Hz-30 kHz for the logarithmic pixel AC outputs of the BVIPS2 IC, comparing measured results with post-layout simulations. The shape of the noise responses here is similar to that seen for the DC output shown in Figure 5.51, but the simulated noise levels here are approximately one third of those for the DC output. This reduction is due to the voltage gain of the HDA, which means that a similar measured voltage noise at the IC output corresponds to a lower input referred current noise (or alternatively that an increase in measured output voltage noise does not correspond to a proportional increase in input referred current noise). This is further demonstrated by the output voltage noise responses shown in Figures 5.53 and 5.54.

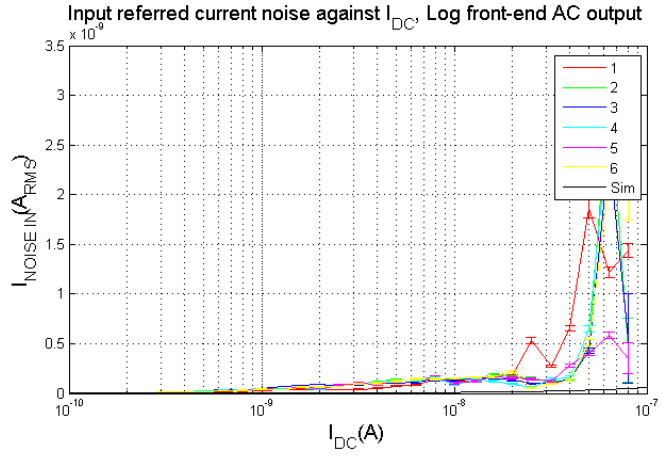

Figure 5.52: Measured and simulated input referred RMS current noise against DC photocurrent of the logarithmic pixel AC outputs on the BVIPS2 IC

Figure 5.53 shows RMS voltage noise measured at the output over 100 Hz-30 kHz for the

logarithmic pixel DC outputs of the BVIPS2 IC, comparing measured results with post-layout simulations. The simulations show voltage noise falling with increasing photocurrent, which is predicted from logarithmic pixel theory (see Section 3.2.1.4). The measured RMS noise follows a similar trend, but with approximately an order of magnitude increase over the simulated results..

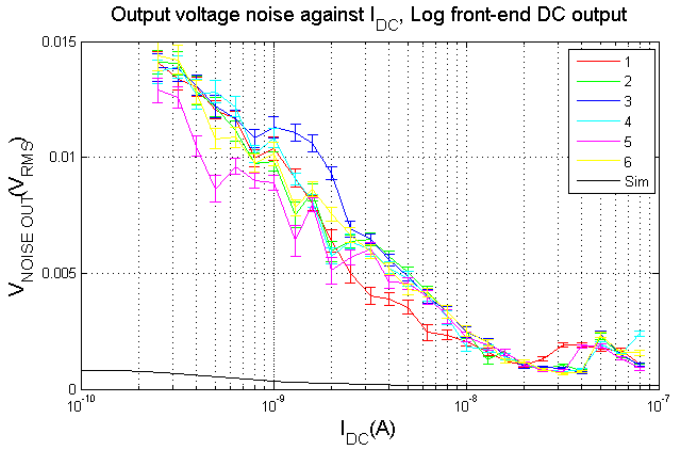

Figure 5.53: Measured and simulated output RMS voltage noise against DC photocurrent of the logarithmic pixel DC outputs on the BVIPS2 IC

Figure 5.54 shows RMS voltage noise measured at the output over 100 Hz-30 kHz for the logarithmic pixel AC outputs of the BVIPS2 IC, comparing measured results with post-layout simulations. The simulated noise response has a similar shape to that of the DC outputs, but with a magnitude 10 (at high DC photocurrents) to 50 (at low DC photocurrents) times greater than the simulated DC output noise level. This is due to noise from the HDA, and noise amplification by the HDA, increasing the noise present at the output. However, the additional gain here also increases the overall transimpedance (and hence signal), so the signal-noise ratio is improved and the input referred current noise falls, as shown in Figure 5.52.

The measured voltage noise at the AC output is similar to that expected from simulation, although the increase in measured noise at the lowest photocurrents (< 300 pA) is greater than shown by simulation. The noise levels are approximately equal across all alternative front-end circuits. The similar noise levels between simulation and measurement here reflect the similarities between the measured and simulated transimpedance of this output channel (shown in (Figure 5.44) - for the DC output, the measured output voltage noise (Figure 5.53) is higher than found from simulation, and again the transimpedance (Figure 5.45) is higher

than that expected.

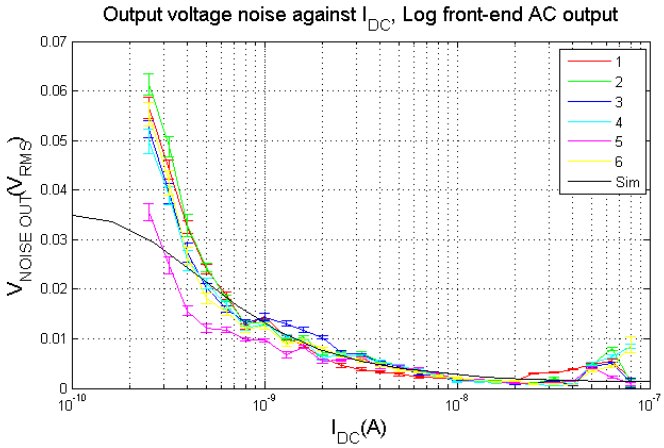

Figure 5.54: Measured and simulated output RMS voltage noise against DC photocurrent of the logarithmic pixel AC outputs on the BVIPS2 IC

5.8.4.3 Linear Pixel Noise Spectra

Figures 5.55 and 5.56 show the input referred current noise spectrum and output voltage noise spectrum of the opamp based pixels at 10 nA DC photocurrent, as found from pre-layout simulations. Again, the input referred current noise spectra has a similar shape to those seen previously, with increases at high and low bandwidth limits caused by the fall in gain at these points caused by high or low pass filters (where a constant output voltage noise combined with a reduced transimpedance causes an increased input-referred noise level). At low frequencies 1/f noise is also a significant noise source.

The noise density in the pass band is lower here than for the logarithmic pixels shown in Figure 5.49, at around 40 fA/$\sqrt{\text{Hz}}$ for all channels. This could be due to the improved power supply and common mode rejection ratio of the opamps used in the front-end here. These features would normally only be expected to reject external sources of noise rather than reduce the theoretical noise levels, but this does include rejecting noise caused by circuits on the IC other than the opamp detectors, which are included in the noise simulations. As the input referred noise takes into account the transimpedance of the circuit, the results shown suggest that the linear pixels may have a better SNR than the logarithmic pixels.

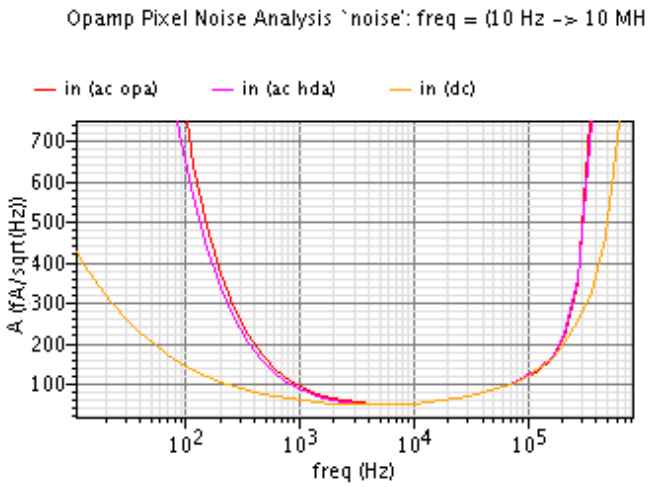

Figure 5.55: Input referred current noise response from pre-layout simulation of the opamp pixels on the BVIPS2 IC

The output voltage spectra here show higher noise on the opamp based AC channel, followed by the HDA AC channel then the DC channel. This is as seen for the logarithmic pixels, where reduced gain on a channel results in reduced noise.

A major difference between the spectra is the increase in output noise on the opamp based AC gain channel at lower frequencies. This could be due to the low-pass filter used here, which may result in increased 1/f noise, although the shape of the spectrum is not the 1/f shape which would be expected.

The output noise density here is higher than that seen for log pixels in Figure 5.50. This is due to the higher signals seen here, as shown by the higher gain in the pass band in Figure 5.46 compared to Figure 5.40. The input referred current noise spectra are a better means of comparing the noise of the circuits, as these compensate for the AC gain of the pixel.

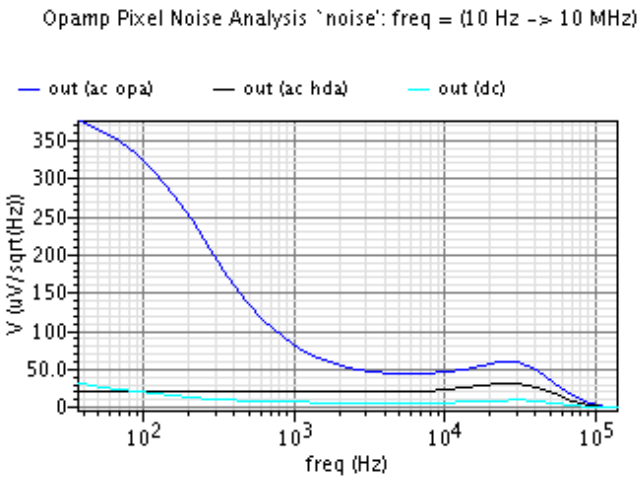

Figure 5.56: Output voltage noise response from pre-layout simulation of the opamp pixels on the BVIPS2 IC

5.8.4.4 Linear Pixel Noise against DC Photocurrent

Figure 5.57 shows input referred RMS current noise measured at the output over 100 Hz-30 kHz for all outputs of the linear pixel on the BVIPS2 IC, comparing measured results with post-layout simulations.

For the AC outputs (Figures 5.57a and 5.57b), the simulated noise response for both HDA and opamp based AC gain circuits is relatively flat (although there is a slight slope below 10 nA for the HDA gain channel output shown in Figure 5.57b), suggesting that the dominant noise source is thermal noise, possibly in the large feedback resistor which sets transimpedance. The noise level therefore does not vary widely with DC photocurrent, as the simulated transimpedance of the linear circuit is constant, unlike the logarithmic front-end where the transimpedance (and hence the noise level) is related to DC photocurrent.

The measured noise of the opamp based AC output (Figure 5.57a) initially matches the simulated behaviour, while the HDA AC output shows a lower noise level than that seen in simulation. Both rise slowly with increasing photocurrent, with the increase being approximately proportional to $\sqrt{I_{DC}}$ as expected for shot noise relative to I_{DC} . There may also be some increase due to laser noise, which would cause a linear increase in input referred noise with rising I_{DC}.

The noise levels seen here are significantly lower than that seen for the logarithmic pixels, with noise levels of approximately 3×10^{-11} A at 10 nA DC photocurrent here compared to

2.5×10^{-10} A for the logarithmic pixel AC output (see Figure 5.52).

The simulation results for the DC output (Figure 5.57c) show noise remaining constant up to a DC photocurrent of approximately 3 nA before dropping to a similar level to the AC outputs, and then rising sharply after 20 nA. This may be related to the non-saturating operating region of the circuit, although the photocurrent limits here do not correspond to values seen in the DC response of the opamp pixel DC output. The measured DC noise response is lower than the simulated value at low DC photocurrents, then rising steadily. The measured noise rises more rapidly once the DC photocurrent rises above 20 nA, although not as sharply as the simulated response rises.

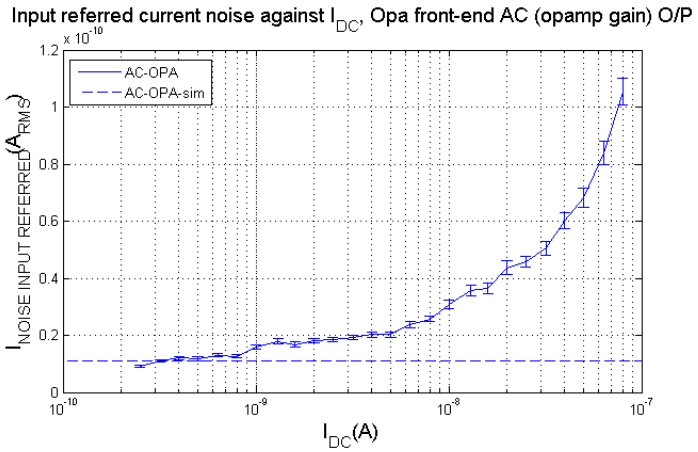

(a) AC (opamp gain) channel

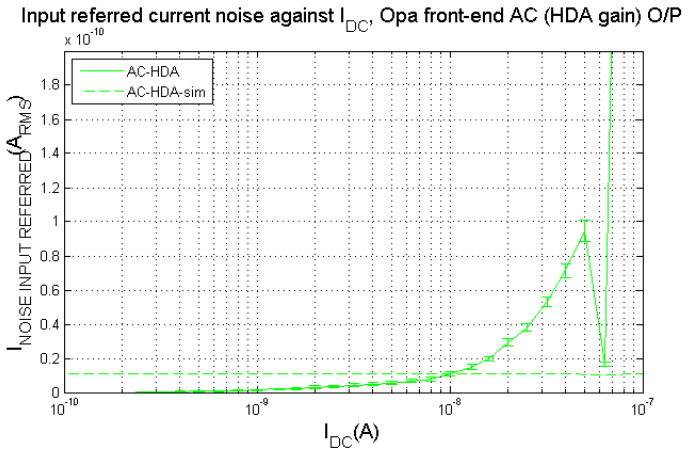

(b) AC (HDA gain) channel

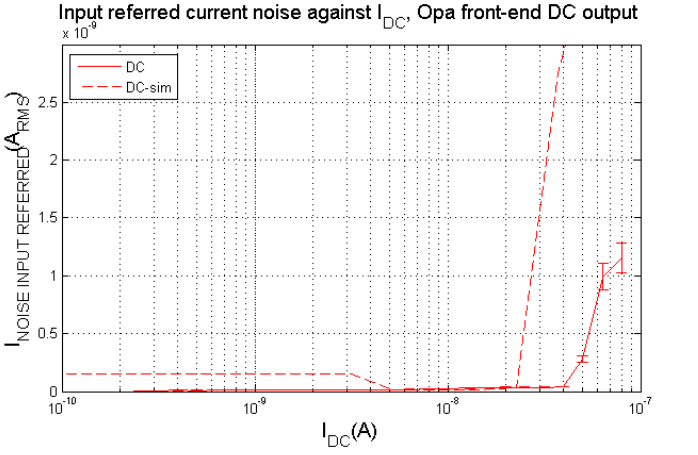

(c) DC channel

Figure 5.57: Measured and simulated input referred RMS current noise against DC photocurrent of the linear pixels on the BVIPS2 IC

Figure 5.58 shows RMS voltage noise measured at the output over 100-30 kHz for all out-

puts of the linear pixel on the BVIPS2 IC, comparing measured results with post-layout simulations.

The simulated responses of the AC outputs (Figures 5.58a and 5.58b) show similar flat noise responses to those seen in the input referred case (Figures 5.57a and 5.57b), although here the opamp based gain stage gives a higher noise level, of approximately 10 mV RMS compared to 5 mV RMS for the HDA version. The similar input referred noise levels with different output voltage noise shows the different gains of each circuit result in similar signal-noise ratios.

The measured noise results for both AC outputs show the same shape as the input referred results, although here the noise levels remain lower than expected from simulation results for a greater range of DC photocurrents than in the input referred case. This suggests that while the absolute noise levels are lower than expected from simulation, the transimpedance/gain is also lower, leading to similar or higher measured input referred noise compared to simulated results.

For the DC output (Figure 5.58c), the simulated results approximately mirror the input referred case shown in Figure 5.57c, with a constant and relatively low noise density at low or high DC photocurrents, and an increased noise density over 5-40 nA. This suggests that the simulated AC gain is at its highest in this region, where the output noise is not suppressed by saturation, but where the increase in gain causes a higher input referred noise level. The measured output voltage noise for the DC output shows an output noise level of approximately 0.4 mV in the saturated region at lower DC photocurrents, rising to a maximum of 1.4 V in the intended operating region (above 10 nA), before falling to 0.2 mV when saturated at high DC photocurrents.

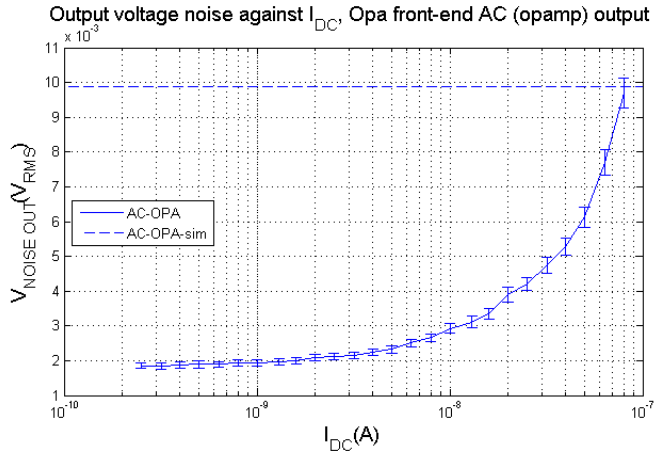

(a) AC (opamp gain) output

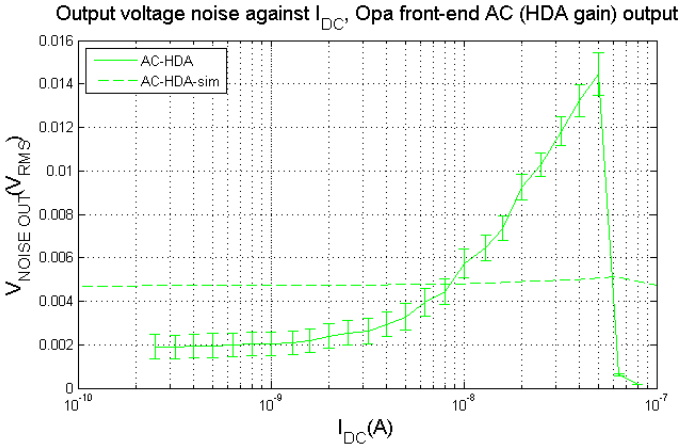

(b) AC (HDA gain) output

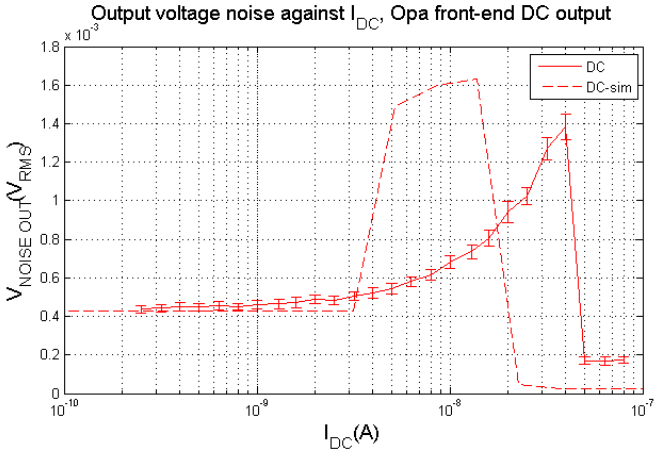

(c) DC output

Figure 5.58: Measured and simulated output RMS voltage noise against DC photocurrent of the linear pixels on the BVIPS2 IC

5.8.5 Transient Response

5.8.5.1 Logarithmic Pixel

To check the general operation of the pixel and multiplexers prior to fabrication, transient simulation was performed to observe the output of the IC at a range of different frequencies (1 kHz-10 kHz) and photocurrents (500 pA-50 nA), verifying that the output signal was as expected in all cases. This also verifies that the multiplexers in the IC design are working as expected, and that the buffers used to drive capacitive outputs are large enough to achieve the required switching times.

Figure 5.59 shows the output voltage from post-layout simulation over several address changes. The changes at 0.8, 1.6 and 2.4 ms are between two pixels in one half of the IC. The change at 3.1 ms is from one block of the IC to another (one block being half of the 64 pixels, connecting to a single ADC). The changes at 3.8 and 4.6 ms are changing between two pixels in the new block.

It can be seen that signals at the four different frequencies all give an undistorted output. The normalisation expected from the logarithmic pixels is also seen, as all pixels have a similar AC magnitude.

The changes in DC level can be seen on the AC and DC channels, but it can also be seen that the changes in DC level are clearer on the DC channel. Conversely, an AC signal can be seen on the DC channel, although with a reduced AC magnitude (as the DC channel does not have the existing AC removed, but does not include any additional AC amplification).

One possible concern shown in Figure 5.59 is the falling DC level of the AC channel output between 0.15 and 0.8 ms. However, as the multiplexer returns to this pixel from 1.6-2.4 ms, during which this artifact is not seen, this can be identified as a start up transient, due to the slow DC response of the HDA caused by its low cut-off frequency (200 Hz, or a time constant of 5 ms). It can also be seen that this effect is not seen on the DC output which does not include the HDA in the signal path. The HDA is included at pixel level, so settling is only required on IC power-up rather than after multiplexer switching, and therefore this is not a problem with the device.

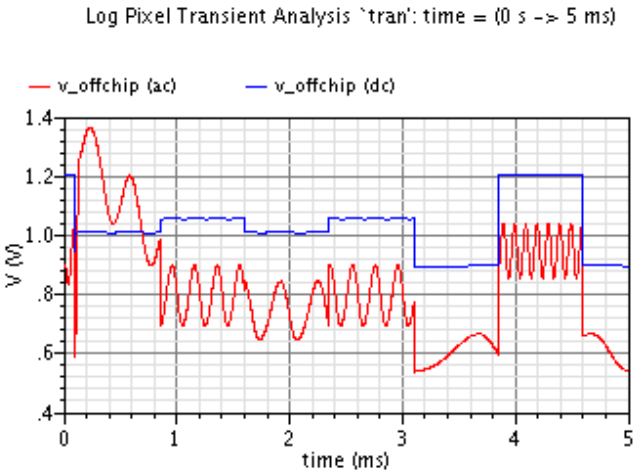

Figure 5.59: Transient response of the logarithmic pixels, switching between four different pixels with different photocurrents and frequencies (500 pA-50 nA, 1 kHz-10 kHz)

Figure 5.60 shows a close-up of the address change at 2.35 ms. It can be seen that the output values from either channel settle at new values within 150 ns. The requirement for device switching speed is set by the ADC sampling frequency of approximately 1.2 MHz required for 40 kHz sampling of all channels by two ADCs:

$$f_{Sample-ADC} = f_{sample-Pixel} \times N_{pixels/ADC} = 40 \times 10^3 \times 32 = 1.28\,\text{MHz} \tag{5.7}$$

This represents an ADC sampling period of:

$$T_{sample} = \frac{1}{f_{sample-ADC}} = 781\,\text{ns} \tag{5.8}$$

However, the sample period from equation 5.8 does not consider data setup and hold times, and therefore the specification requires the multiplexer outputs to have settled by the halfway point between switching, i.e. within 390 ns of switching. The 150 ns settling time shown in Figure 5.60 suggests this requirement is met with margin for increases caused by parasitics larger than expected or not included in the simulations.

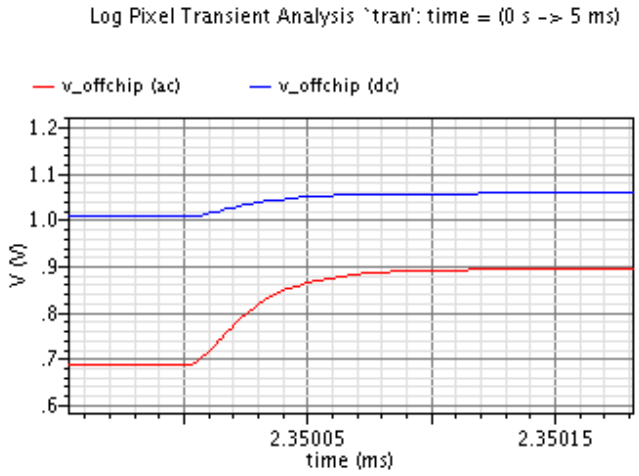

Figure 5.60: Transient response of the logarithmic pixels, showing switching between pixels at 2.35 ms

5.8.5.2 Linear Pixel

Figure 5.61 shows the output voltage from the linear pixels during the same input conditions as used above for the logarithmic pixels, from post-layout simulation. It can be seen that different input signals in most cases give an undistorted output. The output signal between 3.1 and 3.8 ms on the DC channel is saturating at near V_{DD}, as the photocurrent here is too large for the region of operation set. This affects the AC channels to the lesser extent, where some distortion can be seen.

On the opamp based gain channel, the DC level in all cases remains constant, showing that the high-pass filter used here is correctly blocking the DC input voltage. On the other channels the change in DC voltage between different pixels is as expected. Higher DC photocurrent pixels (with lower frequency) have a higher DC output voltage.

As for the logarithmic circuits, the opamp DC channel output in Figure 5.61 shows a significant AC component. This is because the DC channel has reduced AC gain compared to the AC channels rather than having the AC signal removed. Also, as the simulated transimpedance of the linear front-end does not fall with rising DC photocurrent, higher photocurrents in particular are likely to result in significant AC signal levels on the DC channel. The DC channel is intended to have the most useful range of DC output levels (i.e. operating over a wider range of DC photocurrents), particularly compared to the opamp based AC channel where the original DC level is fully removed, but obtaining the DC level itself relies on

averaging performed, either off-chip or by on-chip digital processing.

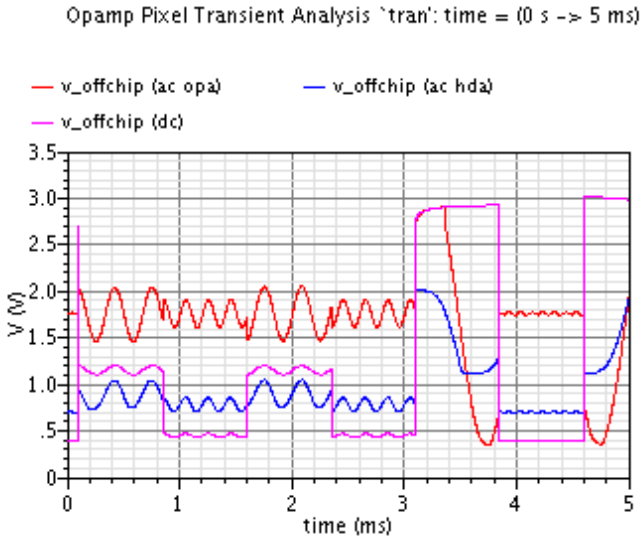

Figure 5.61: Transient response of the opamp pixels, switching between four different pixels with different photocurrents and frequencies (500 pA-50 nA, 1 kHz-10 kHz)

Figure 5.62 shows a close-up of the address change at 2.35 ms. Again it can be seen that the output values from all three channels settle at new values within 150 ns, meeting the requirement described in Section 5.8.5.1.

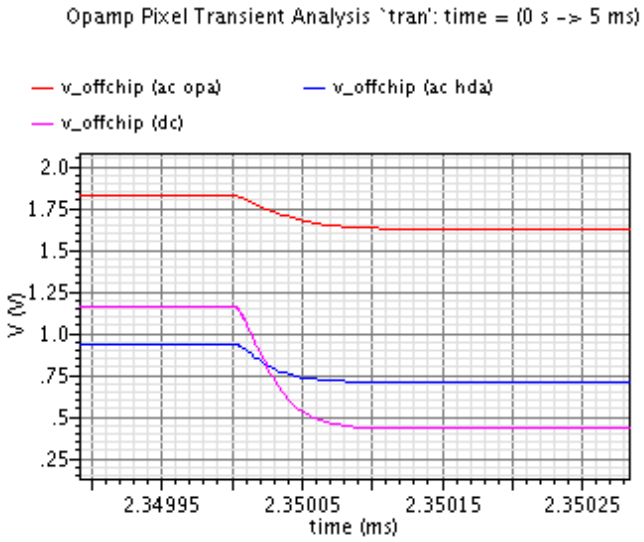

Figure 5.62: Transient response of the opamp pixels, showing switching between pixels at 2.35 ms

5.9 Flow Images

To determine whether the performance of the IC is sufficient for successful blood flow measuring, Doppler imaging was performed as for the previous IC. As this thesis concerns the analogue parts of the IC, the flow was calculated by off-chip sampling of the analogue output signal using the same data acquisition system used for characterisation (shown in Figure 5.34), with processing performed on a PC using Matlab.

5.9.1 Doppler Imaging Setup

The imaging setup was unchanged from that used for BVIPS1, shown in Chapter 4, Figure 4.31, using a 5 mW, 633 nm HeNe laser to produce a 40 mm line on the target. The IC was used to measure blood flow in the fingers during occlusion and release of flow, with the line being projected across three fingers. The system uses fixed line illumination, such that the figures are built up by repeated sampling of the linear array, giving a map of pixel response over time. Pixels were sampled sequentially, such that a line is built up from 1024 samples of pixel 0, 1024 samples of pixel 1 etc. Pixels were sampled at 40 kHz, giving a time per pixel of ~25 ms, and a time per line of ~2 s (i.e. 1.6 s of sampling, with some overhead for saving data and multiplexer switching). The laser line fell across three fingers, centred on the middle finger. The width of the middle finger was approximately 50% of the total line length. This arrangement should mean that pixels 0-15 correspond to the top finger, 16-47 correspond to the middle finger, and pixels 48-63 correspond to the bottom finger. Some variation of this positioning may have occurred due to movement during or between tests, although this was minimised where possible by the use of fixed supports.

5.9.2 Logarithmic Pixel Section

Figure 5.63 shows the flow across the array and over time for each AC output of the logarithmic section of the BVIPS2 IC. Dark red colours represent a pixel/time with high flow, while dark blue shows low flow. Figure 5.64 shows line plots of the flux measurement from a single pixel and across the array for the same outputs.

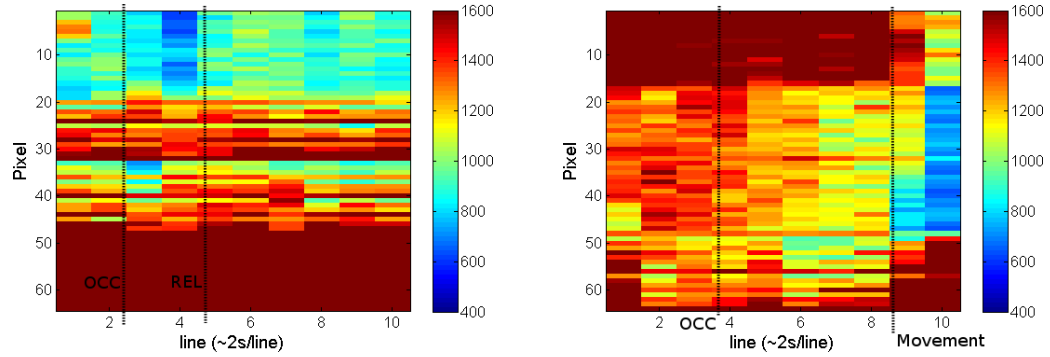

(a) Log AC O/P 1

(b) Log AC O/P 2

(c) Log AC O/P 3

(d) Log AC O/P 4

(e) Log AC O/P 5

(f) Log AC O/P 6

Figure 5.63: Flow across array and over time during occlusion and release of blood flow in fingers, Logarithmic pixel AC outputs 1 of BVIPS2 IC

In all cases a change in flux level over time can be seen at some or all pixels. The most clear example of this is when using front-end #4, where the flux at most pixels varies from ~1500 for high flow to ~1300 at low flow, a fall of 13%. Front-end #3 also shows a significant change between high and low flow, with the change being from ~1100 to ~800, approximately a 30% drop.

The pattern of these flow measurements over time, however, is different to that expected - occlusion then release should give a pattern of normal flow, then a drop in flow during occlusion. On release of occlusion, post-occlusive reactive hyperaemia can be seen - there is a brief spike in blood flow is seen as blood returns to the tissue in the occluded region, after which flow returns to the original level [Shepherd and Oberg, 1990]. Here flow starts at a low level, then increases from lines 3-5. At line 6 the flow reduces, before returning to a higher level at line 8. This is consistent with blood flow being initially occluded, released at row 3, re-occluded at row 6 and released again at line 8. As the output from all alternative circuits was recorded immediately after the measurement from the previous circuit, it is possible that the flow was initially still occluded from a previous measurement - note that the flow at the end of the circuit #2 measurement is similar to that at the start of the circuit #3 plot.

All images show a consistent spatial pattern, with high flow readings from around pixel 50-63. Front-end #2 does not show this pattern throughout, with lower flux readings at these pixels in rows 1-8. However, the region of flux values < 1500 as seen from the other circuits does appear in rows 9 and 10. This high flux region could be caused by a reduction in DC level in this region, caused by this point corresponding to the gap between fingers - the low light level causes high transimpedance, and therefore high flux readings due to noise. The positioning of the laser line on the fingers should result in similar flux patterns for pixels 0-15 as for 48-63, and also in similar flux readings at the top and bottom of the plot as those in the centre, as these points (i.e. pixels 0,32 and 63) correspond to the centre of the fingers. However, if the ends of the laser line extend slightly further than the ends of the detector array then the top and bottom pixels will correspond to gaps between fingers, or at least to parts of the fingers where the skin is not perpendicular to the laser illumination, reducing the effective illumination intensity and hence reflected light level and DC photocurrent. Similarly, if the power in the laser line at the ends is less than that in the centre, then the top and bottom regions of the plot are likely to have different flux readings to that from central pixels due to lower DC photocurrent in these pixels. This latter case explains the changes in flow from pixel 0-15 and 48-63 in Figure 5.63.

Comparing the flow data from these plots with the DC light level of the same pixels would allow 'background' pixels to be marked as such, although this cannot be done here as the DC channel was sampled after the full plots shown in Figure 5.63 were captured, rather than being interleaved with the AC sampling.

It should also be noted that the plots in Figure 5.63 do not show the significant horizontal lines seen in similar plots from the BVIPS1 IC (shown in Chapter 4, Figure 4.33c). This shows that the BVIPS2 IC has lower fixed pattern noise than BVIPS1 IC before the front-end selection technique is applied, suggesting that the other changes made to the pixel design (mainly the alternative HDA design) have been effective.

While the spatial/temporal flux plots in Figure 5.63 show successful detection of temporal and spatial flux changes, the changes from high to low flow are relatively small relative to the average readings, and the noise level at high and low flux levels is high. The flux line plots shown in Figure 5.64 in particular show that the changes from point to point when a constant flux would have been expected are approximately half the magnitude from occluded to free flow, giving an SNR of approximately 2.

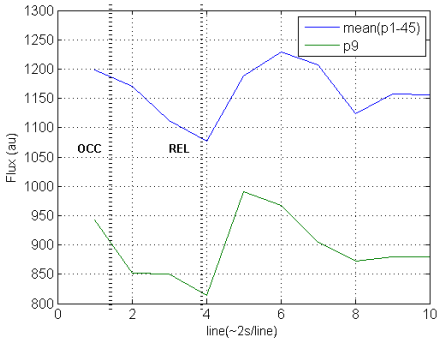

(a) Log AC O/P 1

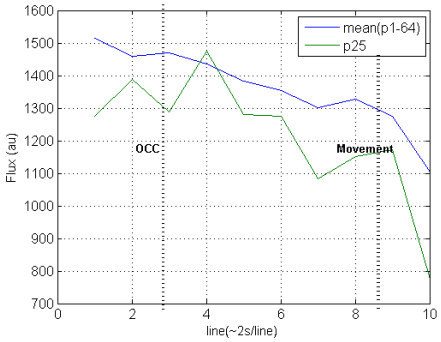

(b) Log AC O/P 2

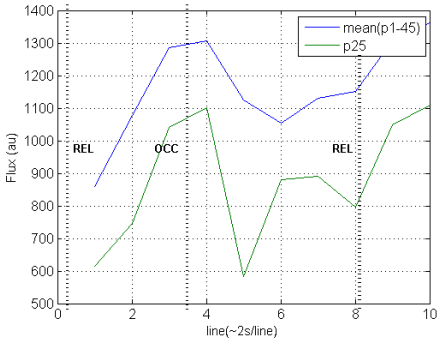

(c) Log AC O/P 3

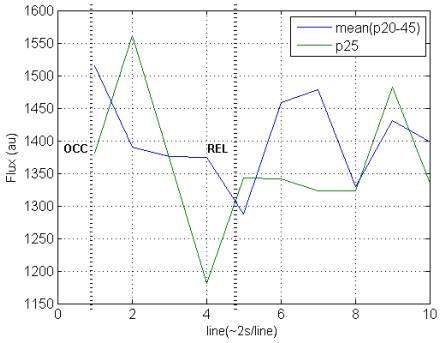

(d) Log AC O/P 4

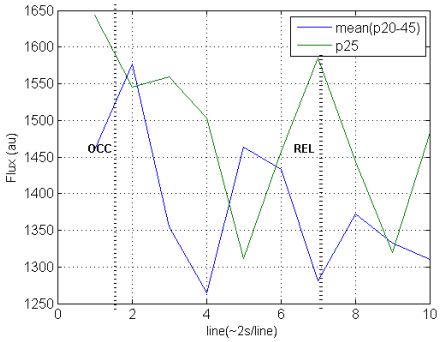

(e) Log AC O/P 5

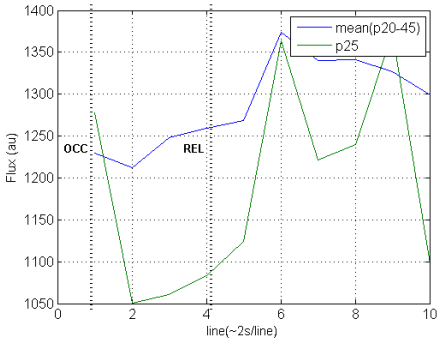

(f) Log AC O/P 6

Figure 5.64: Plots of flux from a single pixel and average flux from all pixels over time, Logarithmic pixel AC outputs of BVIPS2 IC

5.9.3 Opamp Pixel Section

Figure 5.65 shows flux measurements across the array and over time for both AC outputs of the opamp section of the BVIPS2 IC. Dark red colours represent a pixel/time with high

flow, while Dark blue shows low flow. Figure 5.66 shows corresponding line plots.

Both channels show the expected flux pattern, with initial high flow at most pixels, and a reduced flux region from lines 3-7 for the opamp based AC channel and from lines 4-7 for the HDA based AC channel. For the opamp gain channel the drop in flux with occlusion is approximately 180-150, a 17% change, while for the HDA gain channel the drop is from ~280-200 (29%) for pixels 16- 44, and ~200-170 (15%) for pixels 0-15. The drop in flux readings at pixel 15 may be because this part of the image is from another part of the target (i.e. a different finger). The abrupt nature of the transition indicates this may be an electrical effect, although if the problem was related to multiplexers it would have been expected to affect some other outputs (as the signal path for the HDA gain channel uses the same multiplexers as the opamp channel, being selected by sets of transmission gate switches within each pixel, controlled by a global select pin).

As in the logarithmic flux plots in Figure 5.63, a spatial pattern can be seen with a transition around pixel 48. However, for the opamp pixels the flux for pixels from this point to pixel 63 has a low value (< 50). This is consistent with this region of the plots corresponding to a gap between fingers, as for the opamp pixels the gain does not increase sharply at low DC photocurrents, and therefore the low DC light level does not result in high flux readings.

The HDA gain channel shows some horizontal lines similar to those seen in BVIPS1 (shown in Chapter 4, Figure 4.33c), suggesting that while the HDA modifications have reduced the pixel-pixel variation of this circuit, it is still susceptible to process variation, leading to FPN or low-sensitivity pixels.

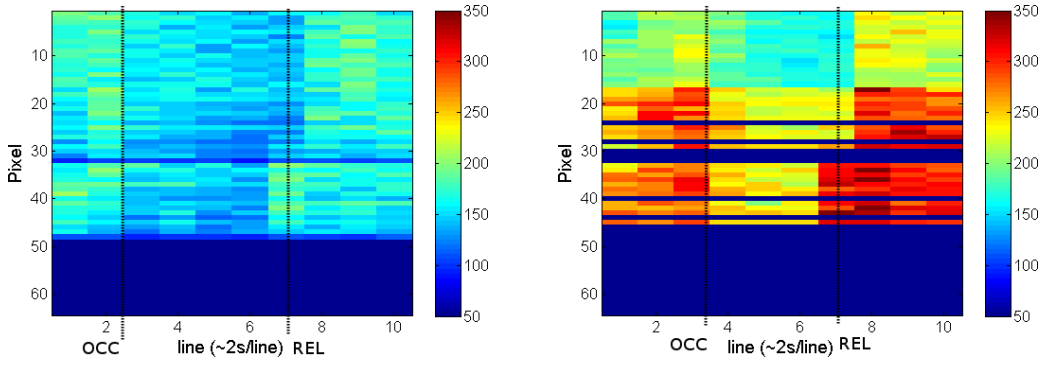

(a) Opamp AC (Opamp gain) O/P

(b) Opamp AC (HDA gain) O/P

Figure 5.65: Flow across array and over time during occlusion and release of blood flow in fingers, Opamp pixel AC outputs of BVIPS2 IC

The flux line plots for the opamp pixels shown in Figure 5.66 show that the opamp pixels have improved SNR compared to the log pixels, with the single pixel results being a closer match to the mean results (except for a constant offset on the HDA gain channel). The variation in flux readings while flow is expected to be constant is ~10 for both outputs, compared to a total range between high and low flux of ~40 for both outputs, giving an SNR of 4. As this represents the difference between biological high and low flow, rather than the full range seen from the static reflector to motility standard seen for BVIPS1, the true SNR is potentially slightly higher than this, although these targets were not available for the testing of BVIPS2.

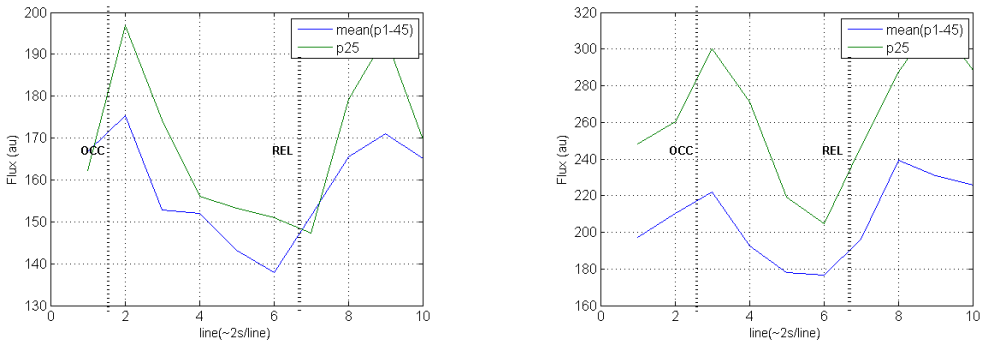

(a) Opamp AC (Opamp gain) O/P

(b) Opamp AC (HDA gain) O/P

Figure 5.66: Plots of flux from a single pixel and average flux from all pixels over time, Opamp pixel AC outputs of BVIPS2 IC

5.10 Summary

The changes made between the BVIPS1 and BVIPS2 ICs have been shown. This includes changes to solve problems found on the first prototype, such as poor signal quality due to problems with the HDA AC amplifier circuits, and high fixed pattern noise. The main change to address the HDA faults seen on BVIPS1 was a switch to an alternative HDA design that has been previously fabricated on a different project and shown successful operation. While simulations suggest the new HDA design has lower gain and a higher low-frequency cut-off,

the discrepancies between simulated and measured performance seen on BVIPS1 suggest the simulation results are not reliable. Changes made to the layout of the HDA should reduce the susceptibility of the HDA to external interference such as stray light.

A further significant change made here, partly enabled by the lower physical size of the new HDA design, is the use of multiple front-end circuits on each pixel. These can be individually selected by a digital control section, meaning that well matched detectors can be used. This should significantly reduce fixed pattern noise, and the need for off-chip calibration, although the individual front-end selection feature was not tested here.

The linear pixels tested on BVIPS1 have also been modified for BVIPS2, with a more compact opamp being designed to achieve a more useful pixel size, and a modified AC channel provided to give higher AC gain.

The greater degree of integration of the second prototype IC has also been shown, and the advantages of this technique both in terms of circuit performance and system flexibility have been discussed.

Pre- and post-layout simulations were performed to verify the IC performance before fabrication, and the results of characterisation tests performed on the IC were compared with some of these results to determine the actual performance of the IC. Some aspects of measured device performance were not as expected, such as the transimpedance of the opamp pixels not being constant over the full range of DC photocurrents and increased noise levels. As limited time was available for testing of the BVIPS2 IC, further testing to identify the causes of these discrepancies was not possible. The large number of external bias adjustments and output configurations of BVIPS2 means that behavour closer to the simulated response may be achieved by adjustment of the external analogue bias circuits. Some effects will also be the same as for BVIPS1, such as power supply noise and interference from stray light. However, most aspects of the IC were shown to operate as intended and the discrepancies seen here are reduced compared to those seen for BVIPS1, suggesting that the design changes made have been at least partially succesful in improving performance.

The additional features of the IC such as front-end selection, dual detector arrays and more complex multiplexing have not been shown to cause adverse behaviour.

Finally, spatial/temporal flux measurements of blood flow in fingers during occlusion and release of flow were shown in Figure 5.63 (for the logarithmic pixels) and Figure 5.63 (for

the opamp pixels). These plots show that the IC is capable of measuring changes in blood flow using the linear and logarithmic pixels. A clear spatial pattern could be seen in the flow images, suggesting that spatial resolution has been improved, although this has not been conclusively shown given that location of veins, the intended purpose for this system, has not been demonstrated.

The fixed pattern noise in the plots produced was reduced compared to similar plots from the BVIPS1 IC, as shown by the reduction in horizontal lines in Figures 5.63 and 5.65.

However, while these features show improved performance, the noise seen on the processed flux signals is high, giving an SNR of ~2 for the logarithmic pixels and ~4 for the opamp pixels, which does not show a significant improvement over BVIPS1. This is partly due to the flux targets used, as the static target and motility targets used to give constant high/low flow readings were not available when this IC was tested so some actual flux variation may have added to the true noise.

Chapter 6

Conclusions and Summary

6.1 Introduction

This chapter will summarise the work presented in this thesis, and will discuss the main points raised with regard to the aims of this project.

6.2 Thesis Summary

Chapter 1 introduced laser Doppler blood flow imaging as a technique, including single point scanning systems and imaging systems. The advantages of CMOS sensors for these systems were considered, in particular the possible benefits of a line scanning system using an integrated sensor with on-chip processing. The physics behind the Doppler effect as used for blood flow measurement were described, along with the fundamentals of systems used to detect these signals. The magnitude and frequency spectrum of the photocurrent to be expected given the sensor size required for this application was also estimated from measurements of existing systems.

Chapter 2 explored the development of existing LDBF systems, including currently available systems and the advances being made using various LDBF techniques. Alternatives to LDBF that may be better suited to a clinical instrument for vein location are considered, such as laser speckle contrast imaging, back-illumination and infra-red image processing. Issues raised by the various devices that need to be addressed before a widely used clinical

instrument can be produced were discussed. The need for a line scanning device, which can offer high acquisition speed, good accuracy and safe laser power was demonstrated.

Chapter 3 investigated various circuits with regards to their suitability for use in a CMOS linear array sensor. The first part of the chapter focused on options for a front-end circuit, particularly logarithmic pixels, which were shown to offer compact size, sufficient gain and bandwidth, as well as inherent normalisation required for LDBF systems. A buffered logarithmic pixel was selected for use on the first prototype IC. Active pixel sensors, as commonly used in CMOS cameras, were dismissed as the light levels and diode capacitances found here result in very small signals. A linear pixel design based on opamps suitable for on-chip implementation was also shown, and found to be suitable for inclusion on an IC despite having drawbacks of large size and output saturation.

The chapter then discussed processing circuits, looking at two alternative designs for a high-pass filter/AC amplifier. A design of hysteretic differentiator AC amplifier was selected for use on the IC - using an 'inverted-inverter' feedback network to give a very low cut-off frequency. For a low-pass/anti-aliasing filter, a g_mC low-pass filter design was simulated to ensure its suitability for this application.

Chapter 4 described the system and pixel structure of the first prototype IC, BVIPS1. Characterisation of the IC was performed using a modulated laser, and testing results were compared with simulation results. Significant discrepancies were shown between measured and simulated behaviour, including higher noise than that predicted in all pixels, and logarithmic pixels in particular, making flow changes harder to detect due to the lower SNR. Efforts made to reduce the noise levels were shown, although these had limited effect. Other artefacts shown during testing suggested a more complex fault with the logarithmic pixels, and work was done to identify and characterise the issue, although the root cause of the fault was not definitively found.

Despite these problems, both the logarithmic pixels and linear pixels were shown to be sufficiently sensitive to detect blood flow changes, although the SNR of processed flux values was low, at around 4-8.5 for the logarithmic pixels compared to at least 10-17 for commercial systems. Spatial resolution of the 64x1 logarithmic array was also poor, with no success in vein location. The opamp pixels appeared to show better SNR for flux outputs, but with unstable readings (showing a gradual drift of flux), no calculation of SNR was made.

Chapter 5 described the second prototype IC, BVIPS2, including changes made to the opamp

pixel design to reduce its physical size, as well as two approaches to avoid the saturation problem associated with pixels with no inherent normalisation. For the logarithmic pixel, the compact size is used to allow multiple individually selectable alternative front-end circuits, allowing the best matching front-ends to be used to reduce fixed pattern noise.

This chapter also briefly described the integrated system built on this IC, using on-chip analogue to digital converters and digital processing. While the digital parts themselves are not the work of the author, the complete system that could be constructed on an IC of modest size represents one of the advantages of the linear detector array technique.

Testing of the BVIPS2 IC compared actual IC behaviour with that predicted from simulation, and found that while discrepancies existed between the measured and simulated responses, these discrepancies were reduced compared to those seen for BVIPS1.

While project delays meant time available for testing of the BVIPS2 IC was limited, initial use of the IC to measure flow changes in tissue were shown, which demonstrated that both the logarithmic and linear detectors on the BVIPS2 IC were capable of detecting flow changes. Furthermore, both arrays demonstrated reduced fixed pattern noise (fewer 'dead' pixels) and improved (though still crude) detection of spatial features compared to the BVIPS1 IC.

Chapter 6 presents conclusions from the project, summarising the work undertaken and it's potential for use in clinical systems. Further work using the ICs and circuits developed is suggested, as well as developments to the systems designed.

6.3 Further Discussion

The prototype ICs have successfully demonstrated the potential of a linear array CMOS detector for blood flow imaging with a line scanning system. Several possible front-end circuits have been shown to successfully measure changes in blood flow, and it has been shown that it is possible to implement several arrays of front-end circuits that share a single photodiode. This approach could potentially be expanded to more than two detector arrays, such that the array could be selected for various expected applications, light levels, or post-processing methods.

6.3.1 Reduction in Size of 1D over 2D Integrated Sensors

The reduction in physical size of a CMOS sensor combining integrated processing with a linear detector array, as opposed to a 2D sensor array, has been demonstrated. The BVIPS2 IC, with two 64x1 sensor arrays, analogue-digital conversion and integrated processing is approximately $4 \times 4\,\text{mm}$. A 2D LDBF imaging IC built on a parallel project used a 64 x 64 array of similar pixels to the logarithmic design used in the BVIPS ICs. This IC had a size of $6 \times 6\,\text{mm}$, offering broadly similar imaging capabilities, albeit with higher frame rates (DOP3 IC, fabricated as part of DTI NEAT project [He et al., 2009, Hoang, 2009]). Figure 6.1 shows the two ICs side-by-side, with the smaller BVIPS2 IC on the left, and the larger 2D array on the left. The area ratio of these ICs is 4:9, so the linear array IC will be less than half the cost of the 2D sensor.

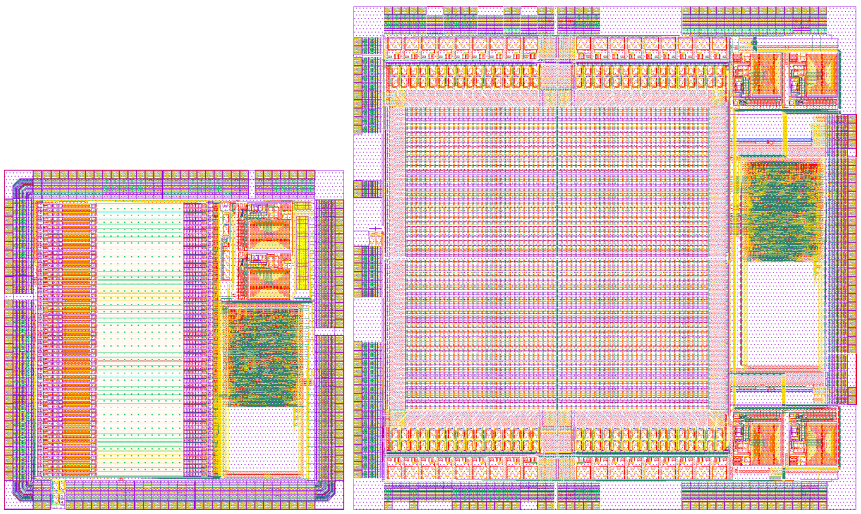

Figure 6.1: Side-by-side comparison of 64x1 (BVIPS2, left) and 64x64 (DOP3, right) LDBF sensors using similar pixels and processing techniques

Furthermore, the smaller linear-array IC has a variety of detector circuits and configurations, with the choice of logarithmic or linear detectors, six alternative front-end/HDA circuits on the logarithmic detectors and two alternative AC channels in the linear pixels.

For this application, the multiple array/multiple channel design was used to test an alternative front-end array design, to improve fault tolerance, and to achieve reduced FPN without off-chip calibration. However, a similar principle could be applied in different ways - for example, the alternative front-end circuits could be designed for different light levels, or a different frequency range depending on the requirements of the application.

It should also be noted that the size of the photodiodes used by the linear array ICs (consuming approximately a quarter of the width of the BVIPS2 IC, as seen in Figure 6.1) could be reduced for some applications. While the large detector width is required here to allow for the beam wandering during mechanical scanning, if the system could be designed to allow tighter tolerances, the individual sensors could be reduced in size. This could allow greater area for processing circuits or other integrated system components, or could allow the alternative front-end circuits to use separate photodiodes. This would allow both arrays to be used simultaneously (effectively a 2×64 2D array), or different photodiode types could be used, for example using shallow well photodiodes to have increased sensitivity to light at shorter wavelengths [Moini, 2000].

6.3.2 Measurement of Flow Changes

Both sensor types developed for this IC have successfully demonstrated detection of flux changes, although the increased noise and reduced gain of the sensors compared to the designed values reduces the quality of the processed flux signal.

The logarithmic pixel array on the BVIPS1 IC has been used to detect change in blood flow across 90% of the pixels, using line illumination with a 5 mW visible red laser, as shown in Figure 6.2. It should be possible to improve on the performance of BVIPS1 by the use of either a higher power laser, or the use of an infra-red laser for increased skin penetration depth. It may be that veins could not be detected as insufficient light was reaching the veins, rather than being limited by the accuracy of the sensor itself.

Figure 6.2 shows flux measured by all pixels of the BVIPS1 logarithmic pixel array over time. The y-axis shows time, while the x-axis shows flux at different points along the illuminated line. The vertical bands show periods of occlusion and release, while the pattern of horizontal stripes shows fixed pattern noise, with approximately 10% of pixels showing no change in flux over time.

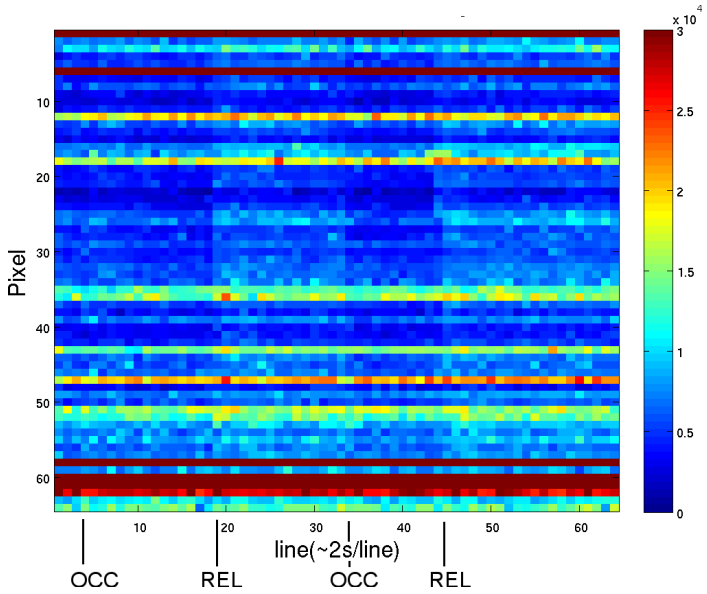

Figure 6.2: Detection of occlusion and release of blood flow in a hand using logarithmic pixels on BVIPS1 IC

Figure 6.3 shows line plots of mean and single pixel flux for various targets. This shows the low SNR achieved for flux readings, with variation of flux readings during no actual change being approximately 1000 for all targets, compared to a minimum flux value of 4000 and a maximum of approximately 8500, giving an SNR value of between 4 and 8.5 (depending on the actual flow level). This demonstrates successful detection of flux changes, but with poor flux resolution and poor repeatability.

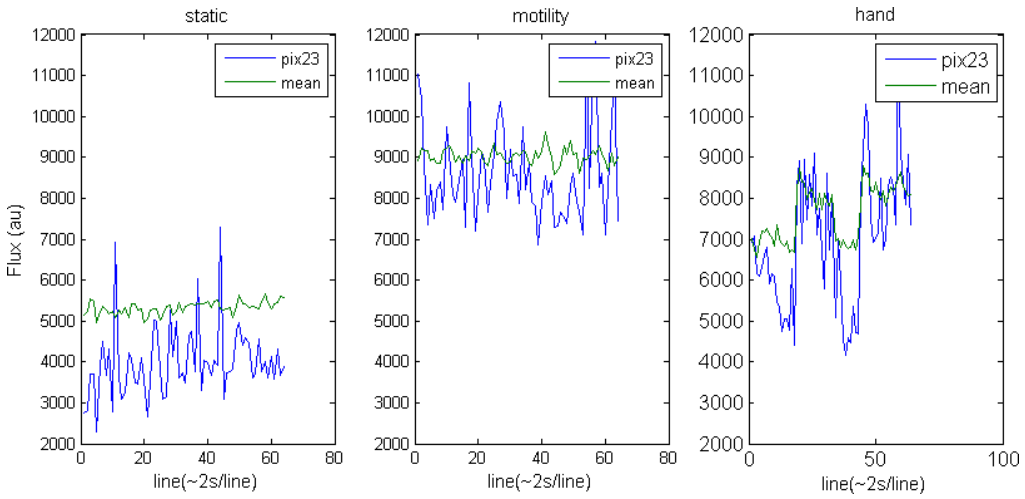

Figure 6.3: Plots of flux from a single pixel of BVIPS1 IC logarithmic array, for static (no flow) target, motility (high flow) target, and a hand during occlusion and release of blood flow

The linear pixels implemented on the BVIPS2 IC attempted to address the problems found on BVIPS1, resulting in reduced fixed pattern noise and fewer low-sensitivity pixels, as demonstrated by the lack of horizontal lines caused by 'dead' pixels in Figure 6.4 and Figure 6.7). These figures also show a spatial pattern caused by the edges of the fingers used as a target, although the spatial pattern is not as clear as would be expected from a commercial system.

Figure 6.4 shows flux measured by all pixels of the BVIPS2 logarithmic pixel array over time, using the AC channel of circuit #3. The y-axis shows time, while the x-axis shows flux at different points along the illuminated line. The vertical bands show periods of occlusion and release (with flow being initially occluded). The dark red area at the bottom of the plot shows a gap between fingers, where the low light level causes high noise and hence high and constant flux readings. The change in flux readings (darker colours) at the top of the plot show the same effect, although the drop is smaller than for the lower pixels.

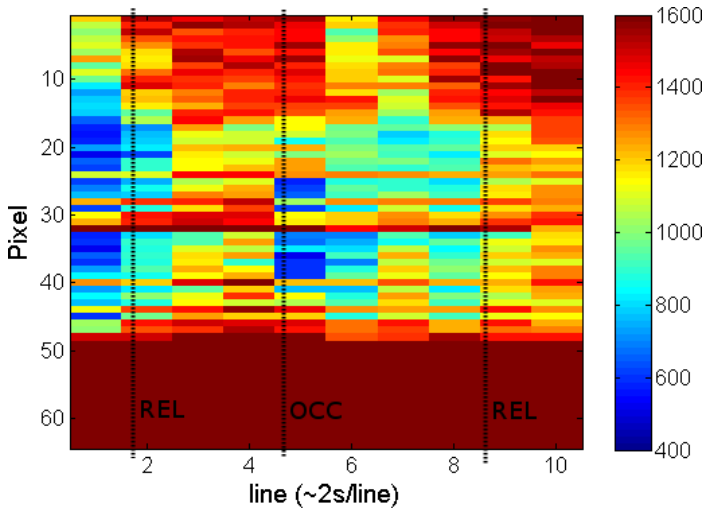

Figure 6.4: Detection of occlusion and release of blood flow in a hand using logarithmic pixels (front-end #3) on BVIPS2 IC

Figure 6.5 shows line plots of mean and single pixel flux on the same log pixel on BVIPS2, showing a change on one pixel from ~1100 to ~700, approximately a 36% drop, but the variation between points where no flux change is expected gives an SNR of around 2.

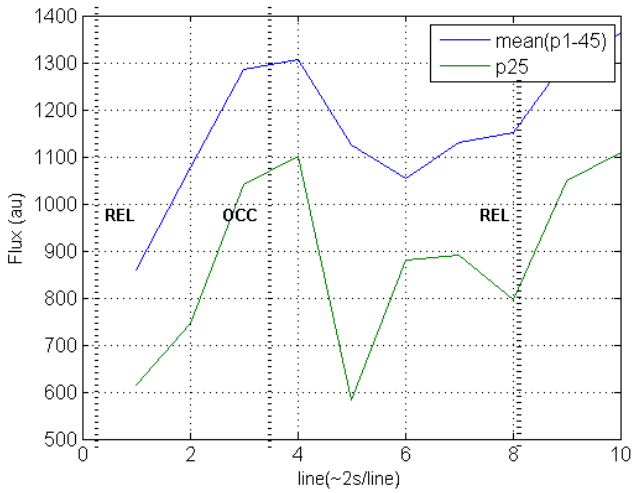

Figure 6.5: Plots of flux from a single pixel and average flux from logarithmic pixel AC outputs (front-end #3) of BVIPS2 IC

The opamp pixels used on both BVIPS1 and 2 showed potentially more promising results, being capable of detecting changes in blood flow during occlusion and release, as shown in Figures 6.6 and 6.7, and appearing to show higher SNRs for processed outputs.

The main issue with the opamp pixels is the physical size of the circuits used, but their inclusion as a 64 x 1 array on BVIPS2 shows that this problem does not prevent the use of this type of pixel on a CMOS integrated LDBF sensor.

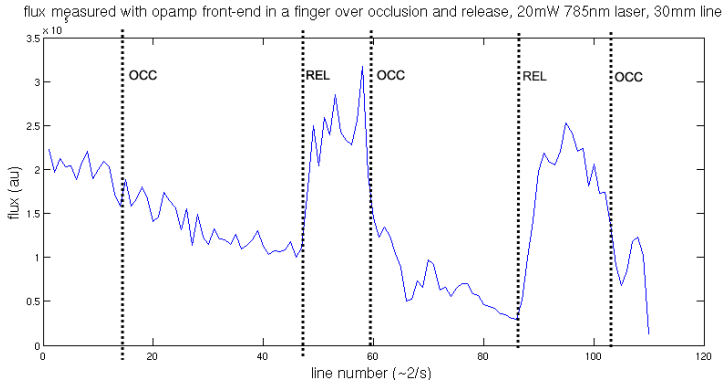

Figure 6.6: Flux signals over occlusion and release measured using an opamp pixel on the BVIPS1 IC

The array flow plot produced using BVIPS2 (shown in Figure 6.7) shows that this version of the opamp pixel, using a compact opamp design produced for use on the front-end, high-pass filter and opamp gain stage, is also also capable of detecting blood flow changes. It also demonstrates that the fixed pattern noise of this front-end design is comparable to that of the BVIPS2 logarithmic pixels, and that the array can detect spatial changes in blood flow as well as those over time.

Figure 6.7 shows flux measured by all pixels of the BVIPS2 opamp pixel array over time, using the opamp based AC channel. The y-axis shows time, while the x-axis shows flux at different points along the illuminated line. The vertical bands show periods of occlusion and release. The dark blue area at the bottom of the plot shows a gap between fingers, where the low light level and the linear response of the front-end result in low flux levels. The blue line along the centre shows some FPN, although as this is pixel 32 this may be caused by slow switching of the final multiplexer stage.

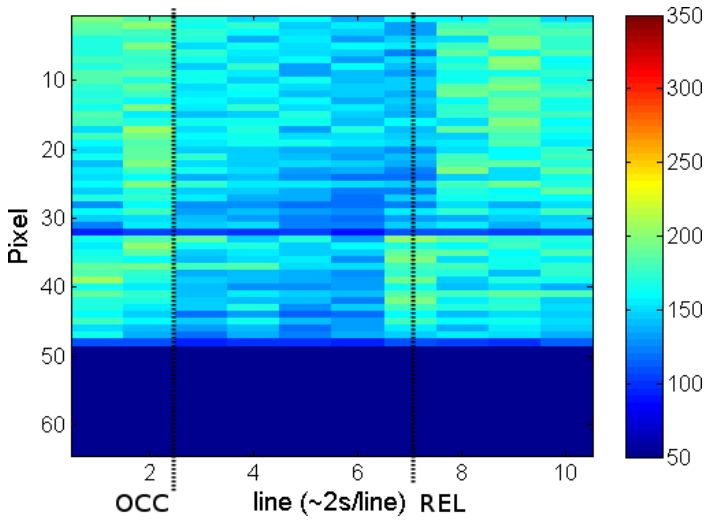

Figure 6.7: Detection of occlusion and release of blood flow in a hand using opamp pixels (opamp AC gain channel) on BVIPS2 IC

Figure 6.5 shows line plots of mean and single pixel flux on the same opamp pixel on BVIPS2, showing improved SNR compared to the log pixels. The variation in flux readings while flow is expected to be constant is ~10 , compared to a total range between high and low flux of ~40, giving an SNR of 4. This represents the difference between biological high and low flow, rather than the full range seen from the static reflector to motility standard seen for BVIPS1, so the true SNR is potentially slightly higher than this. The artificial targets used to determine a full range were not available during testing of BVIPS2.

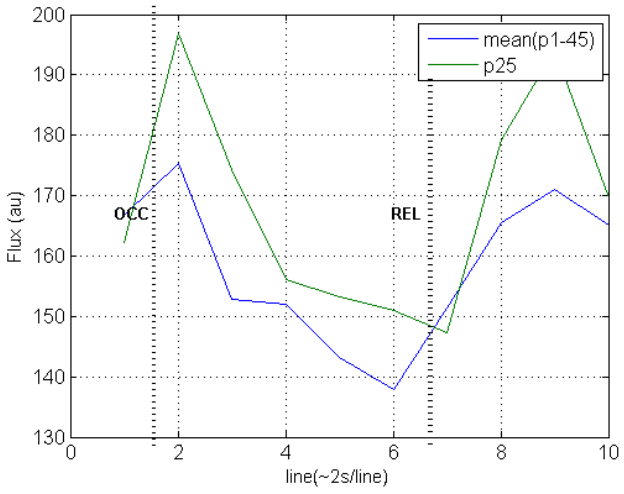

Figure 6.8: Plots of flux from a single pixel and average flux from all pixels over time, Opamp pixel (opamp gain) AC outputs of BVIPS2 IC

6.3.3 Discrepancies between Simulated and Measured Performance

A problem that affected all results presented in this thesis is that behaviour predicted by simulation often differed significantly from that found following fabrication and testing of an IC. This varied from minor variations in a response, such as a DC offset (shown in Figure 6.9 for the BVIPS1 log front-end DC response) which are to be expected to some extent for all analogue ASICs, to major changes in slope and magnitude of responses, as shown in Figure 6.9 for the output noise of the BVIPS1 log front-end.

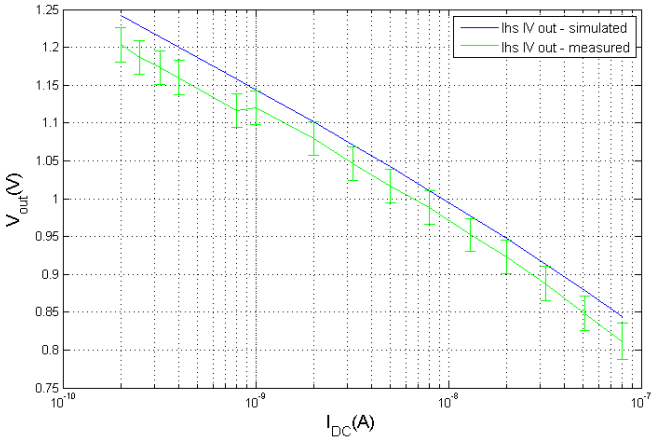

Figure 6.9: Measured and simulated V_{DC} against photocurrent, logarithmic front-end output of BVIPS1 IC.

There are several causes for these discrepancies, although no definitive source of the problem has been found. External sources of noise such as power supply noise (which log pixels are particularly susceptible to) were investigated and reduced (see Section4.4.1), but this had little impact on IC performance, and the work done identified no definitive causes of the noise seen.

Other artefacts seen from the IC during testing of BVIPS1 suggested a more complex fault, with occasional spikes or level shifts seen on output signals. Further testing, described in Section 4.6.3, identified this behaviour as being caused by the HDA AC-amplifier, with the problem being particularly apparent when a signal with little or no AC component was input to the HDA. While this presented a possible explanation (the HDA appearing to have similarities to a switched-capacitor filter circuit which requires an AC clock signal and can

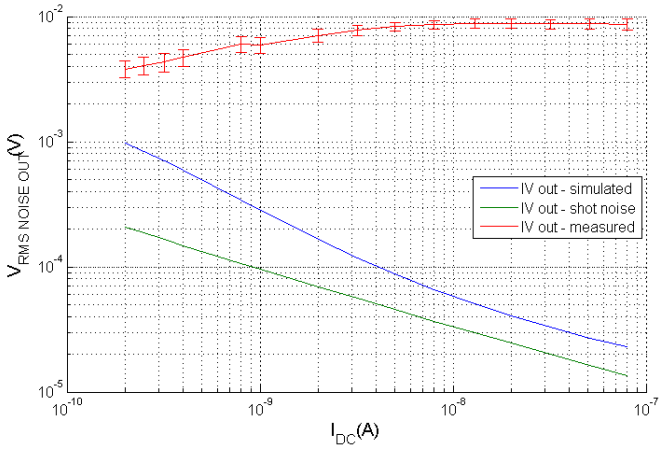

Figure 6.10: Measured and simulated output voltage noise and theoretical shot noise limit against DC photocurrent, logarithmic front-end output of BVIPS1 IC

only pass AC signals), the fault could not be replicated by simulation and no definitive cause could be found.

However, having identified the faulty circuit, an alternative design of HDA was available, which had been previously tested on another comparable IC, and hence this option was taken to address the faults. The new design (discussed in Section 5.5.1) addressed concerns about susceptibility to stray light (i.e. an external interference source) by adding a guard ring to the layout, and a change of capacitor size partially addressed concerns about the possible switched-capacitor filter similarity (with a significant increase in capacitor size being the main design change between previous HDA circuits and that designed for BVIPS2).

Further changes applied to BVIPS2 to address discrepancies included wider use of filters to remove high frequency signals, attempting to prevent noise and interference from propagating through the IC as well as reducing any aliasing caused by slow filter roll-off. The use of alternative front-end circuits was also implemented with the aim of reducing variation in performance, but this step is aimed at reducing variation between two fabricated parts of the IC, rather than between the fabricated IC and simulation results.

The implementation of a full 64x1 array of opamps on BVIPS2, being a relatively uncommon approach for integrated detector arrays, could also be considered as a technique aimed at reducing these discrepancies. Here the approach is to switch to an entirely different type of detector to the logarithmic pixel, accepting that no clear method of reducing noise on logarithmic pixels is available and that the very high resistances and small currents required by HDA amplifiers will remain susceptible to manufacturing variations.

This approach appears to have been partially vindicated, as the characterisation tests of the opamp pixels on BVIPS2 (shown in Section 5.8.2) show reduced discrepancies between measured and simulated results, and the flow readings produced with this array (discussed earlier in this chapter) appear to have better SNR and higher stability than that shown from other arrays used during this work. However, the performance achieved is still not on the same level as that seen from commercial systems.

6.3.4 Large Photodiodes and Speckle Size

Conventional thought would suggest that the large photodiodes used here are an advantage for signal detection, resulting in higher light levels. However, changes in lens focal lengths to focus light to the size of the sensor means that smaller sensors should receive the same light level. The main reason for the large photodiodes here is the requirement for the mechanical scanning to not result in the beam wandering off the sensors.

It is possible that the large sensor actually reduces the Doppler ratio, hence reducing the SNR. This is because the larger pixels may result in more speckles being imaged onto each photodiode.

Speckle size for LDBF is given by [Steenbergen, 2004]:

$$Size_{speckle} = \frac{1.22 \times \lambda_{laser}}{NA} \tag{6.1}$$

Where NA is the numerical aperture of the system. For the imaging setup used here, the NA can be found based on the size of lens used to focus light onto the IC (25 mm) and the distance from the lens to the target (200 mm).

$$NA = \frac{r_{lens}}{\sqrt{D_{target}^2 + r_{lens}^2}} = 0.062 \tag{6.2}$$

From this, and knowing the wavelength of the laser used to be 633 nm, we can find the speckle size of the imaging setup:

$$Size_{speckle} = \frac{1.22 \times 633\,\mathrm{nm}}{0.062} = 12.4\,\mu\mathrm{m} \tag{6.3}$$

The photodiode size is $1000 \times 50\,\mu$m, giving an estimate of 325 speckles per pixel. Compared to previous CMOS detectors with diode size around $50 \times 50\,\mu$m [He et al., 2009,

Gu et al., 2008], giving approximately 16 speckles per pixel. As discussed in Section 2.4.1, the AC photocurrent is due to changes in intensity at each pixel caused by the speckles appearing and disappearing. If the number of speckles on a pixel is sufficiently large, the change in overall intensity as the speckles change will average out to a constant intensity. This means that the larger pixels used on the BVIPS ICs will result in a smaller AC signal, not a larger signal.

This gives several options for increasing the speckle size and hence Doppler ratio. The use of an infra-red laser, already expected to help with Doppler ratio and vein location due to the higher penetration depth, would increase speckle size, as would decreasing the NA, which could be achieved by increasing the distance to the target or reducing the lens size. However, reducing the NA is limited by the need to image the line projected onto the target back onto the sensor array, and the constraints on the overall system such as what distance can be allowed between the target and the other optics elements.

6.3.5 Fixed Pattern Noise on Linear Arrays

The horizontal lines produced by pixels with performance different to neighbouring pixels seen in BVIPS1 (shown in Figure 6.2) and the HDA AC channel of the BVIPS2 opamp pixels (shown in Figure 5.65) highlight a potential issue with the general technique of imaging using a linear array detector and mechanical scanning . This issue is that with a linear array, a single pixel with different performance to most results in a line across the resultant image, whereas the same variation of a pixel in a 2D array would result in only a single pixel 'spot'. While many such spots are likely to occur in the final image in this case, the random nature of these points may be preferable to the lines seen here from a device users perspective.

However, it should be noted that the majority of outputs of the BVIPS2 IC do not show these lines to any obvious extent, showing that careful IC design and development can reduce this fixed pattern noise to acceptable levels, even before the circuit selection techniques demonstrated on BVIPS2, or any calibration routines, are applied.

6.4 Further Work and Possible Alternatives

6.4.1 Clinical Instrument Development

Testing of both BVIPS ICs has shown that this sort of device is capable of detecting flow changes, regardless of the problems from speckle size shown above and the unexpected behaviour seen on BVIPS1. The design of BVIPS2 attempts to addresses the limitations of BVIPS1, and while not totally successful in removing discrepancies between simulated and measured results, the limited testing of BVIPS2 suggests that with refinement of the system (such as more precise analogue bias setup, improved system/PCB design and a more powerful laser), it should be possible to produce clinically useful results. This should mean that BVIPS2 is well suited to use in a clinical instrument such as the Moor Laser Doppler Line Scanner. An obvious follow-up to this work would be to build the BVIPS2 IC into an instrument for clinical use.

6.4.2 Detection of Depth of Blood Flow

There are also possibilities raised by the lower light power required for a linear array system, combined with the small physical size of the integrated sensor. For example, two sensors could be used, with a beam-splitter to simultaneously measure flow from the same line on the target. This could be combined with polarisation on one channel to separate the flow in the top levels of the skin from that in deeper tissue. This technique relies on the polarisation state of light being lost as it is scattered [Belcaro et al., 1994]. If the target is illuminated with polarised light, then a sensor with a polarising filter will only detect light that has not undergone a large number of scattering events, i.e. that has only interacted with the top levels of skin. Light that penetrates deeper into the tissue is scattered more and loses its polarisation state, and is therefore mostly blocked by the polarising filter. A second sensor can either be left un-filtered, or can use a polarising filter at 90^{o} to the polarisation of the illuminating light. This sensor detects all blood flow (if no filter is used), or only flow in deeper tissue (if a filter at 90^{o} is used).

This technique can be applied to any LDBF sensor, but the splitting of light between two sensors, and the blocking of un-polarised light incident on any polarising filter means that the light level incident on the sensor as a fraction of the total illuminating power is reduced.

The lower total light power of a line scanning system compared to a 2D system means that the illuminating light power can be increased without increasing the laser safety risk to an unacceptable level.

6.4.3 Increase of Resolution

The use of a linear array means a doubling of pixel count only doubles the demands on processing and data acquisition, while the resolution of the resultant image is increased by four (assuming the mechanical scan steps are also halved in size). To give the same improvement of resolution for a 2D sensor requires a four times increase in data acquisition and processing capacity, as well as the increase in array area. The resolution could be increased by reducing pixel size rather than increasing IC size, without a significant reduction of fill-factor - the 1D array allows the front-end circuit to be placed to the side of the array, such that more pixels can be added without significantly reducing the overall light sensitive area. For a 2D imager, the front-end circuits have to be interleaved with the photodiodes. The front-end circuits themselves have a fixed size, so reducing the pixel size has to be done by reducing the photodiode size, reducing the fill factor of each pixel. This represents a significant advantage of the line scanning detector, giving greater potential for further improvements of image quality.

6.4.4 Limiting Factors to Further Development

While the results shown here have shown that clinical use of an integrated sensor such as BVIPS2 is feasible with some further development, there are clear drawbacks given the performance reductions (in terms of flux detection) between existing commercial systems and integrated sensors. While the use of an integrated sensor can reduce system size through reduced component count, the size of LDBF systems remains limited by the size of the laser (and power supply/temperature controller), optical components such as lenses/mirrors and the back-end processing/display system (e.g. a tablet PC). This potentially means that the gains offered by reducing the sensor size cannot be translated into an instrument that is significantly more compact. Furthermore, continuous development of commercial CMOS sensors as well as integrated ADCs and low power processors may reduce other advantages of the integrated system presented here such as reducing data bottlenecks and system size.

Bibliography

[Abramowitz, 1970] Abramowitz, M. (1970). *Handbook of mathematical functions : with formulas, graphs, and mathematical tables*. Dover Publications, New York.

[Afshari et al., 2002] Afshari, A., McKinney, W., and Frazer, D. (2002). Design of a novel instrumentation for measurement of blood flow while gripping a handle. In *Bioengineering Conference, 2002. Proceedings of the IEEE 28th Annual Northeast*, pages 215–216.

[Albrecht et al., 2002] Albrecht, H. E., Borys, M., Damaschke, N., and Tropea, C. (2002). *Laser Doppler and Phase Doppler Measurement Techniques*. Springer, 1 edition.

[Allen and Holberg, 2002] Allen, P. E. and Holberg, D. R. (2002). *CMOS Analog Circuit Design*. Oxford University Press, USA, 2 edition.

[austriamicrosystems, 2007] austriamicrosystems (2007). CMOS c35 process parameters. Technical report, austriomicrosystems.

[Belcaro et al., 1994] Belcaro, G., Hoffman, U., Bollinger, A., and Nicolaides, A. N. (1994). *Laser Doppler*. Med-Orion.

[Briers, 1996] Briers, J. D. (1996). Laser doppler and time-varying speckle: a reconciliation. *Journal of the Optical Society of America A*, 13(2):345–350.

[Brown et al., 1998] Brown, R. F. R., Rice, P., and Bennett, N. J. (1998). The use of laser doppler imaging as an aid in clinical management decision making in the treatment of vesicant burns. *Burns*, 24(8):692–698.

[Chen et al., 2007] Chen, L., Zheng, H., Li, L., Xie, P., and Liu, S. (2007). Near-Infrared dorsal hand vein image segmentation by local thresholding using grayscale morphology. In *Bioinformatics and Biomedical Engineering, 2007. ICBBE 2007. The 1st International Conference on*, pages 868–871.

[Clough et al., 1998] Clough, Bennett, and Church (1998). Effects of h1 antagonists on the cutaneous vascular response to histamine and bradykinin: a study using scanning laser doppler imaging. *British Journal of Dermatology*, 138(5):806–814.

[Collins, 2009] Collins, L. (2009). Painting with light. *E&T Magazine*.

[Cutnell and Johnson, 2006] Cutnell, J. D. and Johnson, K. W. (2006). *Physics*. Wiley, 7 edition.

[Dmochowski, 2006] Dmochowski, P. (2006). *CMOS modulated light camera*. PhD thesis, University of Nottingham.

[Dmochowski et al., 2004] Dmochowski, P., Hayes-Gill, B., Clark, M., Crowe, J., Somekh, M., and Morgan, S. (2004). Camera pixel for coherent detection of modulated light. *Electronics Letters*, 40(22):1403–1404.

[Draijer et al., 2009] Draijer, M., Hondebrink, E., van Leeuwen, T., and Steenbergen, W. (2009). Time domain algorithm for accelerated determination of the first order moment of photo current fluctuations in high speed laser doppler perfusion imaging. *Medical and Biological Engineering and Computing*, 47(10):1103–1109.

[Essex, 1994] Essex, T. (1994). Laser doppler imaging of skin blood flow. In *Functional Imaging, IEE Colloquium on*, pages 8/1–8/2.

[Essex and Byrne, 1991] Essex, T. J. and Byrne, P. O. (1991). A laser doppler scanner for imaging blood flow in skin. *Journal of Biomedical Engineering*, 13(3):189–194. PMID: 1870327.

[Fercher and Briers, 1981] Fercher, A. and Briers, J. D. (1981). Flow visualization by means of single-exposure speckle photography. *Optics Communications*, 37(5):326–330.

[Ferrell et al., 1997] Ferrell, W. R., McDougall, J. J., and Bray, R. C. (1997). Spatial heterogeneity of the effects of calcitonin gene-related peptide (CGRP) on the microvasculature of ligaments in the rabbit knee joint. *British Journal of Pharmacology*, 121(7):1397–1405.

[Fossum, 1997] Fossum, E. (1997). CMOS image sensors: electronic camera-on-a-chip. *Electron Devices, IEEE Transactions on*, 44(10):1689–1698.

[Geiger and SÃ¡nchez-Sinencio, 1985] Geiger, R. L. and SÃ¡nchez-Sinencio, E. (1985). Active filter design using operational transconductance amplifiers: A tutorial. *IEEE Circuits and Devices Magazine*, 1:20–32.

[Gschwandtner et al., 1999] Gschwandtner, M. E., Ambrozy, E., Schneider, B., Fasching, S., Willfort, A., and Ehringer, H. (1999). Laser doppler imaging and capillary microscopy in ischemic ulcers. *Atherosclerosis*, 142(1):225–232.

[Gu, 2007] Gu, Q. (2007). *An Analogue Integrated Circuit Design for Laser Doppler Blood Flow Measurement*. PhD thesis, University of Nottingham.

[Gu et al., 2008] Gu, Q., Hayes-Gill, B. R., and Morgan, S. P. (2008). Laser doppler blood flow complementary metal oxide semiconductor imaging sensor with analog on-chip processing. *Applied Optics*, 47(12):2061–2069.

[Halliday et al., 2000] Halliday, D., Resnick, R., and Walker, J. (2000). *Extended , Fundamentals of Physics, 6th Edition*. Wiley, 6 edition.

[He et al., 2009] He, D., Hoang, N., Hayes-Gill, B. R., Crowe, J. A., Zhu, Y., and Morgan, S. P. (2009). Full field laser doppler flowmetry with custom made CMOS sensors. In *Design and Quality for Biomedical Technologies II*, volume 7170, pages 71700J–8, San Jose, CA, USA. SPIE.

[Hecht, 2001] Hecht, E. (2001). *Optics*. Addison Wesley, 4 edition.

[Hoang, 2009] Hoang, N. (2009). *High Speed Processing for Laser Doppler Blood Flow Imaging*. PhD thesis, University of Nottingham.

[Hoang et al., 2010] Hoang, N., Hayes-Gill, B., Morgan, S., Zhu, Y., Boggett, D., Huang, X., and Potter, M. (2010). A Field-Programmable gate array based system for high frame rate laser doppler blood flow imaging. *Journal of Medical Engineering & Technology*, 35(5-6):306–315.

[Holst and Lomheim, 2007] Holst, G. C. and Lomheim, T. S. (2007). *CMOS/CCD Sensors and Camera Systems*. SPIE Society of Photo-Optical Instrumentation Engi, illustrated edition edition.

[Horowitz and Hill, 1989] Horowitz, P. and Hill, W. (1989). *The Art of Electronics*. Cambridge University Press, 2 edition.

[Houdas, 1991] Houdas, Y. (1991). [Doppler, Buys-Ballot, fizeau. historical note on the discovery of the doppler's effect]. *Annales De Cardiologie Et D'angÃ©iologie*, 40(4):209–213. PMID: 2053764.

[Johnston et al., 2009] Johnston, N., Stewart, C., Light, R., Hayes-Gill, B., Somekh, M., Morgan, S., Sambles, J., and Pitter, M. (2009). Quad-phase synchronous light detection with 64x64 CMOS modulated light camera. *Electronics Letters*, 45(21):1090–1091.

[Jung, 2005] Jung, W. G. (2005). *Op Amp applications handbook*. Newnes.

[Khan and Newton, 2003] Khan, F. and Newton, D. J. (2003). Laser doppler imaging in the investigation of lower limb wounds. *International Journal of Lower Extremity Wounds*, 2(2):74–86.

[Klatzky et al., 1985] Klatzky, R. L., Lederman, S. J., and Metzger, V. A. (1985). Identifying objects by touch: an "expert system". *Perception & Psychophysics*, 37(4):299–302. PMID: 4034346.

[Koelink et al., 1994] Koelink, M. H., de Mul, F. F. M., Greve, J., Graaff, R., Dassel, A. C. M., and Aarnoudse, J. G. (1994). Laser doppler blood flowmetry using two wavelengths: Monte carlo simulations and measurements. *Applied Optics*, 33(16):3549–3558.

[Kongsavatsak, 2005] Kongsavatsak, C. (2005). *Full Field Laser Doppler Blood Flow Camera*. PhD thesis, University of Nottingham.

[Kongsavatsak et al., 2008] Kongsavatsak, C., He, D., Hayes-Gill, B. R., Crowe, J. A., and Morgan, S. P. (2008). Complementary metal-oxide-semiconductor imaging array with laser doppler blood flow processing. *Optical Engineering*, 47(10):104401–14.

[Kussman, 2001] Kussman, K. S. (2001). Veinlite - a new tool for sclerotherapy. *Journal fÃŒr Lymphologie*, pages 46–47.

[Le Duff et al., 2004] Le Duff, A., Plantier, G., Valiere, J., and Bosch, T. (2004). Analog sensor design proposal for laser doppler velocimetry. *Sensors Journal, IEEE*, 4(2):257–261.

[Liebert and Maniewski, 1996] Liebert, A. and Maniewski, R. (1996). Influence of probe optical arrangement on biological zero in laser-Doppler perfusion measurements. In *Engineering in Medicine and Biology Society, 1996. Bridging Disciplines for Biomedicine. Proceedings of the 18th Annual International Conference of the IEEE*, volume 1, pages 200–201 vol.1.

[Lindsey, 2005] Lindsey, J. (2005). No more Stab-in-the-Dark IV sticks. *Journal of Emergency Medical Services*, page 90.

[Mbamalu and Banerjee, 1999] Mbamalu, D. and Banerjee, A. (1999). Methods of obtaining peripheral venous access in difficult situations. *Postgraduate Medical Journal*, 75(886):459–462.

[Mead, 1989] Mead, C. A. (1989). *Analog VLSI and neural systems.* Addison-Wesley publishing company, Inc.

[Modha et al., 2008] Modha, K., Stockford, I., Light, R., Clark, M., Pitter, M., and Hayes-Gill, B. (2008). A custom CMOS sensor for pyramidal adaptive optics system. In *Mixed-Signals, Sensors, and Systems Test Workshop, 2008. IMS3TW 2008. IEEE 14th International*, pages 1–4.

[Moini, 2000] Moini, A. (2000). *Vision Chips.* Kluwer Academic Publishers.

[MoorInstruments, 2007a] MoorInstruments (2007a). Moor FLPI. http://www.moor.co.uk/products/imaging/moorflpi (accessed 29 November 2009).

[MoorInstruments, 2007b] MoorInstruments (2007b). Moor LDI2 laser doppler imager specifications. http://gb.moor.co.uk/product/moorldi2-laser-doppler-imager/8/o/14/specifications.

[MoorInstruments, 2009] MoorInstruments (2009). moorLDLS | moor instruments. http://www.moor.co.uk/products/imaging/moorldls (accessed 12 February 2010).

[Morgan and Hayes-Gill, 2009] Morgan, S. and Hayes-Gill, B. (2009). Custom-made cameras for blood-flow imaging. *SPIE Newsroom - Biomedical Optics & Medical Imaging*, http://spie.org/x37664.xml?ArticleID=x37664.

[Morris et al., 1995] Morris, S. J., Shore, A. C., and Tooke, J. E. (1995). Responses of the skin microcirculation to acetylcholine and sodium nitroprusside in patients with NIDDM. *Diabetologia*, 38(11):1337–1344.

[Nilsson, 1996] Nilsson, G. (1996). Laser doppler perfusion imaging. In *Engineering in Medicine and Biology Society, 1996. Bridging Disciplines for Biomedicine. Proceedings of the 18th Annual International Conference of the IEEE*, volume 1, pages 196–197 vol.1.

[Nilsson et al., 1980] Nilsson, G. E., Tenland, T., and Oberg, P. A. (1980). A new instrument for continuous measurement of tissue blood flow by light beating spectroscopy. *Biomedical Engineering, IEEE Transactions on*, BME-27(1):12–19.

[Nixon et al., 1999] Nixon, J., Smye, S., Scott, J., and Bond, S. (1999). The diagnosis of early pressure sores: report of the pilot study. *Journal of Tissue Viability*, 9(2):62–66. PMID: 10480973.

[Oberg et al., 1979] Oberg, P. A., Nilsson, G. E., Tenland, T., Holmstrom, A., and Lewis, D. H. (1979). Use of a new laser doppler flowmeter for measurement of capillary blood flow in skeletal muscle after bullet wounding. *Acta Chirurgica Scandinavica. Supplementum*, 489:145–150. PMID: 158931.

[PerimedAB, 2011] PerimedAB (2011). PeriScan PIM 3 technical specifications | perimed products. http://www.perimed-instruments.com/products/pim3-technical-specifications.

[Powers and Frayer, 1978] Powers, E. W. and Frayer, W. W. (1978). Laser doppler measurement of blood flow in the microcirculation. *Plastic and Reconstructive Surgery*, 61(2):250–255. PMID: 622417.

[Quinn et al., 1991] Quinn, A. G., McLelland, J., Essex, T., and Farr, P. M. (1991). Measurement of cutaneous inflammatory reactions using a scanning laser-Doppler velocimeter. *British Journal of Dermatology*, 125(1):30–37.

[Richards and Briers, 1997] Richards, G. and Briers, J. (1997). Laser speckle contrast analysis (LASCA): a technique for measuring capillary blood flow using the first order statistics of laser speckle patterns. In *Biomedical Applications of Photonics (Digest No. 1997/124), IEE Colloquium on*, pages 11/1–11/6.

[Riva et al., 1972] Riva, C., Ross, B., and Benedek, G. B. (1972). Laser doppler measurements of blood flow in capillary tubes and retinal arteries. *Invest. Ophthalmol. Vis. Sci.*, 11(11):936–944.

[Rudd, 1969] Rudd, M. J. (1969). A new theoretical model for the laser dopplermeter. *Journal of Physics E: Scientific Instruments*, 2(1):55–58.

[Sargent and Scully, 1972] Sargent, M. and Scully, M. (1972). Theory of laser operation - an outline. In *Laser Handbook, edited by Arechhi, F.T. and Schulz-DuBois, E. O.*, volume 1, pages 45–114. North-Holland.

[Serov and Lasser, 2005] Serov, A. and Lasser, T. (2005). High-speed laser doppler perfusion imaging using an integrating CMOS image sensor. *Optics Express*, 13:6416.

[Serov et al., 2006a] Serov, A., Lasser, T., Cohn, G. E., Grundfest, W. S., Benaron, D. A., and Vo-Dinh, T. (2006a). Full-field high-speed laser doppler imaging system for blood-flow measurements. In *Advanced Biomedical and Clinical Diagnostic Systems IV*, volume 6080, pages 608004–8, San Jose, CA, USA. SPIE.

[Serov et al., 2006b] Serov, A., Lasser, T., Cote, G. L., and Priezzhev, A. V. (2006b). Combined laser doppler and laser speckle imaging for real-time blood flow measurements. In *Optical Diagnostics and Sensing VI*, volume 6094, pages 609406–8, San Jose, CA, USA. SPIE.

[Serov et al., 2006c] Serov, A., Nieland, J., Oosterbaan, S., de Mul, F., van Kranenburg, H., Bekman, H., and Steenbergen, W. (2006c). Integrated optoelectronic probe including a vertical cavity surface emitting laser for laser doppler perfusion monitoring. *Biomedical Engineering, IEEE Transactions on*, 53(10):2067–2074.

[Serov et al., 2002] Serov, A., Steenbergen, W., and de Mul, F. (2002). Laser doppler perfusion imaging with a complimentary metal oxide semiconductor image sensor. *Optics Letters*, 27(5):300–302.

[Serov et al., 2005] Serov, A., Steinacher, B., and Lasser, T. (2005). Full-field laser doppler perfusion imaging and monitoring with an intelligent CMOS camera. *Optics Express*, 13:3681.

[Serway and Jewett, 2009] Serway, R. A. and Jewett, J. W. (2009). *Physics for Scientists and Engineers, Chapters 1-39*. Brooks Cole, 8 edition.

[Shepherd and Oberg, 1990] Shepherd, A. and Oberg, P. (1990). *Laser-Doppler blood flowmetry*. Springer.

[Shuwang et al., 2005] Shuwang, C., Mingshi, W., and Rui, Z. (2005). Infrared thermal image analysis in gradual approach for veins location. In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pages 1668–1670.

[Steenbergen, 2004] Steenbergen, W. (2004). Speckle-related scaling effects in laser doppler perfusion imaging. In *Proceedings of SPIE*, pages 35–43, San Jose, CA, USA.

[Steenbergen et al., 2010] Steenbergen, W., Hondebrink, E., Van Leeuwen, T. G., and Draijer, M. J. (2010). Comparison of scanning beam and whole field laser doppler perfusion imaging. pages 75630H–75630H–4, San Francisco, California, USA.

[Stein and Pfeifer, 1969] Stein, H. D. v. and Pfeifer, H. J. (1969). A doppler difference method for velocity measurements. *Metrologia*, 5(2):59–61.

[Stern, 1975] Stern, M. D. (1975). In vivo evaluation of microcirculation by coherent light scattering. *Nature*, 254(5495):56–58.

[Stewart et al., 2005] Stewart, C., Frank, R., Forrester, K., Tulip, J., Lindsay, R., and Bray, R. (2005). A comparison of two laser-based methods for determination of burn scar perfusion: Laser doppler versus laser speckle imaging. *Burns*, 31(6):744–752.

[Tanaka et al., 1974] Tanaka, T., Riva, C., and Ben-Sira, I. (1974). Blood velocity measurements in human retinal vessels. *Science*, 186(4166):830–831.

[Theuwissen, 2001] Theuwissen, A. (2001). CCD or CMOS image sensors for consumer digital still photography? In *VLSI Technology, Systems, and Applications, 2001. Proceedings of Technical Papers. 2001 International Symposium on*, pages 168–171.

[Theuwissen, 1995] Theuwissen, A. J. P. (1995). *Solid-state imaging with charge-coupled devices*. Springer.

[Translite, 2001] Translite (2001). Veinlite.com - finding veins for vein access and imaging varicose veins with transillumination. http://www.veinlite.com/.

[Tuchin, 2007] Tuchin, V. V. (2007). *Tissue optics*. SPIE Press.

[Watkins and Holloway, 1978] Watkins, D. and Holloway, G. A. (1978). An instrument to measure cutaneous blood flow using the doppler shift of laser light. *Biomedical Engineering, IEEE Transactions on*, BME-25(1):28–33.

[Wunderlich et al., 1980] Wunderlich, R. W., Folger, R. L., Giddon, D. B., and Ware, B. R. (1980). Laser doppler blood flow meter and optical plethysmograph. *Review of Scientific Instruments*, 51(9):1258–1262.

[Yeh and Cummins, 1964] Yeh, Y. and Cummins, H. Z. (1964). Localized fluid flow measurements with AN He-Ne laser spectrometer. *Applied Physics Letters*, 4:176–178.

[Zeman et al., 2005] Zeman, H. D., Lovhoiden, G., Vrancken, C., and Danish, R. K. (2005). Prototype vein contrast enhancer. *Optical Engineering*, 44(8):086401–9.

[Zhong et al., 1998] Zhong, J., Seifalian, A., Salerud, G., and Nilsson, G. (1998). A mathematical analysis on the biological zero problem in laser doppler flowmetry. *Biomedical Engineering, IEEE Transactions on*, 45(3):354–364.

[Zhu et al., 2006] Zhu, Y., Hayes-Gill, B. R., Morgan, S. P., and Hoang, N. C. (2006). An FPGA based generic prototyping platform employed in a CMOS laser doppler blood flow camera. In *Field Programmable Technology, 2006. FPT 2006. IEEE International Conference on*, pages 281–284.

[Zimerman, 1991] Zimerman, E. (1991). The landry vein light: increasing venipuncture success rates. *Journal of Pediatric Nursing*, 6(1):64–66. PMID: 1990130.

[Zivanovic and Davies, 2000] Zivanovic, A. and Davies, B. (2000). A robotic system for blood sampling. *Information Technology in Biomedicine, IEEE Transactions on*, 4(1):8–14.

Appendix A

Matlab Code Listing for Front-End Circuit Modelling

A.1 Modelling of Pixel Behaviour

Matlab function plotting DC response, transimpedance, bandwidth of basic and buffered logarithmic pixels, with one or two load PMOS transistors. Includes code to plot comparison of Matlab-modelled performance with Cadence (ie Spice) simulations. Also models capacitance of photodiode:

```
% function to plot bandwidth of log pixels as function of idc

%first plot change in diode capacitance against voltage
function plotlogpixbw_v2()
    v = 0.1:0.1:3.3;
    c = zeros(numel(v),1);
    c_a = c;
    c_p = c;
    for i = 1:numel(v)
        [c(i),c_a(i),c_p(i)] = diodecap(50,1000,v(i));
    end;

    figure('Position',[100   100   700   420],'Name','Diodecap, all');
    plot(v,c);      %plot capacitance - output voltage
    title('Variation of photodiode capacitance with photodiode voltage','FontSize',14);
    xlabel('DC Photodiode (ie output) Voltage (V)','FontSize',12);
    ylabel('Capacitance (F)','FontSize',12);
    h_legend = legend('matlab, basic pixel','FontSize',12);
    set(h_legend,'FontSize',12)
    grid on;
    set(gcf, 'PaperPositionMode', 'auto');
    print(gcf, '-r0', 'diodecap.png', '-dpng');
    saveas(gcf,'diodecap.fig');


    %%%%%%%%%% Matlab modelling of DC output voltage:  %%%%%%%%%%%
    idc_mat = 100e-12:100e-12:30e-9;    % range of expected IDC
```

```
io = 0.5e-12*2;  %PMOS subthreshold leakage from ams docs, 0.5pA/um, 2um gate width
n = 1.2;    %sub-threshold slope factor of NMOS transistor
Ut = 25e-3; %thermal voltage, kT/q = 25mV
vout_basic_mat = 3.3 -  n*Ut*log(idc_mat) + n*Ut*log(io);


%plot output voltage - photocurrent
figure('Position',[100   100   700   420],'Name','VDC, basic matlab');
plot(idc_mat,vout_basic_mat);
title('Variation of output voltage with DC photocurrent','FontSize',14);
xlabel('DC Photocurrent (A)','FontSize',12);
ylabel('DC Output voltage (V)','FontSize',12);
h_legend = legend('matlab, basic pixel');
set(h_legend,'FontSize',12)
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'VDC-basic-matlab.png', '-dpng');
saveas(gcf,'VDC-basic-matlab.fig');

%%%%%%%%%% Read DC voltage results from Cadence):  %%%%%%%%%%%
%Read results from cadence simulation of basic pixel, 1 transistor
A = csvread('vdc_basic_nosf_1.csv');
idc_sim = A(:,1);
vdc_sim_basic_1 = A(:,2);

%Read results from cadence simulation of basic pixel, 2 transistors
A = csvread('vdc_basic_nosf_2.csv');
vdc_sim_basic_2 = A(:,2);

%Read results from cadence simulation of buffered pixel, 1 transistor
A = csvread('vdc_buffered_nosf_1.csv');
vdc_sim_buffered_1 = A(:,2);

%Read results from cadence simulation of buffered pixel, 2 transistors
A = csvread('vdc_buffered_nosf_2.csv');
vdc_sim_buffered_2 = A(:,2);

%plot output voltage - photocurrent, cadence/matlab models
figure('Position',[100   100   700   420],'Name','VDC, all');
plot(idc_sim,vdc_sim_basic_1,idc_sim,vdc_sim_basic_2,idc_sim,vdc_sim_buffered_1,...
    idc_sim,vdc_sim_buffered_2,idc_mat,vout_basic_mat);
title('Variation of output voltage with DC photocurrent','FontSize',14);
xlabel('DC Photocurrent (A)','FontSize',12);
ylabel('DC Output voltage (V)','FontSize',12);
h_legend = legend('Basic, 1 PMOS load, Cadence','Basic, 2 PMOS load, Cadence',...
    'Buffered, 1 PMOS load, Cadence','Buffered, 2 PMOS load, Cadence',...
    'Basic, 1 PMOS load, Matlab');
set(h_legend,'FontSize',12,'Position',[0.55 0.34 0.34 0.25]);
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'VDC-all.png', '-dpng');
saveas(gcf,'VDC-all.fig');

%%%%%%%%%% Matlab modelling of bandwidth & Transimpedance:  %%%%%%%%%%%
```

```
cpd = zeros(numel(vout_basic_mat),1);
cpd_a = cpd;
cpd_p = cpd;
%calculate diode cap over range of output voltages found above
for i = 1:numel(vout_basic_mat);
    [cpd(i),cpd_a(i),cpd_p(i)] = diodecap(47.4,1000,vout_basic_mat(i));
end;

fc_basic_mat = zeros(numel(idc_mat),1);
rac_basic_mat = fc_basic_mat;
%calculate bandwidth from c, v and 1/2piRC, using R = n*Ut/IDC
for i = 1:numel(idc_mat);
    rac_basic_mat(i) = (n*Ut)/idc_mat(i); %Find Rac (transimpedance)
    fc_basic_mat(i) = 1/(2*pi*rac_basic_mat(i)*cpd(i));

end;

%plot matlab modelled bandwidth - photocurrent
figure('Position',[100   100   700   420],'Name','BW, basic matlab');
plot(idc_mat,fc_basic_mat);
title('Variation of cut-off frequency with DC photocurrent','FontSize',14);
xlabel('DC Photocurrent (A)','FontSize',12);
ylabel('Cut-off frequency (Hz)','FontSize',12);
h_legend = legend('matlab, basic pixel');
set(h_legend,'FontSize',12)
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'BW-basic-matlab.png', '-dpng');
saveas(gcf,'BW-basic-matlab.fig');

%plot matlab modelled transimpedance - photocurrent
figure('Position',[100   100   700   420],'Name','Gain, basic matlab');
plot(idc_mat,rac_basic_mat);
title('Variation of transimpedance with DC photocurrent','FontSize',14);
xlabel('DC Photocurrent (A)','FontSize',12);
ylabel('Transimpedance (\Omega)','FontSize',12);
h_legend = legend('matlab, basic pixel');
set(h_legend,'FontSize',12)
ylim([0 1e8]);
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'Gain-basic-matlab.png', '-dpng');
saveas(gcf,'Gain-basic-matlab.fig');

%%%%%%%%%% Read Gain (ie AC resistance results from Cadence):  %%%%%%%%%%%
%Read results from cadence simulation of basic pixel, 1 transistor
A = csvread('gain_basic_nosf_1.csv');
idc_sim = A(:,1);
gain_sim_basic_1 = A(:,2);

%Read results from cadence simulation of basic pixel, 2 transistors
A = csvread('gain_basic_nosf_2.csv');
gain_sim_basic_2 = A(:,2);
```

```
%Read results from cadence simulation of buffered pixel, 1 transistor
A = csvread('gain_buffered_nosf_1.csv');
gain_sim_buffered_1 = A(:,2);

%Read results from cadence simulation of buffered pixel, 2 transistors
A = csvread('gain_buffered_nosf_2.csv');
gain_sim_buffered_2 = A(:,2);

%plot gain/transimpedance - photocurrent, cadence/matlab models
figure('Position',[100   100   700   420],'Name','gain, all');
plot(idc_sim,gain_sim_basic_1,'-x',idc_sim,gain_sim_basic_2,idc_sim,...
    gain_sim_buffered_1,idc_sim,gain_sim_buffered_2,idc_mat,rac_basic_mat);
title('Variation of transimpedance with DC photocurrent','FontSize',14);
xlabel('DC Photocurrent (A)','FontSize',12);
ylabel('Transimpedance (\Omega)','FontSize',12);
h_legend = legend('Basic, 1 PMOS load, Cadence','Basic, 2 PMOS load, Cadence',...
    'Buffered, 1 PMOS load, Cadence','Buffered, 2 PMOS load, Cadence',...
    'Basic, 1 PMOS load, Matlab');
set(h_legend,'FontSize',12)
ylim([0 1e8]);
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'Gain-all.png', '-dpng');
saveas(gcf,'Gain-all.fig');


%%%%%%%%%% Read bandwidth results from Cadence):  %%%%%%%%%%
%Read results from cadence simulation of basic pixel, 1 transistor
A = csvread('BW_basic_nosf_1.csv');
idc_sim = A(:,1);
fc_sim_basic_1 = A(:,2);

%Read results from cadence simulation of basic pixel, 2 transistors
A = csvread('BW_basic_nosf_2.csv');
fc_sim_basic_2 = A(:,2);

%Read results from cadence simulation of buffered pixel, 1 transistor
A = csvread('BW_buffered_nosf_1.csv');
fc_sim_buffered_1 = A(:,2);

%Read results from cadence simulation of buffered pixel, 2 transistors
A = csvread('BW_buffered_nosf_2.csv');
fc_sim_buffered_2 = A(:,2);

%plot bandwidth - photocurrent
figure('Position',[100   100   700   420],'Name','BW, all');
plot(idc_sim,fc_sim_basic_1,idc_sim,fc_sim_basic_2,idc_sim,fc_sim_buffered_1,...
    idc_sim,fc_sim_buffered_2,idc_mat,fc_basic_mat);
title('Variation of cut-off frequency with DC photocurrent','FontSize',14);
xlabel('DC Photocurrent (A)','FontSize',12);
ylabel('Cut-off frequency (Hz)','FontSize',12);
h_legend = legend('Basic, 1 PMOS load, Cadence','Basic, 2 PMOS load, Cadence',...
    'Buffered, 1 PMOS load, Cadence','Buffered, 2 PMOS load, Cadence',...
    'Basic, 1 PMOS load, Matlab');
```

```
    set(h_legend,'FontSize',12)
    grid on;
    set(gcf, 'PaperPositionMode', 'auto');
    print(gcf, '-r0', 'BW-all.png', '-dpng');
    saveas(gcf,'BW-all.fig');


function [c,c_a,c_p] = diodecap(w,l,v)

    cj = 0.08;   %junction capacitance/area (fF/um^2)
    pb = 0.53;      %junction potential (V)
    mj = 0.39;      %area grading coefficient
    cjsw = 0.51; %junction capacitance/perimeter length (fF/um)
    mjsw = 0.27;     %sidewall grading coefficient
    c_a = ((w*l*cj)/(1+((v/pb)^mj)))*1e-15;
    c_p = (2*(w+l)*cjsw)/(1+((v/pb)^mjsw))*1e-15;
    c = c_a + c_p;
```

A.2 Modelling of Noise

Matlab function plotting noise (input and output referred) of basic and buffered logarithmic pixels, with one or two load PMOS transistors:

```
%Function to plot noise in log pixels
%Model noise in basic log pixel (no buffering)
n = 1.2;
Ut = 25e-3;
q = 1.6e-19;
idc_model = 100e-12:100e-12:30e-9;

%Model voltage and input referred current noise:
for i = 1:numel(idc_model)
    vnoise_model(i) = sqrt((n^2+2*n)*(2*Ut^2*q)/idc_model(i));
    inoise_model(i) = sqrt((2+4/n)*q*idc_model(i));
end;


%%%%%%%%%% Read Vnoise results from Cadence):  %%%%%%%%%%%
%Read results from cadence simulation of basic pixel, 1 transistor
A = csvread('vnoise_basic_nosf_1.csv');
idc_sim = A(:,1);
vnoise_sim_basic_1 = A(:,2);

%Read results from cadence simulation of basic pixel, 2 transistors
A = csvread('vnoise_basic_nosf_2.csv');
vnoise_sim_basic_2 = A(:,2);

%Read results from cadence simulation of buffered pixel, 1 transistor
A = csvread('vnoise_buffered_nosf_1.csv');
vnoise_sim_buffered_1 = A(:,2);
```

```
%Read results from cadence simulation of buffered pixel, 2 transistors
A = csvread('vnoise_buffered_nosf_2.csv');
vnoise_sim_buffered_2 = A(:,2);

%%%%%%%%%% Read Inoise results from Cadence):  %%%%%%%%%%%%
%Read results from cadence simulation of basic pixel, 1 transistor
A = csvread('inoise_basic_nosf_1.csv');
idc_sim = A(:,1);
inoise_sim_basic_1 = A(:,2);

%Read results from cadence simulation of basic pixel, 2 transistors
A = csvread('inoise_basic_nosf_2.csv');
inoise_sim_basic_2 = A(:,2);

%Read results from cadence simulation of buffered pixel, 1 transistor
A = csvread('inoise_buffered_nosf_1.csv');
inoise_sim_buffered_1 = A(:,2);

%Read results from cadence simulation of buffered pixel, 2 transistors
A = csvread('inoise_buffered_nosf_2.csv');
inoise_sim_buffered_2 = A(:,2);

figure('Position',[100   100   700   420],'Name','Vnoise, all');
plot(idc_sim,vnoise_sim_basic_1,idc_sim,vnoise_sim_basic_2,idc_sim,vnoise_sim_buffered_1,.
    idc_sim,vnoise_sim_buffered_2,idc_model,vnoise_model);
title('Variation of voltage noise density with DC photocurrent','FontSize',14);
xlabel('Photocurrent (A)','FontSize',12);
ylabel('output voltage noise (V/sqrt(Hz))','FontSize',12);
h_legend = legend('Basic, 1 PMOS load, Cadence','Basic, 2 PMOS load, Cadence',...
    'Buffered, 1 PMOS load, Cadence','Buffered, 2 PMOS load, Cadence',...
    'Basic, 1 PMOS load, Matlab');
set(h_legend,'FontSize',12)
ylim([0 1e-6]);
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'Vnoise-all.png', '-dpng');
saveas(gcf,'Vnoise-all.fig');

figure('Position',[100   100   700   420],'Name','Inoise, all');
plot(idc_sim,inoise_sim_basic_1,idc_sim,inoise_sim_basic_2,idc_sim,inoise_sim_buffered_1,.
    idc_sim,inoise_sim_buffered_2,idc_model,inoise_model);
title('Variation of current noise density with DC photocurrent','FontSize',14);
xlabel('Photocurrent (A)','FontSize',12);
ylabel('input referred current noise (A/sqrt(Hz))','FontSize',12);
h_legend = legend('Basic, 1 PMOS load, Cadence','Basic, 2 PMOS load, Cadence',...
    'Buffered, 1 PMOS load, Cadence','Buffered, 2 PMOS load, Cadence',...
    'Basic, 1 PMOS load, Matlab');
set(h_legend,'FontSize',12,'Position',[0.55 0.13 0.33 0.25])
grid on;
set(gcf, 'PaperPositionMode', 'auto');
print(gcf, '-r0', 'Inoise-all.png', '-dpng');
saveas(gcf,'Inoise-all.fig');
```

Appendix B

Tests of Existing Laser Doppler Flowmetry ICs

The simulations in Chapter 3 were used to select the individual elements to be used for the first prototype IC. In addition to this, as the IC designed here was intended to build on work done on LDF within the University of Nottingham Applied Optics Group, it was possible to test previously fabricated ICs that use the above logarithmic front-end circuits. This section will show experimental test results from these ICs when used to measure Doppler signals.

These tests partially support decisions taken based on simulation results, but concerns about this testing means these results do not form a main part of the thesis. These concerns relate to the quality of the results (in terms of high noise and un-clear differences in performance), and the experimental setup used not being directly comparable to true operating conditions of the line scanning imager.

B.1 Tests Performed

The tests were performed using a 'motility target', which is used as a calibration standard by Moor Instruments. This is a small clear bottle containing a suspension of micro-particles. The Brownian motion of the micro-particles gives a similar Doppler signal to that found from micro-circulation. This provides a useful substitute for testing on actual blood flow as measurements are more repeatable. The tests were performed using elements of a Moor Instruments LDF point scanning system (MoorLDI), mainly a 2.2 mW, 875 nm laser with a 1 mm diameter beam. The system also includes a lower power, coaxial visible red laser which is used for beam alignment and safety reasons. Light was collected through a fibre optic as shown in Figure B.1. The output from the MoorLDI was also acquired to provide a comparison.

The ICs tested were:

- DOP2 - buffered feedback front-end, $50 \times 50\,\mu$m diode. This chip also included, at the pixel level, a hysteretic differentiator amplifier (HDA) using a g_mC low-pass filter in the feedback. The low cut-off frequency of this HDA is around 1kHz [Gu et al., 2008]
- PC3 - CMOS inverter feedback, $18.8 \times 18.8\,\mu$m diode. PC3 includes, at the column level, a HDA with AC gain of around 80, using an inverted inverter in the feedback and an additional operational trans-conductance amplifier (OTA) in the input. This HDA design gives an AC gain of 50 (dependent on bias current and input DC voltage)

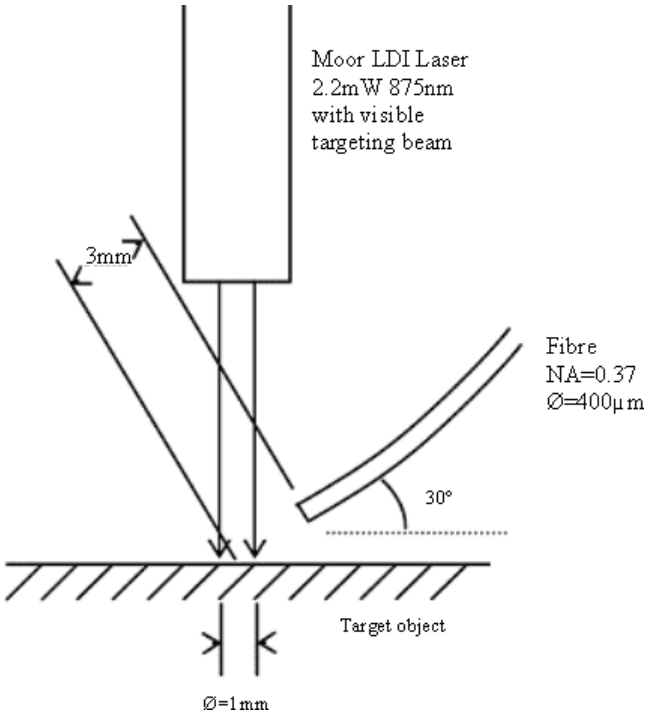

Figure B.1: Optical setup for testing of existing ICs with MoorLDI laser

with a DC gain of unity. The low cut-off frequency of this HDA is 5 Hz. A high-pass filter using a single capacitor and eight transistors as a resistive load, with a designed cut-off frequency of 11 kHz is also present on the chip, but was after the point at which the ADC was connected. This camera was not designed for Doppler blood flow measurements, hence the small diode size and the high cut off frequency of the high-pass filter [Johnston et al., 2009].

- DOP1 - Simple diode-connected PMOS non-feedback front-end. While simulations have shown that this front-end lacks the bandwidth required for this system, it can be used to observe the output signal at lower frequencies for various tests. This IC uses three diode-connected transistors as the active load, with a $100 \times 175\mu$m diode. This chip also included, at the pixel level, a HDA using an inverted CMOS inverter in the feedback, and a g_mC anti-aliasing filter was also included to provide a low pass filter with a 20 kHz cut-off [Kongsavatsak et al., 2008]. While this frequency is below the bandwidth expected from this type of front-end (see Section 3.2.1.3), its inclusion prevents noise above this frequency added by the amplifier, and also restricts the bandwidth should the light power rise above that required for a 20 kHz bandwidth .

B.1.1 Buffered Feedback Front-end Tests on DOP2 IC

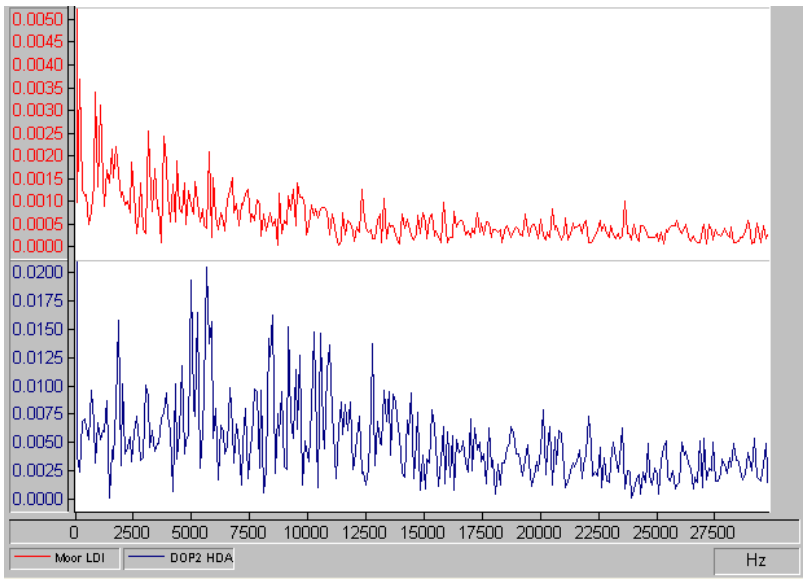

Figure B.2: DOP2 HDA output (bottom) / Moor LDF output (top) (y axis units are Volts)

Figure B.2 shows the output spectrum (calculated by FFT on a PC) of the signal from an existing Moor Instruments LDF system (top) and from the HDA output of DOP2 (bottom). It can be seen that above $\sim 5\,\text{kHz}$ the signals have a similar form. The output from the HDA in DOP2 is larger than that from the Moor Instruments system, being around 20 mV at 5 kHz compared to 2 mV for the Moor Instruments system.

The low signal at lower frequencies is thought to be a problem with the low-cut off frequency of the HDA, rather than being a design problem of the front-end. The HDA design used here is the g_mC feedback variety discussed in Section 3.3.2. This design was shown by simulation to be the better design in terms of controllability, and avoided the high sensitivity of the inverter-inverter feedback arrangement in the alternative design. However, Monte-Carlo simulation showed that it is potentially more prone to manufacturing variation than the inverted-inverter design. The behaviour shown in Figure B.2 confirms that manufacturing variations can cause characteristics of the circuit to be significantly different to the design values. For this reason, the g_mC feedback HDA was not used on the ICs fabricated for the work described in this thesis.

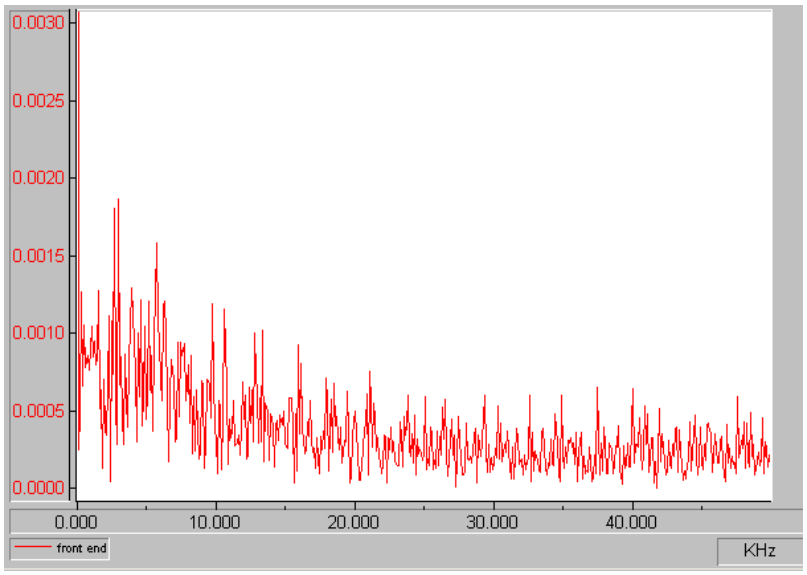

Figure B.3: DOP2 IV output with off-chip filters (y axis units are Volts)

Figure B.3 shows the output spectrum from the DOP2 IC front-end taken off-chip before the HDA, using off-chip anti-aliasing filters instead. The signal level is lower as there is no amplification here, however the HDA low cut-off problem is avoided. It can be seen that the spectrum is similar to that observed from the Moor Instruments system, although without amplification the signal magnitude is lower.

B.1.2 CMOS Inverter Front-end Tests in PC3 IC

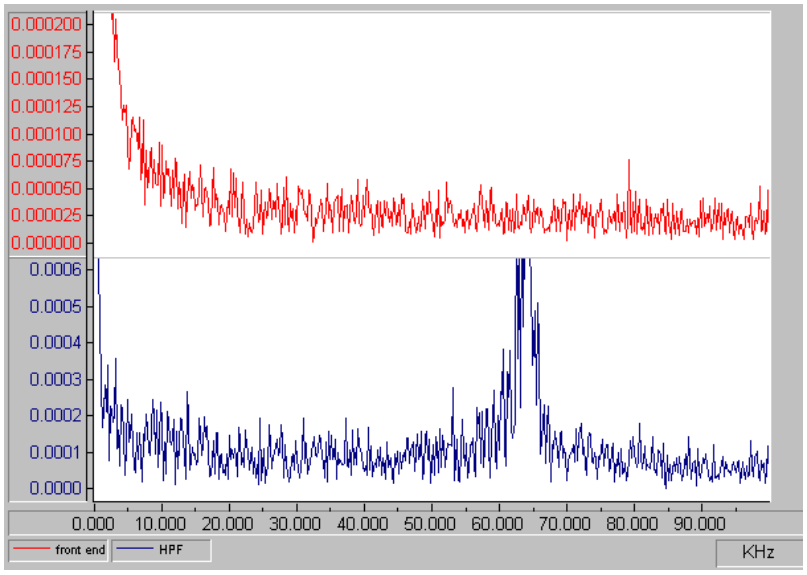

Figure B.4: PC3 IV (top) and HDA (bottom) output (y axis units are Volts)

Figure B.4 shows the output spectrum from the CMOS inverter front-end and the inverted-inverter feedback HDA on IC PC3. The signal appears to have similar shape but, as the output magnitude is considerably lower than expected, this is likely to be due to 1/f noise

rather than a Doppler signal. The signal from the HDA is larger but still not as large as expected, suggesting that this is simply amplified noise rather than the Doppler signal. There is also an unexpected spike at around 65 kHz. This is outside the Doppler bandwidth, but is indicative of a problem with the IC under test. This could have been from a single faulty IC rather than a design flaw, however this IC has not been fully characterised so this is difficult to determine. The smaller diode on this IC means that a lower DC photocurrent would be expected than for the other ICs under test, however the normalising behaviour shown in Section 3.2.1 should mean that a drop in photocurrent due to a smaller photodiode is partially compensated for by an increase in AC gain. This means that a Doppler signal of a similar order of magnitude to that seen on the other ICs would be expected here.

B.1.3 No-feedback Front-end Tests on DOP1 IC

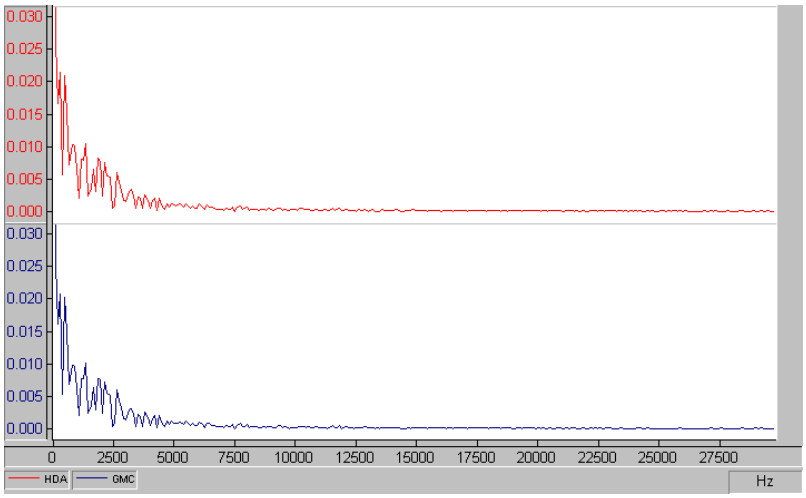

Figure B.5: DOP1 HDA and g_mC output with motility target (y axis units are Volts)

Figure B.5 shows the output from the DOP1 HDA and g_mC when using the motility target. It can be seen that the spectrum shape is similar to that seen from DOP2, but the high cut-off frequency is much lower, at around 5 kHz. This demonstrated the limited bandwidth seen in un-buffered logarithmic pixels as shown in Section 3.2.1.3. The g_mC and HDA outputs appear very similar, which is to be expected as the g_mC has a gain of $\sim$1 for pass band frequencies, and has a cut-off frequency higher than that of the front-end. The g_mC may be having some effect in reducing noise at higher frequencies due to aliasing, however the ADC frequency used here was 100 kHz, considerably higher than the signal bandwidth which should reduce the effect of aliasing.

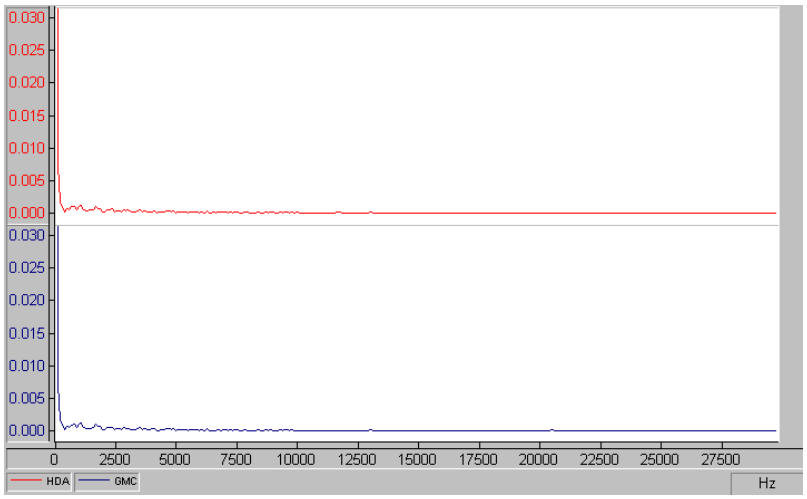

Figure B.6: DOP1 HDA (top) and g_mC (bottom) output with static reflector (y axis units are Volts)

Figure B.6 shows the IC output from the same test configuration as that used for Figure B.5, but with the motility target replaced with a static reflector. This is a smooth cardboard target with colour and finish intended to give a similar DC reflectivity to human skin, but with no Doppler signal. This is used for calibration and noise floor tests. It can be seen that the spectrum becomes much flatter, with only a small amount of noise evident in the spectrum at low frequencies.

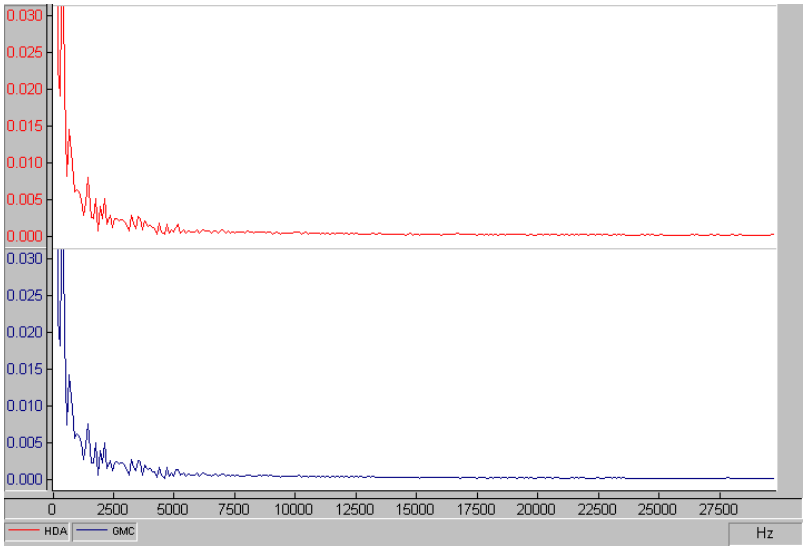

Figure B.7: DOP1 HDA and g_mC output from a finger with normal blood flow (y axis units are Volts)

Figure B.7 shows the HDA and g_mC outputs from DOP1 when used to detect the Doppler signal from a finger. As expected the signal is similar to that from the motility target. However, as this is a test using real blood flow it is possible to observe real-time variations in the signal in response to external stimuli.

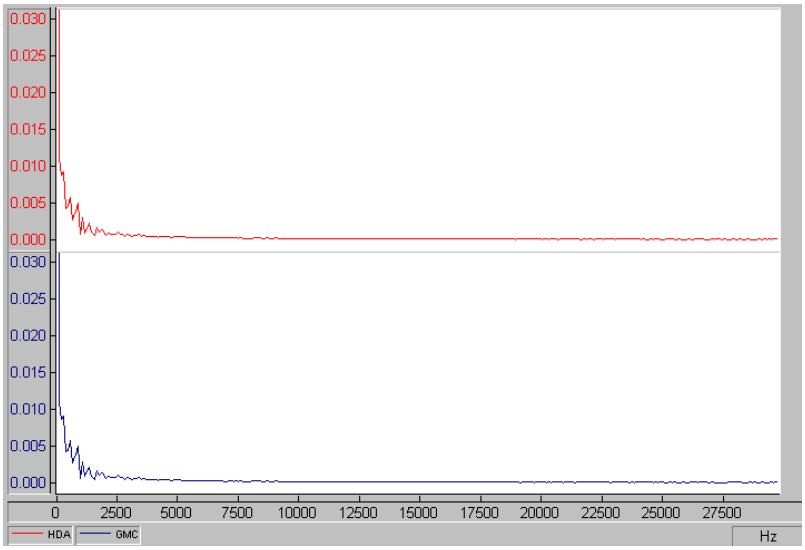

Figure B.8: DOP1 HDA and g_mC output from a finger with occluded blood flow (y axis units are Volts)

Figure B.8 shows the output spectrum from DOP1 from a finger with occluded blood flow. This was achieved by using a blood pressure cuff on the arm of the test subject. A clear difference can be seen between the occluded and free blood flow images, confirming that this type of detector is suitable for blood flow imaging.

B.2 Summary

Tests performed on previously fabricated ICs show that the active load logarithmic front-ends described here are capable of detecting the Doppler signal from blood flow. The need for a feedback buffered front-end has been confirmed by the lower than required bandwidth of the basic front-end. Of the feedback front-ends, the tests on the CMOS inverter feedback front-end were inconclusive, probably due to a fault on the IC. While this design of front-end may be well suited to laser Doppler blood flow applications, that was not confirmed by these tests. For the buffered feedback design, the tests show that the Doppler signal measured from a motility target was as expected.

As well as confirming the choice of the buffered feedback front-end, the tests have shown that the design of the filters and amplifiers after the front-end are critical. The DOP2 IC tests show that while the front-end used on that IC is a suitable design, the design of the HDA used to amplify the Doppler signal is potentially flawed, with variations between the designed and actual cut-off frequencies blocking a large part of the Doppler signal. This suggests that the inverted-inverter feedback design is a more suitable option for the AC amplifier.

Appendix C

Design of a Differential Pixel for Laser Doppler Flowmetry

C.1 Introduction

All of the detectors shown in this thesis are single-ended detectors, that is the signal is measured between one point and ground. This is a common approach, and has advantages in terms of simplicity of circuit design and data acquisition. An alternative approach in various areas of electronics is to use differential signals, measuring between two non-ground nodes.

This appendix describes the design of a detector that applies this approach to LDF. This entails measuring the photocurrent from two separate photodiodes, then measuring the difference between the two. A pixel is designed and simulated using the same CMOS technology as the previous ICs. While this design is a single pixel prototype, it is designed with scaleability in mind, such that it is suitable for use in an array system.

This work is included as an appendix only, as no testing was performed on this design, and the approach taken is significantly different to the other prototype ICs.

C.2 Advantages of Differential Measurement

The reason that differential sampling is of interest for LDF is its potential for noise reduction, and therefore an improvement in signal to noise ratio. Improving this allows for better flux resolution, and ultimately allows smaller changes in blood flow to be accurately measured. The application of differential measurements to LDF is not new, having been used in some early LDF systems[Nilsson et al., 1980] for similar reasons that make the technique potentially useful here.

The main improvements are in terms of Common Mode Rejection Ratio (CMRR) and Power Supply Rejection Ratio (PSRR). Common mode noise refers to sources of noise that are present on both inputs to a differential detector. For early LDF systems the main source of CM noise was laser fluctuations - as the laser fluctuations generate the same signal on both channels, measuring the difference between the channels means that these changes are largely suppressed.

Power supply rejection refers to the ability of circuits to stop noise on the power supply appearing at the output. While steps can be taken to reduce noise on power supplies (separate

analog/digital power supplies, bypass and decoupling capacitors on the power supply on circuit boards etc), some noise is still to be expected. This means that a circuit with superior PSRR will have lower noise at the output than a circuit with poor PSRR.

Differential detection is not in common use in current commercial systems due to improvements made in various areas since early systems. Laser fluctuations were a major source of noise [Nilsson et al., 1980, Sargent and Scully, 1972], but increased laser stability has made this less of an issue. Improvements in other areas such as lower noise op-amps and improved analogue-digital converters also mean that signal-noise ratios are generally sufficient with single-ended detection. However, not all of these developments can be applied to on-chip integrated systems.

C.2.1 Power Supply Rejection of Logarithmic Pixels

The logarithmic pixels used here have been shown to have considerable advantages over opamp based circuits in terms of silicon area required. However, their simplicity means that they have a poor PSRR. As the circuit is effectively a resistor in series with a reverse biased photodiode, the current through the resistor cannot change, as this is set by the photocurrent. The voltage drop across the resistor is therefore constant, so any perturbation of V_{DD} will appear at the output. This means any noise signal on the power supply will appear at the output with very little attenuation.

Figure C.1 shows a simulation of power supply noise applied to a logarithmic pixel of the type used on the previous BVIPS ICs. The bottom plot shows the noise on the power supply, a 10 mV peak-peak signal at 10 kHz. The middle plot is the input photocurrent, a 1 nA DC signal with a 150 pA modulated component at 1 kHz, representing a Doppler ratio of 15%. This is a lower DC photocurrent than the 10 nA used for typical DC photocurrent for BVIPS2 simulations, as the differential pixel has smaller photodiodes due to the requirement for two photodiodes per pixel, and to keep overall circuit size small in order to allow prototyping on un-used space of an IC built for a different project.

The top plot of Figure C.1 shows the output from the front-end. It can be seen that the noise signal is present in the output, and has an amplitude comparable to that of the Doppler signal itself. The power supply noise is attenuated slightly, with a 10 mV power supply noise signal resulting in an 8 mV output noise signal. However the expected output signal at this point is less than 10 mV, so this represents a significant noise source. The noise signal is also within the bandwidth, so will be amplified rather than rejected by subsequent signal conditioning circuitry.

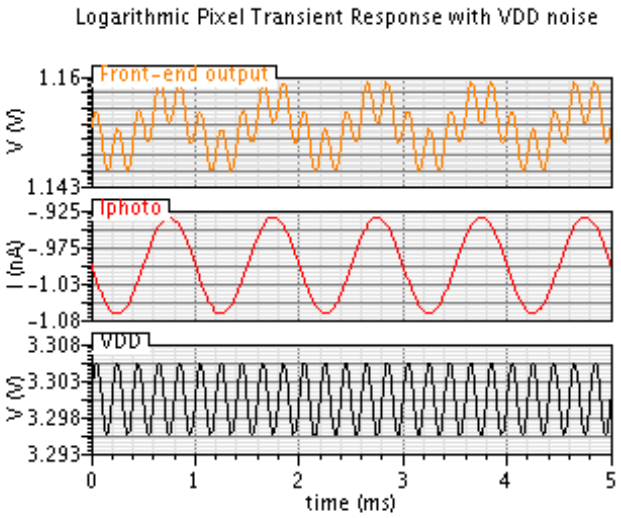

Figure C.1: Photocurrent, power supply (with noise signal) and output signal from a logarithmic pixel

Figure C.2 shows a similar situation on an opamp-based single-ended pixel. Again, the bottom plot shows the power supply, with the same noise signal. The middle plot shows the photocurrent, which is the same as that used for the logarithmic pixel simulation. The top plot shows the output signal, which does not show any of the noise signal applied to the power supply.

This is due to the common mode and power supply rejection of the opamp. While this circuit is single-ended, the opamp itself uses elements of differential detection, as it amplifies the difference between the inverting and non-inverting inputs. As the opamp based circuit is in common use in commercial systems (such as those made by Moor Instruments), this represents a significant disadvantage of the logarithmic pixels used here.

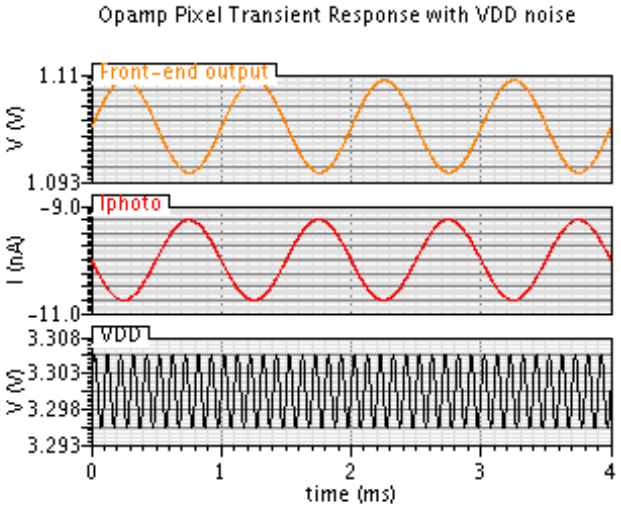

Figure C.2: Photocurrent, power supply (with noise signal) and output signal from an opamp pixel

C.2.2 Common Mode Rejection of Logarithmic Pixels

Both the logarithmic and linear (opamp) pixels shown so far cannot reject common mode signals, as the term is meaningless for detectors with a single input photocurrent. Differential detectors use a pair of photodiodes to give two input photocurrents. There are then various approaches to making a differential detector. The front-end itself could measure the difference between the two photocurrents and output a single-ended voltage. Alternatively, both photocurrents can be converted to voltage in the same manner as for single-ended systems, using the same logarithmic front-end circuit as that used for the BVIPS1 and 2 ICs. The difference between these two output voltages can then be amplified and measured. This approach is more easily applicable here, as it can use the same current-voltage converter circuits as the other BVIPS ICs. A potential drawback is that the initial stage is still single-ended, which might limit the common mode rejection ratio. Ideally, both front-ends will be subject to the same sources of noise, which will be present as duplicate signals on both of the initial front-end outputs. The following differential stage will then remove these common mode signals. However, any mismatch between the two front-ends will mean that the noise signals will not be the same and so will not be removed as effectively by the differential stages.

There are two alternatives for sampling the signal from such a system. Firstly, the difference voltage can be amplified and converted into a single-ended output voltage. This can then be sampled in the same way as the output from the BVIPS1 and 2 ICs. The single-ended signal will be subject to the sources of common mode and power supply noise that this approach is intended to suppress. While this is not ideal, the signal at this point is amplified such that it less susceptible to additional noise than the front-end output.

Another approach is to use a differential ADC, allowing the analogue signal path to be fully differential, minimising the noise added by all amplification stages and interconnects. This should result in the best signal-noise ratio, although it does require that the output is measured using a differential ADC, as well as requiring additional circuitry such as doubling the number of buffers used to drive off-chip outputs, which are often very large due to the need to drive high capacitance loads.

C.3 Differential Detection of a Doppler Signal

Differential detection as used here can seem counter-intuitive, as each pixel is closely located and therefore exposed to the same light level. This would appear to suggest that differential detection will cancel out the signal from each photodiode. The differential technique relies on the speckle size produced by the target and imaging optics being smaller than the photodiode size. The modulated signal detected is effectively a change in the averaged intensity of all speckles. As this change depends on scattering through tissue and and the flow of blood within that tissue (see Chapter 1), the speckle pattern is random, such that the signals produced by separate speckles are uncorrelated, ie having unrelated phase, frequency and amplitude.

This means that a pair of photodiodes that are physically close such that they have a Doppler signal giving equal values of flux do not have the same input signal. While the spectra of the output signal from both detectors may look the same in the amplitude domain, their output in the phase domain may be completely different. This means that a differential signal can be detected.

C.4 On-Chip Differential Pixel Design

C.4.1 Block Diagram of Pixel

Figure C.3 shows the design of the pixel fabricated for this project. It uses a pair of buffered logarithmic pixels, the outputs of which then drive a differential load. It is this stage that will perform the majority of the common mode signal rejection. The outputs of this stage are themselves connected to a differential gain stage, which is used to ensure the effectiveness of the common mode rejection of the differential load, as well as providing additional gain.

Anti-aliasing filters are still required here, so a pair of g_mC filters are used. These are the same as those used in the BVIPS ICs. Ideally this stage would use a differential filter, but at this stage the signal has already been amplified to be $>100\,\text{mV}$, and the g_mC filter does not add large amounts of noise, partly due to having a unity gain. The use of existing filters also reduces the time required to design the circuit.

After the filtering there is a final differential buffering stage, which should remove any common mode noise added by the single-ended filters. Finally, two separate buffers are used to drive the off-chip outputs. These are the same op-amp based buffers as those used to drive off-chip signals on BVIPS1 and 2.

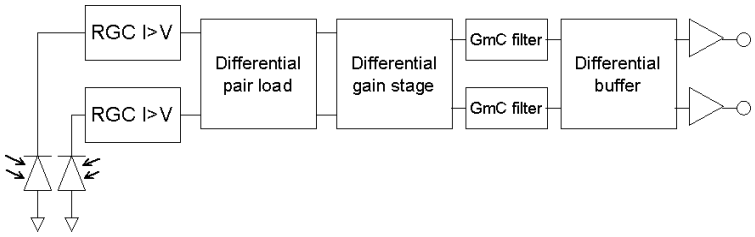

Figure C.3: Block diagram of the on-chip differential pixel

The photodiodes are implemented using four separate photodiodes, allowing a common centroid design [Allen and Holberg, 2002] to be used. This connects the diagonal pairs of diodes in parallel, such that process variations during IC manufacturing affect both diodes equally. This is required to ensure matching between the two channels. If one channel had a larger signal, the change in the operating point of that channel would mean that common mode sources of input noise would be greater on the output of one channel than the other, reducing the effectiveness of the noise cancellation.

This approach is also used by some of the following circuitry. The technique is common for differential pairs, where pairs of transistors have to be closely matched for correct operation. If transistors in a differential pair do not match, the current steering principle does not work as intended, causing distortion of the output signal or reduced gain [Allen and Holberg, 2002].

A potential weakness of the design is that common-centroid design of the pairs of single-ended front-end circuits is not easily achievable. In this case, it may be possible to combine the layout of two front-ends, interleaving the designs in a similar manner to common centroiding of transistor pairs. However, as this circuit is not dependant on close matching of a pair of transistors, it may be that this is not critical for correct operation.

C.4.2 Differential Load Stage

Figure C.4 shows the design of the differential load stage. This is a standard differential pair, resembling the OTAs used in other circuits described in this thesis, but does not have

the connection between the gates of the load transistors MP1 and MP0, which is required for differential inputs and differential outputs. The low width/length ratio of MP1 and MP0 means that these transistors have relatively high AC impedance, increasing the voltage gain of this stage.

The left hand section of the schematic sets the bias current through the main load differential pair. This would not be part of the pixel in an array system, only being required once per IC, with common bias voltage connections to all pixels. A cascode bias configuration is used here, which means the bias current is less susceptible to AC signals on the circuit inputs. Keeping this bias current constant ensures reliable rejection of common mode signals.

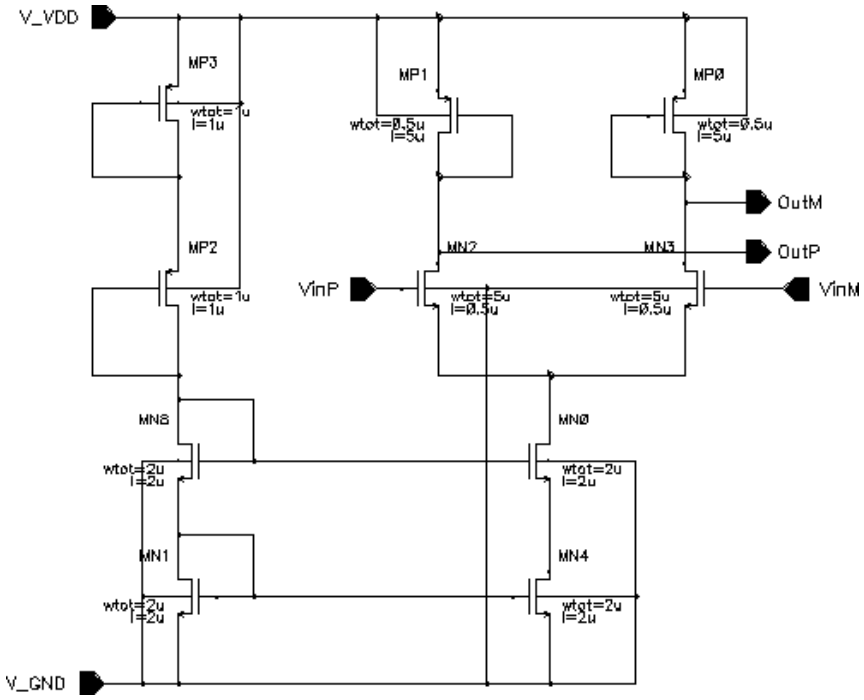

Figure C.4: Differential load stage

C.4.3 Differential Amplifier

Figure C.5 shows the design of the differential amplifier used here. This is based on a circuit designed by Dmochowski for a general purpose modulated light camera using the same type of buffered logarithmic pixels as those used here[Dmochowski, 2006]. The amplifier produced by Dmochowski did not give fully differential outputs, although it was itself based on a fully-differential amplifier. That design was larger and more complex, so the non-differential design was modified for use here. The circuit is based on a standard differential pair, with a pair of current mirrors added (transistors MP3, MP4, MN2 and MN3) used to increase the voltage gain.

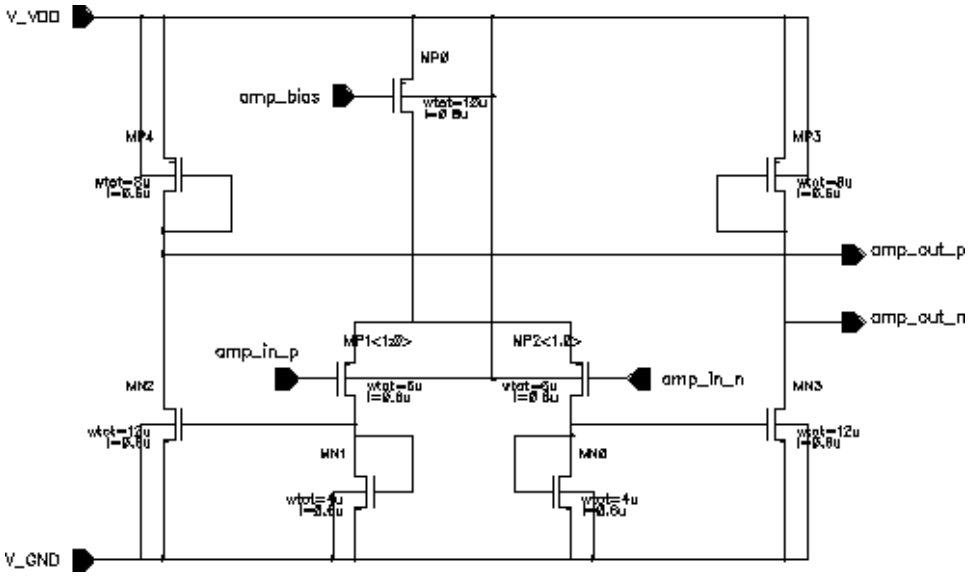

Figure C.5: Differential amplifier stage

C.4.4 Differential Buffer Stage

Figure C.6 shows the schematic of the differential buffer. This is not intended to provide any gain, but isolates the output of the preceding g_mC filters from the large input capacitance of the output buffers. The buffer is designed to have a near unity open loop gain, as this cannot be ensured by using negative feedback due to the requirement for differential inputs and outputs. This is done by setting the sizes of the input and load transistors.

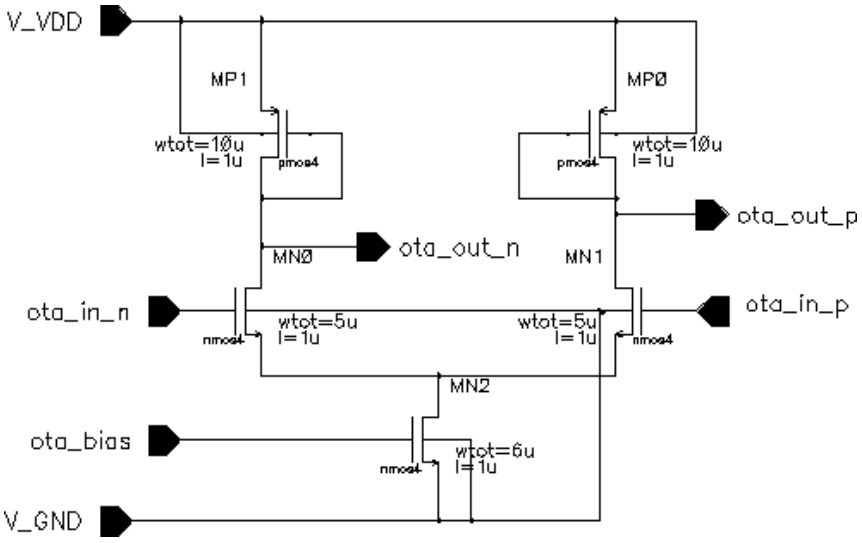

Figure C.6: Differential buffer stage

C.5 Layout of the Differential Pixel

To demonstrate the size of the circuit elements used above, the layout of the pixel is shown in Figure C.7. The photodiodes are shown on the left hand side of the plot, with diagonal

pairs being connected together for common centroiding. Each of the four photodiodes has an active area of 25 μm × 75 μm. The large rectangular devices on the right of the plot are the compensation capacitors in the opamps used to drive off-chip signals, the active circuitry of which is between the capacitors and the main power supply track down the middle of the plot.

The active circuits in the log pixel are in a column between the photodiodes and the power supply track, in an area roughly 25 μm × 150 μm (not including the photodiode). This is suitable for an array system, comparing well to the pixel design used in BVIPS2 which has a pitch of 50 μm and are 250 μm long, although roughly 250 μm of that length is the duplicate front-end/amplifier circuits used for FPN reduction.

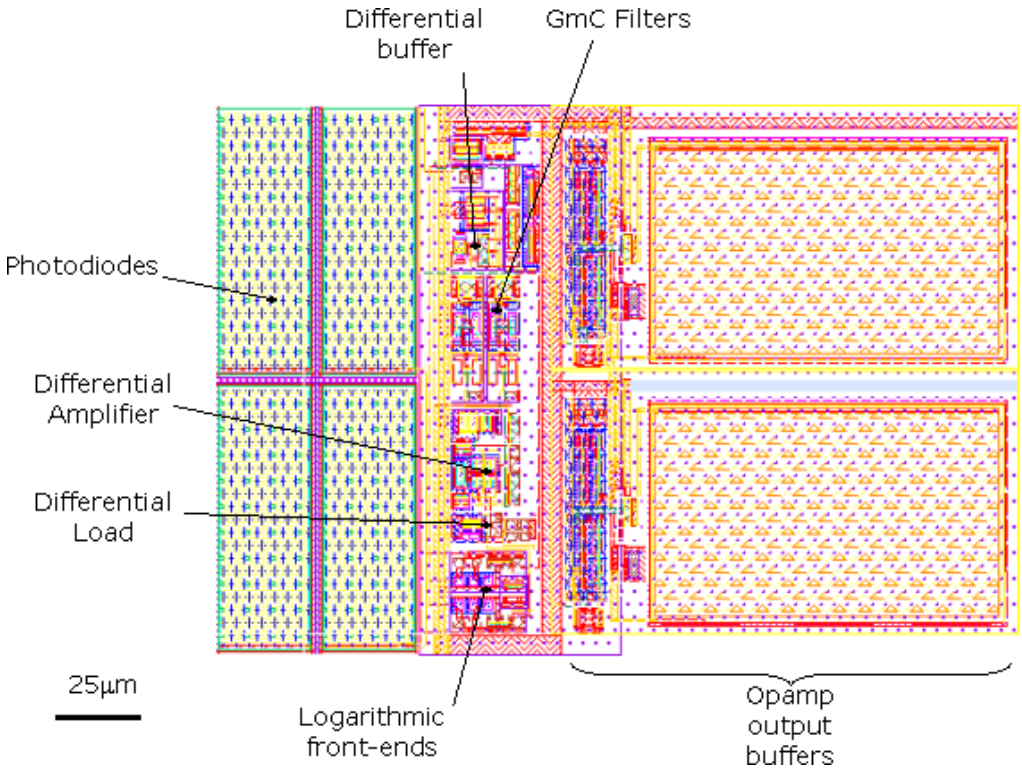

Figure C.7: Layout of the differential pixel

C.6 Simulation of the Differential Pixel

C.6.1 Transient Simulation

C.6.1.1 Ideal Case - No Noise

Figure C.8 shows output signals from the differential pixel with a 1 nA input photocurrent on each pixel, with each photodiode having a 1 kHz, 150 pA modulated signal. The AC component input to each channel has opposite phase, representing the signal described in Section C.3.

The top axes show the output signals from both N- and P-channel front-end circuits. In this situation, with no added noise, these signals are the same as seen for the single-ended circuits.

The lower axes show the outputs after differential load and gain stages, filters and buffers. The signals are as expected, with both channels having a similar signal but with opposite phase. The signals have also been amplified sufficiently for sampling with a standard ADC, with 100 mV amplitude on both channels.

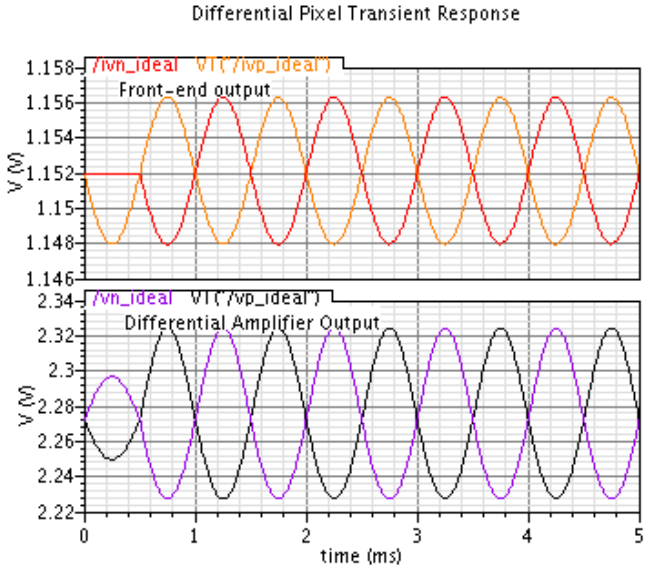

Figure C.8: Output signals from differential pixel, ideal case

C.6.1.2 Noise on Power Supply

Figure C.9 shows the signals as above, but in the case of a noise signal being present on the pixel power supply. This can clearly be seen in the top plot, and some ripple is still evident in the lower plot. However, the main signal is dominant on both channels, and differential sampling of these channels would remove the majority of the remaining ripple.

The power supply noise shown here is larger than that expected, as this level of noise would make successful detection of Doppler signals as shown in previous chapters very difficult. This suggests that single-ended sampling of a single channel may give enough power supply rejection to improve the signal-noise ratio. This provides an extra option for data acquisition from this design.

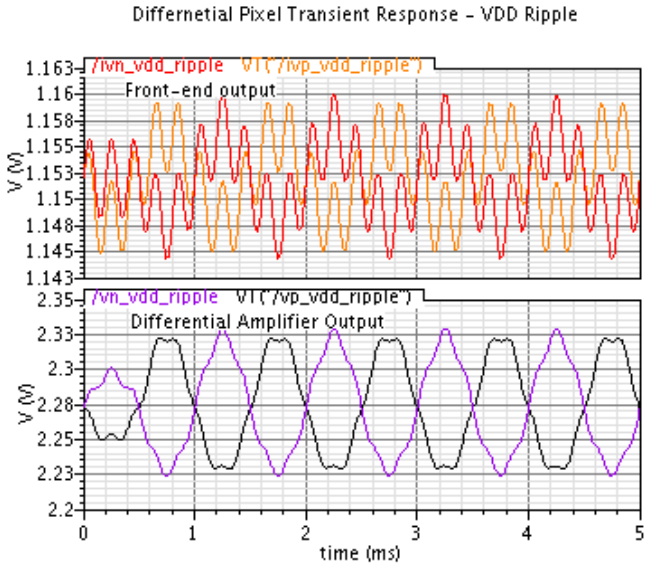

Figure C.9: Output signals from differential pixel, noise on power supply

C.6.1.3 Common Mode Photocurrent Noise

Figure C.10 shows the same signals as Figure C.9, but in the case of a common mode current noise signal being added to the input of both channels. Again, the top plots show the effect this has on the single-ended front-end outputs. The noise signal here has the same amplitude as the power supply noise added above, but at a higher frequency. Despite the noise amplitude being the same, the outputs of the differential pixel show much lower additional noise than in the power supply noise case. This suggests that the pixel will reject noise sources such as laser power fluctuations effectively.

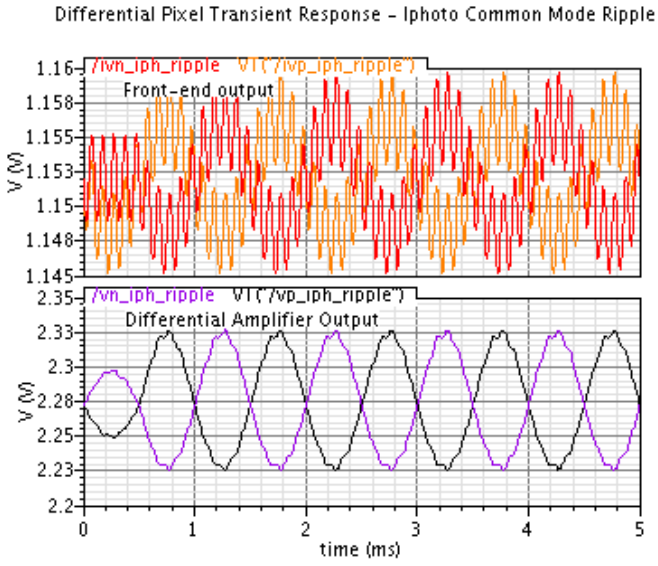

Figure C.10: Output signals from differential pixel, common mode noise on photocurrent

C.6.1.4 Common Mode Photocurrent and Power Supply Noise

The final case of noise sources being present in the photocurrent and on the power supply is shown in Figure C.11. In this case the main 1 kHz signal is masked by several noise sources, being difficult to identify on the single-ended outputs. The differential outputs show a significant amount of noise, but compared to the original signals the expected 1 kHz output is very clear. Again, this is without differential sampling, which will remove some of the remaining noise.

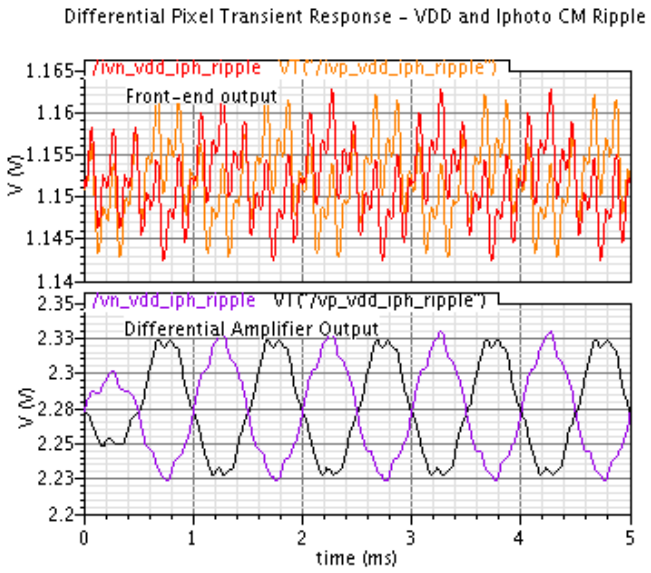

Figure C.11: Output signals from differential pixel, power supply noise and common mode noise on photocurrent

C.6.1.5 Differential Sampling of Pixel Outputs

The differential pixel works most effectively when using a differential ADC, which measures the difference between the two channels rather than the absolute value of either channel. The effect of this sampling is shown in Figure C.12, which shows the difference voltage between both channels for all cases shown above. It can be seen that there is very little noise in either case.

Figure C.13 shows a close-up of the signals at 2.25 ms. This plot shows that power supply noise is blocked more effectively than common mode photocurrent noise. This might be because the effect of photocurrent noise on the output signal will vary as the signal current changes. This is similar to a change in DC bias point - at higher photocurrents (the trough of the output sine wave), the AC gain is slightly lower, so the input noise signal causes a slightly smaller output noise signal. As the difference signal is being measured between the signal peak on one channel and the trough on the other, the mismatch of noise amplitude on each channel will cause a small amount of noise signal to remain. While the noise rejection is not perfect, the CMRR is still a considerable improvement over the single-ended detector.

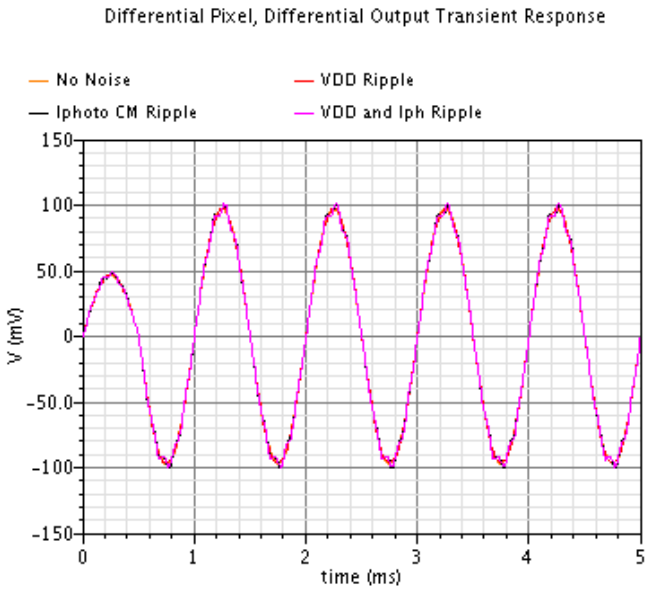

Figure C.12: Output signals from differential pixel, measured differentially, with various noise sources

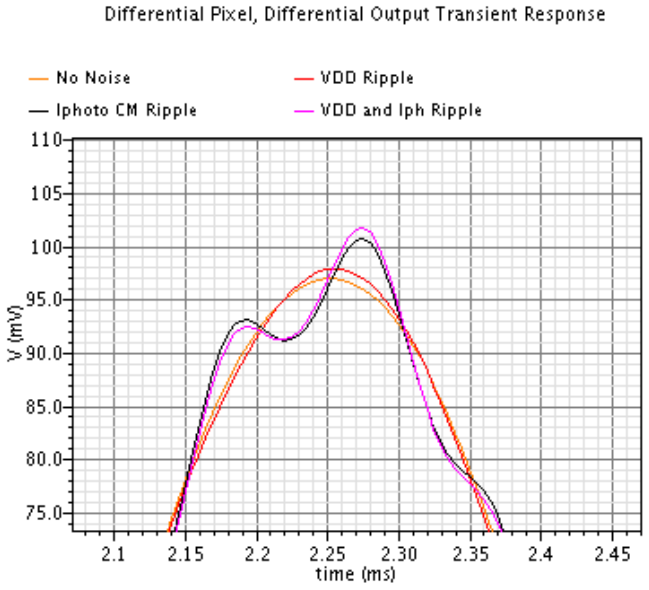

Figure C.13: Close-up of output signals from differential pixel, measured differentially, with various noise sources

C.6.2 DC Simulation

Figure C.14 shows the DC response of the differential pixel. Both channels have the same DC response. This is to be expected, as while both front-ends have the same DC response as the standard single-ended log pixels, the differential stages treat the change in DC output voltage as a common mode signal, and hence reject the input level. The DC response is therefore fixed at around 2.2 V, with no significant dependence on photocurrent.

This means that the circuit used here cannot be used to obtain a DC image. This means that normalisation cannot be performed off-chip, as no separate measure of DC is available. The inherent normalisation that occurs in logarithmic pixels should mean that this is not an issue. An additional problem may be that thresholding cannot be performed to identify image foreground/background.

While this would be a significant problem in an imaging system, for an initial test of this circuit in a single point system this is not a major issue. If a DC channel was required then a separate output from one or both of the front-end circuits could be taken to provide a DC channel. This is one advantage of using separate detectors, as after the front-end buffers the signals can be connected to multiple processing circuits without reducing the signal levels at either - this contrasts with a differential front-end stage, where such an approach may require the photocurrent to be split between several loads.

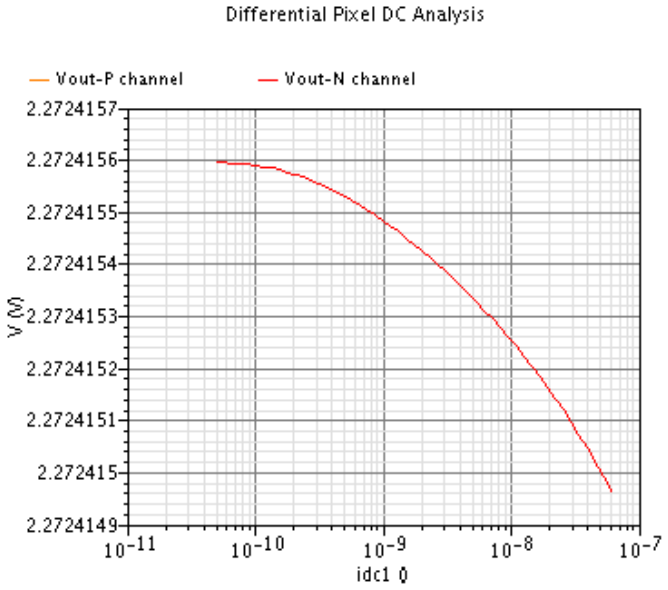

Figure C.14: DC response of both channels on differential pixels

C.6.3 AC Simulation

Figure C.15 shows the frequency response of both channels of the differential pixel. Both channels have the same gain and bandwidth, and the bandwidth is over 100 kHz. This can be reduced to the 20 kHz bandwidth used elsewhere, as the low-pass filter used is the same as in the previous ICs. This simulation shows that the differential circuit designed has sufficient bandwidth for the LDF application.

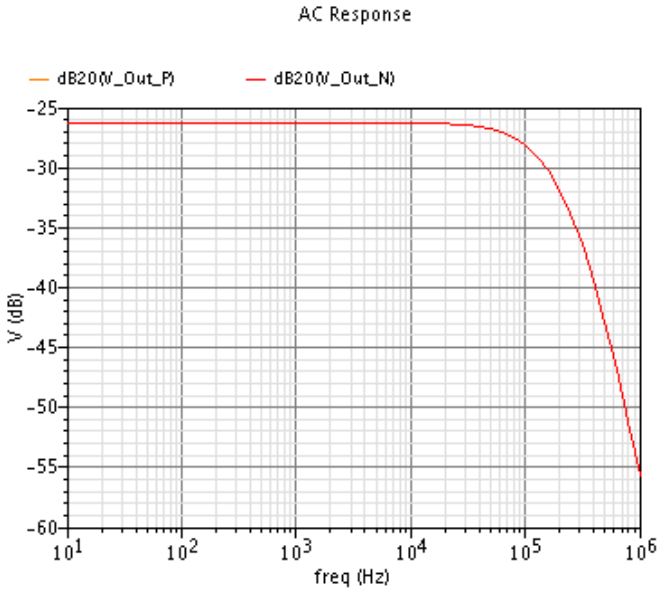

Figure C.15: AC response of both channels on differential pixels

C.6.4 Noise Simulation

Figures C.16 and C.17 show the input-referred and output noise spectra for the differential pixel. The input referred spectra is similar in shape and magnitude to those seen for the single-ended detectors. This is to be expected given the similarities in some of the major circuit elements. While this circuit topology is intended to reduce noise, this reduction is due to the rejection of external noise sources rather than a reduction in the theoretical noise floor. This means that while simulation is unlikely to show a significant reduction in noise density, the measured results should be closer to the simulation results than for the single-ended pixels. For the single-ended circuits, the noise measurements are generally two orders of magnitude higher than the simulated results.

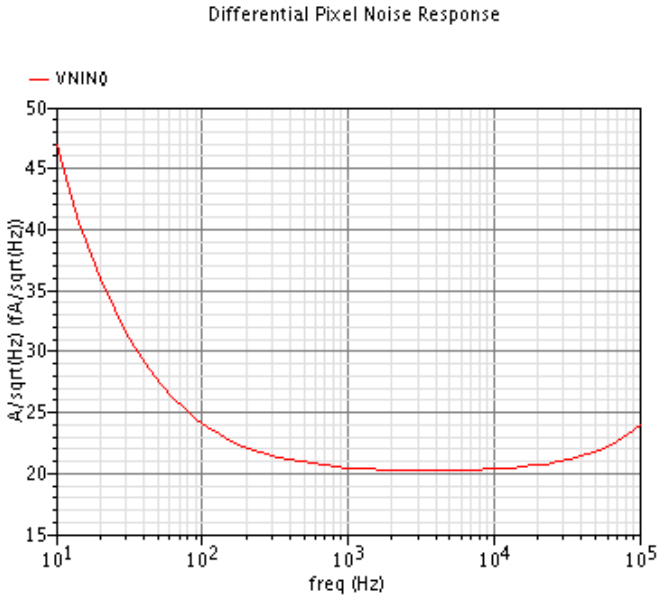

Figure C.16: Input referred noise response of differential pixel

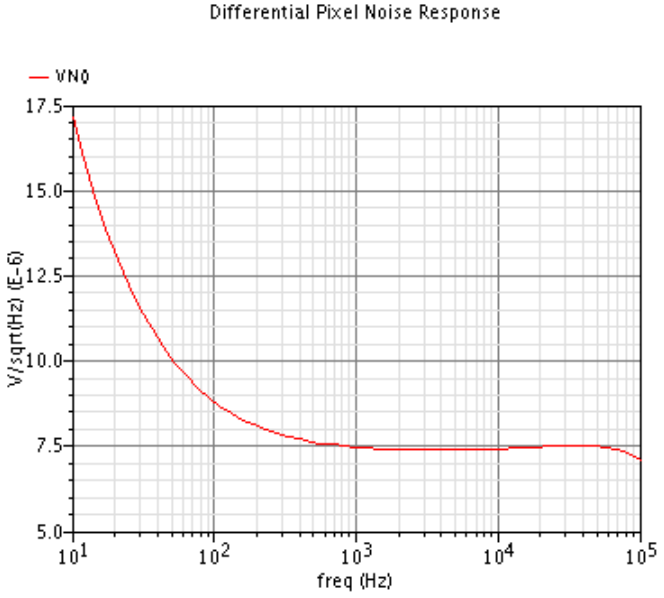

Figure C.17: Output referred noise response of differential pixel

C.7 Effect of Mismatch on a Differential Detector

The results shown above suggest that the differential pixel has a response suitable for detecting Doppler signals, and has good common mode and power supply noise rejection. However, the simulations shown assume that the photocurrent is equal in both channels. Problems with fixed pattern noise have been shown in previous chapters, so the effect of FPN on this pixel design should be considered.

This has been done in simulation in two ways. Firstly, simulations are performed with a specific mismatch set, in this instance by increasing the photocurrent in one detector by 20%. Secondly, Monte-Carlo statistical analysis is performed to test the circuits performance when subject to random variations in device size and behaviour.

The former simulations replicate the effects photodiode variation causing un-even DC and AC photocurrents, and therefore uneven gain of the logarithmic front-end stages (which have a gain set by the DC photocurrent), while the differential aspects of the pixel still operate as expected with no mismatch. The Monte-Carlo simulations vary all aspects of the circuit, giving an indication of the overall effect of mismatch and process variations.

Some of these sources of mismatch could be avoided if an alternative current-voltage converter was used - for example, if a single front-end measured the photocurrent from both photodiodes differentially, then problems caused by uneven gain of the front-end circuits would be reduced. However, the extra time required in the design process to design such a circuit rather than re-using an existing and successfully tested front-end circuit means this option was not explored further here.

C.7.1 Transient Simulation with 20% Photocurrent Mismatch

To verify the basic operation of the detector under these conditions, a transient simulation was performed. Figure C.18 shows the transient response of the differential pixel with input photocurrents having DC levels of 1 nA and 1.2 nA. The output voltages from the front-ends on each channel can be seen to have different DC and AC output levels. The difference voltage between the detector outputs is shown in the top plot. While the noise rejection is not as good as that seen in the ideal case (shown in Figure C.12), the majority of both noise sources is still rejected, giving a fairly clear output signal.

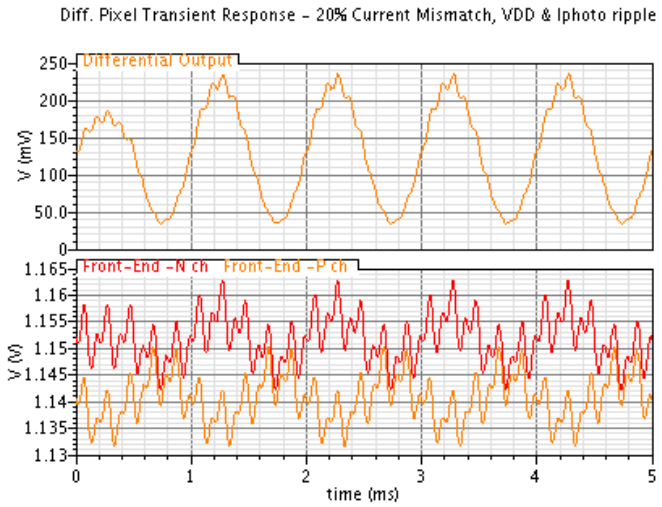

Figure C.18: Output signals from differential pixel, with 20% mismatch in photocurrent, with added common mode and power supply noise sources

C.7.2 DC Simulation with 20% Photocurrent Mismatch

Figure C.19 shows the DC response of the pixel with the same 20% current mismatch. This shows that when the input currents differ, the output voltages from each channel are no

longer equal. The slope on each response with rising photocurrent is also larger, although the change is still much lower than that of a single-ended detector, with an output change of 30 mV as the input rises from 500 pA to 60 nA.

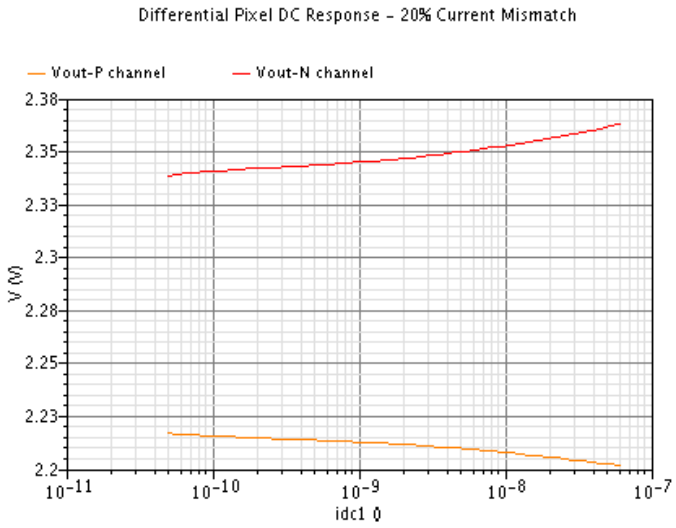

Figure C.19: DC response of both channels of differential pixel, with 20% mismatch in photocurrent

While the DC response has changed compared to the ideally matched case, the change is not significant enough to prevent successful operation of the detector - the DC voltage is still centred on a similar point, and still has little dependence on input photocurrent.

C.7.3 Monte-Carlo Simulation of Differential Pixel

Monte-Carlo simulation is widely used to investigate the impact of manufacturing variation on CMOS circuitry. The effect of these variations is particularly relevant here, as small changes in DC output level or AC gain in the early stages of the detector could cause significant problems in the operation of the later stages. The simulations shown above appear to show that the circuit operates correctly when the photocurrent is significantly different in each channel. However, this still assumes that the differential pairs used at various stages of the detector are perfectly matched. While common centroid design is used to optimise device matching, some mismatch is to be expected, so the design must still operate correctly with non-perfect matching.

The Monte-Carlo tool used can simulate two categories of variation, process and mismatch. Process variation refers to larger scale variation across a wafer during fabrication. This is generally relevant when considering chip-chip variation. Mismatch refers to variation between devices within the same circuit (ie those on a smaller physical scale), and is the relevant source of device variation here.

C.7.3.1 Transient Simulation

Figure C.20 shows the output signal from one channel of the differential pixel for 20 Monte-Carlo iterations. The most significant feature of these plots are the plots at lower DC voltages which have very little or no AC signal at the 1 kHz signal frequency. This would appear to show that this design is not suitable for reliable manufacture. However, the plots

in Figure C.20 only show the output from one of two channels. The detector is designed for differential measurement between two channels.

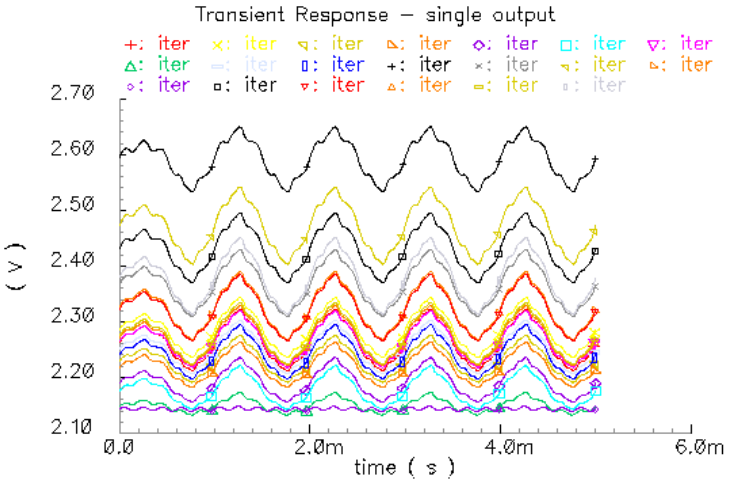

Figure C.20: Transient output from one channel of differential pixel, 20 MC iterations

Figure C.21 shows the differential output voltage from the same simulation. When measured differentially, all iterations show a similar AC magnitude, and all have good power supply and common mode noise rejection. From this it can be seen that while device variation means that the signal from one channel may have very little signal content, this will be compensated for by an increase in gain on the other channel.

This means that using a single channel as a single-ended output may not be a reliable approach, however if this was required a differential-single ended converter could be implemented as part of the output buffering on a PCB for the IC. This could be implemented on-chip, but this function would be best performed by a high quality instrumentation amplifier. There are a range of such devices available if conversion can be done off-chip, whereas an on-chip design would require more silicon area than is feasible for this prototype circuit, as well as adding complexity and time to the design process. This does not contradict the aim of making the design scaleable, as the off-chip instrumentation amplifier could be specified to have sufficient bandwidth for it to be used after an on-chip multiplexer. Alternatively, if on-chip analogue-digital conversion is required, the ADC used in the BVIPS2 IC is capable of differential sampling, but uses a fixed reference for sampling of the single-ended signals present on that IC.

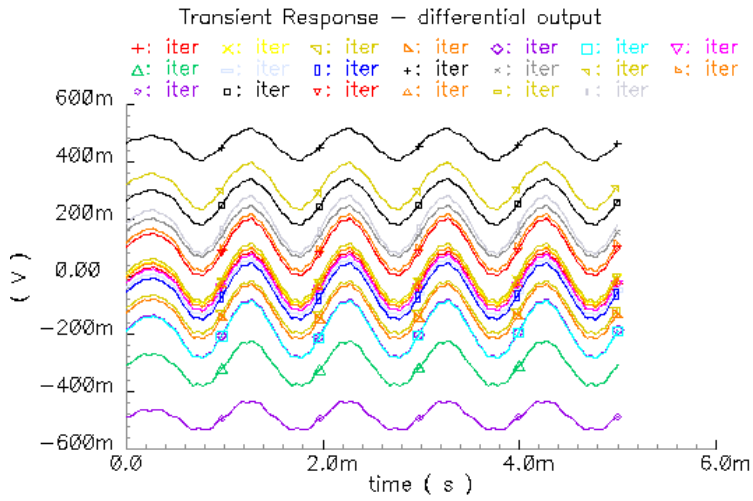

Figure C.21: Transient output, differential measurement of differential pixel, 20 MC iterations

C.7.3.2 DC Simulation

Figure C.22 shows the DC response of one channel of the detector over 20 iterations. The results correspond closely to those shown from the transient response, with the DC levels in each iteration below being the same as that shown in the same iteration in Figure C.20. In all cases, the DC voltage does not vary significantly with photocurrent, and all results are within a range of 0.55 V, which is not large enough to make signal measurement difficult.

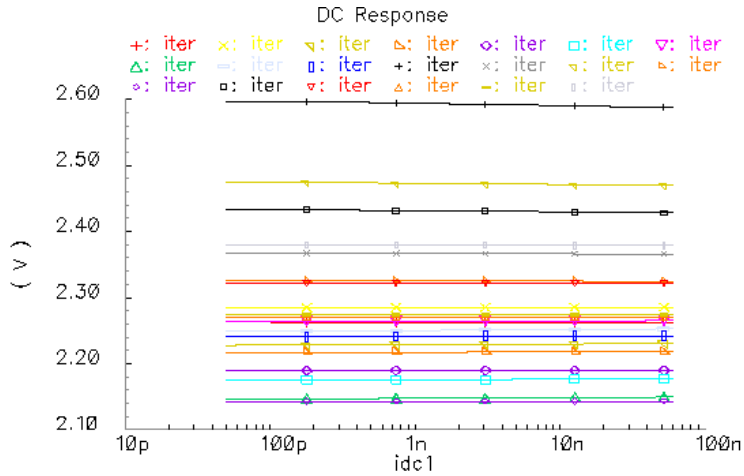

Figure C.22: DC response of differential pixel, 20 MC iterations

C.7.3.3 AC Simulation

Figure C.23 shows the frequency response of one channel of the detector over 20 iterations. Again, the results correspond closely to those seen in the transient simulations. Most iterations give a response similar to the nominal response. There are two iterations that are cause for concern, but these correspond to signals which are shown in Figures C.20 and C.21

to have low AC gain on one channel, but with sufficient gain on the other channel to give a differential output with magnitude similar to the nominal case.

There is relatively little variation in the bandwidth of the circuit, with even the low gain iterations having a bandwidth well over the required 20 kHz.

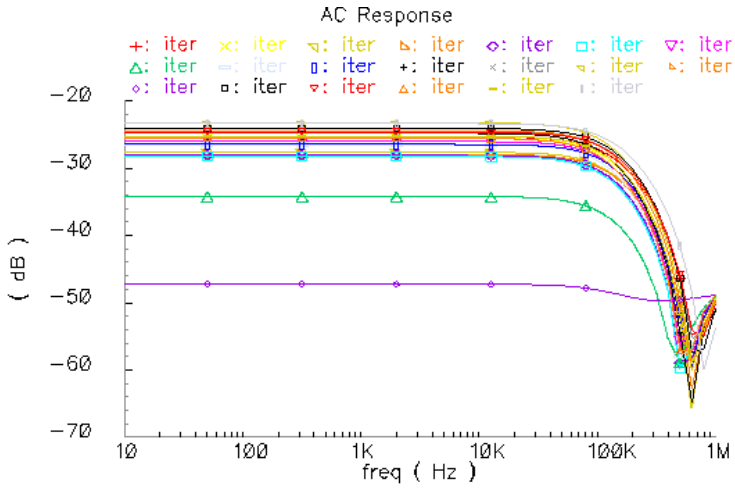

Figure C.23: DC response of differential pixel, 20 MC iterations

C.7.3.4 Variation of Power Supply and Common Mode Rejection Ratio

As well as requiring the basic functionality of the circuits to be unaffected by device variation, it is important that the main advantages of noise rejection are not significantly affected by mismatch. While this appears to be the case from the transient simulations shown, a more precise measure of noise rejection is required to quantify the effect of device variation.

A suitable measure of noise rejection and distortion of the signals is total harmonic distortion. This allows a measure of the signal distortion to be given for each iteration, allowing probability density functions (PDFs) to be shown.

Figure C.24 shows three histograms showing the PDFs of THD for three signals - one output channel measured as a single-ended output, the single-ended output from a single IV, and the difference signal between both output channels.

It can be seen that the noise signals added cause a THD on the first stage output voltage of around 115%, which is as expected given that the noise sources cause an AC noise signal on the output with similar magnitude to the Doppler signal.

After the differential stages, the mean THD has dropped to around 20%, a considerable improvement on the original signals. There is an outlying point with THD of around 80%, which corresponds to the iteration shown to have no Doppler signal component by the transient analysis.

When the outputs are measured differentially, the average THD drops to below 5%, and in this case the THD from all iterations is relatively closely grouped. This again suggests that the design will reject common mode and power supply noise regardless of device variation, and the importance of differential measurement of both output channels has been clearly shown.

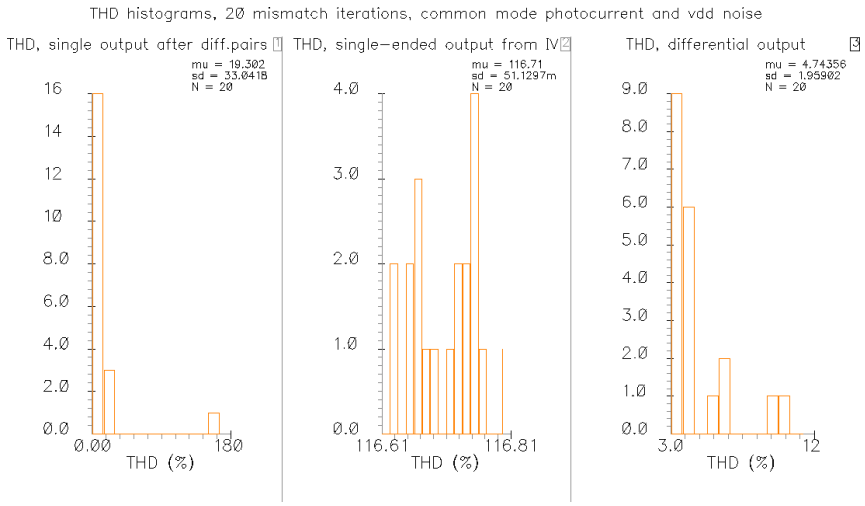

Figure C.24: Histograms showing THD at various points over 20 MC iterations

C.8 Testing Methods for the Differential Detector

The design presented here was fabricated and a PCB was built for testing purposes, however problems in initial testing and debugging meant time did not allow for the intended tests to be performed. This section describes possible methods for testing the differential pixel.

C.8.1 Characterisation of Differential Detectors

Testing of the differential pixel presents additional challenges to the array ICs produced here. The main limitation is that if the pixel shows ideal behaviour, then the output of the circuit when uniformly illuminated with a modulated source will be a fixed DC voltage - the total detector size here is $150\,\mu\text{m} \times 50\,\mu\text{m}$, which is divided into four quarters. To use a modulated light source to characterise the detector requires that the light be focused to a spot which illuminates one quarter of this array (as each half is composed of two photodiodes, one connected to each channel) more powerfully than the rest. This is possible using a high quality lens to focus a laser to a small spot, in combination with a stage to precisely position the IC such that the spot falls entirely (or mostly) on a single photodiode.

Further complications are caused by the need for a reference channel, as the results of the procedure described can only be quantified if the DC photocurrent at each pixel can be independently measured. This could be achieved by using a lower power wide beam (produced from the same laser to give equal wavelength) to uniformly illuminate all photodiodes, the intensity of which can be measured using the same procedure as that used in the characterisation of BVIPS1 and 2 (shown in Chapter 4, Figure 4.8). The intensity of the focused spot can be measured in the same way using a separate beam splitter and photodetector. Both beams are then combined at the IC, producing an area of uniform illumination with a single brighter spot that can be positioned on one of the four photodiodes.

Precise positioning of the focused spot entirely on a single photodiode can be achieved by modulating the laser, and adjusting position until the AC output magnitude is at its greatest - showing the maximum differential photocurrent, which occurs when the additional photocurrent from the small spot is all incident on a single channel.

The overall DC photocurrent level can then be set as for BVIPS1 and 2 characterisation, while a different DC photocurrent can be applied to each channel by varying the power in

the focused spot independently, using neutral density filters. For AC testing, using the same laser for both beams ensures a constant modulation depth, which can be adjusted using neutral density filters to vary the modulation depth if required.

C.8.2 Doppler Testing

Testing using either blood flow in skin or the artificial motility and static phantoms used for testing of BVIPS1 (which closely replicate the situation found during blood flow measurements), results in AC photocurrent components at a wide range of frequencies, which are all incoherent, hence causing a differential photocurrent. This means that testing using these targets can be performed using the same experimental setup as that used for Doppler testing with BVIPS1 and 2, as shown in Chapter 4, Figure 4.31). These tests can be used to compare the performance of the differential pixel with the single-ended circuits used on BVIPS1 and 2, allowing the change in flow readings to be seen between zero flow (static phantom), low flow (occluded finger), normal flow (non-occluded finger) and high flow (motility phantom). The relative change between these readings is a useful figure of merit for an LDBF system, giving an indication of the sensitivity of the pixel to blood flow changes.

C.9 Summary

This appendix has explored the use of an alternative type of detector, using circuit elements taken from existing integrated LDBF sensors. This approach intends to offer improved performance, particularly reduced noise by improved rejection of external noise sources, while remaining suitable for use in integrated LDBF sensors where pixel size must be kept low.

Simulations showed the poor power supply rejection performance of the logarithmic pixels used for the other ICs described in this thesis, and demonstrated the ability of a differential detector to reject this source of noise. Additionally, the common mode photocurrent noise rejection of a differential detector was demonstrated, allowing such a circuit to reject AC photocurrent signals other than those caused by flow in the target.

The design of a CMOS differential pixel was shown, and designs of the new or modified circuits required for this pixel were given. The layout of the complete differential pixel was also shown. Simulations were performed to show the performance of this pixel in the ideal situation in terms of DC, AC and noise response.

The effect of process variations and mismatch on a differential detector were discussed, and simulations were performed to demonstrate that the performance of the design does not deteriorate to an un-usable level in the presence of normal levels of device variation.

Finally, possible methods of testing the pixel were discussed, including the additional challenges presented by the differential detection method. Testing using known illumination to characterise the pixel was considered, along with using the pixel for Doppler blood flow measurements .