Imperial College London

Department of Electrical and Electronic Engineering

Final Year Project Report 2016

Project Title:	**A Low-power Low-noise IC for Electrochemical Impedance Spectroscopy of Biosensors**
Student:	**Kristian Lending**
CID:	**00728904**
Course:	**EEE4**
Project Supervisor:	**Dr Pantelis Georgiou and Dr Sara Ghoreishizadeh**
Second Marker:	**Dr Konstantin Nikolic**

0.1 Abstract

This report describes the principle, design process, and simulation of a low power monolithic Electrochemical Impedance Spectroscopy (EIS) system for application to measuring performance of an amperometric electrochemical sensor. The whole system is designed and simulated using simulink to demonstrate efficacy in principle. A mixed-mode IC is designed to stimulate the sensor in voltage-mode and readout the current. The analogue parts of the circuit are then implemented in Cadence to demonstrate the sensitivity and power consumption of the circuit. In order to be fit for purpose, the circuit can detect impedances with magnitude up to 1GΩ, and up to 100kHz in frequency, while consuming less than 1mW of power from 3.3V power supply. The final circuit has an input dynamic range of 100dB, and can measure across 2 decades of frequency. This is achieved using a translinear current amplifier, with a variable gain of between 0.001 and 1000 at 6 intervals, and an undersampling heterodyne converter, to transform the frequency to a fixed value in order to simplify processing. The report demonstrates mathematically how this signal can be integrated and sampled from this the magnitude and phase data calculated. The implementation of integration and digitisation are combined through use of two integrate and fire circuits with asynchronous up-down counter.

Contents

1 Introduction

With advances in low power, integrated electronics has come the substantial expansion of research into wearable and implantable systems, particularly with applications in bio-medicine[1][2]. To allow for seamless interaction between electronic systems and a patient's biology it is necessary to develop methods for automatic, real-time, in-situ observation[3][4].

One of the most well documented examples of this are attempts to develop artificial pancreases which can supply the insulin requirements of patients with type I diabetes. Currently treatment for this condition requires frequent glucose testing by taking a blood sample, usually through a prick to the finger, and an external device. From this the required insulin dosage can be calculated. This however involves places significant inconvenience, responsibility and discomfort for the patient. Development of an in-situ sensor which allows for continuous glucose readout without the above complications has been discussed since the 1980s[5], however there still remain hurdles to them being completely independent devices.

One of these is that any implanted glucose sensor will gradually degrade over time due to a variety of reasons, such as buildup of peptides and small protein molecules, which limit the surface area of the electrode. Currently the only method of accommodating for this is occasional measurement of blood glucose by traditional means, and comparing the two, limiting the advantages of the new system[6]. Electrochemical impedance spectroscopy (EIS) is a method which can be used to measure the changes in the electrode surface's impedance spectrum, making it possible to track sensor degradation. Although EIS has existed for a number of years and employed for numerous and varied applications, including sensor degradation[7], there currently exists not solution which is suitable to be applied to an implantable system.

Existing systems available rely on external devices, such as signal generators[8], or consumes power beyond that available from the power sources available[9]. It is therefore the aim of this project to design and implement an EIS system which meets the specifications required both to measure the impedance of a glucose sensor (around 1GΩ of impedance up to 100kHz) while allowing it to be implantable (system is monolithic and consumes less than 1mW of power).

This report will be divided into 7 further sections.

- Background: Explains the principle of electrochemical impedance spectroscopy and discusses the application of EIS to measure sensor degradation, using this to outline the basic specification.
- Literature Review: Considers current, similar works and discusses further specification of the final device, relative to them.
- Analysis and Design: Breaks down the objective into constituent problems and discusses the possible solutions to each, with reference to the literature. Completes a top level system design and simulation in Simulink.
- Current Readout Circuit Implementation: Describes the implementation of the analog circuit in Cadence.
- System Simulation and Results: Outlines the process by which the implemented design is simulated with reference to the specifications which need to be fulfilled, and presents the final results.

- Evaluation: Discusses the merits of the design, comparing the results to the specification.
- Conclusion and Further Work: Reflects on the final evaluation of the design and the success of the project, and suggests areas for this work to progress into in the future.

2 Electrochemical Spectroscopy and its Application

Electrochemical Impedance Spectroscopy (EIS) is a method of characterising materials, either organic or in-organic through measurement of their complex impedance.[10] Impedance is the response the electrical current travelling through a material to the potential difference across it, and is therefore defined as $Z = \frac{V}{I}$. By expressing the impedance as a complex number, the frequency effects of the system can be expressed, either in terms of the magnitude ratio and phase difference between the voltage and current ($Z = |Z|\angle\phi_Z = \frac{|V|}{|I|}\angle(\phi_V - \phi_I)$) or as a combination of resistive and reactive parts ($Z = R + iX = \frac{V_{Re}+iV_{Im}}{I_{Re}+iI_{Im}}$). The real-complex and magnitude-phase notations can be interchanged using $|A| = \sqrt{A_{Re}^2 + A_{Im}^2}$ and $\angle A = \phi_A = \arctan\frac{A_{Im}}{A_{Re}}$.

In reality, this impedance will not be fixed for frequency. While ideal resistive components have a fixed resistance, $Z = R$, reactive components have a dependence on frequency, f. The two ideal examples of this are the inductor where $Z = i\omega L$, and the capacitor for which $Z = \frac{1}{i\omega C}$ where $\omega = 2\pi f$. Therefore to gain a full characterisation of a material, it is necessary to measure the impedance across a frequency range.

In biological materials, the impedance spectrum can be associated with a number of possible structures within the tissue. Within a cell, ions carry charge, rather than electrons. When a potential difference is supplied, these ions will flow in the electric field, similarly to electrons in a conductor. The rate at which they do this is dependant on the diffusion constant and concentration of the ions, which can be compared with an equivalent electrical resistance. Where a membrane, which inhibits the flow of the ions, is present, this is similar to a series capacitor, as it will prevent DC flow of the ions, but will pass high frequency AC currents. By this methodology an equivalent circuit model can be developed for any biological sample, however for this the frequency response is needed, and this is the purpose of an electrochemical impedance spectrometer. [10]

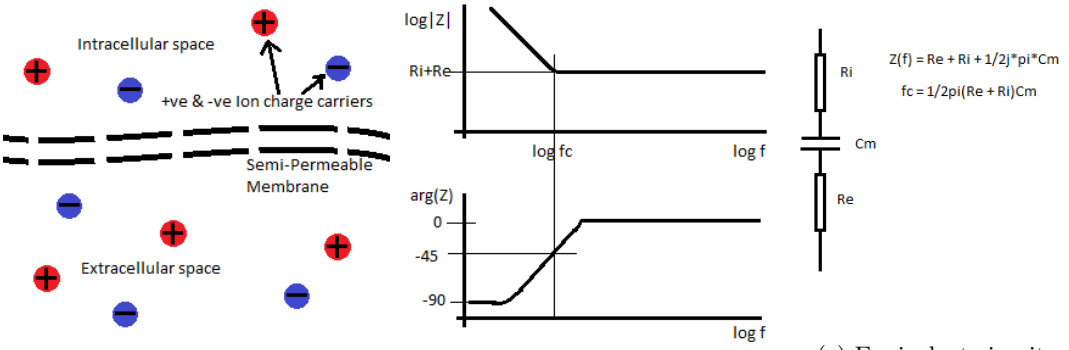

(a) Cell membrane (b) Frequency Response (c) Equivalent circuit

Figure 1: 1a shows a biological structure with a membrane, surrounded by a liquid electrolyte. 1b shows a likely frequency response for this structure. 1c shows a plausible equivalent circuit, where each component is directly associated with each part of the tissue.

EIS has been used in a variety of fields, such as determining the quality of petroleum[11] or characterisation of other oils[12]. In agriculture is has applications in monitoring soil quality[13], as well as root growth and fruit ripening, and others[10]. Specifically in bio-medicine, it has been applied to measuring the molecular characteristics of various parts of the body, from blood[14] to spinal fluid[15] to breast tissue[16], with the objective of detecting or monitoring

disease. It can also be applied to DNA recognition[17], studying proteins[18][19][20], and glucose detection[21][22].

It is particularly useful in the field of implantable glucose sensors. The purpose of these sensors is to measure the blood glucose levels of a patient, most often with diabetes, without requiring the patient to take frequent blood samples. However, during their life time, these sensors are prone to degradation due to the electrodes absorbing peptides and small protein molecules, which pollute them, reducing their surface area. [6] Currently the only method for determining degradation is to take a blood sample and use an external meter, and comparing the two measurements of blood glucose. However, given that degradation of the sensor also affects its impedance spectrum, this is a new possible application of EIS. Naturally measuring the impedance spectrum of an implanted sensor requires a single, low power chip which can perform the entire process without external help, such as from a computer or signal generator. As such the characteristics of the sensor give the basic specifications of this project.

3 Literature Review on on-chip EIS systems

The primary objective of this project is to produce a fully monolithic EIS system, which consumes less than 1mW of power and will measure impedances up to a magnitude of 1GΩ at frequencies up to 1kHz. Therefore these are the primary requirements for the success of the project.

It is important at this stage to consider current work, and the approaches they have taken to achieve similar targets, as well as using them to determine secondary specifications, such as the dynamic range of the system, which would make this project ambitious yet achievable.

The oldest paper considered is [23] which follows the most intuitive method, using a transimpedance amplifier to convert the current response to a voltage, before mixing it with both an in-phase and quadrature wave (which it derives from the stimulus), before applying a low pass filter to each output to retrieve the real and imaginary components of the response. These are presented as the output, from which the impedance can be calculated by taking the reciprocal and multiplying by the stimulus. The block diagram and mathematics for this are shown in Figure 2.

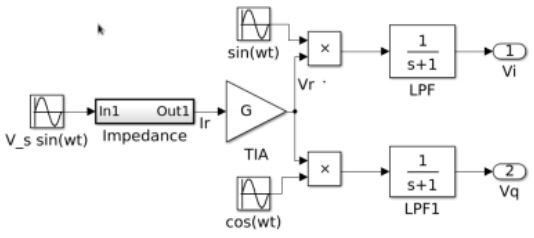

$$
\begin{aligned}
I_r &= \frac{v_s \sin(\omega t)}{Z_r} = i_r \sin(\omega t + \phi) \\
v_i &= G i_r \sin(\omega t + \phi) \sin(\omega t) \\
LPF &\rightarrow \frac{G i_r}{2}(\cos(\phi)) \\
v_q &= G i_r \sin(\omega t + \phi) \cos(\omega t) \\
LPF &\rightarrow \frac{G i_r}{2}(\sin(\phi)) \\
GI_r &= G i_r(\cos(\phi) + i \sin(\phi)) \\
&= 2(v_i + i v_q) \\
Z &= \frac{V_s}{I_r}
\end{aligned}
\tag{1}
$$

Figure 2: Block-diagram and associated arithmetic for a quadrature multiplication scheme

Particularly notable in this implementation, is how the stimulus, in-phase and quadrature signals are all derived from the same $4 \times \omega$ clock, such that all are locked in phase. This enables the system to cover a wide bandwidth from 10Hz to 50MHz. The negative side of this is that a square wave is rich in harmonics, and therefore needs a large quantity of filtering in order to become a useful sine wave. The requirements for a high quality filter with a centre frequency variable across this bandwidth means it must be performed off-chip.

An alternative to this method is given by [9], in which the phase is measured separately using a set of XOR gates, which can compare the response signal directly with the stimulus, therefore removing the need for phase matching. The response is then separately rectified and filtered to give a DC response which is proportional it its amplitude. Although this has been demonstrated to function up to 100kHz, it can only measure phase shifts of up to 90^o. Also interesting is the choice to provide a current stimulus, giving a voltage response. This removes the need for the transimpedance amplifier, and gives a response which is directly proportional to the impedance, simplifying later calculations. It is important to notice here however that this device was only tested on an impedance of $5k\Omega$. For it to measure a $1G\Omega$ sample in a 3.3V regime would require a good quality current mode sinusoidal stimulus of around 1nA which is not discussed in [9].

A second alternative is given by [2]. Here, the same principle is used as in [23], however instead of attempting low distortion mixing, the system integrates the signal and samples it twice in each cycle. This method notes that for a response of $v_r \sin(\omega t + \phi)$, integration will yield $\frac{v_r}{\omega}(\cos(\omega t - \phi) + \cos(\phi)$. By sampling this at half the period, and three quarters of the period we get samples with values $V_{\frac{T}{2}} = -\frac{2v_r}{2\pi f_s}\cos(\phi)$ and $V_{\frac{3T}{4}} = \frac{v_r}{2\pi f_s}(\sin(\phi) - \cos(\phi))$. We can then rearrange these to give values of ϕ and v_r, as in equation 2. Note that these are the magnitude and phase of the current response, and therefore we find $|Z| = \frac{1}{v_r}$ and $\angle Z = -\phi$.

$$\begin{aligned} \phi &= \arctan \frac{V_{\frac{T}{2}} - 2V_{\frac{3T}{4}}}{V_{\frac{T}{2}}} \\ v_r &= \pi f_s \sqrt{V_{\frac{T}{2}}^2 + (V_{\frac{T}{2}} - 2V_{\frac{3T}{4}})^2} \end{aligned} \tag{2}$$

There are two primary disadvantages to this. Although it removes the need for high quality mixers to prevent distortion, it still requires accurate synchronising of the samples to the stimulus frequency. This is not too difficult if the stimulus is generated by a DAC, however at higher frequencies it becomes difficult and will introduce errors. The work itself only demonstrates functionality up to 8kHz. The other difficulty is the introduction of an integrator means that any small offset in the signal could result in it saturating. Even if the integrator were made leaky, the offset would still be amplified, resulting in incorrect mathematics.

At this point a helpful contribution comes from [24]. Noting the challenges discussed with a large frequency range, as well as others including a large power demand arising from high bandwidth, the idea of a heterodyne converter becomes interesting. As is demonstrated in the paper, if the response signal at frequency ω_s is mixed with a signal at $\omega_s + \omega_i$ we receive a signal with components at a sum and difference of these. This can be filtered to give an output with the same phase and magnitude, but with its frequency transformed to ω_i (see Equation 3 for the complete arithmetic). By fixing the frequency, it becomes very easy to filter out noise, as well as simplifying synchronising any further sampling required by the system.

$$\begin{aligned} &\rightarrow v_r \sin(\omega_s t + \phi) \times \sin((\omega_s + \omega_i)t) \\ &= \frac{v_r}{2}(\cos(\omega_i t + \phi) + \cos((2\omega_s + \omega_i)t + \phi)) \\ LPF &= \frac{v_r}{2}(\cos(\omega_i t + \phi)) \end{aligned} \tag{3}$$

A second contribution made by [24] is the moving of the transimpedance amplifier to after the mixing, using a front end comprising of a current mirror with gain x1 or x100 such that the smallest signals can be amplified with the addition of very little noise before anything else is done. As is pointed out, a direct trans-impedance amplifier usually requires a large power to settle the capacitance at the working electrode, whereas this implementation removes that by using a current buffer.

Although [24] provides some useful suggestions, it does not claim to be a complete system, only a front end to an EIS system. It digitises the data after heterodyne conversion and amplification, applying a digital bandpass filter externally to the chip. The lowest frequency resolvable

in a heterodyne system is determined by the bandwidth of the filter, as it is necessary to maintain the signal at ω_i while rejecting the spurious component at $2\omega_s + \omega_i$. This method allows the work to achieve a lower frequency bound of 10Hz by applying a filter with bandwidth <1Hz, implying a Q-factor (for stated intermediate frequency of 2kHz) greater than 2000. Clearly this is not a monolithic solution and it seems implausible that an equivalent on-chip solution will be possible, however a similar solution with relaxed requirements may be achievable.

By careful choosing of the intermediate frequency, ω_i the advantages of this can be optimised. On the one hand, low frequencies suffer from flicker noise, although this generally becomes negligible above 1kHz or so. Furthermore, filters with a lower cut-off or centre frequency typically require larger passive elements, making them consume more area. On the other hand, for a given stimulus frequency, the required Q-factor of the filter is proportional to ω_i, making the the filter harder to design, while other circuit elements must consume more power to achieve the required bandwidth. [24] chooses a compromise of 2kHz.

The question remains as to how best achieve this filter. One possible filter for this could be provided by [25]. This work describes a 4th order filter, which could be set to a frequency of, for example, 2kHz, and have a Q-factor of up to 50 - i.e. giving a theoretical minimum bandwidth of 40Hz. The successful employment of such a filter could mean a lower measurement frequency of perhaps under 100Hz, an order of magnitude lower than the intermediate frequency.

On the other hand, were the signal allowed to remain current mode, it would be possible to implement a log-domain filter[26][27]. These demonstrate strong potential for implementation of high Q-factor filters, while removing the need for a transimpedance amplifier before this stage. Note that current mode amplification would still be desirable before this stage. A challenge may exist in the need for low bias currents in order to achieve the low center frequency, however should these be achievable, the effect is a significant reduction in the size of capacitors needed.

By considering the works above it is possible to outline specifications secondary to those required by the application. It is clear that other designs typically achieve a lower frequency bound of 10Hz – 1kHz [24][9][2]. Therefore for this circuit to be comparable to current technology, the minimum frequency should be aim to be below 1kHz.

With respect to lower impedance bound, it is easy to find examples which achieve measurements down to kΩs[2] or even Ωs[28] however these tend not to achieve the same 1GΩ upper bound as is required here. When this is taken into account we find that a lower bound of 10kΩ would make this circuit competitive[24], providing 100dB of dynamic range, which is the highest found. The maximum impedance determines the maximum noise which is tolerable, as this must be sufficiently low to resolve the signal response, which will likely be in the pico-Amp range for a 1GΩ impedance.

Lastly, previous designs have achieved errors ranging from 2% (RMSE) to 8%, [2][24][9] the upper end of which we will take as our desirable maximum for this design.

Property	min spec	reference
Power Consumption	1mW	Project Brief
Max Impedance	1GΩ	Project Brief
Min Impedance	10kΩ	[24]
Max Frequency	100kHz	Project Brief
Min Frequency	5kHz	[24][9][2]
Max Relative Error	8%	[2][24][9]

Figure 3: Summary of requirements

4 Analysis and Design

In order to begin design of this project it is necessary to separate it into different challenges and derive a solution to each, from which a block diagram of the system can be created and

simulated to show functionality in principle. From the literature review it seems fairly clear that there is a significant advantage in using an heterodyne converter to expand the frequency range of the system, however this involves creation of a second signal with a frequency locked to, but different from the stimulus. This will be discussed first. Second will be considered an approach to automatic gain control, to maximise the dynamic range of the system. Then will be discussed the method for determining the complex phasor of the response wave from the signal. From these a whole system will be presented.

4.1 Heterodyne conversion

The advantages of using a heterodyne converter, such as that used in [24] were discussed in the literature review. In short doing so reduces the bandwidth required of subsequent sections of the circuit. Fixing the frequency to predefined value also simplifies synchronisation of any sampling, and can improve noise rejection by moving the signal out of the range of flicker noise and simplifying design of filters. Ideally in a heterodyne, the signal would be mixed with a pure sinusoid, which was the intention here early in the design process. It was determined that using two DACs, in-phase and quadrature stimulus frequency and intermediate frequency waves could be generated. By mixing these together as shown in figure 4, a wave with the sum of those frequencies could be accurately generated, which could be used to transform the response signal to the intermediate frequency.

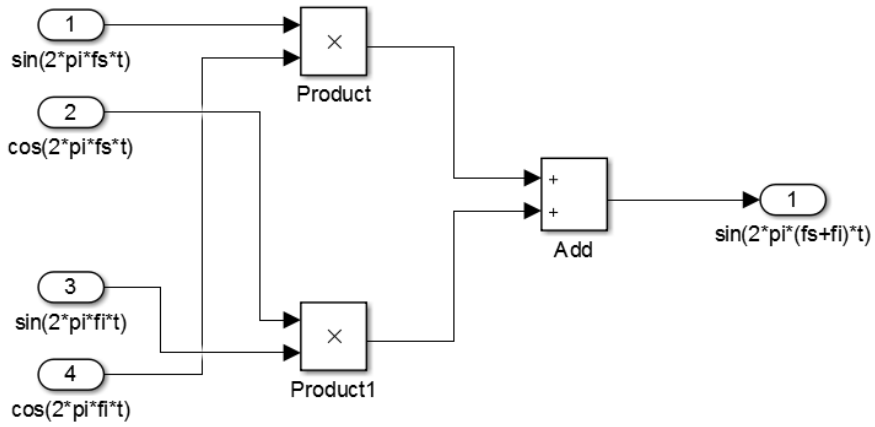

Figure 4: Method for generating $\omega_s + \omega_i$

There are a few difficulties with this method, however. The first is the requirement for the stimulus DAC to have twice the bandwidth of the maximum stimulus frequency, in order to generate the two waves. This then requires 4 high quality sample-hold structures, placed in parallel sampled one after the other to segregate the four signals as they are produced by the DAC. The sum of all this is a larger power consumption. The second main difficulty is the need for 3 good quality, 4 quadrant multipliers, as non-idealities here will introduce distortion, which can become a problem when trying to detect the signal phase.

An alternative is to use under-sampling instead of ideal multiplication. Doing this removes the need for low distortion multipliers, and only requires a single DAC producing a single sinewave, for the stimulus, thus removing all the sample hold circuits. Under-sampling uses the concept that a string of samples is essentially the sum of an infinite number of sinusoids at constant frequency intervals. Therefore these can equally be used to transform the frequency of the signal, so long as a sufficiently low bandwidth filter is available to reject the unwanted harmonics. The arithmetic for this is demonstrated in equation 4. On the left handside is shown the time domain formulae for a string of delta functions, the current response signal (with amplitude i_r and frequency f_s) and their product, respectively. By taking the fourier transform of each, as shown on the right we see how this results in signal which is the sum

of sinewaves at frequency $fi = nf_{samp} - fs$. Given that sampling is essentially multiplying the signal with a string of delta functions, the result will be multiple, superimposed sinusoids separated in frequency by f_{samp}. In order for one of these harmonics to always be at a fixed intermediate frequency we need $f_i = kf_{samp} + f_s \rightarrow f_{samp} = \frac{1}{k}(f_i - f_s)$. In order to remove all other harmonics, the bandwidth of the filter must be less than f_{samp}, meaning that the simplest solution is $k = 1 \rightarrow f_{samp} = f_i + f_s$.

$$\begin{gathered} y(t) \Leftrightarrow \mathcal{F}(y(t))(f) \\ \sum_{n=-\infty}^{\infty} \delta(t - \frac{n}{f_{samp}}) \Leftrightarrow \sum_{n=-\infty}^{\infty} \delta(f - nf_{samp}) \\ i_r \sin(2\pi f_s t) \Leftrightarrow \frac{i_r}{2}(\delta(f - f_s) + \delta(f + f_s) \end{gathered} \tag{4}$$

$$\begin{gathered} i_r \sin(2\pi f_s t) \times \sum_{n=-\infty}^{\infty} \delta(t - \frac{n}{f_{samp}}) \Leftrightarrow i_r \sum_{n=-\infty}^{\infty} \delta(f - nf_{samp} + f_s) \\ i_r \sum_{n=-\infty}^{\infty} \sin(2\pi(nf_{samp} - f_s)t) \Leftrightarrow i_r \sum_{n=-\infty}^{\infty} \delta(f - nf_{samp} + f_s) \end{gathered} \tag{5}$$

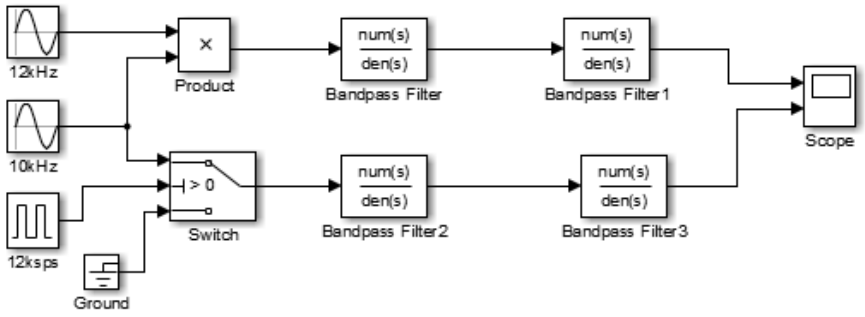

(a)

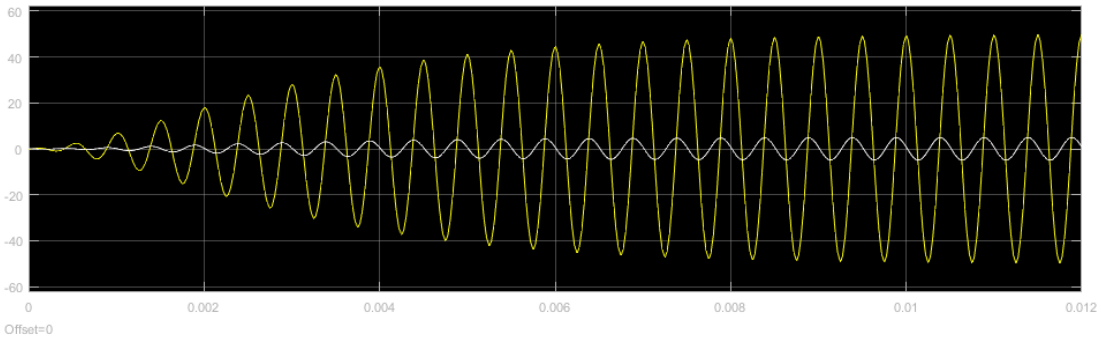

(b)

Figure 5: Comparison between mixing(yellow) and sampling(white) response to intermediate frequency

The effectiveness of this can be demonstrated using MATLAB. The system shown in Figure 5a is used to compare the proposed undersampling with direct multiplication, with the result in Figure 5b. All of the filters used in the system are 2nd order biquad filters with center frequency 2kHz and Q = 20. From the result it is clear the only differences is the sampled signal is a factor of 10 smaller and does not have a phase shift of 90^o with respect to the response signal, unlike

with direct multiplication. While it is true that in this case, the multiplied signals would cope with a substantially simpler filter, it is important to note that removing harmonics caused by sampling is not the only function of the filter. It is also necessary to remove the image frequency which sits at $f_i + 2f_s$ meaning that the bandwidth of the filter must still be less than twice the minimum stimulus frequency. If this stage is placed within and automatic gain loop (to be discussed) then no further consideration is necessary about the size of the signal.

4.2 Automatic gain control

In order to effectively measure a large range of impedances, the system must be able to cope with both very small signals and relatively large ones. In order to avoid the noise floor while also avoiding saturation, it is therefore necessary to include a variable gain amplifier (VGA), preferably with some kind of automatic gain control (AGC). While the VGA as a circuit will be considered in the implementation, the control system surrounding it will considered here.

Given the magnitude of the signal is important information, it is sensible to amplify the signal in a logical way. Perhaps the most obvious is to apply gains which are a power of 10, such that the magnitude of the response is represented by two values - a significand and an exponent. Any external circuit would then have to treat the output as $Sx10^E$. Another advantage of this the output resolution is kept approximately constant relative to the amplitude. If this were not the case, the output would have very poor resolution for large impedances.

The response magnitude can be measured using a peak detector, the output of which can be compared to an upper and lower threshold. Should it cross either of these thresholds, the exponent can be increased or decreased appropriately, adjusting the gain of the VGA at the same time. A system which does this is shown in figure 6. In this diagram the input signal is the response taken directly from the front end. This is then amplified by a gain which is a power of 10. The peak detector detects the the amplitude after amplification, and if this is outside of the acceptable window (between 0.9 and 11 arbitrary units), then the exponent is adjusted. Given that this needs give time for the peak detector to adjust to the new amplitude, the gain can not be adjusted very quickly. In the final circuit, it is likely that the heterodyne and its filters would also be placed with the amplifying block, such that this control circuit would have to give time for these filters to settle. As such the gain is only likely to be changable once every millisecond or so.

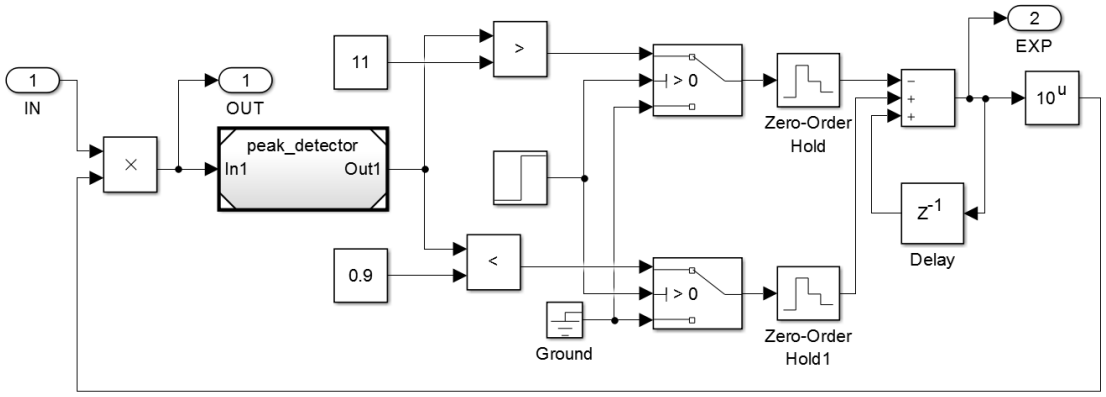

Figure 6: Diagram of the automatic gain control system

4.3 Phasor calculation

It is notable that none of the works considered included a process for actually calculating the impedance on-chip. The closest any achieved was calculating the complex phasor of the response signal. Likely the reason for this is that this calculation inherently involves highly

non-linear functions, such as arc-tangents and square-roots. Given this is necessary in this case, accurate calculation will require digital processing, which while area consuming needn't be power consuming. By using the method of sampling twice per cycle to calculate the phase and magnitude [2], applied after the heterodyne converter, this system only needs recalculate once per cycle, i.e. 2ksps. At this rate even very considerable calculations could be performed at low power. Equally by using the heterodyne to fix the frequency, it is relatively easy to sample effectively. Although the principle of the mathematics from [2] was adopted in this work, the method of implementing the integrator and ADC was different and will be addressed in the next chapter.

Given that the first part of this mathematics involves integrating the signal, here it is worth addressing the susceptibility of the integrator to input offsets, which will tend to cause it to saturate. Even if the integrator is leaky, or resettable, the output offsets will cause substantial errors in the subsequent calculations. This is clear if we consider the arithmetic, including the offset, such that our current response is given by $i_r \sin(\omega t + \phi) + i_o$. Integrating this gives $\frac{i_r}{\omega}(\cos(\omega t + \phi) - \cos(\phi)) + i_o t$. If we then sample at $t = \frac{\pi + n\pi}{\omega}$ and $t = \frac{3\pi + 2n\pi}{2\omega}$ we get $V_{\frac{nT}{2}} = -\frac{2i_r}{\omega}\cos(\phi) + i_o\frac{\pi+n\pi}{\omega}$ and $V_{\frac{3nT}{4}} = \frac{i_r}{\omega}(\sin(\phi) - \cos(\phi)) + i_o\frac{3\pi+2n\pi}{2\omega}$. Equation 6 shows what happens when the phasor is calculated from each sample, n, of this biased input.

$$\begin{aligned}
\phi_n &= \arctan\left(\frac{V_{\frac{Tn}{2}} - 2V_{\frac{3Tn}{4}}}{V_{\frac{Tn}{2}}}\right) \\
&= \arctan\left(\frac{-\frac{2i_r}{\omega}\sin(\phi) - i_o\frac{\pi}{2\omega}}{-\frac{2i_r}{\omega}\cos(\phi) + i_o\frac{\pi}{\omega}(1+n)}\right) \\
i_{rn} &= \frac{\omega}{2}\sqrt{V_{\frac{Tn}{2}}^2 + (V_{\frac{Tn}{2}} - 2V_{\frac{3Tn}{4}})^2} \\
&= \frac{\omega}{2}\sqrt{\left(\frac{2i_r}{\omega}\right)^2 - i_o D}
\end{aligned} \tag{6}$$

$$where \quad D = \frac{4i_r\pi(n+1)}{\omega^2}\cos(\phi) + i_o\left(\frac{\pi(n+1)}{\omega}\right)^2 - 2i_r\sin(\phi) + i_o\left(\frac{\pi}{2\omega}\right)^2$$

It is clear from that if $i_o = 0$ then the arithmetic in equation 6 decomposes to $i_r = i_r$. However if it is non-zero, the effect is significant, very non-linear, and time variant, which self-evidently problematic. The proposed solution in this report is to consider possible samples around the nth cycle. These are listed in equation 7

$$\begin{aligned}
V_{\frac{nT}{4}} &= -\frac{i_r}{\omega}\big[\sin(\phi) + \cos(\phi)\big] + i_o\frac{\pi(1+4n)}{2\omega} \\
V_{\frac{nT}{2}} &= -\frac{2i_r}{\omega}\cos(\phi) + i_o\frac{\pi(1+2n)}{\omega} \\
V_{\frac{3nT}{4}} &= \frac{i_r}{\omega}\big[\sin(\phi) - \cos(\phi)\big] + i_o\frac{\pi(3+4n)}{2\omega} \\
V_{\frac{5nT}{4}} &= -\frac{i_r}{\omega}\big[\sin(\phi) + \cos(\phi)\big] + i_o\frac{\pi(5+4n)}{2\omega}
\end{aligned} \tag{7}$$

From these it is possible to construct two equations, which represent the real and imaginary parts of the response phasor. These are shown in equation 8.

$$\begin{aligned}
2V_{\frac{nT}{2}} - V_{\frac{nT}{4}} - V_{\frac{3nT}{4}} &= 2\frac{i_r}{\omega}\cos(\phi) \\
\frac{1}{2}V_{\frac{nT}{4}} + \frac{1}{2}V_{\frac{5nT}{4}} - V_{\frac{3nT}{4}} &= 2\frac{i_r}{\omega}\sin(\phi) \\
\rightarrow \quad \frac{2i_r}{\omega}(\cos(\phi) + i\sin(\phi)) &= \frac{2}{\omega}I_r
\end{aligned} \tag{8}$$

If the response signal has already been digitised, this is then a relatively easy system to implement, as the sampling can be performed using a series of registers, clocked at $f_{samp} = 4\frac{\omega}{2\pi}$. If the outputs of these registers are combined and sampled once per cycle, they will then show the correct result. The system is shown in figure 7.

In this case we have a 2kHz response signal, with an unknown offset. This is integrated and sampled at 8ksps. By using a series of delay cells (which could be implemented using registers) the samples at $t = \frac{5T}{4}, \frac{3T}{4}, \frac{T}{2}, \frac{T}{4}$ can be extracted. These are then sampled at 2ksps, to give a set of data for each cycle. When the weighted sum of these is taken as discussed the result is the real and imaginary part of the current response. This can then be converted to magnitude and phase form, if necessary. Simulation of this system in Simulink demonstrates complete

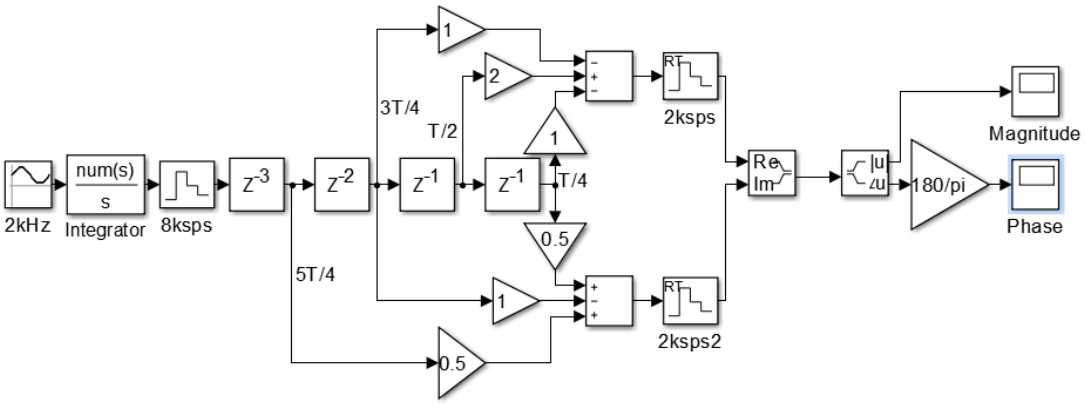

Figure 7: Diagram of the sampling and arithmetic system

immunity to offset in both the magnitude and the phase, as shown in figure 8. By comparison the original method, reacts very poorly to even an offset 1 tenth of the size of the signal amplitude. Also worth noting here is that because this method returns the current response phasor, if no further arithmetic logic is added, this system will match all the rest of the literature. Although calculating the magnitude and phase of the impedance from this is area consuming, were it implemented this would take it beyond any current monolithic device.

Non-idealities in this system may come from two places. The first is that in reality the integrator will overflow at some point, or it must be reset at intervals. This creates times when the output is not correct while it recalculates. To avoid this happening too often the integrator needs to a sufficient number of bits to hold the large numbers requiring it to either be large, or for the resolution of the actual signal to be reduced. The second is that the sampling my not happen exactly in phase with the signal. However this only results in a small offset in the final phase measurement. Non-linear effects will come from jitter, however when sampling at 8ksps, this will have a negligable effect on the final result.

4.4 The complete system

The entire readout system is shown in figure 9. The heterodyne block, and maths block are the same as above, while the automatic gain control have been divided into the control circuit and VGA elements, in order to ensure that the peak detector is measuring the amplitude of a 2kHz wave. If it remained with only the VGA in the control loop, then the peak detector time constant could be too short for low frequency stimuli, causing instability.

The graphs in figure 10 demonstrate that this system works in general. The first shows the significand output (i.e. amplitude, without the exponent) and phase output, as the frequency of the stimulus is swept. Clearly the frequency has no effect on these measurements, except for around 2kHz. The second shows the accuracy of the phase measurement, which is clearly very linear, albeit with a small offset. The third shows the accuracy of the magnitude measurement,

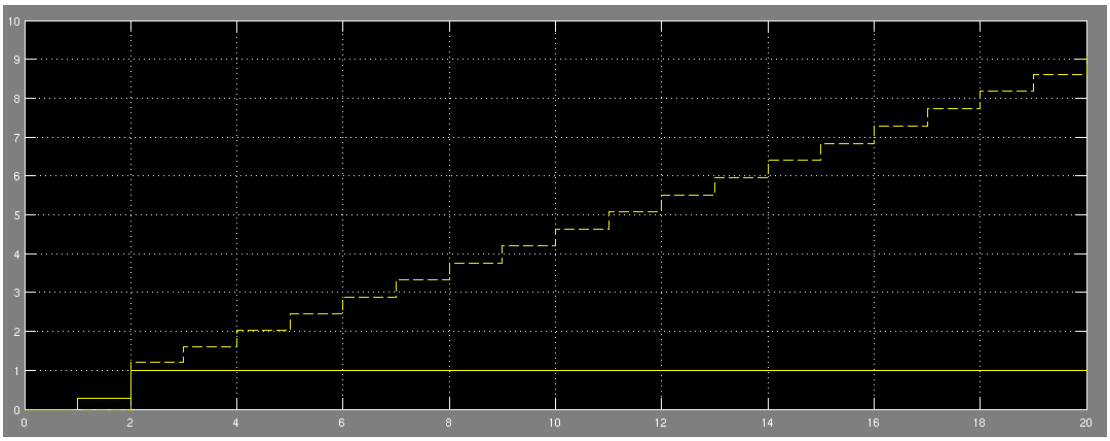

(a) Comparison of the original system (dotted) with the improved system(full) when measuring the amplitude of a response

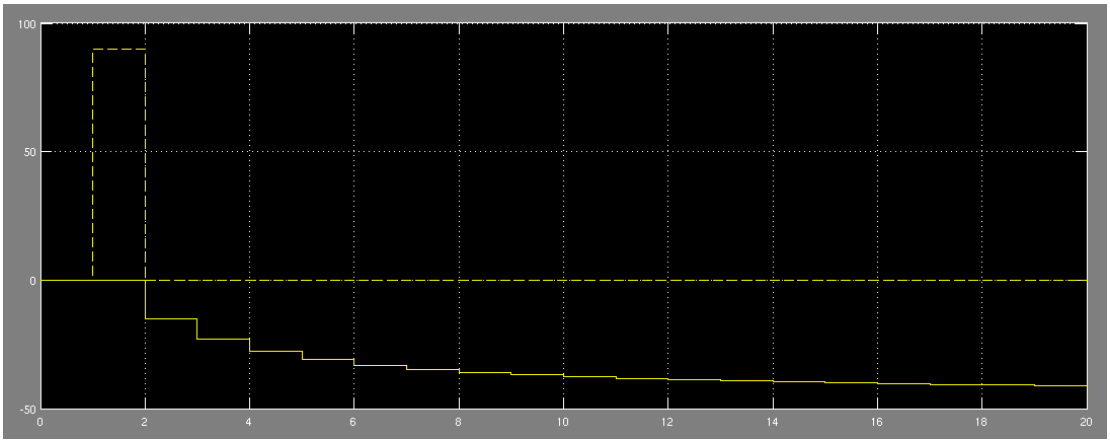

(b) Comparison of the original system (full) with the improved system(dotted) when measuring the phase of a response

Figure 8: Comparison of the two systems, for a response signal with amplitude 1, phase 0, and offset 0.1

for many response amplitudes. As the AGC unit maintains the amplitude of the signal within one order of magnitude, the output is divided between the significand (i.e. the measured amplitude output) and the exponent, which was extracted by the AGC. The final magnitude output is calculated as $SIG \times 10^{EXP}$.

There are a two non-idealities to note, both caused by the heterodyne. The first is a 9 degree phase offset caused by the filters in the heterodyne as visible in figure 10b. The second occurs at a 2kHz frequency, as shown in figure 10a. As this matches our intermediate frequency, it is therefore half of, and exactly in phase with the sampling, such that every sample is zero. Clearly from this it is impossible to measure the magnitude and phase of the signal and the output will display a random value derived from the measured noise at that time. It exists for frequencies approximately 200Hz either side of 2kHz. This is a problem inherent in this method. Arguably this could be taken as the lower frequency band for the design, given this method works down to frequencies of 100Hz, for the filters described (which could be reduced further with higher order filters), depending on application this frequency could be avoided. Besides this the graphs in figure 10 demonstrate that this system works performs the task well.

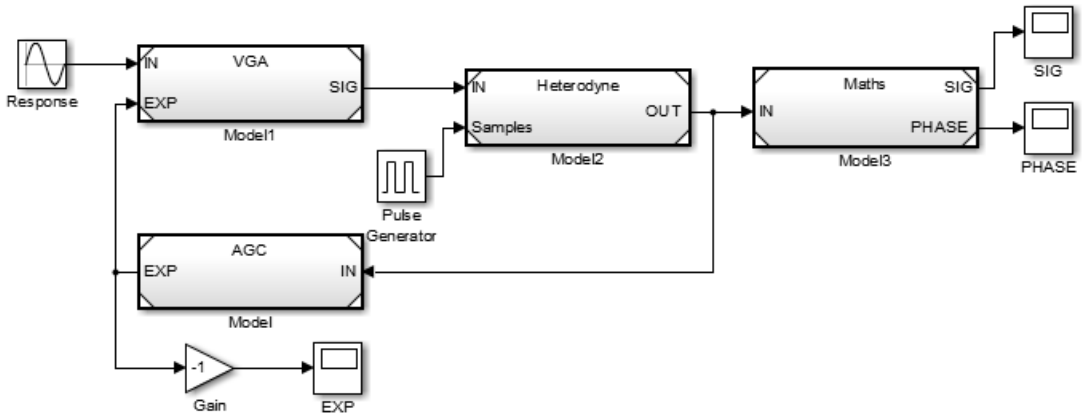

Figure 9: Diagram of entire readout system

5 Implementation

Given that the results from the design stage point to digitisation being most appropriate at the calculation stage, it was concluded that the analogue parts of the readout system were most important to implement. The reasoning for this is that this only leaves calculation and the control state machine to be designed in digital, neither of which will require substantial power (remembering the response is sampled at 8ksps and the arithmetic performed at 2ksps), nor will either introduce any further error into the result. On the other hand, the analogue sections had the potential to be power consuming, and add further distortion, bias or other errors to the response before it could be digitised. Therefore, given that this project was time limited, these parts were the priority to implement on an electronic level, and are the parts shown in this section of the report.

The final design takes advantage of a number of well known circuits, applying them in a unusual application to achieve the desired result. Given the large dynamic range of the response, it was clearly desirable to keep the signal in current mode until after the variable gain amplifier (VGA).

However, given that the system includes an integrator just before digitisation, and the basic integrating component is a capacitor, it makes more intuitive sense to try and keep the signal current mode until this point. As sampling is just as straight forward for either, and we have seen options for either current mode or voltage mode filters in the literature review, it was clear that keeping the signal in the current domain was sensible.

The obvious approach to digitisation would then be to follow the integrator with a voltage mode analog to digital converter (ADC) however, this again depends on a linear, low power converter. Instead this project combines integration and digitisation into a single process, using an integrate and fire circuit with an asynchronous up-down counter. As it is very easy to make integrate and fire circuits linear this should result in a highly accurate low-power circuit analogue to digital integrator.

This section is divided up, in order to explain and demonstrate each circuit block individually, before showing the analog readout circuit simulation in full in the next section.

5.1 Front End

The first task of this system must be to stimulate a sample impedance with a stimulus wave, in such a way as to make it possible to read the current response. Using [24] as a reference, the circuit in figure 11 is used. In this the Op-Amp and transistor combination are used as a high gain voltage follower, while allowing the current to be drawn through the cascode current mirror. This current mirror provides a low input impedance, high output impedance current

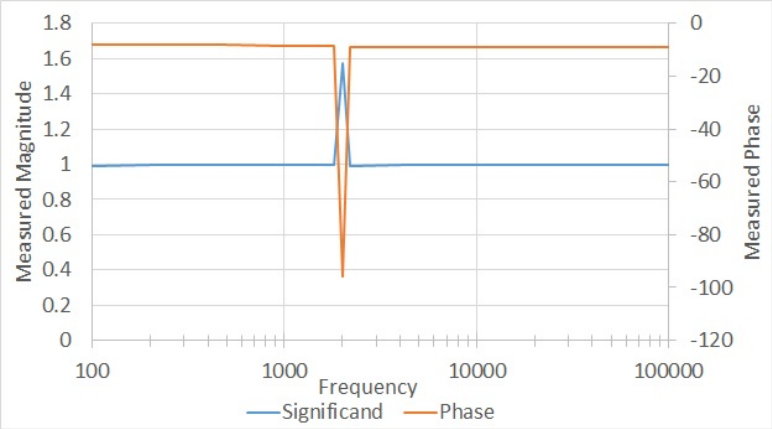

(a) Measured significand and phase across frequencies for response with magnitude 1 and phase 0

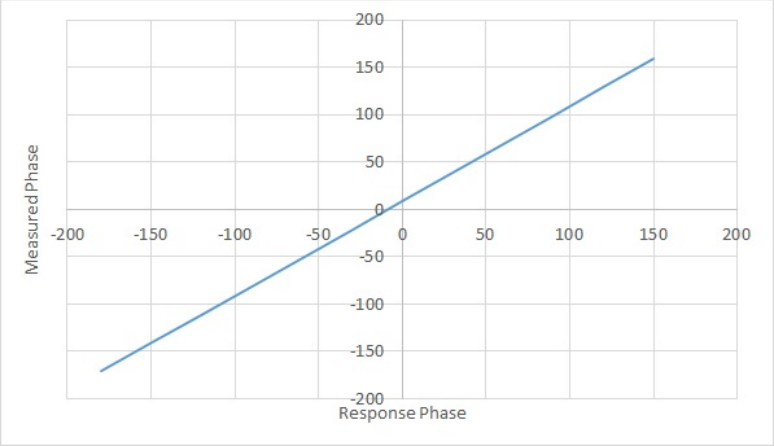

(b) Measured phase and vs response phase for response magnitude 1 and frequency 10kHz

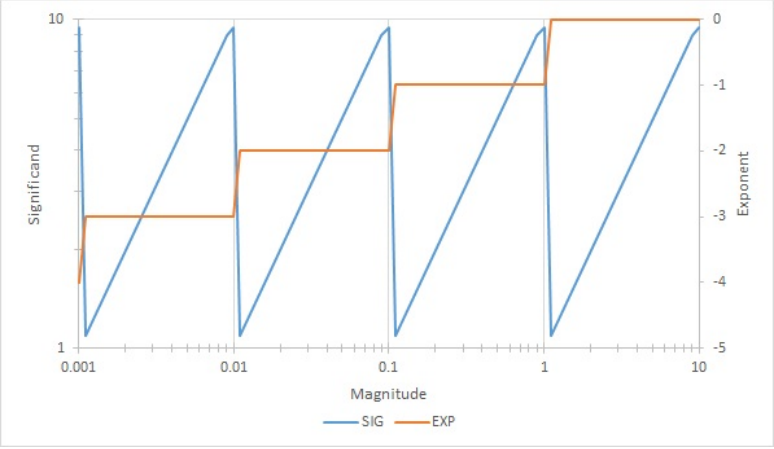

(c) Measured significand and exponent vs response magnitude for response phase 0 and frequency 10kHz

Figure 10: Results from simulink simulations of the readout system, for

follower. This has advantages over connection directly to a transimpedance amplifier (TIA), as in many similar works, as the capacitance in the TIA can effect the response of the measured impedance, as can its non-zero input impedance. Using a current mirror avoids these issues.

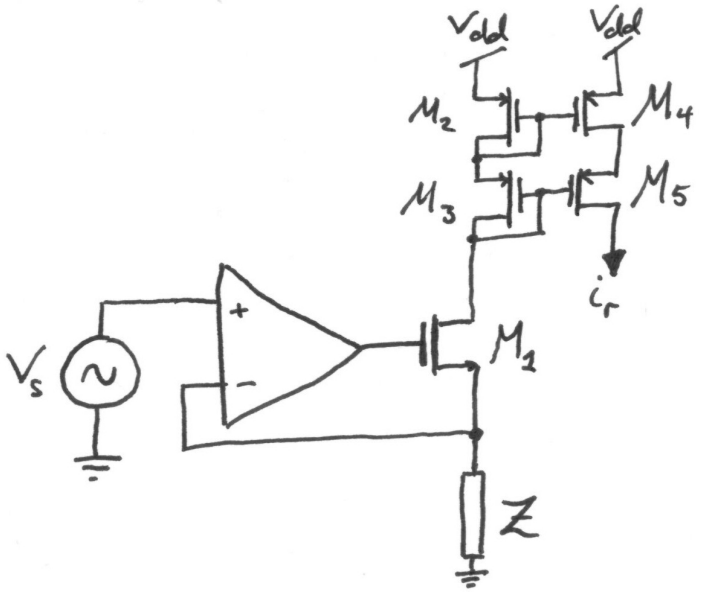

Figure 11: Front End of circuit

The Op-amp is a simple two stage folded cascode amplifier, with a corner frequency of around 50kHz, and a gain above 1000 for all frequencies under 100kHz. The primary disadvantage with this opamp is the parasitic capacitance at its input. Even with small input transistors, at 100kHz the gate capacitance carries approximately 20pA of current. Given that this system measures impedances up to 1GOhm with a stimulus of 10mV, the current read-out is in the range of 10pA. Therefore there is a significant error here. In many applications where the aim is to use the current response data to generate an equivalent circuit model, this capacitance could be found and removed to give the correct circuit model. Given that this Op-amp is not part of the current readout circuit only a simple op-amp was used. Although in a revision of this project, this is something that should be considered for improvement.

5.2 Translinear amplifier

Given that highly variable, low power current mode amplifiers are difficult to find, this block takes advantage of a circuit which is typically used as a multiplier to amplify the signal.

Using the translinear principle, we can deduce that for the circuit in figure 12, $I_1 I_2 = I_3 I_4$ [29]. If we consider I_1 and I_4 as a single differential signal and I_2, I_3 as another, such that $I_1 = I_A(1 + X)$ and $I_4 = I_A(1 - X)$, while $I_2 = I_B(1 + Y)$ and $I_3 = I_B(1 - Y)$ we can then deduce that $Y = -X$, irrespective of their bias currents. As a consequence, if I_A can be fixed, then we have a variable gain amplifier, where the gain is $\frac{I_B}{I_A}$. With a BJT circuit, this works for gains less than the β of the transistors. In theory this does not limit MOSFETs, however generally it is easier to limit it to around 100, and cascade two amplifiers to achieve the required gain. Figure 14 shows an ac simulation against the gain value, demonstrating how accurate this is.

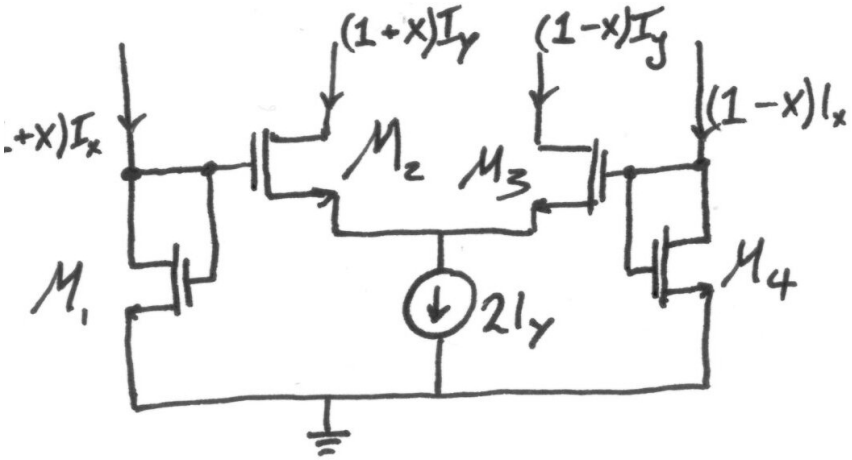

Figure 12: Translinear multiplier cell

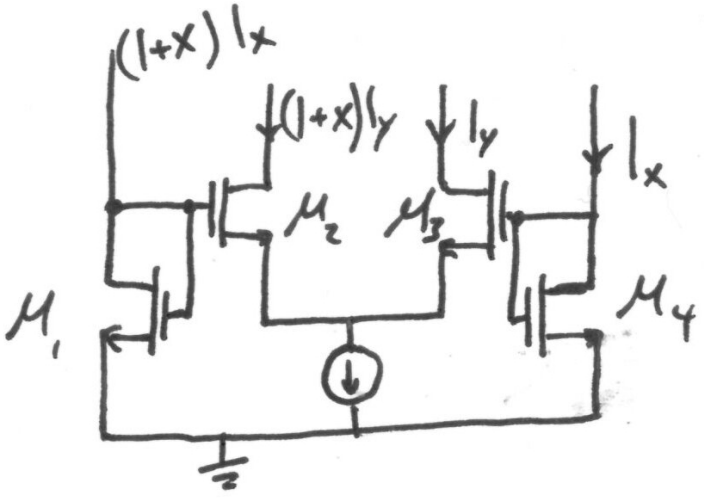

Figure 13: Single Ended adaptation of amplifier

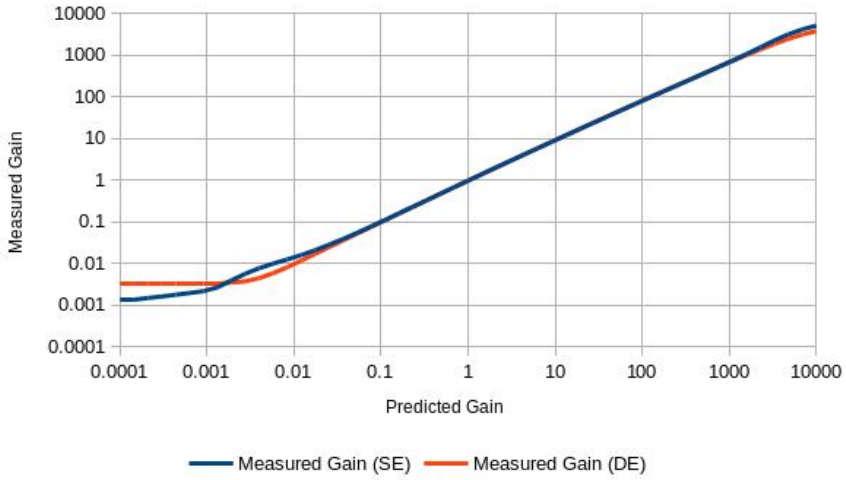

Figure 14: AC simulation of single ended vs differential amplifier

The problem with using this circuit directly like this is that the current response is a single ended, not differential signal. If it were converted, it would be essential that this process does not result in differential bias, as this offset will become a gain error in the amplifier. Given that this system must cope with response amplitudes as low as 10pA, it is nearly impossible to achieve an acceptably low offset.

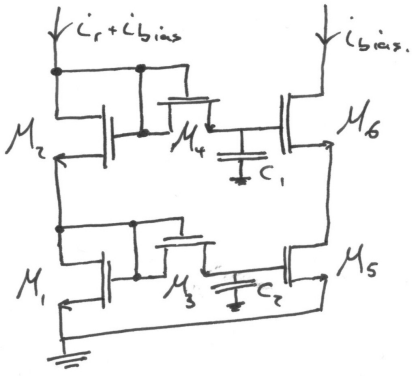

Figure 15: Crude Low-pass filter

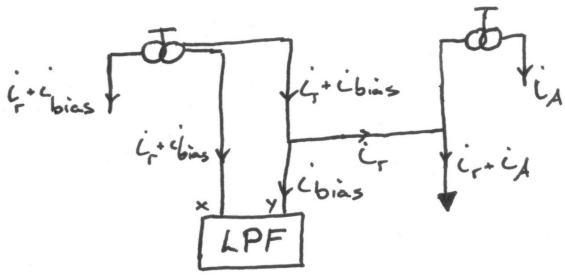

Figure 16: Signal Offset replacement circuit

The alternative is to fix one input and one output, such that our trans-linear equation becomes $I_A(1+X)I_o = I_B(1+Y)I_o$, such as in figure 13. The AC simulation can be repeated as shown in 14, showing that this system actually has slightly better linearity at higher gains, although slightly worse at lower gains. Both are equally good for gains in the region $0.1 < G < 100$ Note that gain is now half the ratio of bias currents. Now the system is single ended, a new way is needed to fix I_A to a known value. This cannot be done perfectly, as the offset to the signal is unknown, and determined by the DC response of the measured impedance. In this case the offset does not lead to a gain error at the output, only a larger offset due to it being amplified also. Therefore, the offset does not have to be zero, only small. The offset can be approximated by a low-pass filter, with a very low corner frequency. This has been implemented as part of a current mirror, using a transistor as a pseudo resistor with a very large resistance, and a capacitor as shown in figure 15. As this is purely for for finding offsets the filter can be crude, as signal distortion, or the precise cut-off frequency are not important.

5.3 Sampler

The sampler forms the core of the heterodyne, and is required to sample the input up to 102ksps. This is even more simple than with a voltage mode circuit, as the current can be simply diverted to ground using a second transistor, which is switched. The circuit diagram is shown in figure 17. In figure 18 this is applied to a 10kHz signal, sampling at 12ksps, such that the 2kHz output component is clearly visible. This method is susceptible to very large current spikes at each transient, however this can be avoided be keeping the rise and fall times of the sampling signal to around 300ns.

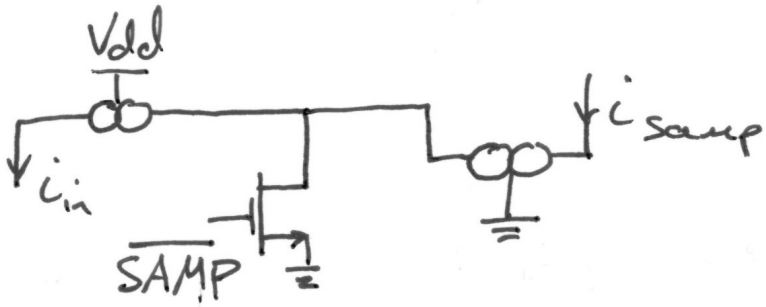

Figure 17: Sampling circuit

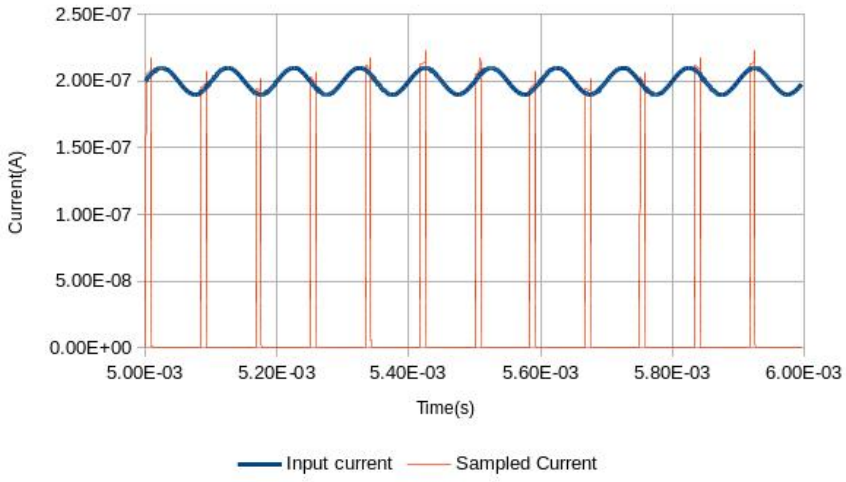

Figure 18: Output of sampler with input signal to show accuracy.

5.4 Log-domain filter

This filter is necessary to remove the unwanted harmonics from sampling. It's principle is not immediately intuitive from the circuit in figure 19 because of the effect of log-domain compression of the signal, for the duration of processing. However its principle can be easily derived from a standard state space equation.

As described in the design section of this report, this system requires a 4th order low-pass filter, with relatively high Q factor. This is more simply constructed from two cascaded 2nd order filters, for which the state space equation and transfer function are:

$$\dot{X_1} = -\omega_0 X_2 + \omega_0 U$$
$$\dot{X_2} = \omega_0 X_1 - \frac{\omega_0}{Q} X_2 \quad (9) \qquad \frac{Y}{U} = \frac{\omega_0^2}{s^2 + \omega_0 s + \omega_0^2} \quad (10)$$
$$Y = X_2$$

Given that a MOSFET in weak inversion has a characteristic given by $I_{ds} = I_s \exp(\frac{V_{gs}}{nU_t})$ in which n is the slope factor and U_t is the thermal voltage, We can transform this state-space description into a circuit description using the following state transforms [27]:

$$X_1 = I_1 \exp\left(\frac{V_1}{nU_t}\right) \qquad X_2 = I_2 \exp\left(\frac{V_2}{nU_t}\right)$$
$$\dot{X_1} = \frac{I_1}{nU_t} \exp\left(\frac{V_1}{nU_t}\right)\dot{V_1} \quad (11) \qquad \dot{X_2} = \frac{I_2}{nU_t} \exp\left(\frac{V_2}{nU_t}\right)\dot{V_2} \quad (12) \qquad U = I_U \exp\left(\frac{V_U}{nU_t}\right) \quad (13)$$

If we substitute this into the state equations we get:

$$\begin{aligned} \frac{I_1}{nU_t} \exp\left(\frac{V_1}{nU_t}\right)\dot{V_1} &= \omega_0 I_u \exp\left(\frac{V_U}{nU_t}\right) - \omega_0 I_2 \exp\left(\frac{V_2}{nU_t}\right) \\ \frac{I_2}{nU_t} \exp\left(\frac{V_2}{nU_t}\right)\dot{V_2} &= \omega_0 I_1 \exp\left(\frac{V_1}{nU_t}\right) - \frac{\omega_0}{Q} I_2 \exp\left(\frac{V_2}{nU_t}\right) \end{aligned} \quad (14)$$

We can rearrange these to describe capacitor currents

$$\begin{aligned} C\dot{V_1} &= \frac{I_\omega I_U}{I_1} \exp\left(\frac{V_U - V_1}{nU_t}\right) - \frac{I_\omega I_2}{I_1} \exp\left(\frac{V_2 - V_1}{nU_t}\right) \\ C\dot{V_2} &= \frac{I_\omega I_1}{I_2} \exp\left(\frac{V_1 - V_2}{nU_t}\right) - \frac{I_\omega}{Q} \end{aligned} \quad (15)$$

where $I_\omega = CnU_t\omega_0$. If we set $I_1 = I_2 = I_\omega$ and $I_U = I_0$ then we get

$$\begin{aligned} C\dot{V_1} &= I_0 \exp\left(\frac{V_U - V_1}{nU_t}\right) - I_\omega \exp\left(\frac{V_2 - V_1}{nU_t}\right) \\ C\dot{V_2} &= I_\omega \exp\left(\frac{V_1 - V_2}{nU_t}\right) - \frac{I_\omega}{Q} \end{aligned} \quad (16)$$

By applying E+ and E- cells to form the positive and negative exponential relationships, we can derive from these two equations the circuit in figure 19. In theory this circuit allows for a wide range of Q-factors, however in practice a Q-factor of 2 was achieved. Nonetheless, this provided a good quality filter, for which the AC transfer function for a single stage and 2 stages is demonstrated in figure 20

In figure 20 is shown the comparison of the sampled current and filtered current, highlighting how smoothly the 2kHz component can be extracted. Although a phase shift is present, this is constant with frequency, and therefore can easily be compensated for later.

5.5 Peak detector

The peak detector is the primary analogue block required for the automatic gain circuit. Its circuit diagram is shown in figure 21. In this the left and right transistor cascodes effectively create a current mirror, between which is placed a voltage-mode peak detector. The op-amp tries to follow the signal at the input, however the diode-connected transistors at its output only permit it to do so when its input is greater than its output. As a consequence it remains approximately at the signal peak. Because this is located after the heterodyne it has been optimised for 2kHz operation. The result is shown in figure 22

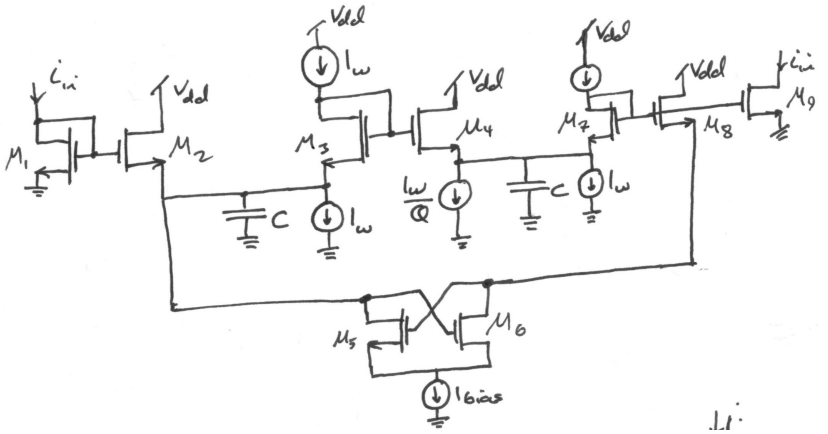

Figure 19: Log domain Biquad Low pass filter circuit

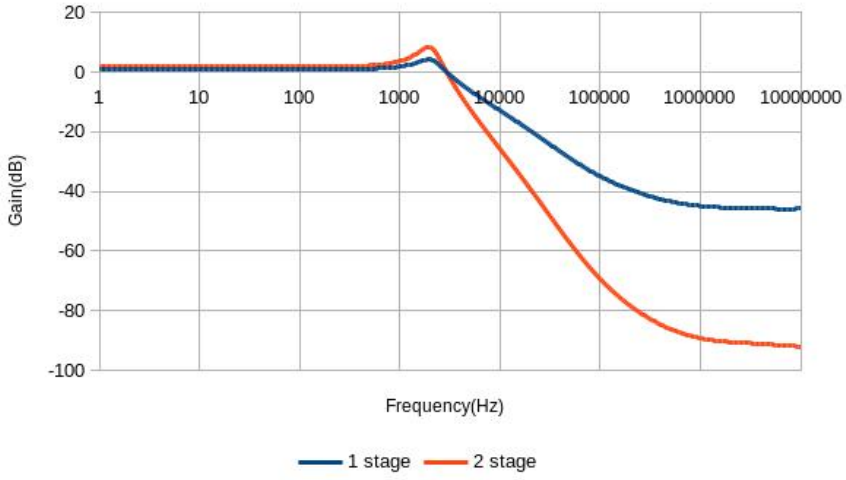

Figure 20: Transfer function of log-domain biquad filter, both on its own and with two in cascade

5.6 Integrator and digitisation

The combination of an Axon-Hillock integrate and fire circuit with an asynchronous up-down counter takes advantage of the need to integrate the signal immediately before digitising in order to digitise the data in an unusual way. The integrate and fire circuit produces a voltage spike for every unit of charge which enters the circuit through the input, and therefore the frequency of its output is proportional to the input current. By connecting this to the input of a counter, the value of the counter will increase at a rate proportional to the input current: in other words to integrate it. Given that an integrate and fire circuit cannot handle negative currents, two can be used to handle the signal differentially to provide the UP and DOWN inputs to an asynchronous up/down counter.

The 'traditional' 'Axon-Hillock' integrate and fire circuit is formed by two capacitors and a series of inverters or a comparator[30][31], as shown in figure 23a. Here the input current charges C_A until the voltage on the C_A reaches the threshold V_th comparator. At this point the circuit fires. As the comparator causes the output voltage to rise fast, which feeds back to the input node through C_fb, reinforcing the spike temporarily. However the high output then turns on the reset transistor, causing the input node to discharge and the circuit to start again. The primary advantage with this circuit is that power consumption is relative to information,

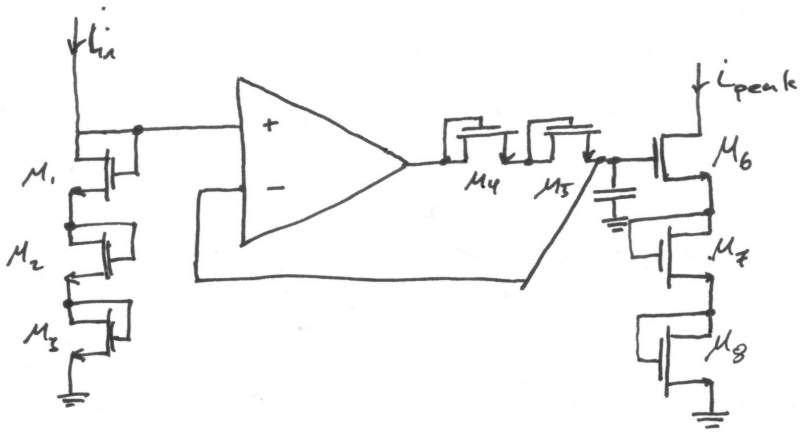

Figure 21: Peak detector Schematic

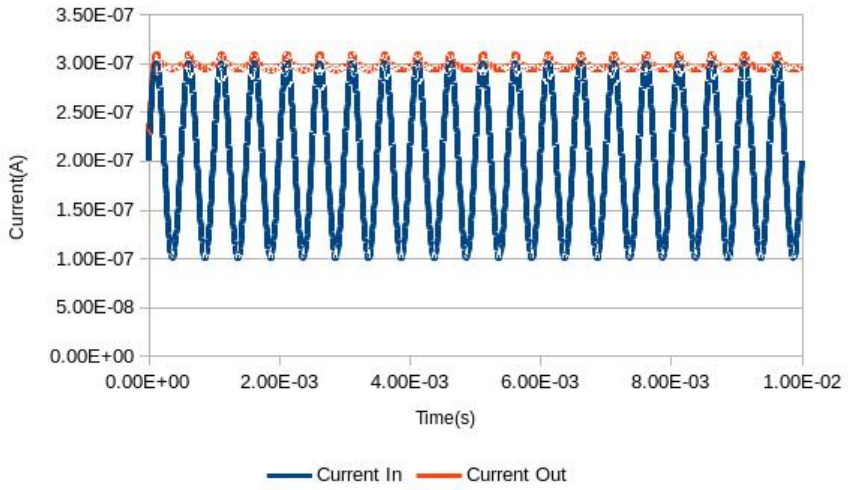

Figure 22: peak detector results

unlike synchronous ADCs. However given that this circuit can only accept a positive signal, to use two, separately integrating and firing each half of a differential signal is clearly a waste of power. A solution to this is to include two comparators, one which represents a negative pulse and the other a positive pulse[32] as shown in 23b. Now a single ended, bi-directional current signal can be applied, and the frequency of firing will be zero when the input signal is zero.

As shown in figure 24 this is fed to an asynchronous up-down counter, which is formed from a synchronous up-down counter, with the clock signal formed by the OR of the two signals. Therefore the value of the counter will change every time one of the two outputs of the integrate and fire circuit fires. As only the first controls the up/down nature of the counter, the counter will count up if this one fires, or otherwise count down. The inverters between the OR gate and the clock port of the counter are to delay the signal. This is because the UP/DOWN input of the counter has a non-zero setup time which must be observed before the clock is triggered. Given that both of this inputs are controlled by the same signal, the inverter string is used to generate a delay, in order to ensure the setup time is satisfied.

From figure 25 we can see that although the amended integrate and fire circuit has good linearity away from zero, there is a central "dead-zone", during which there is no signal from

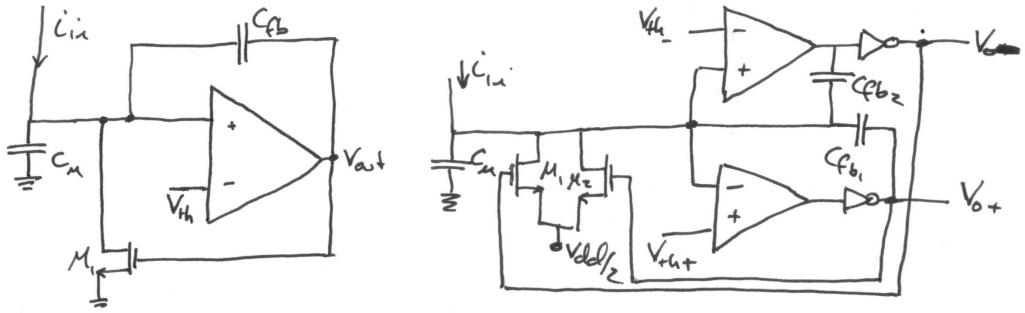

(a) 'Traditional' Integrate and Fire circuit

(b) Revised Integrate and Fire circuit

Figure 23: Integrate and Fire circuits

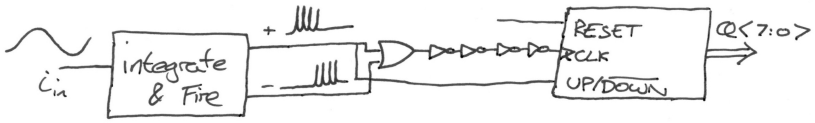

Figure 24: Diagram of the integrate and fire and UP/DOWN counter circuit

either output. This is non-ideal, is relatively small compared to signal amplitudes.

Figure 26 shows the final output after digitisation. This is for a 10nA sine wave on top of a 20nA bias. As is evident form the graph the bias has been almost completely stripped away, leaving an integrated sinewave. Given that there is a small bias remaining, this counter will still need to be reset on occasion, after which there will be a single-cycle time gap during which the output is invalid. Alternatively the counter could be allowed to overflow, which would also result in a single-cycle time gap, however the timing of this would be less predictable. In order to increase the time between resets it is either necessary to increase the size of the counter to hold more bits, or reduce the frequency of the integrate and fire circuits. The former of these options will take up more space and power, while the latter reduces the resolution of the data. As we can see from 26 the signal is currently relatively small, such that for the bias used the frequency possibly could have been increased. However as shall be seen in the testing section, this is a fairly modest bias.

5.6.1 Offset detector

Given that integrators are particularly sensitive to biases, it is sensible to remove as much of the bias as possible prior to integration. Due to the impact a bias will have on the digitised signal, this needs to be more precise than was the case before amplification, and is worth using extra area and power than the crude low pass filter. This improved offset detector is made using two peak detectors, one effectively detecting the signal maximas and the other detecting the signal minimas before the two are averaged. Given that detecting minimas is relatively difficult, this method is adapted slightly to give the circuit in figure 27

The principle is that if we subtract an input signal $i_{off} + i_r \sin(\omega t)$ from a larger bias i_c and use this as the input to a peak detector we get $i_c - i_{off} + i_r$. If we also put the input through a second peak detector and subtract this from the same offset we get $i_c - i_{off} - i_r$. Summing these two signals, and then subtracting twice the offset gives $2i_c - i_{off} + i_r - i_{off} - i_r - 2i_c = -2i_{off}$. Inverting and halving this gives the offset with a ripple twice the ripple on a single peak detector, which can then be subtracted from the signal before integrating.

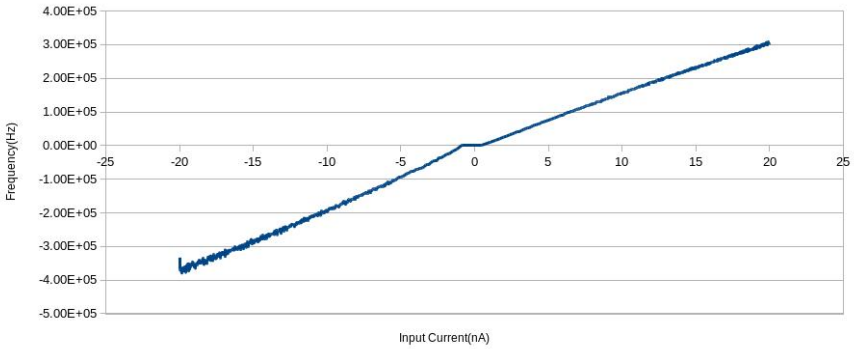

Figure 25: Difference in frequency between the positive and negative outputs of integrate and fire circuit compared to input current

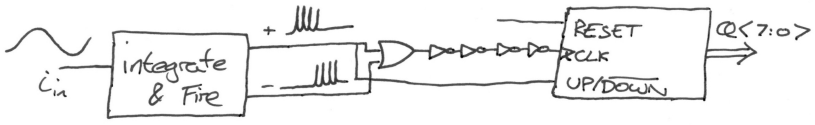

Figure 26: Counter value with time

5.6.2 Integrate and Digitisation results

To perform a basic test on the functionality of this circuit two input signals are tested. The first is a response with amplitude 10nA on a 20nA bias and 0^o phase offset, and the second with 20nA amplitude on 30nA bias and 45^o phase. Given that the value of an LSB in the digital output is arbitrary, comparing two responses, with the second having an amplitude twice the size of the first, it will be possible to see that the circuit is functioning correctly if the digital output is also twice as large. Given that the phase is calculated as the arctangent of a ratio of the real and complex parts, only one value is necessary to show it is is correct, however two adds to the confidence.

28 shows the digitised output of the counter for the second case, with a 45 degree offset. The first cycle must be ignored as during this the offset is still being calculated. However considering the second cycle, starting at t=0.5ms, we can clearly see a sinusoid with some kind of phase shift. However calculation is required to determine exactly what this is. As was demonstrated during the analysis section, this signal would be sampled at 8kHz, and delays applied to extract quarter-phase data. Although this digital part of the circuit has not been implemented, looking up the data shows that the values (in LSBs) at each point are as in figure 29.

In each case the in phase component could be calculated as $V_I = -V_fracT4 + 2V_fracT2 - V_frac3T/4$ and quadrature phase component as $V_Q = \frac{1}{2}(V_fracT4 + 2V_frac3T4 + V_frac3T4$. These have also been calculated and shown in figure 29. We can then check the calculated magnitude and phase of this response using $|I_r| = \sqrt{(V_I^2 + V_Q^2)}$ and $\angle I_r = \arctan(\frac{V_Q}{V_I})$. This information is also included in 29. From this we can see the accuracy of these measurements. Although the amplitudes are measured in arbitrary LSBs, the 20nA signal is clearly measured to have an amplitude of twice the 10nA signal, to within 1 LSB accuracy, giving (taking the 10nA value as a reference) an error of 1.1% when quantisation is taken into account. Equally the phases are both measured to within 1^o accuracy, representing an error of 0.28%.

5.7 Summary

This section has described the circuit level implementation of the analog parts of the current

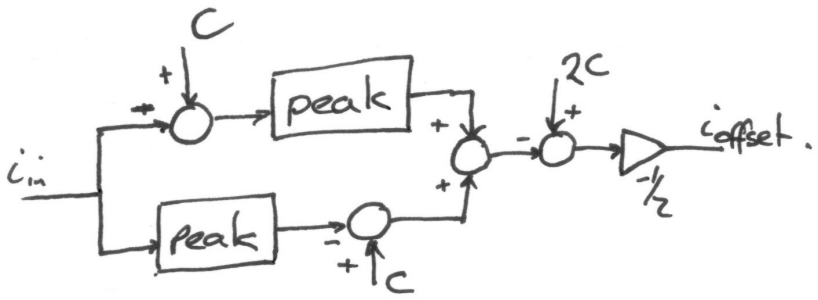

Figure 27: Circuit for determining the offset of a signal

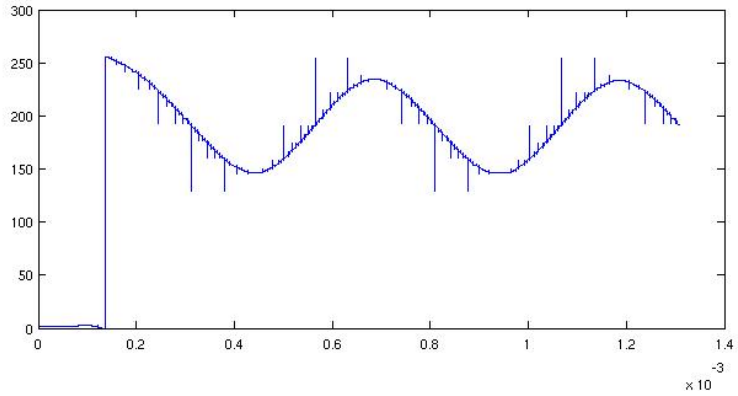

Figure 28: Output of counter for the 20nA wave with a 40^o offset

readout circuit for an EIS system. Given that it has been demonstrated that this system functions in principle, this section has demonstrated that it is possible to implement at a transistor level. Testing of the complete system will be required to understand the interplay between various components.

6 System Simulation and Results

Due to the complexity of this project and its time constraints, implementation of the system was confined to the analogue sections of the circuit. As such a complete digital state machine to control it has not been implemented, nor has the arithmetic which is required after digitisation. As a consequence in the following simulations the gains for the VGA have been set manually.

A transient simulation has been performed for a wholly resistive impedance of 100MΩ, with a 10mV, 10kHz signal on a 500mV bias to demonstrate functionality at each stage of the process, for the circuit as a whole.

6.1 Front End

The purpose of the front end is simply to apply a determined sinusoidal voltage across the sample impedance, and receive a current response which is dependant on it. Figure 30 demonstrates that the front end used results in a clean sine wave with the expected amplitude of 100pA, and 0^o phase. Given that this is a single ended system, this voltage has to be provided atop a

Sample	Time (ms)	10nA, $\angle 0^o$	20nA, $\angle 45^o$
$V_{\frac{T}{4}}$	0.625	120	220
$V_{\frac{T}{2}}$	0.75	142	221
$V_{\frac{3T}{4}}$	0.875	120	159
$V_{\frac{5T}{4}}$	1.125	220	120
V_I	-	44	63
V_Q	-	0	61
$\lvert I_r \rvert$	-	44	87.7
$\angle I_r$	-	0^o	44.026^o

Figure 29: Counter value in LSBs at each sample required for arithmetic

bias. The voltage bias is 500mV, and this results in a bias of just over 50nV, which is useful information for quantifying the bias rejection in later stages.

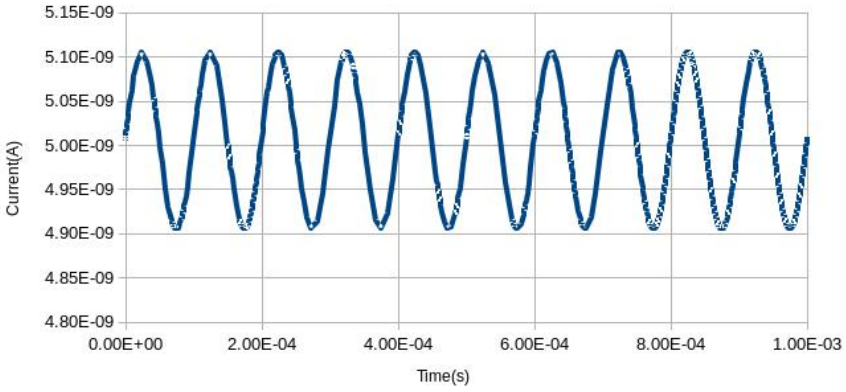

Figure 30: Current response in front end.

6.2 Amplifier

The purpose of the amplifier is to multiply the signal by a power of 10, such that when the magnitude output of the overall system represents the significand of the current response and the exponent is the gain of the amplifier. In order to have a signal around 10nA, which is most appropriate for the filter section, this amplifier should provide a gain of around 100. From this the actual magnitude can be calculated as $SIG \times 10^E XP$. Figure 31 demonstrates that this amplifier is amplifying linearly, albeit with a gain somewhat less than 100. This means that the output signal may not be correctly calculated, although this can easily be adapted by adjusting the bias currents in the amplifier. If the error is small enough, it can be ignored, as the value of an LSB in the digital output is arbitrary. Once a reference is determined, all other measurements will be linear around this.

6.3 Heterodyne

The input and output of the heterodyne are demonstrated in figure 32. From this it is clear that the samples accurately follow the input signal. Equally, it is visible here the effect that undersampling has, with the lower frequency component also visible.

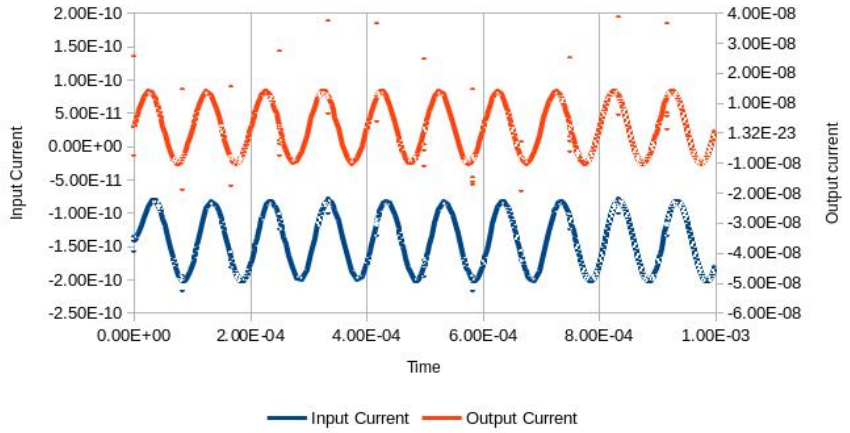

Figure 31: Current at input and output of variable gain amplifier

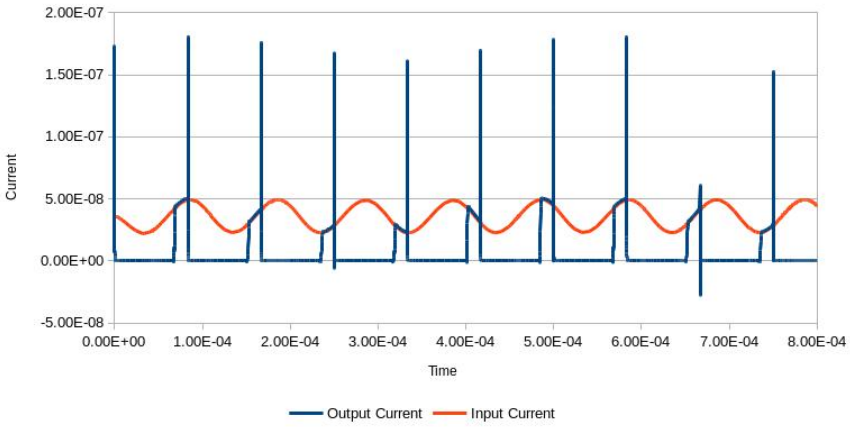

Figure 32: Current at input and output of heterodyne sampler

6.4 Filter

The filter is perhaps the most important component in terms of recovering an accurate measurement. Without sufficient rejection of harmonic and image frequencies, the output of this will not be a pure sine wave which will substantially compromise the mathematics later in the system. Despite the artifacts present from the later integrate and fire circuits which are now visible in the signal, the 2kHz sine wave is clearly visible in figure 33. When compared to the input signal which has come from the sampler it is also clear that these two are in phase with each other.

The primary point to note here is the large settling time at the beginning of the simulation. This has a notable effect on the integrate and fire circuit, as will be discussed.

6.5 Integrate and fire circuit

The integrate and fire circuits have been demonstrated to function well, however with the whole circuit it is necessary to cope with the startup settling time from the filter. Because of this initially high input, the offset finder requires a comparatively long time to adjust to the correct value, during which there is a significant bias in the signal. During part of this time the current input signal is out of the circuit's range, and results in a mid-rail output voltage on the negative output, with neither output firing as is visible in figure 34. This results in a significant power consumption and the large bias negatively affects the integration stage.

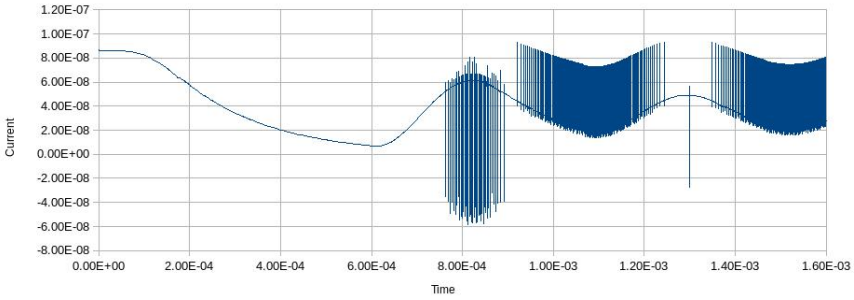

Figure 33: Current at output of filter

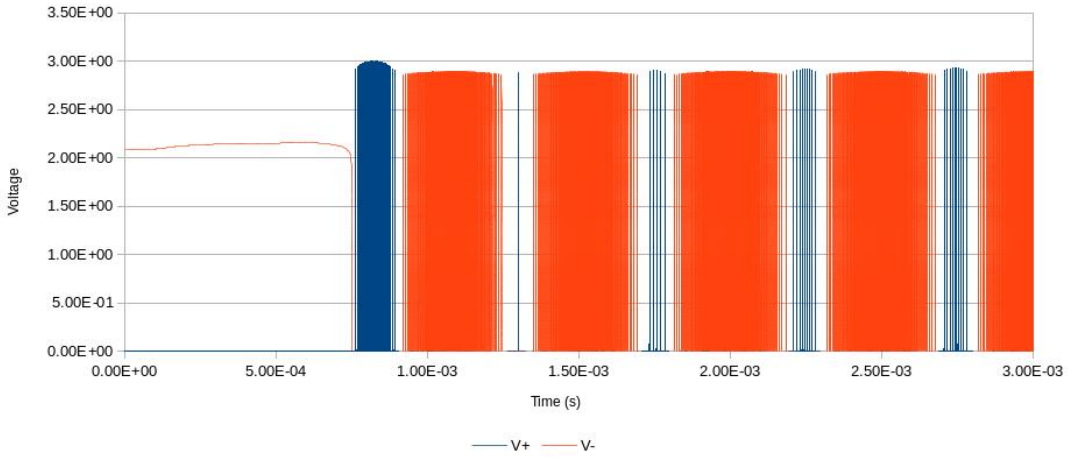

Figure 34: Voltage output of integrate and fire circuit

6.6 Digital output

As has been discussed before, a bias in the signal results in a rapidly increasing output when the signal is integrated. This is visible in the digital output data in figure 35. The effect of this is that the counter overflows at least once within each cycle making it impossible to calculate any data from it. However if the peak-detector is given time to settle, it does become possible to make a reading.

Another imperfection which is important to note is how the counter tends to glitch resulting in multiple bits to change erroneously and the output to jump. This is probably the effect of an insufficient set up time, and the probability of this happening can be decreased with the addition of further inverters on the clock input of the counter. If these occur too often then it becomes very difficult to extract a reliable reading. In all, some further adjustments need to be performed to make this entirely usable.

6.7 Phase and magnitude calculation

The first case where a measurable cycle occurs is at 2ms, as the data between 2.125ms and 2.625ms is all contained without an overflow. There is a second at 4m. At this point the data received and the results calculated are shown in 36. With only one example wave, it is difficult to discern whether the amplitude calculation is correct, (note that the inclusion now of the preceeding amplifier, and filter etc. means that the reference determined in the implementation section is no longer valid). However, it is notable that there is a non-negligible difference between the two samples.

On the other hand we can interpret the phase calculation. If we acknowledge that using an

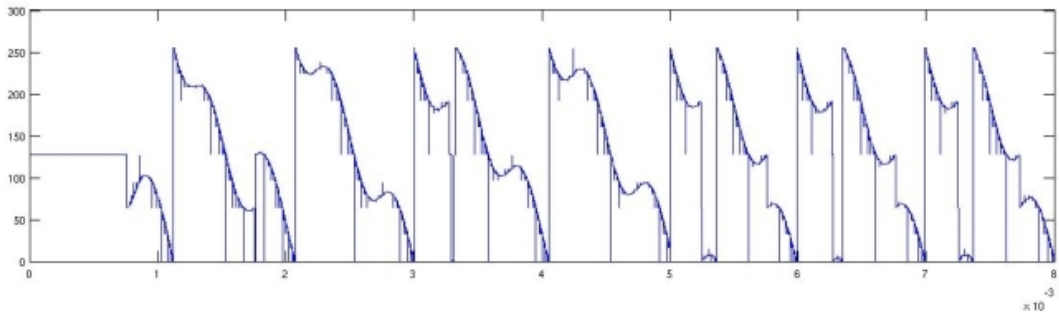

Figure 35: Digital output of circuit

undersampling heterodyne will result in a 90 degree phase shift and we accommodate for this, then the calculated values become 3.12^o and 5.5^o. This implies an error of 0.87% and 1.52% respectively which is a good accuracy.

Sample	S1 time (ms)	S1 Values	S2 time (ms)	S2 Values
$V_{\frac{T}{4}}$	2.125	234	4.125	224
$V_{\frac{T}{2}}$	2.25	228	4.25	225
$V_{\frac{3T}{4}}$	2.375	219	4.375	218
$V_{\frac{5T}{4}}$	2.625	82	4.625	87
V_I	-	3	-	6
V_Q	-	-55	-	-62.5
$\|I_r\|$	-	55	-	62.79
$\angle I_r$	-	-86.88 ^{o}	-	-84.5^o

Figure 36: Counter value in LSBs at each sample required for arithmetic

6.8 Overall power consumption

Besides the functionality of the system, another requirement of this project was that the device consume less than 1mW of power. Therefore the transient of the current leaving the vdd voltage source is shown in figure 37. As was discussed earlier, after startup the filter causes an indeterminate state in the integrate and fire circuit, which results in increased power consumption. This can be quantified now as being around 725uA, representing a power consumption of 2.393mA, which is clearly substantially over the power budget. However, once normal operation begins, at around 750us, although the current drawn from the power supply is very spikey (due to the integrate and fire circuit) it averages to 58.3uA, which represents a power consumption of around 193uW on average. Unlike during startup, this is substantially below the power target.

How much the start up current matters is very much dependant on application. If the intention is that the circuit will remain on relatively continuously then it is a minor issue. On the other hand in applications where the circuit will be turned on in order to make a few readings, before being switched off again, this high startup power becomes more problematic.

7 Evaluation

The criteria for success for this project as outlined at the end of the literature review, are shown again in figure 38. The basic requirements of the project were that the circuit could measure impedances up to 100kHz and 1GΩ while consuming less than 1mW of power. To an extent these have all been met, in that it has been demonstrated that a 100kHz signal, of 10pA amplitude

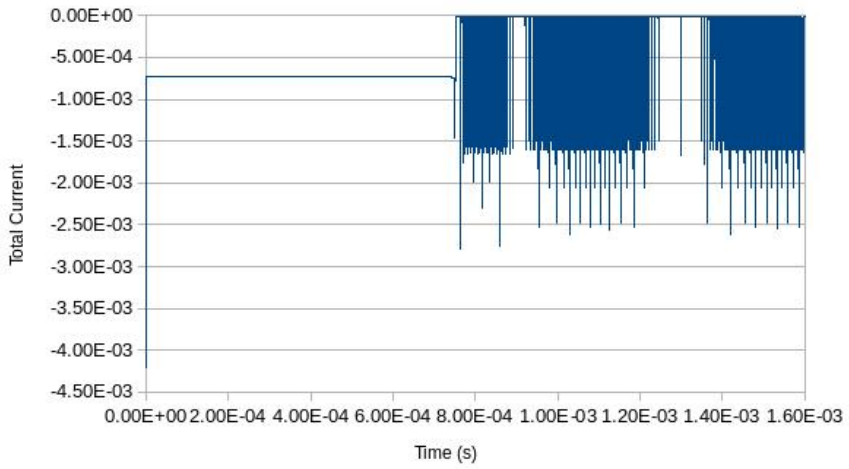

Figure 37: Total current consumed by circuit

(the amplitude expected of the response of a 1GΩ impedance) can be amplified to the working amplitude of 10nA and frequency 2kHz by the VGA and heterodyne respectively. Therefore in theory these requirements can be met. The caveat to this is that, as discussed during the circuit implementation, the Op-amp which controls the stimulus has an input capacitance which is non-negligable at this frequency and impedance. This capacitance appears in parallel with the measured impedance, and distorts the result. Consequently it can be argued that while the stimulus circuit requires some improvement, the current readout circuit does in fact succeed in meeting these basic requirements. This is part of the reason that the overall system test was performed at 10kHz and 100MΩ, along with the fact that these values were seen as representative of a typical measurement, rather than extremes, which is more appropriate in lieu of a complete simulation across the whole field of possible conditions. There is also further work needed in decreasing the offset of the signal before integration and reducing the number of glitches in the UP/DOWN counter in order to make the results more readable.

With regard to the power consumption, the circuit was shown to substantially over perform, therefore providing confidence that the remaining digital peripherals as well as the stimulus DAC can be added later without significant risk of compromising the power budget.

Property	min spec	reference
Power Consumption	1mW	Project Brief
Max Impedance	1GΩ	Project Brief
Min Impedance	10kΩ	[24]
Max Frequency	100kHz	Project Brief
Min Frequency	5kHz	[24][9][2]
Max Relative Error	8%	[2][24][9]

Figure 38: Summary of requirements

With regard to the secondary aims, for lower impedance, it has been shown that the VGA can achieve a sufficient attenuation to bring the amplitude down sufficiently to be usable by the log-domain filter (which requires inputs to be less than 100nA). Equally it has been demonstrated in simulink that the heterodyne should be able to attain frequencies as low as 2.5kHz perfectly, and frequencies even below that are measurable if the data point at 2kHz itself can be ignored. There is therefore no obvious reason why this shouldn't function in the circuit itself, however more comprehensive testing is required to confirm this. The only dependancy is the bandwidth of the filter. It has been shown that this is sufficient for the higher frequency signals to be

accurately recovered, however as it is not quite to the specification desired, it is possible this may compromise the ability to measure low frequency signals. However with a measured Q-factor greater than 1, recovering data from a 5kHz stimulus should still be possible.

Equally it has been shown that the error in calculation can be very low, however in order to demonstrate that this is constant across all operating requirements a much greater number of tests will need to be run. Given the complexity of this system each simulation can take a substantial amount of time to run, which is why this comprehensiveness has not been applied for this report. However the results shown so far do give hope that they would be successful.

8 Conclusions and Further Work

In hindsight this project was an ambitious undertaking, which required not only a large number of well designed analogue systems, but relatively substantial digital and control circuitry as well. Although the digital design does not require any particularly novel solutions, it is not an insubstantial part of the task. Consequently given the time available, it was sensible to focus on the intended system as a whole, but then to only implement the analogue blocks given that it is on the performance of these that the capability of the device as a whole depends.

Ultimately this was done relatively successfully, given that all the sections of the readout circuit individually have been shown to function successfully, and can meet all the specifications. Although some problems are faced with the circuit as a whole, the correct results can still be read, albeit with some manual discretion about which time samples to use. Were the offset and glitching issues solved, however, it seems probable that this circuit could provide the expected results reliably.

Although there are a number of system topologies which can, theoretically, calculate the magnitude and phase of a wide range of responses, a major part of this project was determining which of these are simple to implement. The best example of this was in the heterodyne, in which the most intuitive topology was to multiply whole sinusoids. However at the implementation stage, it was clear that sampling would be a much simpler approach and lead to less distortion, assuming a high quality output filter. The gradual change in topology, after initial system design was first finished, was a major part of the time costs of this project. With greater foresight this time could have been reduced.

One of the effects of having a deadline is the need to definitively decide to stay with the design as it is at a certain time, in order to complete the remaining work in time. This does not preclude from further thoughts on how the project could be improved. Beyond simply fixing the remaining issues in the design, and adding the necessary control and arithmetic circuits, there are a few adaptions which would be worth investigating.

The first of these is considering whether the integrator is necessary at all. Although it is present in [2], on which this circuit is based, the principle simply involves sampling the response at exactly the frequency necessary to gain the information required to determine its magnitude and phase. Given integration does not add or remove information, this should be possible without. If this is the case then an alternative to the integrate and fire and UP/DOWN counter implementation would be needed for digitisation, which has an equal or better linearity. One option for this would be to use a successive approximation ADC. These require an internal DAC to back-convert the result to compare with the input; one option for which being a current steering DAC. The primary source of non-linearity in these is the transimpedance amplifier, however given in this case the signal is current mode, the amplifier can be dispensed with, such that the input current is compared with the total steered current. Assuming reasonable quality current sources, this should be highly linear and very low power, as if the input 100nA then the DAC will only need consume a further 300nA. This would also remove the high frequency components introduced by spiking. Although the propagation of these backwards along the circuit does not seem to have had any negative effects in a real circuit they would capacitively couple with parts of the circuit which may be completely unrelated yet are geographically close on the silicon, which may cause unexpected problems.

While this report outlines significant progress towards the aims of this project, and can achieve most if not all of the specifications, further work is required primarily in improving the integrity of the integrating and digitising circuits. Besides that, an improved amplifier in the stimulus circuit would be desirable, and for actual application the appropriate arithmetic logic and a suitable control state-machine should be added. From this the circuit should be fully functional, although there are also multiple areas which could be investigated with an aim to improving the system.

9 Acknowledgements

Firstly I would like to acknowledge my supervisors, Dr Georgiou and Dr Ghoreishizadeh for their help and technical support throughout this project. Also Julian Tyrrell from Dialog Semiconducto for lending me his copy of "Analogue IC design", which has both valuable and interesting during this project and in my general university study. Lastly and definitely not leastly, my fiancee, Malyha Ahmed, for her constant encouragement and for keeping me sane as I bumped into the various technical dead-ends and pushed through long days of typing.

References

[1] A. Lymberis, "Wearable smart systems: From technologies to integrated systems," presented at the IEEE Annual Conf (2011), Boston, MA, p. 503.

[2] J. Gu, H. Yao, K. Wang, B. Parviz, and B. Otis, "A 10ua on-chip electrochemical impedance spectroscopy system for wearables/implantables," presented at the Solid-State Circuits Conference (A-SSCC), 2014 IEEE Asian (2014), KaoHsiung, pp. 309–312.

[3] S. Narasimhan, J. Park, and S. Bhunia, "Digital signal processing in bio-implantable systems: Design challenges and emerging solutions," presented at the Quality Electronic Design (ASQED), 2010 2nd Asia Symposium on (2010), Penang, pp. 223–229.

[4] G. L. Faycal Mounaim and M. Sawan, "Long-term monitoring of electrochemical parameters from stimulated neural tissues," presented at the Biomedical Circuits and Systems Conference, 2007. BIOCAS 2007. IEEE (2007), Montreal, Que., pp. 25–28.

[5] E. Wilkins and M. Wilkins, "Implantable glucose sensor," *Journal of Biomedical Engineering*, vol. 5, no. 4, pp. 309–315, 1983.

[6] N. Yang, R. Gautham, B. Liang, R. Shah, C. Szyman, M. Miller, J. Wang, Y. Li, W. Morgan, P. Chen, *et al.*, *Application of electrochemical impedance spectroscopy in sensor systems, devices, and related methods*, WO Patent App. PCT/US2013/042,680, Dec. 2013. [Online]. Available: `http://www.google.com/patents/WO2013184416A2?cl=en`.

[7] B. Pejcic and R. de Marco, "Impedance spectroscopy: Over 35 years of electrochemical sensor optimization," *Electrochimica Acta*, vol. 51, no. 28, pp. 6217–6229, 2006.

[8] V. Valente, D. Jiang, and A. Demosthenous, "Design of a wideband cmos impedance spectroscopy asic analog front-end for multichannel biosensor interfaces," presented at the EMBC, 37th IEEE Annual Conf. (2015), Milan, pp. 885–888.

[9] P. Kassanos, I. Triantis, and A. Demosthenous, "A cmos magnitude/phase measurement chip for impedance spectroscopy," *Sensors Journal, IEEE*, vol. 13, no. 6, pp. 2229–2236, 2013.

[10] E. Azzarello, E. Masi, and S. Mancuso, "Electrochemical impedance spectroscopy," in *Plant Electrophysiology*, A. Volkov, Ed., Berlin: Springer Berlin Heidelberg, 2012, ch. 9, pp. 205–223.

[11] K. Folgerø, "Broad-band dielectric spectroscopy of low-permittivity fluids using one measurement cell," *IEEE Transactions on Instrumentation and Measurement*, vol. 47, no. 4, pp. 881–885, 1998.

[12] A. A. S. Akmal, H. Borsi, E. Gockenbach, and V. Wasserberg, "Dielectric behavior of insulating liquids at very low frequency," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 13, no. 3, pp. 532–538, 2006.

[13] V. L. Mironov, "Spectral dielectric properties of moist soils in the microwave band," presented at the Proceedings of the International Geoscience and Remote Sensing Symposium (2004), vol. 5, pp. 3474–3477.

[14] S. Abdalla, "Low frequency dielectric properties of human blood," *IEEE Transactions on Nanobioscience*, vol. 10, no. 2, pp. 113–120, 2011.

[15] A. Lonappan, V. Thomas, G. Bindu, C. Rajasekaran, and K. Mathew, "Analysis of human cerebro spial fluid at the ism band of frequencies," *Journal of Electromagnetic Waves and Applications*, vol. 20, no. 6, pp. 773–779, 2006.

[16] M. Thill, K. Röder, K. Diedrich, and C. Dittmer, "Intraoperative assessment of surgical margins during breast conserving surgery of ductal carcinoma in situ by use of radiofrequency spectroscopy," *The Breast*, vol. 20, pp. 579–580, 2011.

[17] P. Michaels, S. Ciampi, C. Y. Yean, and J. J. Gooding, "Target dna recognition using electrochemical impedance spectroscopy," presented at the International Conference on Nanoscience and Nanotechnology (2010), Sydney, NSW, pp. 282–284.

[18] Y. Hayashi, N. Miura, J. Isobe, N. Shinyashiki, and S. Yagihara, "Molecular dynamics of hinge-bending motion of igg vanishing with hydrolysis by papain," *Biophysical Journal*, vol. 79, pp. 1023–1029, 2000.

[19] A. Oleinikova, P. Sasisanker, and H. Weingärtner, "What can really be learned from dielectric spectroscopy of protein solutions? a case study of ribonuclease a," *Journal of Physical Chemistry B*, vol. 108, no. 24, pp. 8467–8474, 2004.

[20] B. A. Mazzeo and A. J. Flewitt, "Observation of protein-protein interaction by dielectric relaxation spectroscopy of protein solutions for biosensor application," *Applied Physics Letters*, vol. 90, 2007.

[21] A. Caduff, M. S. Talary, M. Mueller, F. Dewarrat, J. Klisic, M. Donath, L. Heinemann, and W. A. Stahel, "Non-invasive glucose monitoring in patients with type 1 diabetes: A multisensor system combining sensors for dielectric and optical characterisation of skin," *Biosensors and Bioelectronics*, vol. 24, pp. 2778–2784, 2009.

[22] A. Caduff, F. Dewarrat, M. Talary, G. Stalder, L. Heinemann, and Y. Feldman, "Non-invasive glucose monitoring in patients with diabetes: A novel system based on impedance spectroscopy," *Biosensors and Bioelectronics*, vol. 22, pp. 298–604, 2006.

[23] A. Manickam, A. Chevalier, M. McDermott, A. Ellington, and A. Hassibi, "A cmos electrochemical impedance spectroscopy (eis) biosensor array," *Biomedical Circuits and Systems, IEEE*, vol. 4, no. 6, pp. 379–390, 2010.

[24] J. Guo, W. Ng, J. Yuan, and M. Chan, "A 51fa/hz0.5 low power heterodyne impedance analyzer for electrochemical impedance spectroscopy," presented at the VLSI Circuits (VLSIC), 2013 Symposium on (2013), Kyoto, pp. C56–C57.

[25] Y. Sundarasaradula and A. Thanachayanont, "A 1-v, 6-nw programmable 4th-order bandpass filter for biomedical applications," presented at the New Circuits and Systems Conference (NEWCAS), 2015 IEEE 13th International (2015), Grenoble, pp. 1–4.

[26] D. R. Frey, "Log-domain filtering: An approach to current mode filtering," *IEE Proceedings G - Circuits, Devices and Systems*, vol. 160, no. 6, pp. 406–416, 1993.

[27] E. M. Drakakis, A. J. Payne, and C. Tomazou, "Log-domain state-space: A systematic transistor-level approach for log-domain filtering," *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, vol. 46, no. 3, pp. 290–305, 1999.

[28] L. M. Alexander Renner and A. Koenig, "Architecture and cell design of the iisic cmos-chip serving as a front-end for integrated impedance spectroscopy," presented at the Sensors and Measuring Systems 2014; 17. ITG/GMA Symposium (2014), Neuremberg, Germany, pp. 1–6.

[29] C. Toumazou, F. J. Lidgey, and D. G. Haigh, "Analogue ic design, the current-mode approach," in. Peter Peregrinus Ltd, 1990, ch. 2, p. 19.

[30] M.-W. Kwon, H. Kim, J. Park, and B.-G. Park, "Integrate-and-fire neuron circuit and synaptic device with floating body mosfets," *Journal of Semiconductor Technology and Science*, vol. 14, no. 6, pp. 755–759, 2014.

[31] R. Sarpeshkar, L. Watts, and C. Mead, "Refractory neuron circuits," California Institute of Technology, Tech. Rep., 1992.

[32] M. Rastogi, A. S. Alvarado, J. G. Harris, and J. C. Principe, "Integrate and fire circuit as an adc replacement," presented at the 2011 IEEE International Symposium of Circuits and Systems (ISCAS) (2011), Rio de Janeiro, pp. 2421–2424.