

MASTER

Robust detection platform and algorithms for magneto-resistive biosensors

de Boer, B.M.

Award date:
2007

[Link to publication](#)

Disclaimer

This document contains a student thesis (bachelor's or master's), as authored by a student at Eindhoven University of Technology. Student theses are made available in the TU/e repository upon obtaining the required degree. The grade received is not published on the document as presented in the repository. The required complexity or quality of research of student theses may vary by program, and the required minimum study period may vary in duration.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

9302

Eindhoven University of Technology
Department of Electrical Engineering
Signal Processing Systems

Vertrouwelijk

TU/e technische
universiteit
eindhoven

Robust Detection Platform and Algorithms for Magneto-resistive Biosensors

by B.M. de Boer

Master of Science thesis

Project period: 08 2006 – 08 2007

Report Number: 16-07

Commissioned by: Philips Research

Supervisors:

Prof. dr. ir. J.W.M. Bergmans (TU/e)

Dr. ir. J. Veen (Philips Research)

Additional Commission members:

Dr. ir. P.J.M. Cluitmans (TU/e)

Contents

- 1. Introduction..... 7**
 - 1.1. Biosensing and its Applications..... 7
 - 1.2. Immunoassays..... 9
 - 1.2.1. Sandwich Assay..... 9
 - 1.2.2. Inhibition Assay..... 10
 - 1.3. Magnetic Detection..... 11
 - 1.4. Magnetic Actuation..... 14
 - 1.5. Problem Formulation and Report Outline 15

- 2. State of the Art 17**
 - 2.1. Description of Current Platform 17
 - 2.1.1. Excitation and Sensing..... 17
 - 2.1.2. Performance 22
 - 2.2. Limitations of Current Platform 25

- 3. New Platform..... 27**
 - 3.1. Functional Description 28
 - 3.2. Analog Front-end Architecture..... 30
 - 3.2.1. Excitation Current Source..... 30
 - 3.2.2. Sense Current Source..... 31
 - 3.2.3. Pre-amplifier..... 33
 - 3.3. Analog Front-end Measurement Results..... 37

- 4. Combined Sensor and Front-end Model 42**
 - 4.1. Model Description..... 42
 - 4.2. Model Parameters 44
 - 4.3. Magnetic Signal and Noise Model..... 47

- 5. Robust Detection Algorithms 49**
 - 5.1. Description of Robust Detection Algorithms 49
 - 5.1.1. Algorithm 1: Reference Sensor Scheme..... 49
 - 5.1.2. Algorithm 2: Vector Decomposition Scheme 51
 - 5.2. Joint Phase and Amplitude Estimation 53
 - 5.2.1. Gradient Based Least Squares Adaptation..... 53

5.2.2.	Gradient Based Closed Loop Phase Estimation.....	54
5.2.3.	Gradient Based Closed Loop Amplitude Estimation.....	56
5.2.4.	Joint Phase and Amplitude Estimation	58
5.2.5.	Discretisation.....	59
5.2.6.	Dynamic Behaviour of the Adaptation Loops.....	60
5.2.7.	Parameter Estimate Gradient Noise.....	62
5.2.8.	Estimate Reliability	65
5.3.	Algorithm Simulation Results	66
6.	Digital Back-end Implementation and Verification.....	70
6.1.	Digital Back-end Architecture	70
6.1.1.	Waveform Generator	71
6.1.2.	Joint Phase and Amplitude Estimator.....	73
6.1.3.	Digital Pre-amplifier.....	77
6.2.	Digital Back-end Simulation Results.....	78
6.2.1.	Waveform Generator	78
6.2.2.	Joint Phase and Amplitude Estimator.....	81
6.3.	Digital Back-end Measurement Results.....	84
6.3.1.	Waveform Generator	84
6.3.2.	Joint Phase and Amplitude Estimator.....	86
6.3.3.	Digital Pre-amplifier.....	88
7.	Experimental Results	91
7.1.	Measurement Methods	91
7.2.	Heating Experiment Measurement Results	93
7.2.1.	Heating of the Sensor.....	93
7.2.2.	Heating of the Sensor and Platform.....	96
8.	Summary, Conclusions and Recommendations	98
8.1.	Summary and Conclusions	98
8.2.	Recommendations.....	101
	References	104

1. Introduction

Molecular diagnostics, i.e. diagnostics based on the detection of specific molecular markers for certain predispositions or conditions (contrary to conventional diagnostics based on symptoms), is a field that is very much alive, with new disease markers being discovered on a regular basis. It covers a wide variety of tests, ranging from simple pregnancy tests in the form of a dipstick for home use, to high-end large DNA arrays used in hospital laboratories.

In the field of molecular diagnostics a trend can be observed towards more and more rapid testing at point-of-care, e.g. patient bed-side, rather than in the centralised laboratories where the more sensitive tests are conventionally performed. Advancing molecular diagnostic technologies from a laboratory setting to near-patient tests requires the availability of portable and sensitive bio analytical equipment. Ideally, such devices are easy to use and low-cost and combine a low detection limit with short analysis time. In order to meet these requirements, a high degree of integration and miniaturization is essential.

Magnetic biochips are promising candidates for sensitive detection of low concentrations of target molecules in body fluids [1,2]. Combined with suitable biochemistry to selectively attach labels, magnetic biosensors allow for miniaturized molecular diagnostics in an array format. Within the biosensor project, running at Philips Research, magneto resistive biosensors and a complementary readout platform are developed to measure the concentration of specific molecules in bodily fluid samples. The biosensor platform is specifically targeted at point-of-care applications.

The following sections can be considered to be a brief introduction on biosensing in general and on the Philips magnetic biosensor. This chapter concludes with an overview of the topics that are discussed in the remainder of this report.

1.1. Biosensing and its Applications

In general, with the exception of label-free biosensors, it can be said that biosensors implement two subsequent steps to derive the concentrations of relevant molecules:

1. A biochemical assay, in which relevant molecules are labelled and bound to the sensitive area of the biosensor.
2. A detection procedure to estimate the label density in order to determine the concentration of the relevant molecule in the sample.

The following sections briefly discuss how both steps are implemented in the case of the Philips magnetic biosensor.

Applications that are currently a topic of research within Philips include, but are by no means restricted to, the measurement of the concentration of proteins in blood, and testing for drugs-of-abuse in saliva.

More specifically, the blood test may be used to detect cardiac markers, i.e. proteins that provide information on the condition of the heart. A number of cardiac markers are known and can already be detected in centralised laboratories. It would however be beneficial, both in terms of patient convenience and in terms of cost reduction for the healthcare system, if a patient experiencing chest pains could be tested in a matter of minutes at the General Practitioner office, as opposed to being sent to a centralised laboratory. Such an application requires more or less handheld analytical equipment in combination with low cost, disposable biosensors, preferably in array format to be able to test for a set of different markers, also known as a panel, simultaneously. Figure 1 shows an example of the concentration of three cardiac markers in blood as a result of time from the moment of occurrence of a heart attack. It can be seen that certain markers can be detected at an earlier stage than others. It can furthermore be seen that the ratios between the different marker concentrations convey information on the time of occurrence of the heart attack.




Figure 1: The development of the concentration (a.u.) of three different cardiac markers over time from the moment that the first symptoms of a heart attack occurred (source: National Academy of Clinical Biochemistry).

A second, more advanced application of the blood test is the detection of tumour markers. The concentration of these markers is generally much lower than the concentration of cardiac markers, but can indicate the development of tumours in the body, even before the patient develops symptoms, such that treatment can be provided at an earlier stage. This is important since treatment of tumours is in general more effective if started at an earlier stage of development. Fast and low-cost tests would allow for routine screening for these markers.

The ability to test for drugs-of-abuse in saliva allows for a road-side drug screening application. The concentrations to be measured here are in general much higher than for the blood test applications, but the measurement time should also be much smaller, i.e. smaller than one minute to allow through-the-window testing similar to alcohol tests. The

required analytical performance is therefore similar. Again, the goal is to simultaneously test for a set of for example 5 different drugs.

1.2. Immunoassays

In this section the immunoassays are described that can be used for the different applications described above. Immunoassays can be defined as biochemical assays, or laboratory techniques, that make use of the binding between a specific molecule of interest and its antibody in order to identify and quantify the molecule of interest in a sample. Proteins are generally very large molecules consisting of hundreds or even thousands of atoms, while drugs-of-abuse typically consist of much smaller molecules. Since both applications test for very differently sized molecules, a different type of assay is required in both cases. Both assays are briefly discussed in the following subsections.

1.2.1. Sandwich Assay

The sandwich assay can be employed for tests involving large molecules, such as proteins, that provide multiple binding sites for antibodies. Figure 2 depicts a cross sectional view of a biosensor. It consists of a silicon chip containing current conductors and a giant magneto-resistive (GMR) element, which will be discussed later. A gold surface on top of the sensor is used to immobilize antibodies that are specifically sensitive to the molecules of which the concentration will be determined, i.e. the analyte. If a fluid sample that is applied to the sensor contains the analyte, binding occurs between the analyte and the antibodies. Subsequently, superparamagnetic particles, i.e. the labels, coated with a second antibody are applied to the sensor, which will also bind to the analyte. By using this so-called sandwich assay the labels, also referred to as beads, are specifically bound to the surface of the sensor, and the density of bound beads is related to the concentration of the analyte in the sample. The assay shown in Figure 2 is typical for the detection of proteins.




Figure 2: Cross sectional view of a biosensor using a sandwich assay. The photo inset shows the magnetic labels that are typically used in the assays.

1.2.2. Inhibition Assay

Since the molecules of drugs-of-abuse are in general too small to allow binding of two antibodies simultaneously, a sandwich assay cannot be used for the drug screening application. Instead, an inhibition assay is used, which is schematically depicted in Figure 3. The relevant molecules e.g. morphine, are in this case bound to the sensor surface. Upon application of the fluid sample to the sensor, it is mixed with labels coated with antibodies against morphine. If morphine is not present in the sample, the labels will bind to the sensor surface, such that they will be detected. However, if morphine is present, it will bind to the antibodies on the labels. This causes the labels to saturate, such that no free binding sites are left and binding to the morphine on the sensor surface is inhibited.




Figure 3: Cross sectional view of a biosensor using an inhibition assay.

For the inhibition assay the response as a function of drug concentration is therefore characterised by a $1-x$ behaviour, until all labels are fully saturated and no beads bind to the surface. This behaviour can clearly be observed in Figure 4, which shows the response of the sensor as a result of morphine concentration in a buffer solution.




Figure 4: Change in GMR signal as a result of the concentration of morphine in a buffer solution [3].

1.3. Magnetic Detection

The use of superparamagnetic particles as labels for molecular detection has a number of distinct advantages. Firstly, biological samples contain hardly any magnetic material;

therefore the detection background is intrinsically low. Secondly, magnetic sensors for label detection can be integrated on-chip, thus enabling extreme miniaturization of diagnostic systems. Thirdly, magnetic labels are already applied widely to speed up immunoassays, which are otherwise diffusion-limited, by actively concentrating targets within a sensing zone [4]. Mainly for this purpose, magnetic particles are available commercially. In addition, magnetic actuation techniques for particle transportation [5-7] and magnetic washing [8] can also be integrated on-chip.

Giant magneto-resistive (GMR) sensors have already been identified as suitable devices for detecting magnetic labels [2,5,7,9,10]. GMR sensors have good sensitivity to small magnetic fields, such that high signal-to-noise ratios can be realized with submicron-sized beads. Moreover, stacks of appropriate materials can be deposited on silicon substrates with micron scale surface dimensions. These properties make GMR sensors very well suited for application in fully integrated magnetic biosensors.

Magneto-resistance is the property of some magnetic materials to change their electrical resistance in the presence of an external magnetic field. Figure 5 shows a simplified schematic representation of a giant magneto-resistance element. The giant magneto-resistance (GMR) effect, which occurs in magnetic multilayer stacks with nanometer-scale layer thicknesses, is caused by the spin-dependent scattering of electrons between a layer with freely moveable magnetization direction, i.e. the free layer, and one with a fixed magnetization direction, i.e. the pinned layer. The resistance of the GMR depends on the angle (φ) of the direction of magnetization of the free layer with respect to that of the pinned layer, and is given by [11]

$$R(\varphi) = R_{\varphi=0} + \frac{\Delta R}{2}(1 - \cos \varphi) \quad (1)$$

where $R_{\varphi=0}$ represents the resistance when the direction of magnetisation of the pinned and free layers is the same and ΔR represents the total change in resistance when φ changes from 0 to 180 degrees.




Figure 5: Simplified schematic representation of a GMR element. The preferred magnetisation of the free layer is along the y-axis.

Figure 6 depicts a typical GMR curve, i.e. the GMR resistance vs. applied magnetic field. It can be seen that for small applied magnetic fields the GMR resistance depends linearly on the applied field strength, while for larger field strengths the sensor saturates. For small field strengths, the resistance can thus be approximated by

$$R = R_0 + S \cdot H \quad (2)$$

in which S is called the sensitivity of the GMR and corresponds to the slope of the GMR curve, while R_0 is the resistance in the absence of a magnetic field. In this case the free layer will be magnetised along the y-axis, such that

$$R_0 = R_{\varphi=0} + \frac{\Delta R}{2} \quad (3)$$




Figure 6: Typical GMR curve. The sensitivity of the GMR is the slope of the curve.

Figure 7 schematically depicts the excitation of beads bound to the surface of the sensor. Superparamagnetic beads are used that consist of a biocompatible polymer matrix containing iron oxide (usually Fe_2O_3 or Fe_3O_4) grains. The beads may, upon magnetization, be considered as a single dipole. After connection of the sensor to the measurement platform, an excitation current is applied to the current conductors in order to generate a magnetic field. This excitation field induces a magnetic moment in the beads. The field of the magnetised beads subsequently induces a resistance change in the GMR. Because the GMR has the property that it is only sensitive to magnetic field components in the horizontal (in-plane) direction, it is highly sensitive to the field of the magnetized beads, while being insensitive to the primary field of the current conductors. The resistance of the GMR is finally determined by measuring the voltage across the GMR upon application of a constant sense current.




Figure 7: Excitation of a biosensor. The field lines shown are those of the magnetic field from the current carrying conductor and the detected stray field of a magnetized bead.

1.4. Magnetic Actuation

Magnetic actuation can be employed to speed up biochemical assays that are otherwise diffusion limited. This greatly reduces the required analysis time. Reduction of the analysis time is important especially for time critical applications, such as through-the-window drugs-of-abuse testing, where analysis times longer than one minute are unacceptable. Several methods of actuation are possible, however, a convenient method is to use on chip conductors to attract beads by means of magnetophoresis. The magnetic force acting on a bead, which can be considered as a dipole, as a result of an inhomogeneous magnetic field is given by

$$\vec{F} = \nabla(\mu \vec{H} \cdot \vec{M}) \quad (4)$$

in which \vec{H} is the applied magnetic field, \vec{M} is the magnetic moment induced in the bead, and μ is the magnetic permeability. Since the magnetic moment of the bead aligns with the applied magnetic field, the force acts in the direction of the gradient of the applied field. In case of a current carrying conductor, the gradient is towards the conductor, such that beads are drawn towards the conductor.

Apart from reducing analysis time, actuation can be used to concentrate labels in the most sensitive area of the sensor in order to increase the signal per bead. This is, as will become clear in Subsection 2.1.2, the area between the excitation conductors and the GMR. Attraction conductors that are located under the most sensitive area of the sensor, can be used to actuate beads to this area.

Furthermore, magnetic actuation can be used to remove non-specifically bound labels from the sensor surface prior to detection of the surface bound beads, i.e. magnetic washing.

1.5. Problem Formulation and Report Outline

This report discusses the work that was done on the development of a flexible and robust biosensor platform and detection algorithms. This platform is the successor of the platform that is currently used within the project. Although the current platform achieves sufficiently high signal-to-noise ratio (SNR), the corresponding low limit of detection can only be achieved in laboratory settings, since the platform lacks sufficient robustness against environmental variations like variations in temperature and (external) magnetic fields. The envisioned application of the biosensor at points-of-care, however, means that there is no prior knowledge of the environment in which the platform is used. This requires a more robust system that can be used even in highly varying environmental conditions, which can for instance be expected for the road-side drug screening application. The new platform should finally be more flexible, such that newly developed and more complex excitation and detection methods can be implemented and evaluated in a short time span.

A description of the state-of-the-art, i.e. the currently used magnetic biosensor platform, is given in Chapter 2. This chapter should give the reader insight in the general method of excitation of the Philips magnetic biosensor and the detection scheme that is used. It also discusses the performance of the current platform in terms of the electronic SNR and the limit of detection of the system. Chapter 2 ends with a description of the limitations of the platform and the main causes of the current lack of robustness.

In Chapter 3 a functional description is given of the new robust hardware platform. Furthermore, the architecture of the part of the system that is most critical in terms of SNR and robustness, i.e. the analog front-end, is discussed in more detail. This chapter should give the reader insight into the most important design considerations that were made in order to optimise SNR and robustness. Chapter 3 ends with some measurement results obtained with the new front-end and a discussion on its performance.

In order to be able to test the robust detection algorithms in an environment that allows simulation with short iteration times and easy simulation parameter adaptation, a model of the combined new analog front-end and biosensor was developed and implemented in Matlab Simulink. A description of this model and a discussion on how the model parameters were obtained is given in Chapter 4.

Chapter 5 describes two robust detection algorithms and discusses the results of simulations performed in Simulink with the combined sensor and front-end model of Chapter 4. It also discusses the joint phase and amplitude estimation that provides input for both algorithms.

Chapter 6 discusses the digital back-end. The architecture and implementation of its key components are described. Furthermore, the results of simulations of its key components are discussed. It ends with some measurement results of the digital back-end obtained with the new platform.

Chapter 7 discusses the results of heating experiments that were conducted with the new platform in order to test if the combination of new hardware and detection algorithms have improved the robustness of the platform output signal against temperature variations.

Finally, a summary, conclusions and further recommendations are given in Chapter 8.

2. State of the Art

This chapter describes the state-of-the-art of the magnetic biosensor electronics and its limitations. A thorough understanding of the state-of-the-art is necessary to understand what should be improved, as well as to fully appreciate the design considerations that serve as a starting point for the new platform. To this end, a description of the currently used platform is given in the next section, which is followed by a section that discusses the limitations of this platform and the arguments to design a new, more robust platform.

2.1. Description of Current Platform

Figure 8 shows a photograph of the current platform. The platform, which is currently used for most biochemical experiments that are performed in the process of assay development, allows quasi-simultaneous excitation and measurement of 4 on-chip biosensors. It consists of a more or less portable and stand-alone platform that is connected via a USB interface to a PC, which is used merely to display and store measurement results, i.e. it has no post-processing functions.




Figure 8: Biosensor measurement platform.

The next subsection describes the general method of excitation of the magnetic biosensor and the detection principle as used in the current platform and in the new platform. This is followed by a subsection that discusses the performance of the current platform, both in terms of electronics SNR, which is compared to SNR resulting from the largest of the other noise sources in the system, and in terms of sensitivity.

2.1.1. Excitation and Sensing

Figure 9 depicts a simplified block scheme of the platform. It consists of the biosensor that is connected to an analog front-end that is used to excite the sensor and condition the sensor signal. The analog front-end is in turn connected to a digital back-end that is

merely used to store and display measurement results. This back-end is not important to obtain proper insight in the working of the platform and is non-critical for its performance. For these reasons it is not discussed in detail.

During measurements, an excitation current (I_e) is applied to the on-chip current conductors R_c in order to generate the magnetic field (H) that excites the surface bound beads. The resulting change in GMR resistance, induced by the magnetized beads, is determined by measuring the voltage across the GMR (U_{GMR}) upon application of a constant sense current (I_s). This signal is subsequently low-pass filtered (LPF1) to reduce the dynamic range of the signal, amplified by a factor A and demodulated by multiplication by a reference signal U_{ref} . The signal is then further low-pass filtered (LPF2) to remove out-of-band noise and unwanted signal components and finally converted to the digital domain by an analog-to-digital converter (ADC). The considerations that lead to this specific signal path, will now be discussed in more detail.

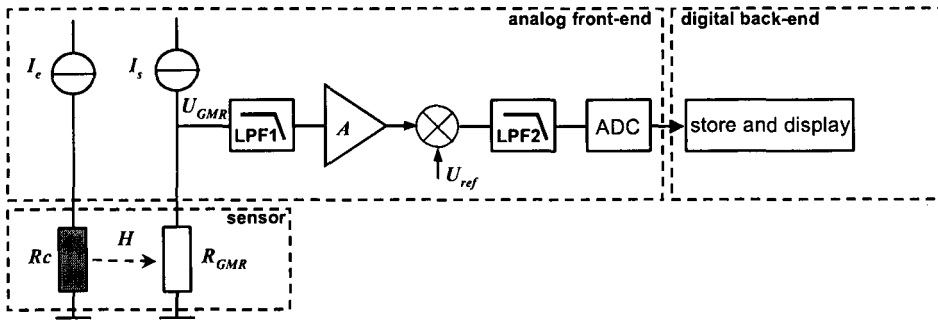


Figure 9: Block scheme of current detection platform.

The voltage across the GMR is given by

$$U_{GMR}(t) = I_s(t)(R_0 + S \cdot H(t)) \quad (5)$$

where $I_s(t)$ is the sense current applied through the GMR, R_0 is the GMR resistance in the absence of a magnetic field, S is the GMR sensitivity, which corresponds to the slope of the GMR curve of Figure 6, and $H(t)$ is the average in-plane magnetic field in the GMR. The magnetic field strength $H(t)$ can in turn be separated into two contributions: Apart from a bead dependent component (H_B), the magnetic field comprises a magnetic crosstalk component (H_{mxt}), resulting from the fact that the current conductors and the GMR are not centred at the same horizontal plane, such that the excitation field is not exactly perpendicular to the sensor plane, but has a small component in its sensitive x-direction. This misalignment results from the difference in height of the conductors (300 nm) and GMR (80 nm) and the fact that they are deposited on the same substrate. The geometrical misalignment giving rise to magnetic crosstalk is schematically depicted in Figure 10.




Figure 10: The centres of the current conductor and the GMR strip are not perfectly aligned in the xy -plane, resulting in magnetic crosstalk H_{mxt} .

Since the transfer from conductor to GMR is constant, i.e. fixed in the geometry of the sensor, we can consider the resulting crosstalk signal to be an offset to the bead dependent part of the magnetic signal, such that

$$U_{GMR}(t) = I_s(t)(R_0 + S(H_B(t) + H_{mxt}(t))) \quad (6)$$

To be able to appreciate how the frequency of the excitation field was chosen, a short discussion on the so-called limit of detection is required. The limit of detection (L_D) of (bio-)analytical equipment, i.e. the lowest amount of analyte in a sample that can be detected, but not necessarily quantitated, is often defined as

$$L_D = 2.71\sigma \quad (7)$$

where σ is the standard deviation of an ensemble of blank measurements obtained with the measurement system [12]. For our platform that detects a surface bead density, which is in turn related to the amount of analyte in the sample, this corresponds to a signal-to-noise ratio (SNR), which is defined as the ratio between the power of the signal generated by the surface bound beads and the noise power, equal to 7.3. Electronics noise contributes to this standard deviation, but also inter-sensor variations e.g. in biochemistry, bead positional noise and counting statistical noise. These contributions and their relative significance will be discussed in more detail later. To obtain a system with a low limit of detection, the electronics noise should be minimised. The resistance of the GMR is characterised by $1/f$ noise, of which the origin is not fully understood, but which is thought to be the result of random flipping of the direction of magnetization of individual magnetic domains in the GMR [13]. Figure 11 shows a typical measured noise voltage density spectrum. To minimise noise, the excitation current is modulated with a frequency such that, for typical sense current amplitudes, the thermal noise of the GMR dominates over the $1/f$ noise. In practice an excitation frequency (f_e) of 1.052 MHz is used.




Figure 11: Noise voltage spectral density at GMR output. $I_s = 1 \text{ mA}$, $R_0 = 540 \, \Omega$. A $1/\sqrt{f}$ character can be observed for lower frequencies, i.e. the noise power shows a $1/f$ character.

However, due to the small on-chip distance between the current conductor and the GMR, which is only $3 \, \mu\text{m}$, as well as the fact that the sense- and excitation current loops on the flex foil connecting the sensor to the detection platform are close together, as is schematically depicted in Figure 12, there is significant capacitive and inductive coupling between the current conductors and the GMR. This coupling gives rise to a cross talk current ($I_{C/I-XT}$) that, multiplied by the relatively large R_0 of the GMR, results in a signal component that has the same frequency as the magnetic signal of interest, yet has a 40 dB larger amplitude.




Figure 12: Origin of capacitive and inductive cross talk and coarse electric model of the combination of sensor chip and flex foil.

The cross talk and magnetic signal are separated in the frequency domain by modulation of the sense current. The currently used sense current frequency (f_s) of 1 MHz results in a magnetic signal at 52 kHz, which is low enough to facilitate easy amplification, yet above the frequency range where $1/f$ noise dominates in the amplifier (A). The GMR voltage now satisfies

$$U_{GMR}(t) = \left(\hat{I}_s \cos(2\pi f_s t) + \hat{I}_{C/I-XT} \cos(2\pi f_e t + \frac{\pi}{2}) \right) \cdot (R_0 + S(H_B(t) + \hat{H}_{mxt}) \cos(2\pi f_e t)) \quad (8)$$

where \hat{I}_s , $\hat{I}_{C/I-XT}$ and \hat{H}_{mxt} represent the peak values of the corresponding currents and fields respectively. Since the crosstalk current $\hat{I}_{C/I-XT}$ and the magnetic field induced GMR resistance changes are both small, their product may be neglected, such that

$$\begin{aligned} U_{GMR}(t) \approx & \frac{1}{2} \hat{I}_s S(H_B(t) + \hat{H}_{mxt}) \cos(2\pi(f_e + f_s)t) \\ & + \frac{1}{2} \hat{I}_s S(H_B(t) + \hat{H}_{mxt}) \cos(2\pi(f_e - f_s)t) \\ & + \hat{I}_{C/I-XT} R_0 \cos(2\pi f_e t + \frac{\pi}{2}) \\ & + \hat{I}_s R_0 \cos(2\pi f_s t) \end{aligned} \quad (9)$$

Figure 13 shows the most relevant components of the spectrum that were measured using a sensor with typical resistance and typical values for the respective current amplitudes. In this figure the magnetic signal, i.e. the signal of interest, corresponds to the second term in the right-hand side of Equation 9, while the capacitive- and inductive crosstalk signal (C/I-XT) and the sense signal correspond to the third and fourth term respectively. For clarity the first term, which appears at the highest frequency in the spectrum, is omitted. As can be seen, the sense signal is the dominant signal as a result of the

relatively large I_s and R_0 . The capacitive- and inductive crosstalk current $I_{C/I-XT}$ is about a factor 100 smaller than I_s , such that the resulting signal is 40 dB smaller. The total GMR resistance change induced by the magnetic field change, which is for small bead surface densities dominated by the magnetic crosstalk H_{mxt} , is 84 dB smaller than R_0 , while the average GMR resistance change induced by a single surface bound bead is even 164 dB smaller than R_0 .

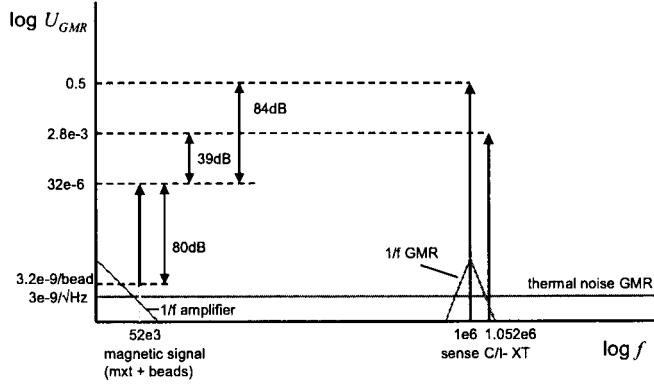


Figure 13: Most relevant components of the GMR output spectrum ($R_0 = 425\Omega$, $I_s = 1.31$ mA, $I_e = 2.24.75$ mA). The signal of interest is 3.2 nV per 300 nm bead and the magnetic crosstalk is 32 μ V.

After low-pass filtering (LPF1) to attenuate the large sense- and crosstalk signals, the signal is amplified by a factor A . The demodulated signal of interest is obtained by multiplication by a 52 kHz reference signal (U_{ref}). After final low-pass filtering (LPF2) to remove out-of-band noise and unwanted signal components, the magnetic signal is obtained, which is (for unity amplitude U_{ref}) given by

$$U(t) = \frac{A}{4} \hat{I}_s S(H_B(t) + \hat{H}_{mxt}) \quad (10)$$

2.1.2. Performance

In order to determine the sensitivity, linearity and limit of detection of the platform, the response of the detection platform to various numbers of beads on a sensor surface was experimentally determined by making use of an inhibition assay as described in Subsection 1.2.2. In the experiment the sensor surface was prepared by incubation with a concentration of 3 μ g/ml Bovine Serum Albumin (BSA)-opiate in a carbonate buffer (pH 9.6). Since BSA binds well to the gold on the sensor surface, it acts as a link between the sensor surface and the opiate. Biotinylated antibodies against opiate, i.e. antibodies with an attached biotin molecule, were subsequently applied to the sensor surface, such that they formed bounds with the surface-bound opiate. The antibodies were incubated for 1 hour, after which the unbound antibodies were washed away. In the second assay step, a sample containing magnetic beads with a streptavidin coating was applied to the sensor. These particles are composed of an ensemble of magnetite (Fe_3O_4) grains in a biocompatible polymer matrix. Commercially available beads were used, namely Ademtech particles, with a diameter of 300 nm and a susceptibility of $4 \cdot 10^{-20} \text{ m}^3$. When the beads came in contact with the sensor surface, they were immobilized as a

result of the high binding affinity between streptavidin and biotin. During the last assay step, the detection platform acquired the bead signal at regular time intervals. A CCD camera in combination with a microscope was used for optical monitoring of the beads on the active sensor surface. During the assay, images were recorded and a time stamp was provided for each image. The instantaneous number of beads on the sensor surface was determined from each image, and then linked to the bead signal at that time-instant. In Figure 14, the signal change is plotted as a function of the number of beads on the sensor [14]. A straight line was fitted through the data by least-squares minimization, showing that the signal increases linearly with the number of beads on the sensor surface, namely 5 nV per bead on average. The total input-referred noise voltage of the detection platform, including the sensor and the electronics, was measured to be 5.5 nV_{RMS} with a 1 Hz detection bandwidth. This is close to the theoretical limit, which is determined by the thermal noise voltage density $e_{n,GMR}$ of the GMR sensor that is given by

$$e_{n,GMR} = \sqrt{4kTR_0} \quad (11)$$

where k is Boltzmann's constant, T is the temperature in Kelvin and R_0 is the resistance of the GMR. The GMR thermal noise amounts to 2.9 nV_{RMS} in a 1 Hz bandwidth. Since the limit of detection is defined to be 2.71 times the noise standard deviation (Equation 7), it follows that a single bead is detectable with a platform having a bandwidth of 0.1 Hz on a sensor surface of 1500 μm^2 .

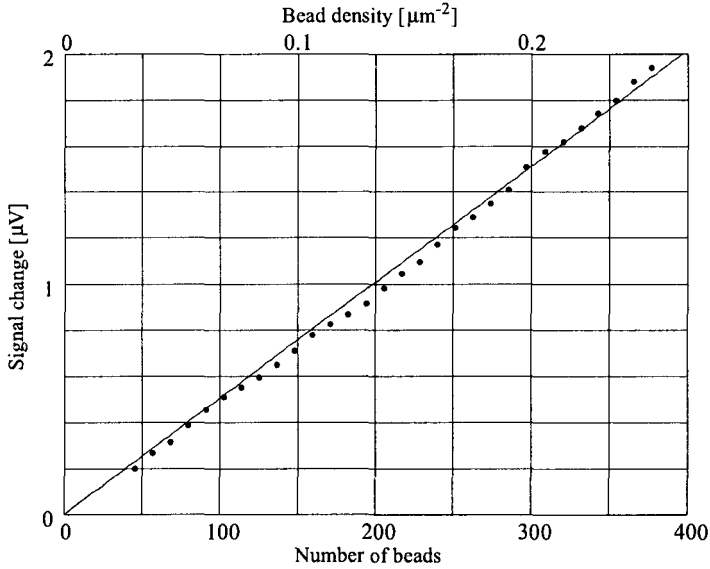


Figure 14: Bead-dose response curve. The change in output signal of the detection platform (referred to the GMR voltage) as a function of the number of 300 nm beads bound to the sensor surface.

As stated before, the detection limit is determined by the standard deviation across an ensemble of measurements. Apart from electronic noise, several other stochastic processes contribute to this standard deviation. Their combined effect, which we will refer to as statistical noise, results from:

1. Non-uniform positional sensor sensitivity: The signal as a result of the binding of a bead to the sensor surface depends heavily on the x-position of the bead on the sensor surface, as can be seen from Figure 15 [15]. Since the bead x-position is a stochastic process, the signal per bead is also stochastic.
2. Non-uniform bead size: The diameter of different beads, and therefore their magnetic moment upon magnetization, may differ as much as a factor two.
3. Non-uniform bead susceptibility: The susceptibility of different beads may vary due to non-uniformity in the magnetic content.
4. Shot noise due to the bead arrival process on the sensor surface.




Figure 15: Signal per 200 nm bead as a result of the position of the bead on the sensor surface. The location of the sensor elements on the x-axis is as follows: conductor 1 0-3 μm , GMR 6-9 μm , conductor 2 12-15 μm [15].

In order to determine if the electronic SNR suffices, it should be compared to other noise sources of the system. Figure 16 depicts an SNR plane. On the x-axis the signal to electronic noise ratio, or electronic SNR, is shown, while on the y-axis the signal to statistical noise ratio, or statistical SNR is shown. The black dot represents the combination of electronic- and statistical SNR of the current platform when measuring a single surface bound bead. Since the GMR voltage increases linearly with surface bead density (B), the electronic SNR scales quadratically with B . It can furthermore be shown that the statistical SNR is linearly related to the surface bead density [16]. The arrow represents the path that is travelled through the SNR plane as the surface bead density increases. The system signal to noise ratio SNR_{system} , is given by

$$SNR_{system} = \frac{P_s}{\sigma_{total}^2} = \frac{P_s}{\sigma_{electronic}^2 + \sigma_{statistical}^2} \quad (12)$$

For clarity, two equi- SNR_{system} lines are shown in Figure 16. It can be concluded from Figure 16 that for any value of B the SNR_{system} is determined by the statistical SNR. The electronic SNR of the current system is thus sufficiently low.




Figure 16: SNR plane. The SNR of the current system is indicated by the dot [16].

Note that a surplus of electronic SNR can be traded for statistical SNR by changing the area of the sensor. For this reason a high signal to electronic noise ratio is always advantageous. This can be achieved by optimising sense- and excitation current amplitude and measurement bandwidth. Alternatively, a surplus in electronic SNR may be traded for more efficient binding by using beads with smaller diameter, which reduces the signal per bead, but also reduces steric hindrance at the sensor surface. Steric hindrance is defined as the physical blockage of a particular binding site within a molecule by the presence of local atoms or groups of atoms. As a result of steric hindrance, which is in this case caused by the relatively large size of the magnetic bead compared to the antibody, a reaction of the antibody with the analyte may be prevented.

2.2. Limitations of Current Platform

The limit of detection of the current (electrical) platform is determined by the thermal noise, which is 5.5 nV for a 1 Hz bandwidth system. To obtain a sufficiently robust system, fluctuations in the output signal resulting from environmental changes such as temperature variations should be smaller than this. The current platform lacks this required robustness, since the following effects are not (or only partly) accounted for:

1. The sensitivity of the GMR is a function of temperature. This temperature dependency is discussed in more detail in Chapter 4 .
2. The sensitivity of the GMR is a function of externally applied magnetic fields. It can be seen from Figure 6 that large non-alternating fields result in a shift of the

working point (around which the field is modulated by the excitation current) towards the non-linear parts of the GMR curve, i.e. close to saturation.

3. The transfer of the analog part of the measurement system is a function of temperature, e.g. due to temperature related shifts of filter transfers.
4. Excitation- and sense- currents sources that are non-ideal and are characterized by a temperature dependency of the output current amplitude.
5. Inherent capacitive and inductive crosstalk introduces errors in the measured signal. This will be explained in more detail in Section 4.3.

All these effects combined result in output signal variations that are 2 orders of magnitude larger than the thermal noise. The large influence of the above-mentioned effects on the output voltage can be explained using Equation 10, which is for clarity repeated here.

$$U(t) = \frac{A}{4} \hat{I}_s S(H_B(t) + \hat{H}_{mxt}) \quad (13)$$

The magnetic crosstalk field H_{mxt} will, for practical surface bead densities, be much larger than the bead signal H_B . As a result of this, small variations in sensitivity S , sense current I_s , excitation current I_e (both H_{mxt} and H_B are proportional to I_e) and amplification A , will result in much larger variations in output voltage compared to what can be expected purely from variations in bead density. The aim of this research is therefore to realise a more robust measurement platform. This can be achieved by reducing the temperature dependency of the current sources and filters, mainly by reducing the number of passive components to a minimum. Furthermore, robust detection algorithms can be employed to compensate for variations in sensitivity, e.g. by normalizing the magnetic signal using a reference signal that depends similarly on the sensitivity, such that the resulting signal is not dependent on the sensitivity. In the following chapters this robust platform is discussed, both in terms of hardware and detection algorithms. It should finally be noted that robustness can also be significantly improved by a reduction of the magnetic crosstalk H_{mxt} , which can be achieved by improving the design of the sensor chip, such that the conductors and the GMR are centred in the same xy-plane. This is, however, beyond the scope of this report.

3. New Platform

The new platform is designed to be used in combination with a newly designed sensor. Since this has consequences for the platform design, e.g. in terms of required maximum sense- and excitation current amplitudes and resulting GMR output voltage, we start this chapter by describing the new sensor, as well as explaining the motivation for the design of a new sensor. The following sections discuss the functional description of the new platform and its analog front-end architecture.

The basic shape of the new sensor is similar to the old sensor, however, the new sensor is composed of 5 parallel rows of sensor sections each in turn composed of 3 sensors in series, as schematically depicted in Figure 17. Moreover, the width of the GMR strip is slightly increased. This results in an increase of the effective sensor area by a factor 15, which reduces statistical noise power by a factor 15, since more beads are measured at the same surface bead density as compared to the old sensor. On the other hand, the maximally achievable signal to electronic noise ratio is reduced by a factor 15. By increasing the area of the sensor, some electrical SNR has thus been traded for statistical SNR in order to obtain a better system SNR. The resistance of the GMR of the new sensors is about $150\ \Omega$, while the resistance of the current conductors is about $12\ \Omega$.




Figure 17: Schematic representation of top view (top) and cross-sectional view (bottom) of old and new sensor.

The new sensor contains so-called compensation conductors that can be used to cancel magnetic crosstalk in the GMR. This function will be discussed in more detail as part of the recommendations in Chapter 8. Moreover, the compensation conductors may be used to attract beads by magnetic actuation towards the sensor surface in order to speed up the immunoassay, as explained in Section 1.4.

3.1. Functional Description

To facilitate fast implementation and validation of new and more robust detection algorithms, a dedicated analog front-end is combined with a flexible digital back-end. This results in a mixed-signal system as depicted in Figure 18. Its digital back-end consists of functionality to store and display measurement results, which is relatively straightforward and therefore not described in further detail, and a digital signal processing (DSP) unit that can produce a sense waveform and an excitation waveform. The analog front-end consists of an excitation current source, a sense current source and a pre-amplifier. The excitation current source uses a digital-to-analog converter (DAC) to convert the excitation waveform generated by the DSP unit to the analog domain. This waveform is subsequently used to drive a current source that supplies an excitation current (I_e) to one of the excitation conductors (R_e) on the sensor chip. The sense current source performs a similar function: it converts the sense waveform to the analog domain and supplies the GMR (R_{GMR}) with a sense current (I_s). A high-pass filter (HPF) is required to reduce the noise generated by the sense current source, in particular at the frequency of the magnetic signal, which is located spectrally at the difference of the sense- and excitation frequency. This will be explained in more detail in Subsection 3.2.2. The pre-amplifier conditions the sensor signal (U_{GMR}) before it can be processed by the DSP unit. The GMR voltage is first low-pass filtered (LPF) to reduce the dynamic range of the signal, i.e. to attenuate the large sense signal and capacitive- and inductive crosstalk. Figure 13 clearly indicates that this dynamic range should be heavily reduced before amplification can take place. The amplification is followed by a conversion to the digital domain using an analog-to-digital converter (ADC). Further processing, e.g. demodulation, is subsequently performed in the DSP unit.

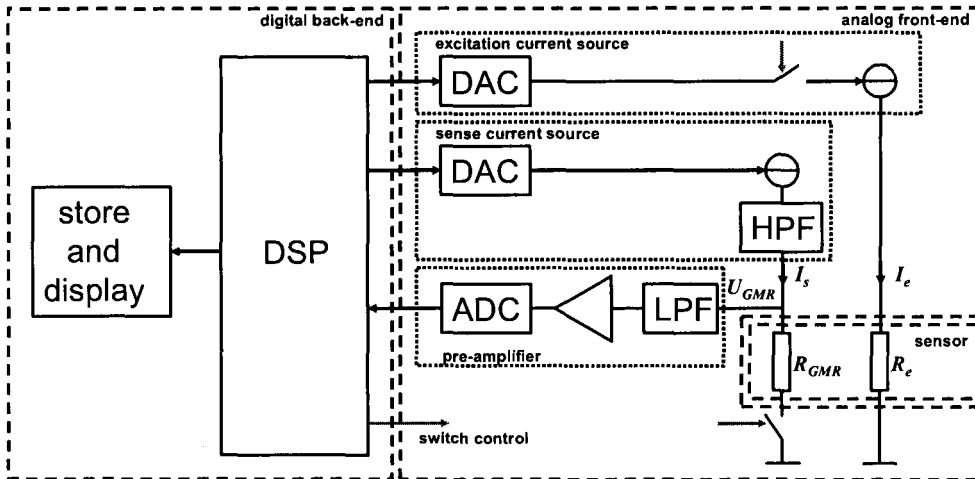


Figure 18: Block scheme of the new detection platform.

Switch control signals are furthermore generated by the DSP unit to activate individual current conductors of a sensor, as well as switch between any of the four sensors of the biosensor chip. Note that for clarity only one of the sensors is shown in Figure 18. By using time multiplexing, i.e. measuring all sensors sequentially by switching between

them, the analog front-end is identical for each sensor. Inter-sensor differences in signal at the output of the platform can therefore only be attributed to differences in sensor signal, i.e. voltage across the GMR, which allows for fair comparison of signals from different sensors. Comparing signals from different sensors is only allowed if the bandwidth of gain variations and phase shifts in the detection electronics, should they occur, is much smaller than the switching frequency. This is in general the case. A drawback of this sequential measurement strategy is that the SNR of the detected signal is reduced by a factor 4, since we measure each sensor effectively one fourth of the total measurement time. However, since it was already shown in Subsection 2.1.2 that the obtained SNR is sufficient to measure even minute surface bead densities, e.g. in the order of several beads, this disadvantage is not of large significance.

Figure 19 schematically depicts the contents of the DSP unit. The demodulation and filtering of the magnetic signal is now performed in the digital domain using an adaptive joint phase and amplitude estimator (ϕ, A estimator). This ϕ, A -estimator will be discussed in detail in Section 5.2 and Subsections 6.1.2, 6.2.2 and 6.3.2. Before being applied to the ϕ, A -estimator, the output signal of the analog pre-amplifier is further amplified and filtered in a digital pre-amplifier (pre-amp). This digital pre-amplifier is discussed in Subsections 6.1.3 and 6.3.3. A waveform generator is employed to generate the excitation- and sense waveform, as well as a reference waveform that is supplied to the ϕ, A estimator. This waveform generator will be discussed in detail in Subsections 6.1.1, 6.2.1 and 6.3.1. The algorithm block controls the waveform frequencies and the sensor selection. The sensor selection block in turn generates the signals to control the switches, based on the sensor select signal produced by the algorithm block. The algorithm block furthermore processes the estimated phase ϕ and amplitude A in order to estimate the surface bead density B according to the methods described in Section 5.1. Finally, the DSP unit interfaces via a USB interface to the store and display functionality of the digital back-end.

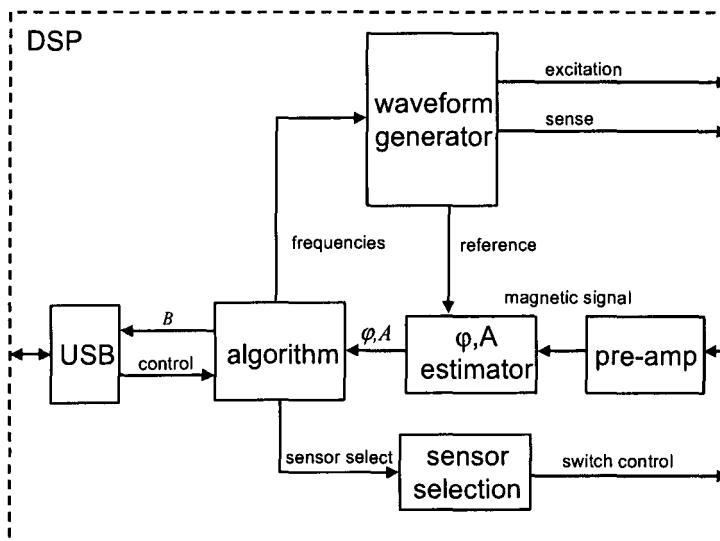


Figure 19: Block schematic view of DSP unit contents.

The DSP unit can be implemented as a digital signal processor, a microcontroller or a field-programmable gate array (FPGA). An FPGA was chosen for several reasons. Firstly an FPGA provides the parallelism required to simultaneously output three high-frequency and synchronous waveforms to the DACs. This would be more difficult with a DSP and virtually impossible with a microcontroller. A DSP could be considered an overkill for the modest computational load generated by the algorithms that run at relatively low sample rates, i.e. no more than a few outputs/second. Since only few or no if-statements are used, which result in parallel hardware branches, an FPGA can perform the limited amount of calculations equally well. The added cost associated with the employment of an FPGA is less of an argument, since it concerns a research project and only a limited amount of measurement platforms will be made. The power consumption of the FPGA is low enough for the intended handheld application, i.e. in the order of 100 mW. And finally, an interface with Matlab Simulink is available from the FPGA manufacturer, providing a rich simulation environment to test the algorithms.

3.2. Analog Front-end Architecture

Since the analog front-end is the most critical part of the system in terms of robustness and SNR, the following subsections discuss all its relevant parts. These subsections discuss the design, both in terms of design considerations made to reduce the temperature dependency of the measurement results, i.e. to improve the robustness, and to obtain a high signal to electronic noise, i.e. keeping the noise close to the limit imposed by the thermal noise of the GMR itself.

3.2.1. Excitation Current Source

Figure 20 depicts the excitation current source. The excitation conductors are driven from a current source that is able to supply a maximum excitation current (I_e) of 160 mA_{pp}. It uses a 10-b parallel current DAC clocked at 40 MHz that converts the excitation waveform to the analog domain. The amplitude of the waveform is controlled by a reference signal (U_{ref}) generated by a serial 8-b DAC, which is in turn controlled via a serial peripheral interface (SPI) from the FPGA. The output of the DAC is connected to four separate transconductance amplifiers that (sequentially) drive the excitation conductors (R_e) of the four sensors on chip. A sensor is selected by enabling the corresponding operational amplifier (op-amp) by providing the appropriated switch control signals. To achieve maximum robustness, the same highly temperature stable feedback resistor (R_f) is shared by all transconductance amplifiers, such that all excitation currents remain equal and highly stable during the measurements.

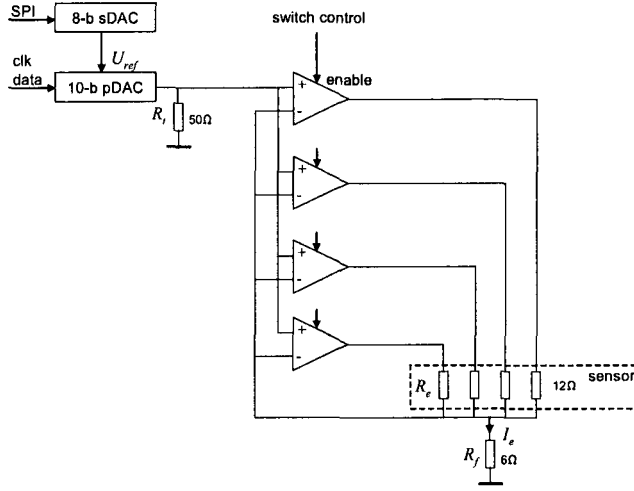


Figure 20: Excitation current source.

The allowable current noise can be determined by comparing the induced voltage noise across the GMR to the thermal noise of the GMR. From the current sensors the transfer from excitation current to GMR voltage (U_{GMR}/I_{ex}) was measured to be equal to $1.6 \cdot 10^{-3}$ V/A using a 425Ω GMR and a sense current of 1.3 mA. The new sensors will be operated at an increased sense current of 9 mA, such that the transfer becomes $11.1 \cdot 10^{-3}$ V/A. Using Equation 11, the thermal noise of the new 150Ω sensors can be calculated to equal $1.6 \text{ nV}/\sqrt{\text{Hz}}$. The excitation current noise density should therefore be lower than $144 \text{ nA}/\sqrt{\text{Hz}}$. This is easily achieved with the current source depicted in Figure 20. The current noise produced by the op-amp can be calculated by multiplying the input voltage noise of each input by the transconductance, which equals 0.167 A/V . The input voltage noise consists of op-amp input voltage noise ($1 \text{ nV}/\sqrt{\text{Hz}}$), op-amp input current noise ($2.6 \text{ pA}/\sqrt{\text{Hz}}$ in 50Ω) and DAC output noise ($50 \text{ pA}/\sqrt{\text{Hz}}$ in 50Ω). The total input-referred voltage noise generated by the op-amp ($e_{i,op-amp}$) is given by

$$e_{i,op-amp} = \sqrt{e_-^2 + e_+^2 + (i_- R_f)^2 + (i_+ R_t)^2} \quad (14)$$

in which e_- and e_+ are the input-referred voltage noise of the op-amp at the negative and positive input respectively, i_- and i_+ are the input-referred noise currents at the negative and positive input respectively, R_t is the DAC output termination resistor and R_f is the feedback resistor. The op-amp generated input-referred voltage noise amounts to $1.4 \text{ nV}/\sqrt{\text{Hz}}$, such that the total input-referred voltage noise becomes $2.9 \text{ nV}/\sqrt{\text{Hz}}$. The total current noise then amounts to $480 \text{ pA}/\sqrt{\text{Hz}}$, which is sufficiently low.

3.2.2. Sense Current Source

Figure 21 schematically depicts the sense current source. The maximum required sense current (I_s) for any of the envisioned sensor types is 26 mA_{pp} , which results in 3.9 V_{pp} across a 150Ω GMR. This means that the employment of merely a current DAC will not suffice, as these in general have a output range of $\pm 1 \text{ V}$. An op-amp current amplifier is

therefore employed to cope with the GMR voltage range. An integrated switch is used to select one of the four GMRs (R_{GMR}) on the biosensor chip. A 10-b parallel DAC clocked at 40 MHz converts the sense waveform to the analog domain. The amplitude of the waveform is controlled by a reference signal (U_{ref}) generated by a serial 8-b DAC that is again controlled via an SPI interface. A filter is employed between the output of the DAC and the op-amp to reduce the noise contributions of the DAC and op-amp in the output U_{GMR} at the frequency of the signal of interest. This is important since we strive to approach the lower limit of the noise, which is the thermal noise of the GMR. In other words, the GMR should preferably be the dominating noise source.




Figure 21: Sense current source.

The noise current spectral density at the output of the DAC is $50 \text{ pA}/\sqrt{\text{Hz}}$, which would result in a noise voltage spectral density of $50 \text{ pA}/\sqrt{\text{Hz}} \cdot 150 \Omega = 7.5 \text{ nV}/\sqrt{\text{Hz}}$ at the output of the op-amp. First-order filtering is thus required to reduce the DAC noise contribution to a level below the thermal noise of the GMR, i.e. $1.6 \text{ nV}/\sqrt{\text{Hz}}$ at room temperature. In practice it was found that, for electro-magnetic compatibility (EMC) related reasons that are beyond the scope of this report, a third order filter is required. A Butterworth filter was used since this type of filter is characterised by a small pass-band ripple, which is important since we strive for minimal temperature induced filter transfer variations at the excitation frequency. Figure 22 shows the frequency response of applied filter, which is depicted in Figure 21.




Figure 22: Frequency response of sense current source filter.

The contribution of the op-amp noise ($e_{o,opamp}$) to the total noise at the output of the current source is given by

$$e_{o,opamp} = \sqrt{\left(e_- \left(\frac{R_{GMR}}{Z_f} \right) \right)^2 + \left(e_+ \left(1 + \frac{R_{GMR}}{Z_f} \right) \right)^2 + (i_- R_{GMR})^2} \quad (15)$$

where e_- and e_+ are the input-referred noise voltages of the op-amp at the negative and positive input respectively, R_{GMR} is the GMR resistance, Z_f is the filter impedance as seen from the op-amp and i_- is the input referred current noise of the op-amp. For the op-amp used, the contributions of the voltage noise sources dominate. These contributions are minimized if Z_f is much larger than R_{GMR} . For the applied filter this is the case, since its impedance Z_f is about 475 Ω at 50 kHz. The total noise generated by the op-amp amounts to 1.3 nV/ $\sqrt{\text{Hz}}$. Since the noise contribution of the DAC after filtering is negligibly small, the output noise of the sense current source is determined by the combined op-amp noise and GMR thermal noise and amounts to 2.1 nV/ $\sqrt{\text{Hz}}$.

3.2.3. Pre-amplifier

Figure 23 depicts the expected spectrum of the voltage across the GMR for the new sensors. The differences with respect to the measured spectrum of the formerly used sensors, as depicted in Figure 13, result from a smaller GMR resistance (150 Ω), a larger sense current (9 mA) and a reduced excitation frequency (1.026 kHz).

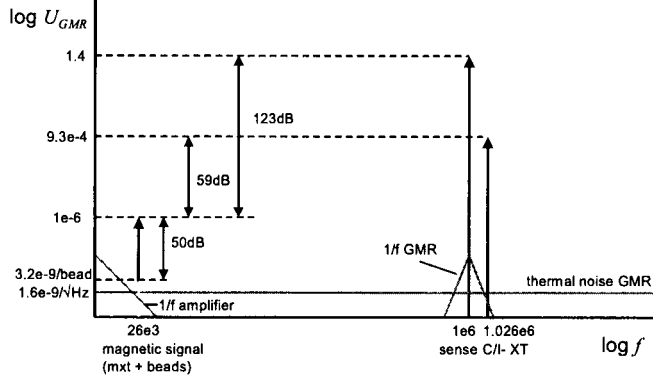


Figure 23: Expected spectrum of GMR voltage adapted to new sensor.

From the spectrum of Figure 23 it can be calculated that the signal to noise density ($S/N(f)$), defined as the ratio between the signal power and the noise power spectral density, is 179 dB. Note that the word signal refers to the largest signal component in the spectrum, which is not necessarily the signal of interest. In this case the sense signal determines the $S/N(f)$. This signal is however not of any interest and may be reduced by filtering in the pre-amplifier, thereby reducing the $S/N(f)$.

The pre-amplifier should be designed such that the $S/N(f)$ of the output signal of the analog-to-digital converter (ADC) is determined by the $S/N(f)$ of the signal applied to its input. Since the full scale input voltage range of an ADC U_{RMS} equals

$$U_{RMS} = \frac{q(2^N - 1)}{2\sqrt{2}} \quad (16)$$

where N is the effective number of bits of the ADC and q is the quantisation step size, the full scale input signal power P_s of the ADC for sinusoidal signals equals

$$P_s = \frac{q^2(2^N - 1)^2}{8} \quad (17)$$

The quantisation noise power density $P_n(f)$ of the ADC equals

$$P_n(f) = \frac{q^2}{12f_{sig}} = \frac{q^2}{6f_s} \quad (18)$$

where f_{sig} is the input signal bandwidth, which is the sample frequency f_s divided by two. The $S/N(f)$ of the ADC thus equals

$$S/N(f) = \frac{P_s}{P_n(f)} = (2^N - 1)^2 \frac{6}{8} f_s \quad (19)$$

Figure 24 depicts for the current system the combinations of number of bits and sample frequency of the ADC that are required to obtain just enough $S/N(f)$, given a certain number of filter orders with $f_c = 50$ kHz used in the analog pre-amplifier. A suitable ADC should be chosen based on application specific requirements in terms of cost, power

consumption and performance. Since an experimentation platform is developed, rather than a product, the focus is more on performance than on cost or power consumption. From a robustness perspective, it is advisable to keep the number of analog filters to a minimum, as analog filters are characterised by a temperature dependent transfer. For this reason, a 14-b ADC (with an effective number of bits of 11) is used that is clocked at a frequency of 40 MHz. From Figure 24 it can be seen that this configuration, which is indicated by the black dot, requires second order filtering in the analog domain. It should be noted that this configuration, where the signal of interest is much smaller than one quantisation step, is only allowed when a sufficiently large dither signal is present. In this case, the relatively large sense signal fulfils this function.




Figure 24: Number of bits and sample frequency of the ADC in order to obtain the required $S/N(f)$ for different orders (N) of filtering in the analog domain. The curves shown are obtained using Equation 19, where $S/N(f) = 179 \text{ dB} - 26 \text{ dB} \cdot N$.

To determine the minimum gain required for the pre-amplifier, the noise voltage spectral density of the ADC is compared to that of the signal. The noise voltage spectral density at the output of the ADC equals $103 \text{ nV}/\sqrt{\text{Hz}}$. Ideally, the noise voltage spectral density of the GMR signal is determined by the thermal noise of the GMR, which for a $150 \, \Omega$ GMR amounts to $1.6 \text{ nV}/\sqrt{\text{Hz}}$. A minimum gain of 64.4 is thus required.

The maximum gain is determined by the GMR signal amplitude, the attenuation as a result of filtering and the full-scale input range of the ADC. The maximum GMR signal amplitude is $3.9 \text{ V}_{\text{pp}}$, which is fully determined by the sense signal at 1 MHz. The pre-amplifier gain should thus be such that, given the second order filtering, the amplitude of the sense signal after amplification does not exceed the ADC input range. The filter attenuation at the sense frequency, i.e. 1 MHz, is ideally a factor 20 per added filter order, assuming a corner frequency of twice the maximum magnetic signal frequency, i.e. 50 kHz. As can be seen from Figure 25, the best practical filters with acceptable passband ripple (0.1 dB), however, attenuate only 47.2 dB, i.e. a factor 229. Given the 2

V_{pp} input range of the ADC, this results in a maximum gain of $(2V_{pp}/3.9V_{pp}) \cdot 299 = 115$, which is still well above the minimally required gain.




Figure 25: Simulated filter characteristic for 2nd order filter as shown in Figure 26.

The resulting pre-amplifier, which complies with the aforementioned requirements, is depicted in Figure 26 and has a gain of 106, i.e. 40.5 dB. The second stage is a general single-ended to differential input stage for the ADC. A 1 μ F capacitor finally introduces a high-pass cut-off frequency of 3.6 kHz to reduce the offset voltage of the sense current source before amplification.




Figure 26: Pre-amplifier.

For applications where form factor and power consumption are the main issues, e.g. a hand held saliva tester, it may be beneficial to trade some SNR for a reduction of the number of op-amps used. For experimentation purposes the front-end printed circuit board (PCB) therefore allows bypassing of the pre-amplifier. For applications where power consumption and cost are key, it may be beneficial to use a less expensive ADC at the cost of some extra gain and filtering in the analog domain. The front-end PCB therefore also allows insertion of an additional gain stage and filtering up to 10 orders. The gain of the additional stage is chosen equal to the attenuation of the sense signal in case of second order filtering, i.e. 47 dB. The complete pre-amplifier design is depicted in Figure 27. The component values shown are the values after optimizing for 5th order filtering per stage. It should however be noted that in the remainder of this research the single-stage pre-amplifier of Figure 26 will be used.

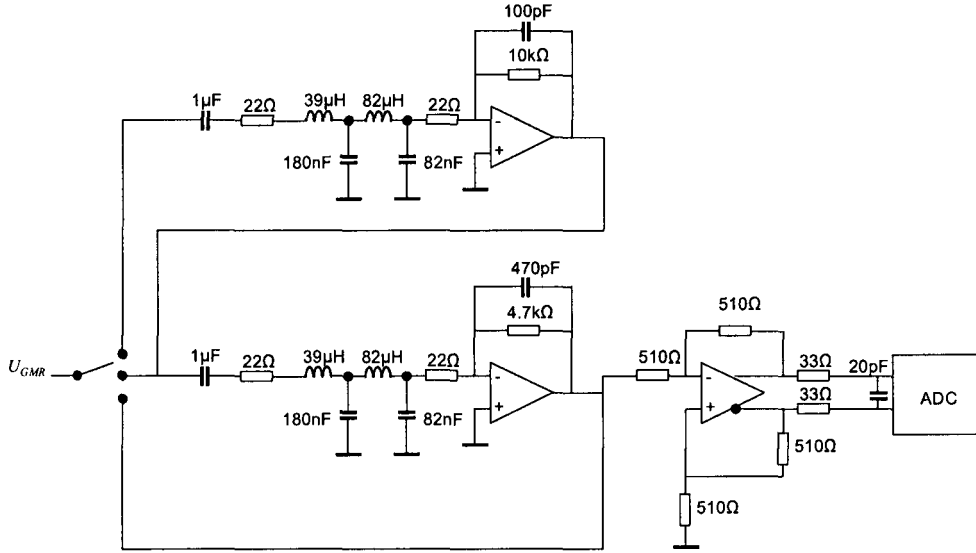


Figure 27: Flexible pre-amplifier.

Finally, it should be noted that in practice, for a system generating signals with such large differences in amplitude, proper PCB lay-out design is of vital importance. Whether the theoretically achievable performance is actually met is thus to a large extent dependent on the skills of the lay-out designer.

3.3. Analog Front-end Measurement Results

In order to analyse the performance of the analog front-end, the transfer of the complete (two stage) pre-amplifier of Figure 27 was measured to determine the exact gain and filter characteristics. Subsequently, an excitation- and sense current were applied to the sensor to measure the signal and noise performance of the front-end. A sensor with a GMR resistance R_0 of 500 Ω was used in this case. Using the pre-amplifier transfer, the input referred noise, i.e. the noise at the output of the GMR, could finally be determined and compared to the thermal noise of the GMR. Apart from this, the transfer of the single-stage pre-amplifier of Figure 26 was determined, since this configuration will be used in the remainder of this research. The input referred noise of this configuration is equal to that determined with the two-stage pre-amplifier. Finally, the robustness of the single-stage amplifier, in terms of temperature dependency, was measured.

Figure 28 depicts the transfer of the complete pre-amplifier as depicted in Figure 27. The gain at the frequencies of interest, i.e. 10 to 50 kHz, should theoretically be about 80 dB, which corresponds well with the measured gain of 79.7 dB.




Figure 28: Transfer of the pre-amplifier. Input and output attenuation was used during this measurement. The actual gain at 25 kHz equals 79.7 dB.

Figure 29 and Figure 30 show the spectrum at the output of the pre-amplifier when a sense current of 1 mA at a frequency of 987 kHz and an excitation current of 166 mA through one conductor at a frequency of 977 kHz are applied. In order not to overload the pre-amplifier, the spectrum was measured through 510 Ω , i.e. the actual values are 10 times higher than shown in the graphs. The noise voltage spectral density in the vicinity of the magnetic signal, located at 10 kHz, is -62 dBV in a bandwidth of 300 Hz. Since the gain of the pre-amplifier stages is 80 dB, this corresponds to -142 dBV (or 79.4 nV) input referred noise. The input referred noise voltage spectral density thus amounts to 4.6 nV/ $\sqrt{\text{Hz}}$, which is close to the theoretical limit imposed by the thermal noise of the GMR (2.9 nV/ $\sqrt{\text{Hz}}$). Note that the magnetic signal at 10 kHz consists solely of magnetic crosstalk, since no beads were applied to the sensor at this stage.

It can thus be concluded that the gain of the pre-amplifier is practically equal to the specified value and that the noise generated by the electronics is such that near-optimal signal to electronics noise is achieved with the new platform. It can furthermore be concluded from the spectra of Figure 29 and Figure 30 that indeed some filter orders may be omitted and substituted for filtering in the digital domain, since the sense signal is now 35 dB lower than the magnetic signal, i.e. the dynamic range reduction of the GMR signal is more than is required.




Figure 29: Spectrum at the output of the pre-amplifier 0-2 MHz.

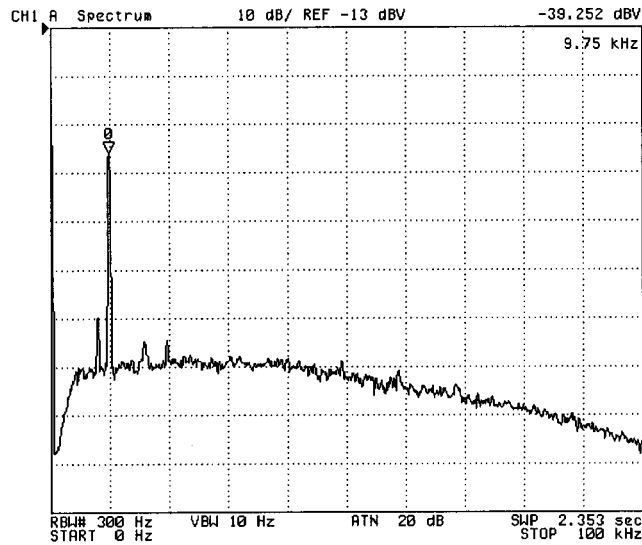


Figure 30: Spectrum at the output of the pre-amplifier 0-100 kHz.

The transfer of the single-stage pre-amplifier was subsequently measured. This transfer is shown in Figure 31. Input attenuation (34.3 dB) and output attenuation (21.0 dB) were used during the measurement, such that the actual gain can be calculated to be 95. The difference between the measured gain and the specified gain of 106 can be attributed fully to the residual impedances of the 1 μ F capacitor and the 68 μ H coil at measurement frequency (20 kHz).




Figure 31: Gain (top) and phase (bottom) of single-stage pre-amplifier.

Finally the temperature dependence of the transfer of the single-stage pre-amplifier was measured. To this end, the analog front-end was heated, while both the temperature of the front-end and the transfer of the pre-amplifier were measured. Figure 32 shows the measured transfer of the pre-amplifier at 10 kHz and the temperature as a function of time during the heating experiment. From this figure it can be concluded that the temperature coefficient of the pre-amplifier gain is about $0.025/35 = 714$ ppm/°C and the temperature coefficient of the phase is 0.13 °/°C.




Figure 32: Gain (top), phase (bottom) and temperature of the single-stage pre-amplifier at 10 kHz as a function of time.

With respect to the single-stage pre-amplifier it can be concluded that the gain is slightly lower than specified, but within the allowed range. Furthermore, although the

temperature dependency has been reduced significantly compared to the previous platform, it is still considerable, even though the number of analog filter orders has been reduced to a minimum. Further improvement of the robustness of the system is therefore required and will be provided by the robust detection algorithms.

4. Combined Sensor and Front-end Model

The following sections describe a model of the combined sensor and analog front-end, as well as a model of the magnetic signal across the GMR. The sole purpose of the combined sensor and analog front-end model is to generate input for the algorithms during the simulations that are discussed in Section 5.3. The model thus facilitates easy performance testing of the algorithms in a rich simulation environment like Simulink. Section 4.1 discusses the model itself, while Section 4.2 discusses the associated model parameters. In order to be able to understand the principles of the algorithms, which are discussed in Section 5.1, it is necessary to discuss the magnetic signal at the GMR. This discussion will be provided in Section 4.3.

4.1. Model Description

Figure 33 depicts the model consisting of a sensor part and an analog front-end part. In the sensor part, the voltage across the GMR (U_{GMR}) is obtained by a multiplication of the current through the GMR (I_{GMR}) and the resistance of the GMR (R_{GMR}). The GMR current results from the addition of the applied sense current (I_s) and a small cross talk current ($I_{C/I-XT}$) that originates from the capacitive and inductive coupling between the GMR and current conductors. This coupling is modelled by multiplying the excitation current (I_e) by a transfer (α) and taking the derivative, where for convenience we assume a dominant capacitive coupling, even though the inductive coupling may in fact be larger. This can be compensated for by taking a negative value for α , although for the algorithms to be tested the absolute phase of the crosstalk signal is of no importance.

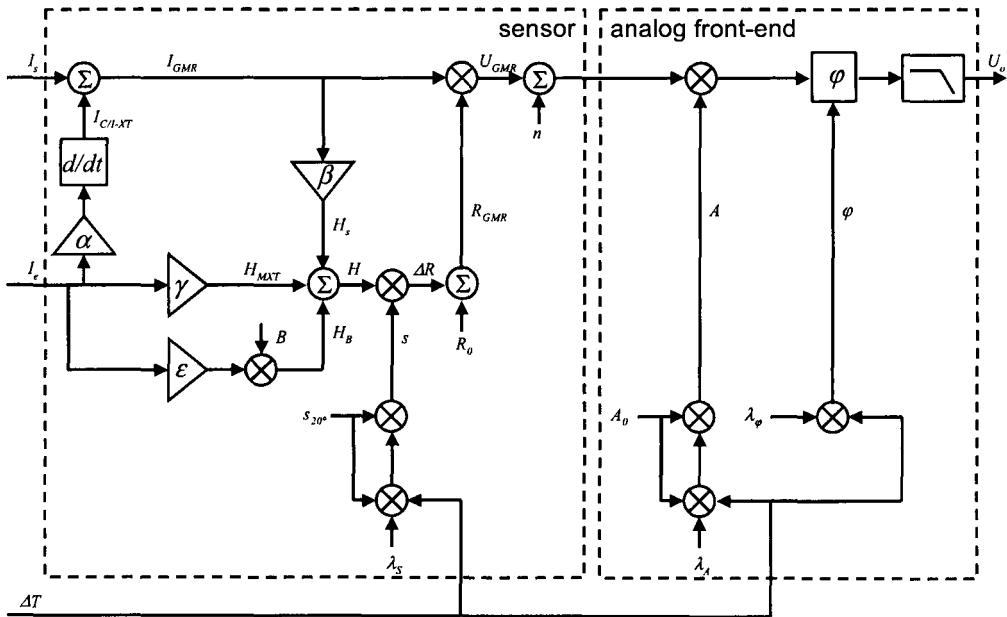


Figure 33: Combined sensor and font-end model

The GMR resistance results from the addition of a constant part (R_0) and a field dependent part (ΔR). The field dependent part is obtained by a multiplication of the magnetic field (H) and the GMR sensitivity (S). The GMR sensitivity is dependent on temperature. This dependency is modelled by means of a sensitivity temperature coefficient (λ_s), such that the sensitivity is given by

$$S = S_{20^\circ} (1 + \lambda_s \cdot \Delta T) \quad (20)$$

where S_{20° is the GMR sensitivity at room temperature and ΔT is the temperature change with respect to room temperature. The magnetic field (H) in turn consists of three contributions. Firstly, magnetic crosstalk results in a magnetic crosstalk field (H_{mxt}), i.e. the horizontal component of the magnetic field in the GMR originating directly from the current conductors. This is modelled by introducing transfer γ . Secondly, a surface bead density dependent field component is modelled by transfer ε , representing the average field contribution per bead, multiplied by the number of surface bound beads B . Finally, a so-called self-magnetisation field is present, which needs some further explanation.

Figure 34 depicts the typical stack of metal layers in a GMR element. Since a sense current that is applied to the GMR will flow primarily through the low-ohmic copper layer of the GMR, an additional magnetic field component will be present in the sensitive free layer, i.e. the NiFe layer of the GMR that causes a corresponding resistance variation R_s with the frequency of the sense current. Multiplied by the capacitive crosstalk current, this results in a component in the GMR output voltage, which is not dependent on surface bead density, having the same frequency as the magnetic signal of interest. The transfer β in Figure 33 accounts for this self-magnetisation effect. Moreover, the capacitive crosstalk current flowing through the GMR produces a magnetic field in the GMR at the excitation frequency. This current, multiplied by the sense current produces another bead surface density independent component in the GMR voltage with an amplitude exactly equal to aforementioned component. For convenience, both fields are combined in the symbol H_s .




Figure 34: GMR stack (not to scale) and magnetic field H_s in the sensitive NiFe layer as a result of a current through the GMR.

The noise voltage spectral density at the output of the GMR consist of a 1/f resistance noise component multiplied by the sense current, as well as a thermal noise component. Since the signal of interest appears, due to the modulation of the excitation current, in the spectrum at a frequency where thermal noise is dominant, 1/f noise is not included in the model. The noise is thus modelled solely as additive white Gaussian noise n .

The analog front-end is finally modelled by a gain, a phase shift and a low pass filter. As discussed in Section 3.3, the temperature dependency of the front-end was measured during a heating experiment. This experiment resulted in a temperature coefficient for the amplifier gain λ_A , and a temperature coefficient for the phase shift λ_φ . These coefficient have been used in the front-end model to obtain temperature dependent gain and phase shift. The temperature dependent gain (A) is given by

$$A = A_0 (1 + \lambda_A \Delta T) \quad (21)$$

where A_0 is the gain of the analog front-end at room temperature. Since we are interested in phase shifts, rather than absolute phase, the temperature dependent phase φ is given by

$$\varphi = \lambda_\varphi \Delta T \quad (22)$$

After final low-pass filtering, the analog front-end output signal U_0 is obtained.

4.2. Model Parameters

Since the availability of the new type of sensors for experimentation purposes was still uncertain at the time of creation of the model, it was decided that all simulations and experiments would be performed with the old sensor type. Therefore, the model parameters discussed here have values corresponding to the old sensor type.

The sensitivity of the GMR corresponds to the slope of the GMR curve. A typical GMR curve is shown in Figure 6. This curve was measured by placing a GMR between two Helmholtz coils that were used to apply a well-defined homogeneous magnetic field. A resistance measurement was subsequently performed on several sensors at room temperature while varying the applied magnetic field [17]. From these experiments, the average sensitivity was determined to be $3.5\text{e-}3 \text{ } \Omega\text{m/A}$.

The temperature dependency of the GMR sensitivity was determined experimentally by heating experiments, in which several sensors were heated in an oven. During the experiments the sensor surface temperature was monitored by means of a thermocouple. Simultaneously, the change in sensor output was measured. From these experiments the temperature coefficient of the GMR sensitivity could be determined, which was found to be $1.8\text{e-}3 \text{ K}^{-1}$.

The parameters α , β , γ and ε can be determined from Figure 13, which shows the relevant spectral components of the voltage across the GMR. This spectrum was measured using a GMR with a R_0 of $425 \text{ } \Omega$, while the amplitudes of the sense and excitation currents were set to $1.31 \text{ mA}_{\text{RMS}}$ and $24.75 \text{ mA}_{\text{RMS}}$ per conductor, respectively. The capacitive crosstalk was measured to be $2.8 \text{ mV}_{\text{RMS}}$ and is given by

$$U_{C/I-XT} = I_{C/I-XT} R_0 = j\omega_e \alpha I_e R_0 \quad (23)$$

Thus,

$$\alpha = \frac{U_{C/I-XT}}{I_e R_0 \omega_e} = 20 \cdot 10^{-12} \text{ (rad/s)}^{-1} \quad (24)$$

Parameter β , corresponding to the self-magnetisation effect, was measured by applying only the sense current and measuring the amplitude of the second harmonic of the sense current frequency U_{2fs} , which was found to have an amplitude of 0.23 V_{RMS}. The signal at this frequency is given by

$$U_{2fs} = I_s R_s = \beta S I_s^2 \quad (25)$$

in which R_s is the GMR resistance as a result of the self-magnetising field of the sense current. Parameter β is thus given by

$$\beta = \frac{U_{2fs}}{S I_s} = 50050 \text{ m}^{-1} \quad (26)$$

The relatively high value for β can be explained from the fact that the magnetic field is generated very close to the sensitive layer of the GMR.

Parameter γ , corresponding to the magnetic crosstalk, can be obtained from the magnetic signal in the absence of beads, i.e. the magnetic crosstalk signal U_{mxt} that was found to have an amplitude of 32e-6 V_{RMS}. This signal is given by

$$U_{mxt} = \gamma S I_e I_s \quad (27)$$

such that,

$$\gamma = \frac{U_{mxt}}{S I_e I_s} = 141 \text{ m}^{-1} \quad (28)$$

Finally, parameter ε corresponds to the signal per bead that is about 80 dB lower than the magnetic crosstalk, such that

$$\varepsilon = 10^{-4} \cdot \gamma = 14.1 \cdot 10^{-3} \text{ m}^{-1} \quad (29)$$

Figure 35 depicts the simulated spectrum of the voltage across the GMR when the model is excited with identical sense- and excitation signals as those used to excite the real sensor of which the measured spectrum is shown in Figure 13. It can be seen that the same spectral components with similar amplitudes are obtained with the model.

The noise voltage was modelled as additive white Gaussian noise by multiplying a zero mean unity variance normally distributed random number by a noise gain of $(5.5\text{nV}/\sqrt{\text{Hz}}) \cdot \sqrt{(f_s/2)}$. In this simulation a sample frequency f_s of 4Mhz was used, such that a noise gain of $7.78\text{e-}6$ was used. During the simulations discussed in Section 6.2, a sample frequency of 40 MHz was used, such that a noise gain of $2.46\text{e-}5$ was required. The noise level of -147 dB shown in Figure 35 is obtained by multiplying the noise voltage spectral density ($5.5\text{nV}/\sqrt{\text{Hz}}$) by the square root of the frequency resolution used in the FFT, which was chosen 61 Hz.




Figure 35: Spectrum of voltage across GMR. The frequency resolution used is 61Hz.

It can be concluded that the same spectral components with similar amplitudes and the same noise voltage spectral density are obtained when exciting the model, compared to when exciting a real sensor, such that the validity of the model is concluded.

Finally, the pre-amplifier gain was chosen 106 and a 2nd order Butterworth filter was selected having a cut-off frequency of 100 kHz in order to obtain 40 dB suppression at 1 MHz. Figure 36 shows the spectrum at the output of the analog pre-amplifier.




Figure 36: Spectrum at output of analog pre-amplifier.

4.3. Magnetic Signal and Noise Model

Figure 37 depicts a vector diagram of a typical magnetic signal, i.e. the components of the voltage across the GMR occurring at the difference between the sense- and excitation frequency, as well as the different components contributing to this signal. These components and their origin will be discussed and expressed in the model parameters of the sensor model of Section 4.1.

The bead signal, i.e. the signal of interest, is the result of the multiplication of the sense current and the GMR resistance change R_B as a result of the field originating from the magnetised surface bound beads. The bead signal U_B is given by

$$U_B = I_s R_B = \epsilon B S I_e I_s \quad (30)$$

Magnetic cross talk due to misalignment of current conductors and GMR results in resistance variations R_e induced by the primary field of the excitation current. Multiplied by the sense current, these resistance variations result in a magnetic crosstalk signal having the same phase as the bead signal. This magnetic crosstalk signal U_{mxt} is given by

$$U_{mxt} = I_s R_e = \gamma S I_e I_s \quad (31)$$

Capacitive and inductive cross talk, which is inherent to the sensor geometry, give rise to a current $I_{C/I-XT}$ through the GMR with a frequency equal to the excitation frequency. Furthermore, the applied sense current gives rise to a self-magnetisation field H_s in the GMR at the sense current frequency. This field subsequently results in a GMR resistance variation R_s at the same frequency. Their product results in a GMR voltage component at the same frequency as the bead signal, yet 90 degrees shifted in phase. Furthermore, the crosstalk current itself will generate self-magnetisation and a corresponding resistance change $R_{C/I-XT}$. The multiplication of this resistance and the sense current I_s results in a

second 90 degrees shifted component. The resulting quadrate signal U_Q can be described using the model parameters of Section 4.1 as

$$U_Q = I_{C/I-T} R_s + I_s R_{C/I-T} = 2j\omega_e \alpha \beta S I_e I_s \quad (32)$$

Finally, $U_{GMR,0}$ and $U_{GMR,t}$ represent the actual measured magnetic signal before the experiment and during the experiment, respectively. The measured magnetic signal is the combination of the three above-mentioned contributions.




Figure 37: Magnetic signal and different contributions in the complex plane.

Parameters α, β, γ and ϵ are all fixed in the geometry of the sensor and will thus not change during measurements. It is not expected that large variations will occur between sensors on the same chip. The parameters may, however, vary between different batches of sensors, as mask misalignment during the lithographic steps of the sensor production will cause parameter variations.

It should be noted that all vectors contain the factor $S I_e I_s$. Therefore, variations in the sensitivity, e.g. as a result of fluctuations in temperature and/or external magnetic fields, variations in the excitation- and sense current, and gain variations in the analog front-end scale all signal components equally. Phase variations in the analog front-end, on the other hand, cause a rotation of the axes around the origin.

As already discussed, the noise is white at the frequency of interest and the noise voltage spectral density is lower bound to the thermal noise of the GMR. The noise density used during simulations was $5.5 \text{ nV}/\sqrt{\text{Hz}}$, while the actually achieved noise density was shown in Section 3.3 to be $4.6 \text{ nV}/\sqrt{\text{Hz}}$.

5. Robust Detection Algorithms

A spectrum comprising the most important components of the voltage across the GMR has already been shown in Figure 13. This figure also shows the noise limit, which is the thermal noise of the GMR. As discussed in Chapter 3, analog pre-processing is used to select and amplify the magnetic signal in order to be able to determine its amplitude. However, although required, measurement of the amplitude of the magnetic signal alone is not enough to accurately determine the bead density on the sensor surface. This results from the fact that the magnetic signal comprises multiple components of different origin that contribute to its amplitude. Moreover, temperature induced gain and phase variations in the analog front-end and magnetic field induced variations in GMR sensitivity may also change the signal. Further processing is needed to extract the actual information-carrying signal that is linearly related to the bead density on the sensor surface and independent of temperature and external magnetic fields. Two algorithms were developed to this end and subsequently implemented in the FPGA as part of the ‘algorithm’ block in Figure 19. These algorithms are discussed in Section 5.1. Since it is such an important subsystem, an extensive discussion of the joint phase and amplitude estimator (ϕ, A estimator) that is employed to produce amplitude- and phase estimates as input for the algorithm block, as shown in Figure 19, is given in Section 5.2. This chapter finally ends with a section on simulations of the algorithms in Matlab Simulink.

5.1. Description of Robust Detection Algorithms

In the previous platform the phase of the demodulation signal, i.e. U_{ref} in Figure 9, was manually adjusted to maximize the output signal of the platform at the start of the experiment. This effectively means that the demodulation axis was aligned with $U_{GMR,0}$ in Figure 37. It can be seen that during the experiment only the amplitude of the projection of the bead vector U_B on the demodulation axis is measured, instead of its real amplitude. The severity of misalignment between the demodulation axis and the bead vector furthermore depends on the a priori unknown crosstalk vector U_Q , i.e. the bead signal is scaled as a result of the amount of crosstalk. Moreover, phase shifts in the analog front-end output signal occurring during the measurement result in non-linear scaling of the measured bead signal. Finally, gain variations in the front-end also scale the bead vector with respect to the demodulation signal. Clearly, methods are needed to extract the bead signal from the magnetic signal.

To this end, two algorithms were developed. The first algorithm, which is discussed in Subsection 5.1.1, makes use of an on-chip reference sensor to extract the bead signal. The second algorithm, discussed in Subsection 5.1.2, does not require a reference sensor, resulting in more efficient chip area usage.

5.1.1. Algorithm 1: Reference Sensor Scheme

As we are only interested in the contribution of the beads, we need to determine the length of bead vector U_B . This vector is obtained by subtracting the magnetic signal at time zero $U_{GMR,0}$, i.e. before the start of the experiment, from the magnetic signal

measured during the experiment $U_{GMR,t}$. These vectors can be determined using the phase and amplitude estimation techniques described in Section 5.2. An arbitrary but constant demodulation vector U_{ref} is chosen as depicted in Figure 38. Phase and amplitude estimation is subsequently used to determine angles φ and θ , as well as the modulus of the vectors $U_{GMR,0}$ and $U_{GMR,t}$, denoted as A_0 and A_t respectively. The modulus of the bead vector A_B , which is given by

$$A_B = |A_0 e^{j\varphi} - A_t e^{j\theta}| = |\epsilon B S I_e I_s| \quad (33)$$

is linearly dependent on the density of beads on the sensor surface.




Figure 38: Magnetic signal vectors and demodulation vector.

However, gain and phase variations in the transfer of the sensor and analog front-end will result in scaling and rotation of $U_{GMR,t}$ with respect to $U_{GMR,0}$. These effects can be compensated for by making use of a reference sensor that has a shielded sensor surface such that beads cannot reach the surface, i.e the reference sensor measures only crosstalk. Figure 39 depicts the reference signals $U_{r,0}$ and $U_{r,t}$ that are obtained from the reference sensor at time zero and time t , respectively. This reference signal is used to cancel the effects of gain and phase variations on $U_{GMR,t}$ during the experiment. The modulus of the compensated bead signal $A_{B,c}$ is given by

$$A_{B,c} = \left| A_0 e^{j\varphi} - A_t e^{j\theta} \cdot \frac{A_{r,0} e^{j\sigma_0}}{A_{r,t} e^{j\sigma_t}} \right| \quad (34)$$

Note that absolute angle and modulus of the reference vector are not important: only their relative change during the measurement is important, as long as they can be determined with sufficient SNR. Since the crosstalk is relatively large compared to the bead signal, this is guaranteed. Note furthermore that this algorithm only works if the reference sensor reacts similarly to environmental variations as the actual measurement sensor. This assumption is fair since they are close to each other on the sensor chip and thus experience practically the same external magnetic field strength and temperature.




Figure 39: Magnetic signal vectors of the measurement sensor and the reference sensor and demodulation vector.

5.1.2. Algorithm 2: Vector Decomposition Scheme

The second algorithm does not require a reference sensor. Instead, it uses the quadrature signal U_Q as a reference signal and relies on a vector decomposition scheme to obtain the modulus of the bead signal.

As can be seen from Equation 32, the amplitude of the quadrature signal is linearly dependent on the frequency of the excitation current. The amplitude of bead signal U_B and magnetic crosstalk signal U_{mxt} , on the other hand, are independent of the excitation frequency, as is shown by Equations 30 and 31. This knowledge, and the fact that the phase of U_Q is 90 degrees shifted with respect to U_B and U_{mxt} , can be used to extract the amplitude of U_Q and the combination of U_B and U_{mxt} . For convenience the addition of the bead signal and the magnetic cross talk signal will be referred to as the in-phase signal U_I , hence

$$U_I = U_B + U_{mxt} \quad (35)$$

In order to determine above-mentioned amplitudes, the sensor is alternatingly excited at two excitation frequencies, $f_{e,1}$ and $f_{e,2}$, where the ratio between the two frequencies is a factor $N > 1$. Figure 40 depicts all signals schematically for both excitation frequencies, where $U_{GMR,1}$ and $U_{GMR,2}$ are the magnetic signals obtained using the first and second excitation frequency, respectively. Similarly, $U_{Q,1}$ and $U_{Q,2}$ are the quadrature signals for the first and second excitation frequency, respectively.




Figure 40: magnetic signal vectors resulting from excitation at two different frequencies.

From the measured amplitudes of the magnetic signal $U_{GMR,1}$ and $U_{GMR,2}$, which will be referred to as A_1 and A_2 respectively, the amplitude of the in-phase and quadrature signal can be obtained by solving the following system of equations:

$$\begin{cases} A_1^2 = A_I^2 + A_Q^2 \\ A_2^2 = A_I^2 + (NA_Q)^2 \end{cases} \quad (36)$$

where A_I and A_Q denote the amplitude of U_I and U_Q , respectively and N is the ratio between the two excitation frequencies used. Solving the system for A_I , we obtain

$$A_I = \sqrt{\frac{(NA_1)^2 - A_2^2}{N^2 - 1}} = |(\gamma + \epsilon B)SI_e I_s| \quad (37)$$

The obtained amplitude of the in-phase signal is, however, still sensitive to gain variations in the sensor and analog front-end during the measurement. These can be compensated for by normalising the amplitude of the in-phase signal by the amplitude of the quadrature signal, which is scaled equally by gain variations. This quadrature amplitude is obtained by solving the system for A_Q , resulting in

$$A_Q = \sqrt{\frac{A_2^2 - A_1^2}{N^2 - 1}} = |2\omega_e \alpha \beta SI_e I_s| \quad (38)$$

After normalisation, an amplitude $A_{I,c}$ is obtained that is compensated for gain and phase variations in the analog front-end, variations in excitation- and sense current amplitude and variations in the sensitivity of the GMR, i.e.

$$A_{I,c} = \frac{\gamma + \epsilon B}{2\omega_e \alpha \beta} \quad (39)$$

Finally, the compensated amplitude of the bead signal $A_{B,c}$ is extracted by subtracting the in-phase amplitude $A_{I,c0}$ measured before the start of the experiment, from the in-phase component measured during the experiment, such that we obtain

$$A_{B,c} = A_{I,c} - A_{I,c0} = \frac{\epsilon B}{2\omega_e \alpha \beta} \quad (40)$$

Since the parameters α, β and ϵ are all fixed in the geometry of the sensor, the amplitude of the bead signal obtained using this algorithm only varies with the amount of surface bound beads B .

Note that by adjusting the sense frequency accordingly, the frequency of the magnetic signal can be preserved while alternating between both excitation frequencies. This guarantees an equal transfer of the magnetic signal by the analog front-end in both excitation cases. The method to generate the corresponding excitation and sense signals is discussed in detail in Subsection 6.1.1.

5.2. Joint Phase and Amplitude Estimation

This section discusses a joint phase and amplitude estimator that is used to track variations in amplitude and phase of the sensor signal during the measurement [18]. The obtained estimates are used as input for the detection algorithms. The estimator uses gradient-based least-squares optimisation to obtain phase and amplitude estimates with maximum likelihood. The now following subsections give a theoretical treatment of the joint phase and amplitude estimator, supported by some simulation results of a joint phase and amplitude estimator that was implemented in Matlab.

5.2.1. Gradient Based Least Squares Adaptation

Least-squares (LS) optimisation is employed to estimate with maximum likelihood the value of a parameter θ of an input signal $r(t)$ in the presence of additive white Gaussian noise (AWGN). To this end, LS optimisation strives to minimise the compound energy $\mathcal{E}(\varphi)$ of an error signal $e(\varphi, t)$, which is indicative of the mismatch between the actual value of the parameter θ and its estimate φ . The compound error energy is given by

$$\mathcal{E}(\varphi) = \int e^2(\varphi, t) dt \quad (41)$$

When it is known a priori that the compound energy as a result of the estimate φ has only a single local minimum in the interval of interest, the minimum of the compound energy can be found by locating the zero-crossing of its gradient $\xi(\varphi)$, which is given by

$$\xi(\varphi) = \frac{d\mathcal{E}(\varphi)}{d\varphi} = 2 \int e(\varphi, t) \frac{de(\varphi, t)}{d\varphi} dt = -2 \int e(\varphi, t) s(\varphi, t) dt \quad (42)$$

in which $s(\varphi, t)$ is referred to as the signature of the mismatch Δ between θ and φ , since the presence of this component in the error signal $e(\varphi, t)$ indicates a parameter mismatch. The term $e(\varphi, t)s(\varphi, t)$ is known as the stochastic gradient η of the error energy.

As will become clear in the now following, a joint phase and amplitude estimator can be constructed by simply combining a phase estimator and an amplitude estimator, without introducing unwanted coupling between the estimators. Both estimator types may therefore be analysed separately. The results obtained in this way are also valid for the joint estimator. Properties that hold for both estimators are furthermore treated in general. In the general case the parameter to be estimated is denoted as θ , while the estimate is denoted as φ .

5.2.2. Gradient Based Closed Loop Phase Estimation

As stated before gradient based estimation requires that the error energy shows only a signal local minimum in the interval of interest. In the absence of noise, the error energy for phase estimation $\mathcal{E}(\varphi)$ is found by integration of the error signal $e(\varphi, t) = r(t) - x(\varphi, t)$. In which the input signal $r(t)$ is given by $r(t) = \sin(\omega t + \theta)$, while reference signal $x(\varphi, t)$ is given by $x(\varphi, t) = \sin(\omega t + \varphi)$. The error energy then satisfies

$$\mathcal{E}(\varphi) = \int_{I_0} e^2(\varphi, t) dt = I_0 (1 - \cos(\theta - \varphi)) \quad (43)$$

The normalised error energy $\mathcal{E}(\Delta)$, which equals $\mathcal{E}(\varphi)/I_0$ and is depicted in Figure 41, is a periodic function of φ . The occurrence of a single local minimum within the interval of interest $[-\pi, \pi]$ guarantees convergence.




Figure 41: Normalised error energy versus phase mismatch Δ .

Figure 42 depicts a gradient based closed loop phase estimator. The estimator loop adapts the phase of the reference signal $\sin \omega t$ to obtain the adapted reference signal $x(\varphi, t) = \sin(\omega t + \varphi)$ such that the amplitude of the signature component $s_p(\varphi, t)$ in the error signal $e(t) = r(t) - x(\varphi, t)$, and therefore the error energy $\mathcal{E}(\varphi)$, is minimised.

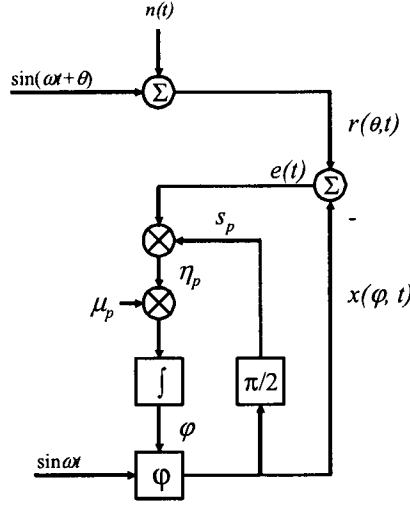


Figure 42: Closed loop phase estimator.

The phase misadjustment signature s_p is given by

$$s_p(\varphi, t) = \frac{dx(\varphi, t)}{d\varphi} = \cos(\omega t + \varphi) \quad (44)$$

In practice, minimisation is achieved using the stochastic gradient η_p of the error energy, i.e. a noisy instantaneous measure of phase mismatch Δ , given by

$$\eta_p(\varphi, t) = e(\varphi, t)s_p(\varphi, t) \quad (45)$$

Here $\eta_p(\varphi, t) = -\frac{1}{2}\sin(2\omega t + 2\varphi) + \frac{1}{2}\sin(\theta - \varphi) + \frac{1}{2}\sin(2\omega t + \theta + \varphi) + n(t)\cos(\omega t + \varphi)$.

The DC component of this signal, i.e. the mean of η_p , represents the phase mismatch and is found by integration over interval I_0 , which is preferably as large as possible. The effective size of interval I_0 is inversely proportional to bandwidth of the adaptation loop, which is controlled via the adaptation gain μ_p . This gain influences the adaptation speed and the final mismatch, as will become clear in the following subsections. The gradient of the error energy $\xi(\varphi)$, which is given by

$$\xi(\varphi) = \int_{I_0} \eta_p(\varphi, t) dt = \frac{1}{2} I_0 \sin(\theta - \varphi) + v_p \quad (46)$$

is a phase mismatch indicator that is used to adapt the phase of the reference signal. Here $v_p = \int_{I_0} n(t)\cos(\omega t + \varphi) dt$. The constant factor -2 used in Equation 42 is omitted here,

since it effects signal and noise equally. To make further analysis less complex, the gradient may in the vicinity of $\theta \approx \varphi$ be linearised using first order Taylor approximation, i.e. $x(\varphi, t) \approx x(\theta, t) - \Delta_p s_p$ and thus

$e(t) = r(\theta, t) - x(\varphi, t) = x(\theta, t) + n(t) - x(\varphi, t) = \Delta_p s_p + n(t)$. In the absence of noise, the

linearised gradient then satisfies $\xi(\phi) = \frac{1}{2} I_0(\theta - \phi)$. Figure 43 depicts the gradient and linearised gradient as a function of the phase mismatch Δ .




Figure 43: Gradient and linearised gradient of phase mismatch around $\theta = \phi$, normalised to I_0 .

Since our interval of interest is $[-\pi, \pi]$, also the extreme values of $-\pi$ and π may occur. Figure 43 shows that for these specific values the gradient is zero and the system is in equilibrium. This equilibrium is however unstable, as can be seen from the error energy function depicted in Figure 41. The LS-algorithm, which is a steepest-decent algorithm, will converge in the direction opposite to the gradient direction, i.e. it will converge to the minimum of the error energy function. However, in theory, it may take the estimator infinitely long to converge. In practice noise will introduce small deviations from $|\Delta| = \pi$ such that convergence will always occur, but this cannot be guaranteed. In practice, measures may be taken to overcome this problem, as will become clear in the following subsections.

5.2.3. Gradient Based Closed Loop Amplitude Estimation

Figure 44 depicts a gradient based amplitude estimator. This estimator adapts the amplitude of the reference signal $x(A', t)$ such that the amplitude of amplitude mismatch signature s_a in the error signal $e(t) = r(A, t) - x(A', t)$ is minimised. Here the input signal $r(t)$ is given by $r(t) = A \sin \omega t + n(t)$, while reference signal $x(A', t)$ that is adapted by the estimator is given by $x(A', t) = A' \sin \omega t$.




Figure 44: Closed loop amplitude estimator.

The signature s_a is given by

$$s_a(A', t) = \frac{dx(A', t)}{dA'} = \sin(\omega t) \quad (47)$$

The stochastic gradient η_a of the error energy, i.e. an instantaneous measure of amplitude mismatch $A - A'$ is given by

$$\eta_a(t) = e(t)s_a(A', t) \quad (48)$$

In this case $\eta_a(t) = \frac{1}{2}(A - A') - \frac{1}{2}(A - A')\cos(2\omega t) + n(t)\sin(\omega t)$. The DC component of this signal, i.e. the mean of η_a , represents the parameter mismatch and is obtained by integration, while the adaptation gain μ_a can again be used to influence the adaptation speed and the final mismatch. The gradient $\xi(A')$, shown in Figure 45 and given by

$$\xi(A') = \int_{I_0} \eta_a(t) dt = \frac{1}{2} I_0 (A - A') + v_a \quad (49)$$

is an amplitude mismatch indicator that can be used to adapt the amplitude of the reference signal $x(A', t)$. Here $v_a = \int_{I_0} n(t) \sin(\omega t) dt$.




Figure 45: Gradient of amplitude mismatch, normalised to I_0 .

5.2.4. Joint Phase and Amplitude Estimation

The joint phase and amplitude estimator is depicted in Figure 46. For proper functioning of this joint phase and amplitude estimator both estimation loops should be uncoupled. This is, however, guaranteed since the signatures used in the phase- and amplitude loops, s_p and s_a respectively, are exactly orthogonal, i.e. the presence of s_a in the error signal will not be demodulated to DC upon multiplication with s_p and vice versa.

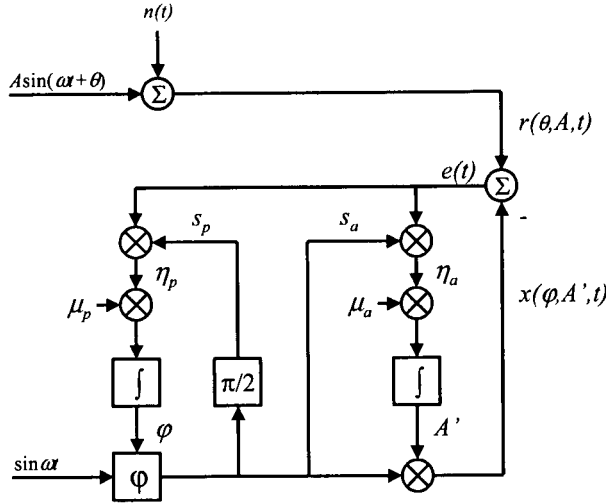


Figure 46: Joint phase and amplitude estimator.

The combined estimates give ambiguous solutions to describe the input signal, since $(A', \varphi + k2\pi)$ equals $(-A', \varphi + (2k+1)\pi)$ for all A' , φ and $k \in \mathbb{Z}$. This poses no problem for algorithm 1 as it uses only the combined estimates, i.e. $A' e^{j\varphi t}$, to obtain the signal of

interest, i.e. the modulus of the bead vector U_B . Algorithm 2 on the other hand uses only the amplitude estimate. However, to ensure correct functioning of algorithm 2 it suffices to guarantee positive values of A' . Equation 50 gives expressions of the stochastic gradients η_a and η_p for the joint estimator, while Equation 51 gives the expressions for the gradients $\xi(A')$ and $\xi(\varphi)$.

$$\begin{aligned}\eta_p &= -\frac{A'}{2} \sin(2\omega t + 2\varphi) + \frac{A}{2} \sin(\theta - \varphi) + \frac{A}{2} \sin(2\omega t + \theta + \varphi) \\ \eta_a &= -\frac{A'}{2} \cos(2\omega t + 2\varphi) + \frac{A}{2} \cos(\theta - \varphi) + \frac{A}{2} \cos(2\omega t + \theta + \varphi) - \frac{A'}{2}\end{aligned}\quad (50)$$

$$\begin{aligned}\xi(\varphi) &= \int_{I_0} \eta_p dt = \frac{A}{2} I_0 \sin(\theta - \varphi) \\ \xi(A') &= \int_{I_0} \eta_a dt = \left(\frac{A}{2} \cos(\theta - \varphi) - \frac{A'}{2} \right) I_0\end{aligned}\quad (51)$$

Where the noise terms v_a and v_p have been omitted. It can be seen that the estimator converges to the stable equilibrium $\theta - \varphi = 0$ if A is positive. On the other hand, if A is negative the gradient function $\xi(\varphi)$, switches sign and $\theta - \varphi = \pi$ becomes the stable equilibrium to which the estimator will converge. In this case the gradient $\xi(A) = (-A - A')I_0$, such that A' will be positive. A positive value for the amplitude estimate A' is thus guaranteed for all values of A . Finally, it can be seen that the value of the mismatch estimate to which the system converges is independent of the initial amplitude estimate A_0 .

5.2.5. Discretisation

In the new platform the joint phase and amplitude estimator will be implemented in the digital domain. Therefore, after having discussed the general principle of phase and amplitude estimation in the time continuous domain, more specific properties like the range of convergence, the time constant of the loops and the estimate gradient noise are analysed using the time discrete representation of the LS algorithm. To this end, a conversion from the time continuous domain to the time discrete domain is required. In practice, the estimator is also amplitude discrete such that quantisation noise will be added upon transformation of the input signal to the digital domain. However, in Subsection 3.2.3 it was already shown that the quantisation noise is smaller than the noise already present in the input signal. The required word length for the estimates is furthermore discussed in Subsection 6.1.2.

In the time continuous case, the estimates are obtained by continuous integration of the stochastic gradient. As can be seen from Figure 46, considering the general case, the estimate φ is given by

$$\varphi(t) = \mu \int_0^t \eta(t) dt = \mu \int_0^t e(t) s(t) dt \quad (52)$$

The time discrete representation of this relation is given by

$$\phi_k = \mu \sum_{i=0}^{k-1} e_i s_i \quad (53)$$

in which e_k and s_k are the time discrete representations of the error signal $e(t)$ and the signature $s(t)$, respectively. These time discrete representations are obtained from the time continuous representations by substituting t by kT_s , where T_s is the sample period of the digital system and k is an integer. Here, it is assumed that the signals are band-limited, such that their maximum frequency is smaller than $1/(2T_s)$.

From Equation 53 it follows that

$$\phi_{k+1} = \mu \sum_{i=0}^k e_i s_i = \mu \sum_{i=0}^{k-1} e_i s_i + \mu e_k s_k \quad (54)$$

which results in the time discrete representation of the LS algorithm for parameter estimation, i.e. the LS update rule:

$$\phi_{k+1} = \phi_k + \mu e_k s_k \quad (55)$$

5.2.6. Dynamic Behaviour of the Adaptation Loops

For the amplitude estimator, the LS update rule is given by $A'_{k+1} = A'_k + \mu_a (s_k A + n_k - s_k A'_k) s_k$. Given the parameter mismatch $\Delta_{A,k} = A - A'_k$ we obtain $\Delta_{A,k+1} = (1 - \mu s_k^2) \Delta_{A,k} - \mu s_k n_k$. For the phase estimator the analysis is more difficult, but can be approximated by linearisation around $\theta \approx \phi_k$ using the Taylor approximation, i.e. $x_{\phi,k} \approx x_{\theta,k} - \Delta_{\phi,k} s_k$ and thus $\phi_{k+1} \approx \phi_k + \mu_p (s_k^2 \Delta_{\phi,k} + n_k s_k)$. As can be seen, the adaptation rule for the phase estimator now has the same form as that of the amplitude estimator, such that the phase and amplitude case can be treated simultaneously using the general parameter mismatch $\Delta_k = \theta - \phi_k$. The parameter estimation mismatch equals $\Delta_{k+1} = (1 - \mu s_k^2) \Delta_k - \mu s_k n_k$. Since the expected value of the noise $E\{n_k\} = 0$ it follows that the expected value of the parameter estimation mismatch $E\{\Delta_{k+1}\}$ is given by $E\{\Delta_{k+1}\} = E\{(1 - \mu s_k^2) \Delta_k\}$. Since variations in Δ_k are in general much slower than variations in s_k this may also be written as $E\{\Delta_{k+1}\} = (1 - \mu E\{s_k^2\}) E\{\Delta_k\} = (1 - \mu \sigma_s^2) E\{\Delta_k\}$, in which σ_s^2 is the power of the signature s_k . Finally, since $E\{\Delta_0\}$ is deterministic it follows that $E\{\Delta_0\} = \Delta_0$, such that

$$E\{\Delta_k\} = (1 - \mu \sigma_s^2)^k \Delta_0 \quad (56)$$

For both estimators we can now determine the region of convergence for the adaptation gain μ . Parameter mismatch Δ_k is guaranteed to converge to zero only if $-1 < (1 - \mu \sigma_s^2) < 1$ and thus $0 < \mu < \sigma_s^2$. Since for our sinusoidal signatures $\sigma_s^2 = \frac{A^2}{2} = 0.5$ we obtain $0 < \mu < 4$ as the range of convergence for the adaptation gain.

Note that the range of convergence does not in any way depend on the amplitude or SNR of the input signal.

The time constant of the loops follows from Equation 56 which is an exponentially decaying function, i.e. $(1 - \mu\sigma_s^2)^k = e^{\ln(1 - \mu\sigma_s^2)^k} = e^{k \ln(1 - \mu\sigma_s^2)} = e^{-kT_s/\tau}$, with time constant

$$\tau = \frac{-T_s}{\ln(1 - \mu\sigma_s^2)} \quad (57)$$

In which T_s is the sample period. The time constant of the loop thus depends solely on the sample period, the adaptation constant and the power of the signature and is independent of the SNR of the input signal. It is, however, dependent on the amplitude of the input signal.

Figure 47 shows the time constant as a result of the adaptation gain μ_ϕ , both calculated and simulated using a phase estimator that was implemented in Matlab. A sample frequency of 40 MHz and a signature amplitude of one were used. It can be seen that the simulation results agree well with the theory when $\tau \gg 2\pi/\omega$, as $\overline{s_k^2}^\tau = \sigma_s^2$ only holds for sufficiently long integration intervals. For larger adaptation constants the simulated values deviate somewhat from the theoretical values.




Figure 47: Time constant (simulated and calculated) vs. adaptation constant μ_ϕ

Similarly, Figure 48 shows the time constant as a result of the adaptation gain μ_a , both calculated and simulated using an amplitude estimator that was implemented in Matlab. Again, a sample frequency of 40 MHz and a signature amplitude of one were used.




Figure 48: Time constant (simulated and calculated) vs. adaptation constant μ_a .

5.2.7. Parameter Estimate Gradient Noise

Although noise does not affect the time constant, it does introduce uncertainty in the estimates in the form of gradient noise. Least-squares optimisation will produce maximum likelihood estimate with smallest RMS estimation error σ_φ , provided that $n(t)$ is additive white Gaussian noise (AWGN). Figure 49 depicts two equivalent representations of the amplitude adaptation loop, assuming for convenience that $s_k \neq 0 \quad \forall k$. In order to analyse σ_φ more easily, in the bottom figure the signature has been transferred to the right of the addition point, such that the input of the loop is now the parameter to be estimated (θ) and the output is the parameter estimation (φ). The noise on the input of the loop now becomes $z_k = n_k / s_k$ with variance $\sigma_z^2 = \sigma_n^2 / s_k^2 = \sigma_n^2 / \sigma_s^2$.




Figure 49: Equivalent representations of the adaptation loop.

The parameter estimate gradient noise σ_ϕ^2 is limited by the bandwidth of the loop, such that

$$\sigma_\phi^2 = \sigma_z^2 B A_{cl} \quad (58)$$

where A_{cl} is the closed loop gain, which is equal to one, and B is the normalised equivalent noise bandwidth of the loop, i.e. the ratio of the equivalent noise bandwidth of a first order loop $B_n = \pi B_1/2 = \pi/2\tau$ and the Nyquist frequency $\omega_{niq} = \pi/T_s$. For small loop gains $\tau \approx T_s/\mu\sigma_s^2$ and the normalised equivalent bandwidth equals

$$B = \frac{B_n}{\omega_{niq}} \approx \frac{\mu\sigma_s^2}{2} \quad (59)$$

The noise variance is thus given by

$$\sigma_\phi^2 = \frac{\mu\sigma_n^2}{2} \quad (60)$$

Note that filtering of input signal in the analog and digital pre-amplifier also limits the bandwidth of the noise and thus the input noise power. This, however, only affects the gradient noise if the bandwidth of the loop is larger than that of the input filter. If, on the other hand, the bandwidth of the loop is smaller than that of the pre-filter, the noise equivalent bandwidth scales inversely proportional to the input noise power, such that the gradient noise remains equal. In our case the bandwidth of the loop is determined by the bandwidth of the variations in parameters A and θ , which are much smaller than the bandwidth of the filters of the analog and digital pre-amplifier. The filtering that occurs in the pre-amplifier therefore has no effect on the SNR of the estimates and is merely employed to limit the dynamic range of the input signal, as explained in Section 3.1.

The quality of the estimate can for the general estimate ϕ be expressed as the signal to noise ratio (SNR) of the estimate SNR_ϕ , defined as

$$SNR_{\varphi} = \frac{\varphi^2}{\sigma_{\varphi}^2} \quad (61)$$

In the case of amplitude estimation of sinusoidal input signals the SNR of the amplitude estimate $SNR_{A'}$ can be expressed in the SNR of the input signal SNR_r by substituting estimate A' for the general estimate φ in Equations 60 and 61 and combining the results.

$$SNR_{A'} = \frac{4}{\mu} SNR_r \quad (62)$$

The quality of the estimation is thus proportional to the SNR of the input signal and inversely proportional to the loop gain. It is thus advisable to use a loop gain as small as allowed by the requirement for convergence time.

Figure 50 depicts the simulated and calculated amplitude estimation quality, in the form of the ratio between input signal SNR and amplitude estimate SNR. The simulation results confirm Equation 62.




Figure 50: Amplitude estimation quality: ratio between estimate- and input SNR vs. adaptation constant μ_A .

In the case of phase estimation the quality of the estimation can be expressed as a relative phase uncertainty range e_{φ} , which is defined as the phase range that included the actual phase with a certainty of 99.7% (i.e. 3σ) divided by the total phase interval $[-\pi, \pi]$.

$$e_{\varphi} = \frac{3\sigma_{\varphi}}{2\pi} \quad (63)$$

Note that this definition requires the gradient noise to have Gaussian probability density. This, however, is the case for small mismatch values for which the gradient noise depends linearly on the input noise, which is AWGN.

Figure 51 depicts the simulated and calculated phase estimation quality, which is here represented by the relative phase uncertainty range, as a function of the adaptation constant μ_ϕ . The simulations agree well with Equation 63.




Figure 51: Phase estimation quality: relative phase uncertainty range vs. adaptation constant μ_ϕ .

5.2.8. Estimate Reliability

Three criteria can be formulated to test the reliability of the estimates obtained with the joint phase and amplitude estimator. Firstly, since it is of no use to measure surface bead densities that are below the detection limit of the system, which was defined in Subsection 2.1.1 as $L_D = 2.71\sigma$, the minimum value of the amplitude estimate is equal to L_D . Hence, the minimum value for the amplitude estimate A_{min}' is given by $A_{min}' = 2.71\sigma_{A'}$. Using Equation 60 and substitution A' for φ , A_{min}' can be expressed in terms of the a priori known noise power σ_n^2 of the input signal and the adaptation constant μ , i.e.

$$A_{min}' = 1.92\sqrt{\mu\sigma_n^2} \quad (64)$$

Secondly, convergence should have occurred. Convergence is defined here as having approached the final value to within $3\sigma_{A'}$. In other words the residual mismatch is comparable to the gradient noise. Since $A_k' = A(1 - e^{-kT_s/\tau})$ and mismatch $\Delta_{A'} = A - A_t'$, it follows that

$$\Delta_{A'} = \frac{e^{-kT_s/\tau}}{1 - e^{-kT_s/\tau}} A_k' < 3\sigma_{A'} \quad (65)$$

where $\sigma_{A'}$ can be calculated using Equation 60.

Finally, for $\theta = \pi$ it may take infinitely long to converge. It is thus important to prevent the system from operating in this area. The occurrence of $\theta \approx \pi$ can be detected by

simultaneously monitoring the means of η_φ and η_A , e.g. by integrating over a reasonably small number of periods of the input signal. If simultaneously the mean of η_φ is small and that of η_A is negative, the system is not converging sufficiently fast, indicating that θ is in the vicinity of π . In this case we should invert the sign of the reference signal to obtain guaranteed convergence within the allowed time span and add π to the phase estimate φ .

5.3. Algorithm Simulation Results

In order to test the algorithms, simulations were performed in Matlab Simulink. Figure 52 schematically depicts the simulation set-up. The combined sensor and analog front-end model is excited using the Simulink model of the exact waveform generator that will be implemented in the FPGA. The waveform generator is discussed in detail in Subsection 6.1.1. The sensor and front-end model was extended to include some digital pre-amplification, by increasing the gain by a factor 100 and adding 8 orders to the lowpass filter. The digital pre-amplifier that will perform this task in the actual platform is discussed in Subsection 6.1.3. The output signal of the sensor and front-end model, comprising the magnetic signal, is supplied to the simulink model of the exact joint phase and amplitude estimator that will be implemented in the FPGA. The estimator is discussed in Subsection 6.1.2. The phase and amplitude estimates, φ and A respectively, were recorded for different values for the sense and excitation frequencies and the temperature T . The recorded estimates were subsequently used to test the algorithm. In order to measure only signal variations as a result of temperature variations, the number of beads on the sensor surface was kept constant during the simulations, namely zero, such that the in-phase component of the magnetic signal consisted solely of the constant magnetic crosstalk signal.




Figure 52: Simulation set-up that was used to test algorithm.

To prevent excessive simulation times the algorithm was tested in the absence of noise in the sensor and front-end model. By doing so, the adaptation gain of the estimator could be chosen relatively large, such that convergence occurs within reasonable times while at the same time accurate estimate values are obtained. The estimator was however also simulated separately with noisy front-end output signals, in order to determine the performance of the estimator. The simulation result of the estimator are discussed in Subsection 6.2.2.

Algorithm 1 could be simulated by alternatingly exciting two separate models of the sensor and analog front-end. Since the parameters of the models are identical, except for the room temperature sensitivity s_{20° and an additional initial phase shift that may be added to one of the models, the simulated relative amplitude and phase variations as a result of temperature would be exactly equal for both models. Moreover since the sense- and excitation waveforms applied to the sensors are identical and noise is not included in the simulations, algorithm 1 would give exact results. Not much insight would therefore be gained from simulating algorithm 1. Actual measurements should be performed to confirm the assumption that the sensitivity of both sensors varies comparably as a result of temperature.

Contrary to algorithm 1, algorithm 2 is worth simulating, since in this case only a single sensor is employed, but this sensor is excited at two pairs of sense- and excitation frequencies. The simulation therefore not only tests the algorithm itself, but also the performance of the waveform generator. During the simulations, the sensor was excited at frequency pairs ($f_s = 1.25$ MHz, $f_e = 1.27$ MHz) and ($f_s = 2.50$ MHz, $f_e = 2.52$ MHz), such that the magnetic signal remained at a constant frequency of 20 kHz.

Figure 53 shows the calculated magnetic crosstalk (U_{mxt}) upon excitation with frequency pair ($f_s = 1.25$ MHz, $f_e = 1.27$ MHz), obtained using Matlab to calculate the standard deviation of the magnetic cross talk related component of the combined sensor and front-end model output signal. Furthermore the magnetic crosstalk upon excitation with frequency pair ($f_s = 1.25$ MHz, $f_e = 1.27$ MHz), obtained using the estimator ($U_{mxt,est}$) is shown, as well as the magnetic crosstalk, compensated for temperature deviations from room temperature by the vector decomposition algorithm ($U_{mxt,decomp}$). From the results shown in Figure 53, it can be concluded that the estimated values for U_{mxt} correspond well with the calculated values over the complete temperature range. Hence, the estimator performs well. Furthermore, it can be concluded that the algorithm is indeed capable of extracting, and referring back to 20° Celsius, the in-phase component of the magnetic signal, for all temperatures. It was found that the maximum deviation of U_{mxt} from its calculated value was 22e-9 V. Given the fact that, for typical sense and excitation currents, the average signal per surface bound bead is about 5 nV, this maximum deviation corresponds to about 5 surface bound beads.




Figure 53: Calculated magnetic crosstalk (U_{mxt}) using Matlab, magnetic cross talk estimates obtained using the ϕ, A -estimator ($U_{mxt,est}$) and magnetic crosstalk compensated for temperature deviations from room temperature by the vector decomposition algorithm.

In Figure 54 the decomposition algorithm is compared to demodulation with a fixed phase. Fixed phase demodulation is employed in the old platform, where the phase of the demodulation signal U_{ref} is fixed at the value that maximises the demodulated magnetic signal before the start of the surface bead density measurements. It can be seen that, contrary to the results obtained with the decomposition algorithm, fixed phase demodulation results in a temperature dependency of the output signal. Moreover, the relation between the output signal and temperature is non-linear. Since in the fixed phase demodulation case the amplitude of the magnetic signal is measured, which is here equal to

$$U_{GMR} = \sqrt{U_Q^2 + U_{mxt}^2},$$

an error is made in the estimation of the signal of interest U_{mxt} . This error is shown in Figure 54 as the difference between the decomposition algorithm result and the fixed phase result for room temperature. Finally, the results obtained with an adaptive phase scheme, where the phase of the demodulation signal U_{ref} is adapted in order to correctly measure the magnetic signal amplitude, are shown in Figure 54. Using the adaptive phase scheme, the output signal is also dependent on temperature, albeit linearly. Furthermore the same error is made, since U_{GMR} is measured instead of U_{mxt} or, in general, U_I . On other words, adaptive phase demodulation corrects for phase variations in the transfer of the analog front-end during the measurements, but it does not compensate for variations in front-end gain or GMR sensitivity and for U_Q . The maximum deviation of the estimates obtained using the adaptive phase demodulation corresponds to about 300 beads.




Figure 54: Estimation of the magnetic crosstalk signal U_{mxl} vs. temperature using different estimation techniques. The calculated value at room temperature is $3.6545e-5$ V.

It can thus be concluded that the application of the vector decomposition algorithm greatly reduces the temperature dependency of the platform, as compared to other non-compensating demodulation schemes.

6. Digital Back-end Implementation and Verification

In this chapter the digital back-end is discussed that, as can be seen from Figure 18, consist of a digital signal processing (DSP) unit and store and display functionality. However, since the store and display functionality are non-specific for this project and of no importance for the robustness of the platform, we omit this functionality from the discussion here. In this chapter, digital back-end therefore refers to the DSP unit of the actual digital back-end. In Section 6.1 the architecture of the parts of the digital back-end that are most critical from a robustness perspective are discussed. Section 6.2 discusses simulation results for the critical parts of the digital back-end, while in Section 6.3 measurement results obtained with the platform are discussed.

6.1. Digital Back-end Architecture

Figure 55 depicts the architecture of the digital back-end, which was implemented in a Xilinx Spartan 3 field-programmable gate array (FPGA). The digital pre-amplifier (pre-amp) comprises a filter to further reduce the dynamic range of the analog pre-amplifier output signal and an amplifier to scale the signal to an appropriate input level for the adaptive joint phase and amplitude estimator (ϕ, A -estimator). Subsequently, the ϕ, A -estimator estimates the amplitude and phase of the magnetic signal. These estimates are used as input for the algorithm block. A waveform generator is employed to generate the excitation- and sense waveform, as well as a reference waveform that is supplied to the ϕ, A estimator. The algorithm block controls the waveform frequencies and the sensor selection. The sensor selection block in turn generates the signals to control the sensor selection switches on the analog front-end, based on the sensor select signal produced by the algorithm block. The algorithm block furthermore processes the estimated phase ϕ and amplitude A in order to estimate the surface bead density B according to the methods described in Section 5.1. Finally, the DSP unit interfaces via a USB interface to the store and display functionality of the digital back-end.

When algorithm 1 is used, i.e. the reference sensor scheme, the frequencies of the produced waveforms are constant, since both sensors are excited at the same sense-and excitation frequency. The sensor select signal will, on the other hand, be such that the sensor selection block alternately selects the reference sensor and the measurement sensor. When algorithm 2 is used, i.e. the vector decomposition scheme, the sensor select signal is constant since only one sensor is used in this scheme, while the ‘frequencies’ signal will be such that the waveform generator alternately produces two pairs of waveforms with differing frequencies.

The principles of both algorithms were already discussed in Section 5.1. Moreover, their implementation in the Very High Speed Integrated Circuit Hardware Description Language (VHDL) is fairly straightforward, therefore the algorithm block is not discussed in more depth here. The same applies for the sensor selection block. Furthermore, the USB interface is not interesting in the context of this work and will thus not be treated here either. The remaining blocks will be discussed in more detail in the following subsections. Subsection 6.1.1 discusses the architecture of the waveform

generator, while Subsection 6.1.2 discusses the ϕ, A -estimator. Finally, Subsection 6.1.3 discusses the digital pre-amplifier.

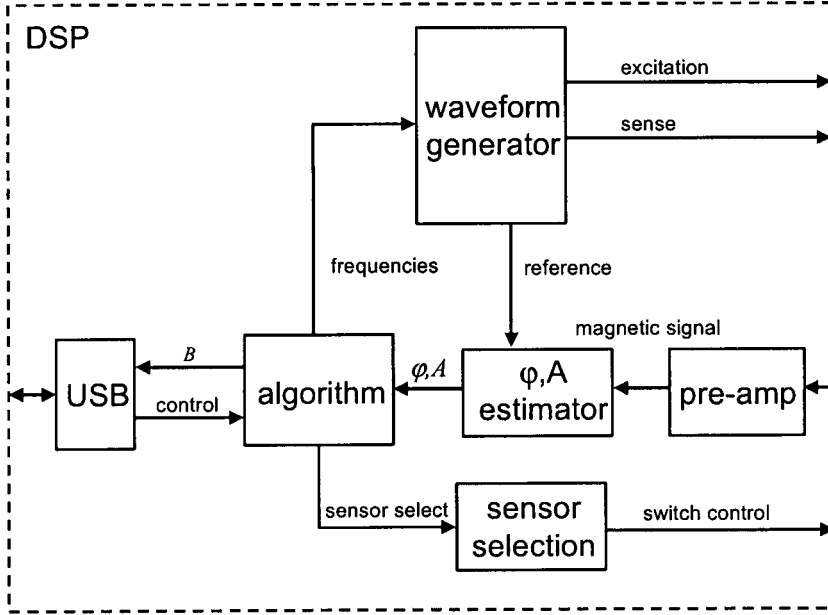


Figure 55: Digital back-end architecture, implemented in an FPGA.

6.1.1. Waveform Generator

The function of the waveform generator is to produce a sense-, excitation- and reference waveform, based on frequency information received from the algorithm block. As stated in Subsection 5.1.2, it is advantageous if the frequency of the magnetic signal can be preserved while alternating between excitation frequencies, since this guarantees an equal transfer of the magnetic signal by the analog and digital pre-amplifiers in both excitation cases. This can be achieved by first generating the sense and reference waveforms, and subsequently employing a single-sideband modulator to generate the excitation waveform. The reference signal, that is used as reference input for the ϕ, A -estimator, has the same frequency as the magnetic signal. Figure 56 schematically depicts a time continuous representation of a single-sideband modulator [19]. It uses the sense waveform $\sin(\omega_s t)$ and the reference waveform $\sin(\omega_r t)$ to produce the excitation waveform $\sin(\omega_e t)$. It can be shown that the frequency of the excitation waveform ω_e is given by

$$\omega_e = \omega_s + \omega_r \quad (66)$$

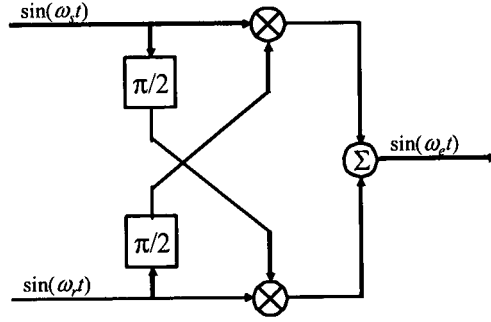


Figure 56: Single-sideband modulator.

Figure 57 depicts the complete waveform generator. It is implemented in Matlab Simulink combined with Xilinx System Generator. System Generator adds a blockset to the Simulink library browser that holds a large variety of Simulink models provided by Xilinx. A system that is built using these models can first be simulated in the rich simulation environment of Simulink and subsequently converted into VHDL that is optimised for Xilinx devices. The generated VHDL can finally, along with other digital back-end VHDL code, be used to generate a program file to be loaded in the FPGA. The FPGA should then, according to Xilinx, behave “bit true and cycle true” to the simulation results obtained in Simulink.

The waveform generator comprises two pairs of read-only memories (ROMs). One pair (ROM, ROM1) stores the excitation waveform and a 90 degrees phase shifted version of the excitation waveform, the other pair (ROM2, ROM3) comprises the reference waveform and its 90 degrees shifted equivalent. All ROMs are addressed by the output of one of the two accumulators (Accumulator1, Accumulator2). The accumulator outputs are set to wrap upon overflow, i.e. once the maximum value is reached that can be represented by the word length of the output, which corresponds to the address length of the ROMs, the accumulator will reset the output to zero and continue accumulating. The value on the input of the accumulators is a constant that can be adjusted by the algorithm block. The ‘frequencies’ signal in Figure 55 comprises solely these two constants. The first constant (f_sense) controls the frequency of the sense waveform, as it effectively determines the difference between the addresses of two consecutively addressed ROM entries. In other words, each clock cycle the ROM is addressed, and the increment in the address that is made each clock cycle equals constant f_sense . It can be easily seen that a larger value for f_sense results in a larger sense frequency f_s . In a similar way, the second constant (f_demod) determines the frequency of the reference waveform. Since the depth of the ROMs is 4096 entries, the generated sense frequency f_s is given by

$$f_s = f_{CLK} \frac{f_sense}{4096}$$

in which f_{CLK} is the clock frequency of the FPGA, which equals 40 MHz. Setting f_sense to 128 thus results in a sense frequency of 1.25 MHz, while an f_sense of 256 results in a 2.5 MHz sense frequency. Similarly, by choosing f_demod to be 2, we obtain a reference frequency f_r of 19.53 kHz.

The 10-b wide outputs of the ROMs are subsequently pair-wise multiplied (Mult, Mult1) and the result added (Add), according to the principle shown in Figure 56, in order to generate the excitation waveform. Subsequently, a shift right operation (Shift) is required to reduce the word length of the now obtained excitation signal to 10 bit. Finally, a constant (Constant2) of 512 is added (Add1) in order to convert the signed excitation waveform to an unsigned signal that can be applied to the excitation DAC inputs. A similar type conversion (Constant1, Add2) is performed on the sense waveform before it is applied to the sense DAC inputs.

Note, finally, that the employment of ROMs to store the waveforms also allow the application of waveforms other then sinusoids, adding to the flexibility of the platform.

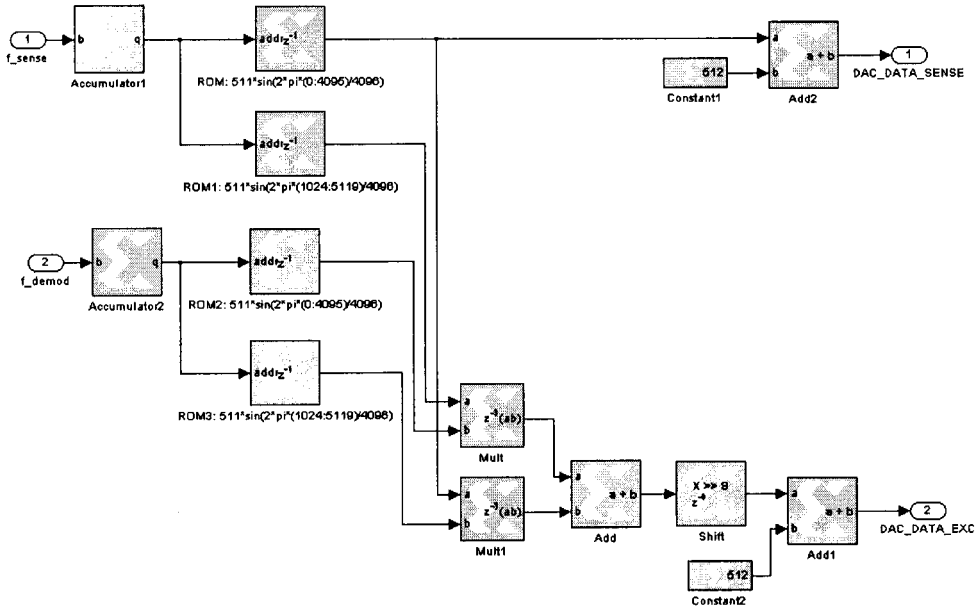


Figure 57: Implementation of the waveform generator.

6.1.2. Joint Phase and Amplitude Estimator

Figure 58 depicts the implementation of the joint phase and amplitude estimator of which a schematic representation is shown in Figure 46. The upper loop corresponds to the phase estimation loop, while the lower loop corresponds to the amplitude estimation loop. In the phase estimation loop, the error signal is multiplied (Mult1) by the phase signature, which is stored in a ROM (ROM2). The resulting signal, i.e. the stochastic gradient is multiplied by the adaptation constant μ_φ (u_p). This multiplication is implemented as a shift left. This fixes μ_φ to powers of 2, but is costless in terms of hardware. The result of the shift left operation is subsequently used as input for a discrete integrator, i.e. an accumulator (Accumulator3). The output of this accumulator represents the phase estimate. The value of this estimate is finally added (Add) to the address input of the ROM, which is furthermore updated each clock cycle in a similar way as already explained for the waveform generator of Subsection 6.1.1, i.e. the constant f_{demod} ,

representing the frequency of the reference signal, is applied to an accumulator that wraps upon overflow. The output of this accumulator is used as a pointer to a ROM address, while the phase estimate gives an offset to this address.

The same pointer that is used to address the ROM of the phase estimation loop is used to address a ROM in the amplitude estimation loop (ROM1). This ROM contains the amplitude signature, which is 90 degrees shifted in phase with respect to the phase signature. The output of this ROM is multiplied (Mult2) by the error signal in order to obtain the stochastic gradient of the amplitude estimation loop. This stochastic gradient is subsequently multiplied by the adaptation constant μ_A (u_a) and supplied to an accumulator (Accumulator2). The output of the accumulator, which represents the amplitude estimate, is used to control the amplitude of the adapted reference signal that is to be subtracted (Subtract) from the input signal (filtered ADC_DATA) in order to obtain the error signal. The amplitude of the reference signal is, before multiplication (Mult4) by the amplitude estimate, reduced to 0.5 as a result of a 14 bit shift right operation (Shift2) on the output of ROM1. The amplitude estimate A' therefore satisfies

$$A' = 2A_p \quad (67)$$

Where A_p is the peak amplitude of the input signal.

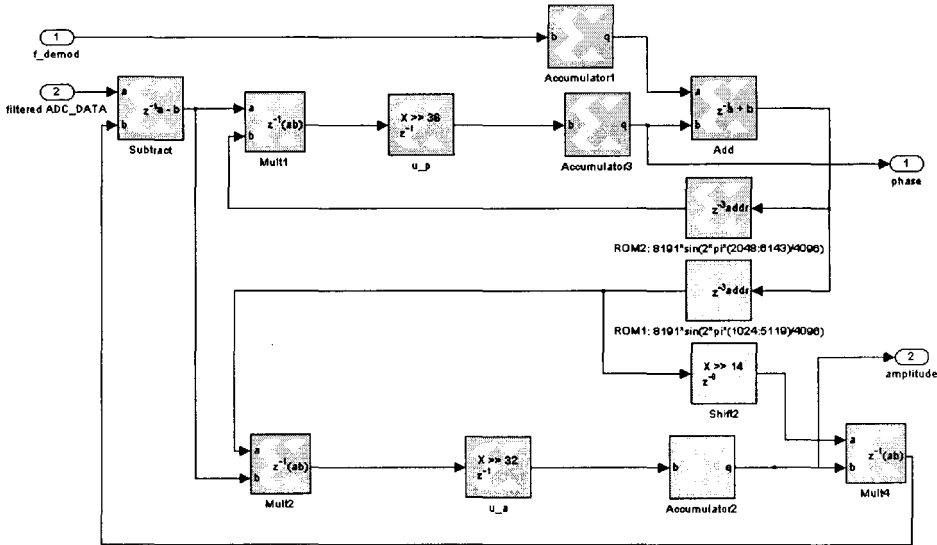


Figure 58: Implementation of the joint phase and amplitude estimator.

Changes in amplitude A and phase θ of the input signal are the result of temperature variations and changes in the number of beads on the sensor surface, which are both low frequency processes with a bandwidth that is assumed to be smaller than 0.5 Hz. The tracing speed should thus be such that parameter variations of 0.5 Hz can be tracked. Moreover, to obtain information about temperature induced changes in GMR sensitivity and the transfer of the analog front-end, as explained in the algorithm descriptions of Section 5.1, either the excitation frequency or the sensor selection is modulated with a frequency of about 0.5 Hz. The speed of convergence should therefore be such that

sufficiently accurate estimates are obtained within less than 1 s. The latter constraint determines the required bandwidth of the loops. The adaptation constant should be chosen as small as possible to minimise gradient noise, with the constraint that convergence should occur within 1 s. The estimator furthermore runs at the clock frequency of the FPGA, which is 40 MHz. In Subsection 5.2.8 the convergence criterion, $\Delta < 3\sigma$ was found, i.e. the loops are assumed to have sufficiently converged if the residual parameter mismatch is smaller than 3 times the gradient noise standard deviation. This criterion will be used to obtain a suitable value for the adaptation constant μ . To that end we first derive an expression for the noise variance, specifically for the implemented amplitude estimation loop.

The noise in the input signal is not white up to the Nyquist frequency, but band limited by the 10th order filter in the front-end model, which has a bandwidth of 100 kHz, i.e. $\omega_f = \omega_{nyq}/200$. Using Equation 59, the normalised equivalent noise bandwidth of the loop now satisfies

$$B = 100\mu\sigma_s^2 \quad (68)$$

In order to obtain an expression for the amplitude estimate gradient noise it is informative to redraw the equivalent representations of the general adaptation loop of Figure 49 here and more specifically for the implemented amplitude adaptation loop of Figure 58. The equivalent loops are shown in Figure 59. In the lower part of Figure 59 it can be seen that $z_k = n_k/0.5s_k$, such that in this case $\sigma_z^2 = \sigma_n^2/0.25$ $\sigma_s^2 = \sigma_n^2/0.125$.




Figure 59: Equivalent representations of the implemented amplitude adaptation loop.

Using Equation 58 and 68 and substituting the general estimate φ by the amplitude estimate A' , we obtain for the noise variance of the amplitude estimate

$$\sigma_{A'}^2 = 400\mu\sigma_n^2 \quad (69)$$

The noise power at the input of the estimator σ_n^2 can be found by multiplication of the GMR noise power spectral density, referred to the output of the pre-amplifier, and the bandwidth of the analog pre-amplifier filters and amounts to $(5.5\text{nV}/\sqrt{\text{Hz}} \cdot 10\text{e}4)^2 \cdot 100\text{e}3 = 302\text{e-}6 \text{ V}^2$. In order to determine the estimated gradient noise $\sigma_{A'}^2$, a suitable value for μ should furthermore be chosen, making use of the convergence criterion discussed in Subsection 5.2.8. The largest possible, and thus most critical, value for A is equal to the ADC input range, which is 1 V_p . Using Equation 57, which may for small values for μ be approximated by $\tau \approx T_s / \mu \sigma_s^2$, Equation 65, where $A_k' = A(1 - e^{kT_s/\tau})$ and $kT_s = t$, and finally Equation 67, the convergence criterion $\Delta_{A'} < 3\sigma_{A'}$ can be written as

$$2Ae^{\frac{-\mu\sigma_s^2}{T_s}} < 3\sqrt{400\mu\sigma_n^2} \quad (70)$$

A value for μ can now be found such that the criterion is met within one second.




Figure 60: $\Delta_{A'}$ and $3\sigma_{A'}$ vs. μ .

A value of $0.95\text{e-}6$ was chosen for the adaptation constant, corresponding to a μ_A in Figure 59 equal to 2^{-32} . The parameter gradient noise variance $\sigma_{A'}^2$ then amounts to $115\text{e-}9 \text{ V}^2$, such that the standard deviation $\sigma_{A'}$ equals $339 \text{ }\mu\text{V}$. GMR referred this amounts to 17 nV . The GMR referred standard deviation can, given that the signal per bead is about 5 nV , be used to calculate the detection limit of the platform. The detection limit of the platform, which can be calculated using Equation 7, amounts to $L_D = 2.71 \cdot 17\text{nV} = 46 \text{ nV}$ which roughly corresponds to 10 surface bound beads. The time constant of both loops can be chosen equal, since they are fully orthogonal and behave equally in terms of noise performance and loop time constant, such that the overall adaptation speed remains in the same order.

The word length for the estimates should be such that gradient noise is dominant over the quantisation noise due to the finite word length of the estimate. In other words, the quantisation step should be smaller than the standard deviation of the estimate. Since the input range of the ADC is 1 V_p , the maximum value for A_{max}' is 2 V . The minimum

required resolution then amounts to $A_{max}' / \sigma_A = 5900$ q or 13 bits. A resolution of 14 bits was chosen. The word length for the phase estimate is given by depth of the ROMs, which is 4096, i.e. 12 bits. This particular depth was chosen, since it is the smallest depth that still allows the frequency of the magnetic signal to be altered in step of 10 kHz. In other words, the smallest frequency that can be produced is 10 kHz, corresponding to an increment of the ROM address in address of one each clock cycle, given the clock frequency of 40 MHz.

6.1.3. Digital Pre-amplifier

To minimise temperature induced phase and amplitude variations in the output signal of the analog pre-amplifier, it was decided to use the single stage pre-amplifier shown in Figure 26, rather than the dual-stage variant shown in Figure 27. This means that additional filtering and amplification is required in the digital domain before the magnetic signal can be applied to the ϕ, A -estimator. To this end, a digital pre-amplifier was implemented that comprises a filter and an amplifier. The filter is employed to further reduce the dynamic range of the output signal of the analog pre-amplifier, more specifically the sense signal and the capacitive and inductive crosstalk signal. The amplifier is used to obtain a sufficiently large input signal for the joint phase and amplitude estimator.

The filter consists of two filter stages in series. Since specific spectral components need to be attenuated, the filter stages can simply be implemented as comb filters, of which the structure is shown in Figure 61 [20].




Figure 61: Structure of the comb filter

The transfer function of the filter is given by

$$H(e^{j\omega}) = \frac{y(e^{j\omega})}{x(e^{j\omega})} = 1 + e^{-j\omega K} = 1 + \cos(\omega K) - j \sin(\omega K) \quad (71)$$

Where $e^{j\omega}$ equals z and ω is the normalised frequency $2\pi f / f_{CLK}$. Here f_{CLK} is the clock frequency of the FPGA. The magnitude response satisfies

$$|H(e^{j\omega})| = \sqrt{2 + 2 \cos(\omega K)} \quad (72)$$

This magnitude response shows notches at $\omega K = 2\pi \frac{f}{f_{CLK}} K = \pi(1 \pm 2i)$, $i \in N$ and since the clock frequency f_{CLK} is 40 MHz, the frequencies of the notches are given by

$$f = \frac{1 \pm 2i}{2K} f_s, \quad i \in N \quad (73)$$

The first stage has its first notch at the frequency of the sense signal, i.e. 1.25 MHz, such that $K = 16$. Since the excitation frequency differs only 20 kHz from the sense frequency, the capacitive and inductive crosstalk signal is also sufficiently attenuated by this filter stage. A second filter stage was required, since the vector decomposition algorithm excites the sensor alternatingly at two different frequencies. The second sense frequency that is used is 2.5 Mhz, which is not attenuated by the first filter. A second filter stage was therefore applied with $K = 8$, resulting in a notch at 2.5 MHz.

Finally, the output signal of the filter is amplified by a factor 16 by shifting 4 bits to the left. Since the gain of the comb filters is approximately 2 for the low frequency of the magnetic signal, the total gain of the digital pre-amplifier amounts to 64.

6.2. Digital Back-end Simulation Results

In order to test their performance prior to the actual hardware implementation, the waveform generator and joint amplitude and phase estimator, discussed in Section 6.1, were simulated in Matlab Simulink. Subsection 6.2.1 reports on the simulation of the waveform generator, while Subsection 6.2.2 discusses the results of simulation of the estimator.

6.2.1. Waveform Generator

The waveform generator was simulated in order to test whether it produces the required frequencies. Furthermore, it was evaluated whether the spurious components, i.e. the generated spectral components, other than those intended, are sufficiently low. Finally, the amplitude of the generated waveforms was compared for different frequencies, to check that the amplitude of the generated waveforms is independent of their frequency.

Figure 62 shows the simulated spectrum of the sense current. The sense current is 20 mA_{pp}, which is equivalent to -43.0 dBI_{RMS}. The frequency f_s is 1.25 MHz. Since the increment in ROM address per clock cycle (128) was chosen to be an integer dividend of the ROM depth (4096), or equivalently, the sense frequency is an integer dividend of the clock frequency, no subharmonics are produced and the spectrum consists of a single component at the sense frequency in the frequency range up to 2 MHz. This is important since components at the frequency of the magnetic signal that are already present in the spectrum of the sense current may, upon multiplication by the relatively large R_θ of the GMR, result in significant voltage across the GMR. This non-magnetic signal can not be discriminated from the magnetic signal generated in the GMR.




Figure 62: Simulated spectrum of sense current for $f_s = 1.25$ MHz.

Figure 63 shows the simulated spectrum of the excitation current. The excitation current is 160 mA_{pp}, which is equivalent to -24.6 dBI_{RMS}. The frequency f_e is about 1.27 MHz.




Figure 63: Simulated spectrum of excitation current for $f_s = 1.25$ MHz.

Contrary to the sense frequency, the excitation frequency is not an integer dividend of the clock frequency, resulting in the generation of subharmonics. It can, however, be seen that these subharmonics are at least 80 dB smaller than the fundamental harmonic. The GMR resistance variations, that correspond to these subharmonics will upon multiplication by the sense current not appear at the frequency of the magnetic signal and, moreover, be smaller than the noise power for practical measurement bandwidths. This can be seen from Figure 64 and Figure 65, which show the simulated spectrum at the output of the combined sensor and front-end model, in the presence and absence of noise, respectively, upon excitation with the above shown excitation current and a sense current of 5 mA_{pp}. The resolution bandwidth used was 610 Hz, therefore the noise shown in Figure 64 would occur about 25 dB lower if a bandwidth of 1 Hz would be used. Even

in the latter case, the noise is still larger than the spectral components due to of the subharmonics in the spectrum of the excitation current shown in Figure 65.




Figure 64: Simulated spectrum at the output of the sensor and front-end model $f_s=1.25$ MHz, resolution bandwidth = 610 Hz.




Figure 65: Simulated spectrum at of the sensor and front-end model $f_s=1.25$ MHz, without noise.

Similar results are obtained if the sense frequency is chosen 2.5 MHz, resulting in an excitation frequency of about 2.52 MHz. The amplitude of the waveforms that are generated by the waveform generator were furthermore compared for different frequencies by recording a large number of periods and having Matlab calculate the variance of the recorded signals. From this it was found that the amplitude varies about 100 ppm across the frequency range of interest, which results in amplitude variations in the magnetic signal that are smaller than those caused by the GMR thermal noise.

6.2.2. Joint Phase and Amplitude Estimator

The performance of the joint phase and amplitude estimator was tested with the simulation set-up shown in Figure 52. During the simulation the temperature (T) was kept at a constant value of 20 °C. The spectrum of the magnetic signal at the output of the combined sensor and front-end model, which in this case includes the digital pre-amplifier, is shown in Figure 66. The amplitude of the magnetic signal, which, as can be seen from Figure 66, is about -10 dB, was calculated more accurately with a noise free signal recording using Matlab. The amplitude was found to be $37.449 \mu\text{V}_{\text{RMS}}$ across the GMR. The noise power of the recorded signal σ_n^2 was determined to be $302\text{e-}6 \text{ V}^2$. Since 10^{th} order low-pass filtering was used, this power is confined to the bandwidth of the pre-amplifier filters (100 kHz). The noise voltage spectral density at the output of the pre-amplifier thus amounts to $55 \mu\text{V}/\sqrt{\text{Hz}}$. By dividing this by the gain of the pre-amplifier (80 dB), we obtain, as expected, a GMR referred noise voltage spectral density of $5.5 \text{ nV}/\sqrt{\text{Hz}}$.




Figure 66: Spectrum of recorded magnetic signal. Resolution bandwidth = 610 Hz.

Figure 67 depicts the results of the simulation. The upper part of Figure 67 depicts the convergence for phase and amplitude, while the lower part shows the amplitude and phase after convergence on a smaller scale. The amplitude estimates are GMR referred. From Figure 67 it can be concluded that the estimator converges well within the specified convergence time, i.e. 1 second. Using Matlab, the noise variance in the amplitude estimation was determined to be $63\text{e-}9 \text{ V}^2$, about half of the value for σ_A^2 obtained in Subsection 6.1.2. The difference can be attributed to the fact that only one simulation was performed for a 2 second period, since more or longer simulations would result in unpractically long simulation times. The number of samples that were used to determine the variance was therefore limited. The GMR referred noise standard deviation σ_A equals 12 nV. The variance of the phase is smaller than the phase resolution given by the depth of the ROMs (10b) that store the reference signal, which is $0.35 \text{ }^\circ/\text{q}$. The estimated amplitude is $37.443 \mu\text{V}_{\text{RMS}}$. Comparison with the amplitude of the input signal reveals a bias of 6 nV, which is lower than standard deviation of 12 nV. In general, a bias always

exists since we simulate a single realisation of the noise process. Moreover, in a real system a bias will occur as well, since only one realisation of the noise is used. However, the bias should be comparable to the standard deviation, if a sufficient number of samples is taken into account.




Figure 67: Phase estimation results (left) and GMR referred amplitude estimation results (right).

In order to confirm the relation between the adaptation constant and the noise variance of the estimates, i.e. Equation 69, more simulations were performed with estimators with different adaptation constants. Figure 68 shows the results of simulations for amplitude estimators with different adaptation constants, where the left, the middle, and the right pair of graphs each correspond to a different adaptation constant. The adaptation constants used were from left to right 488e-6, 122e-6 and 30.5e-6, i.e. a reduction by a factor 4 each time from left to right. It can be seen that for smaller adaptation constants the time constant increases and the standard deviation of the estimates reduces. These simulations were repeated for different noise realisations in order to obtain an average mean and standard deviation for each time constant used.

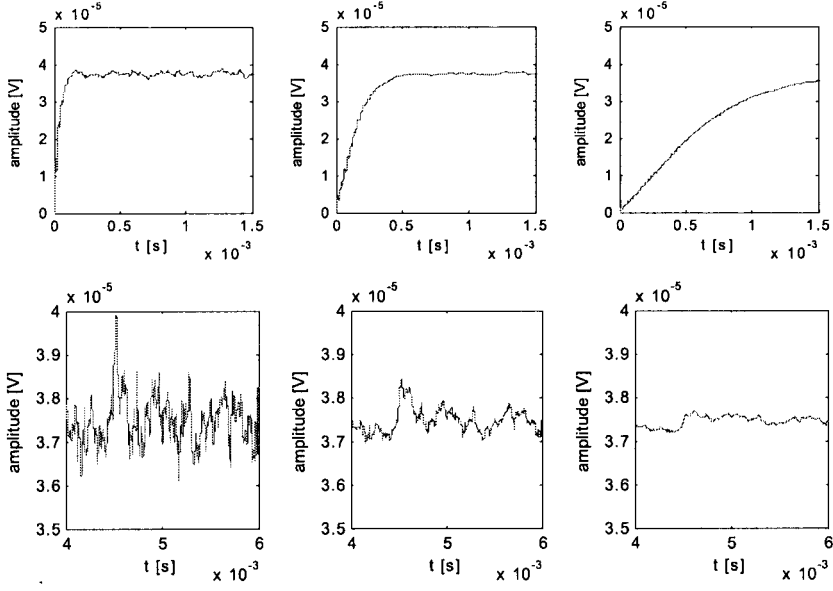


Figure 68: Amplitude estimation results for different time constants of the loops. The top figures show the convergence of the loops, while the bottom figures show the estimates when convergence has occurred.

Figure 69 shows the simulated noise variance and the noise variance calculated using Equation 69 as a function of μ . It can be concluded that the simulations confirm Equation 69.




Figure 69: GMR referred amplitude estimate variance vs. the adaptation constant.

6.3. Digital Back-end Measurement Results

The Simulink models discussed in Section 6.2 were converted to VHDL that was subsequently compiled along with other parts of the digital back-end and loaded in the FPGA, such that heating experiments could be performed in order to test the robustness of the platform. The results of these experiments are discussed in Chapter 7. However, before commencing with the heating experiment, the waveform generator, joint phase and amplitude estimator and digital pre-amplifier were tested. The waveform generator was tested by measuring the generated signals at specific points on the analog front-end. The results of these measurement are discussed in Subsection 6.3.1. The performance of the joint phase and amplitude estimator was measured in terms of estimate standard deviation. This measurement is discussed in Subsection 6.3.2. Finally, the effect of the comb filters of the digital back-end was measured. These measurement are discussed in Subsection 6.3.3.

6.3.1. Waveform Generator

In order to test the excitation waveform that is generated by the actual hardware implementation of the waveform generator, the spectrum of the excitation current was measured at the output of the DAC, i.e. across R_i in Figure 20 using a differential 1:10 probe. The measured spectrum for a sense current of maximum amplitude, i.e. 20 mA_{pp} DAC output current and a frequency of 1.25 MHz is shown in Figure 70. Except for a small interference of which the origin is unclear and that will appear in many of the spectra shown in the following, no additional components are present in the spectrum, other than the 1.27 MHz and its second harmonic. The measured amplitude of the fundamental corresponds well with the expected value, i.e. 29 dBV, of which 20 dB results from the 1:10 probe. It can furthermore be seen that any spurious components are at least 80 dB smaller than the ground harmonic at 1.27 MHz. It should be noted here that in all following spectrum plots, the scale per y-division is 10 dB and the y-value for the top of the plot is shown just above the plot, following the indication REF, i.e. -23 dBV in Figure 70. In order to compare the voltage spectrum of the excitation waveform shown in Figure 70 with the simulated current spectrum of Figure 63, the attenuation of the probe (20 dB) should be subtracted and the gain of the V/I conversion (-15.5 dB) added to the spectrum in Figure 70. By doing so we indeed obtain the -24.6 dB excitation signal shown in Figure 63.




Figure 70: Spectrum at the output of the excitation DAC, before conversion to current. $f_s = 1.25$ MHz.

The spectrum of the sense current could for practical reasons most easily be obtained by measuring the spectrum across the GMR. This means that, apart from the sense signal, also the capacitive and inductive crosstalk signal appears in the spectrum, namely at a frequency 20 kHz higher than the sense signal. Figure 71 depicts the spectrum of the voltage across the GMR. The sense signal results from the multiplication of the sense current and the R_0 of the GMR. Since an old type sensor was used, which had an R_0 of 520 Ω , compared to 150 Ω for the new type sensors, the sense current was reduced in order not to saturate the pre-amplifier. In this case a sense current of 1.25 mA_{pp} was used. In order to compare the voltage spectrum of the sense waveform shown in Figure 71 to the simulated sense current spectrum of Figure 62, the attenuation of the probe (20 dB) should be subtracted from the sense signal in Figure 71 and the result divided by R_0 , such that the sense current amplitude is obtained. Since for Figure 62 a sense current of 20 mA_{pp} was used, the obtained sense current should be multiplied by 16 in order to find the value of -43 dB_{I_{RMS}} shown in the spectrum of Figure 62. It can furthermore be seen that eventual possible spurious components are at least 80 dB smaller than the sense signal.




Figure 71: Spectrum of voltage across the GMR, $f_s = 1.25$ MHz.

It can be concluded that the waveform generator performs as expected. For completeness, Figure 72 depicts the spectrum at the output of the analog front-end.




Figure 72: Spectrum at output of the analog pre-amplifier, $f_s = 1.25$ MHz.

6.3.2. Joint Phase and Amplitude Estimator

In order to verify the simulated value, the GMR referred standard deviation of the amplitude estimate of a sensor signal was measured. To that end, a single sensor was excited and the estimated amplitude of its signal was recorded for a period of 10 minutes.

The recorded signal is depicted in Figure 73. The amplitude estimate variations in this measurement are caused by gradient noise, as well as temperature induced sensitivity changes that even occur in a laboratory setting. Since the temperature variations are in the extreme low frequency range, a subset of the recorded signal was selected having the smallest low frequency variations. By doing so we eliminate the temperature induced variations from the measurement, such that the residual variations can only be attributed to gradient noise.




Figure 73: Sensor amplitude estimate, GMR referred.

The chosen subset of the data shown in Figure 73 is shown in Figure 74 and consists of roughly the points between 515 and 577 seconds. The standard deviation of estimate was determined using Matlab and was found to be 16 nV. It can thus be concluded that the measured GMR referred standard deviation matches the simulated value of 17 nV that was obtained in Subsection 6.2.2.




Figure 74: Subset of data shown in Figure 73 that was used to determine standard deviation.

6.3.3. Digital Pre-amplifier

In order to test if the filter in the digital pre-amplifier sufficiently suppresses the sense signal, as well as the capacitive crosstalk signal, the spectra of the magnetic signal before and after filtering were compared. Specifically for the purpose of examining FPGA internal signals, a 10-b test DAC followed by an operational amplifier was added to the design of the platform. The transfer from the input of the test DAC to the input of the spectrum analyser was measured to be 0.915 mV/q.

Figure 75 depicts the spectrum of the input signal of the digital pre-amplifier. To this end, the 10 most significant bits (MSB) of the output signal of the analog pre-amplifier were, via the FPGA, routed to the input of the test DAC. It should be noted here that apart from the expected spectral components a number of unexpected components appear in the spectrum. These can be attributed to crosstalk from the test DAC to the analog pre-amplifier. In this measurement we are, however, only interested in the effect of the digital pre-amplifier filter on the spectrum. During actual measurements, no signal will be applied to the test DAC. It can be seen that the sense signal, in this case at 1.25 MHz is about 15 dB larger than the magnetic signal at 20 kHz. To prevent the signal from clipping upon amplification in the digital pre-amplifier, this sense signal needs to be further attenuated, such that it is smaller than the magnetic signal.




Figure 75: Spectrum of the signal from the test DAC when the ADC output data is looped back via the FPGA to the test DAC. $f_s = 1.25$ MHz.

Figure 76 depicts the signal from the test DAC when the 10 MSBs of the output of the first filter stage are applied to the inputs of the test DAC. The appearance of a notch at 1.25 MHz can be clearly observed. The sense signal is attenuated by 50 dB, which suffices. The capacitive and inductive crosstalk signal, with a frequency of 1.27 MHz, is attenuated by approximately the same amount.

The second stage of the digital pre-amplifier filter is employed to attenuate the sense signal when a sense frequency of 2.5 MHz is used, since the first filter stage does not attenuate at that frequency, but in fact amplifies 6 dB at that frequency. Figure 77 depicts the spectrum of the signal at the output of the second filter stage. A notch can now be observed at 2.5 MHz. This notch attenuates the sense signal upon excitation of the sensor with a sense frequency of 2.5 MHz. Note that the filter indeed amplifies the magnetic signal at 20 kHz by about 6 dB per filter stage.

From these measurements, it can be concluded that the digital pre-amplifier filter shows notches at the intended frequencies and is capable of sufficiently attenuating the sense and capacitive and inductive crosstalk signals.




Figure 76: Spectrum of the signal from the test DAC when the output of the first filter stage is applied to the test DAC. $f_s = 1.25$ MHz.




Figure 77: Spectrum of the signal from the test DAC when the output of the second filter stage is applied to the test DAC. $f_s = 1.25$ MHz.

7. Experimental Results

In this chapter the experimental results are discussed of heating experiments that were conducted in order to test the performance of the platform in combination with the reference sensor detection algorithm. Section 7.1 discusses the type of experiments that were conducted, as well as the methods used. In Section 7.2 the results of these experiments are discussed. The vector decomposition algorithm did not result in any usable data since the signal obtained when the sensor was excited with a sense frequency of 2.5 MHz appeared to be smaller than the signal obtained when the sensor was excited at a sense frequency of 1.25 MHz. This is opposite to what is expected, since the capacitive crosstalk component in the magnetic signal should increase with frequency. Apparently, signal components other than those included in the signal model of Section 4.3 exist at the frequency of the magnetic signal and hinder the proper working of the algorithm. The exact cause of this is subject of further research.

7.1. Measurement Methods

Two types of experiments were conducted in order to test the reference sensor algorithm. In the first type of experiments, which are discussed in Subsection 7.2.1, only the sensor is heated. This experiment will verify our assumption that the sensitivity of two adjacent GMRs on a sensor chip changes comparably as a result of temperature variations. Furthermore, since the platform itself remains at a constant temperature such that its gain in the analog front-end is constant as well, this experiment can be used to measure the temperature coefficient of the sensitivity. The measured value can be compared to the value used in the model, in order to verify this model parameter. The sensor is heated by means of a soldering iron that is fixed in close proximity, i.e. about 3 mm, of the sensor chip, and heated up to 270 °C. subsequently the soldering iron is removed to allow the chip to release its heat to the environment. Using this method, the temperature of the chip can be increased and reduced relatively fast and in a controlled way. During the heating of the chip, a thermocouple, which is in direct contact with the silicon substrate of the chip, measures the temperature of the sensor. Since the material of the substrate is silicon, which conducts heat well, the GMRs on the sensor chip may be assumed to have the same temperature as the substrate. Finally, the amplitude and phase of the sensors are recorded, alternatigly for the measurement- and reference sensor. These recordings were finally used as input for the algorithm. During the measurements no beads were added to the sensor, i.e. the magnetic signal consists solely of magnetic crosstalk, which is constant during the measurement. The desired output of the platform is constant during all experiments discussed here.

In the second type of experiment the complete platform, including the sensor, was heated in an oven. In this way the reference sensor algorithm could be tested for the complete platform. During this measurement the thermocouple used to measure the temperature was suspended in mid-air in the middle of the oven.

The switching frequency is, for reasons of easy implementation, chosen to be 0.31 Hz, i.e. a different sensor is selected each 1.6 seconds. Figure 78 shows the amplitude

estimate, measured during a typical measurement. It can be seen that estimated amplitude differs slightly for the two sensors used. This is the result of the sensitivity of the two sensors not being exactly equal. This difference in sensitivity is of no importance to the performance of the detection algorithms. Only the relative change in sensitivity as a function of temperature should be comparable. It can furthermore be seen that during switching the amplitude estimate reduces drastically. This results from the fact that during the switching temporarily no GMR is selected, such that the output of the sense current source, i.e. U_{GMR} in Figure 21, saturates. The magnetic signal is therefore zero during switching. Just $2.5 \mu s$ before switching between sensors, the amplitude- and phase estimate produced by the estimator are stored. The estimator, which was designed to converge starting from zero within 1 second, thus has sufficient time to converge.




Figure 78: Typical amplitude estimate output generated by the estimator during experiments where every 1.6 s a different sensor is selected.

For completeness, Figure 79 depicts an example of the phase estimate output of the estimator during a typical measurement. It can be seen that the phase estimate converges sufficiently fast.




Figure 79: Typical phase estimate output generated by the estimator during experiments where every 1.6 s a different sensor is selected.

7.2. Heating Experiment Measurement Results

In this section the result are discussed of the two types of heating experiments described in Section 7.1. In Subsection 7.2.1 the results are shown of experiments where only the sensor was heated, while in Subsection 7.2.2 the heating experiments of the complete platform are discussed.

7.2.1. Heating of the Sensor

Figure 80 depicts a typical result of a heating experiment in which only the sensor was heated. The top graph shows the expected amplitude and the amplitude measured with the platform, while the bottom graph shown the temperature as a function of time. The straight line in the top graph is the GMR referred amplitude that was measured at the start of the experiment, i.e. at $t = 0$ s. It can be seen that the uncompensated sensor signal deviates considerably as a result of a change in sensor temperature. The signal that is obtained using the reference sensor algorithm is much less dependent on temperature. The algorithm thus compensates for a large fraction of the temperature induced variations




Figure 80: Typical sensor output uncompensated and compensated using the reference sensor algorithm vs. time (top) and temperature vs. time (bottom).

Figure 81 depicts the measured amplitude of the magnetic signal as a function of temperature. In the now following discussion, it was assumed that the measurement started after the sensor was heated to maximum temperature and started to cool down, i.e. t_0 was set to 40 s. This starting point was chosen because the cooling process is more gradual than the heating process. Since the sensor is heated very fast the algorithm was not able to accurately compensate for the fast temperature increase. In other words, the bandwidth, or more specifically the slew rate of the temperature variations during the heating process is beyond that of the algorithm. It can be seen from Figure 81 that the uncompensated signal depends linearly on temperature and that the temperature coefficient of the GMR, which causes this temperature dependency, is $-0.16 \text{ } \%/^{\circ}\text{C}$. This corresponds well to the value of $-0.18 \text{ } \%/^{\circ}\text{C}$ that was used in the model of Chapter 4. It can furthermore be seen that the signal after compensation is virtually independent of temperature.

The standard deviation of the error, i.e. difference between the compensated amplitude and the expected amplitude was found to be 18 nV. Given the fact that for the old type sensor the average signal per bead is about 80 dB smaller the magnetic crosstalk, which is measured here, the standard deviation of the error corresponds to about 11 surface bound beads. The maximum error that was observed is 50 nV, which corresponds to about 29 surface bound beads.




Figure 81: Measured amplitude vs. temperature.

Although not of any importance for the measurement of bead surface densities, since it is compensated for, it is informative to investigate the phase of the sensor signal as a result of temperature. Since phase variations are only assumed to result from heating of the filters in the analog front-end, no phase variations are expected during this measurement. From Figure 82, it can, however, be seen that a small and linear dependency of the phase on temperature exists. The observed phase shift is $-0.02 \text{ }^{\circ}\text{C}^{-1}$. This can possibly be attributed to thermal expansion of the sensor chip, resulting in an increased distance between the current conductors and the GMR. This reduces the capacitive coupling and therefore the amplitude of the quadrature signal U_Q .




Figure 82: Measured phase vs. temperature.

7.2.2. Heating of the Sensor and Platform

Figure 83 depicts a typical result of an experiment in which the complete platform, including the sensor, was heated in an oven. The temperature variations observed are considerably slower and smaller as a result of the use of the oven, compared to the experiments in which only the sensor was heated using a soldering iron. The algorithm should thus be able to easily track the temperature induced amplitude and phase variations. However, as a result of the way in which the temperature is measured, i.e. the thermocouple suspended in mid-air, the measured temperature is probably not exactly equal to the temperature of the platform and sensor. When the oven periodically switches on, temperature gradients occur, causing convection. The thermocouple responds faster to this than the platform. This discrepancy between the measured temperature and the platform temperature will also be observable in the graphs where amplitude and phase are plotted against temperature.




Figure 83: Typical sensor output uncompensated and compensated using the reference sensor algorithm vs. time (top) and temperature vs. time (bottom).

Figure 84 depicts the measured amplitude of the magnetic signal as a function of temperature. The temperature dependency of the uncompensated signal is linear and was found to be $-0.47 \text{ } \%/^{\circ}\text{C}$. The temperature dependency of the transfer of the analog front-end filters is found by subtracting the temperature dependency of the GMR sensitivity from this value, and amount to $0.31 \text{ } \%/^{\circ}\text{C}$.

The standard deviation of the error, i.e. difference between the compensated amplitude and the expected amplitude was found to be 59 nV. The standard deviation of the error corresponds to about 24 surface bound beads. The maximum error that was observed is 273 nV, which corresponds to about 109 beads. A residual temperature dependency of about $0.06 \text{ } \%/^{\circ}\text{C}$ can be observed in the compensated amplitude signal.

Figure 85 depicts the measured phase of the magnetic signal as a function of temperature. The observed phase shift temperature dependency is linear and amounts to $0.31 \text{ } ^{\circ}/^{\circ}\text{C}$. The phase shift temperature dependency is thus much larger and of different sign than the

phase shift temperature dependency of the sensor. By subtracting the phase shift temperature dependency of the sensor, which was found to be $-0.02\text{ }^{\circ}\text{C}$, we obtain a phase shift temperature dependency of $0.33\text{ }^{\circ}\text{C}$ for the analog front-end.




Figure 84: Measured amplitude vs. temperature.




Figure 85: Measured phase vs. temperature.

From the experimental results described in this and the previous subsections, it can be concluded that the temperature dependency of the platform output signal is negligible in the case that only the sensor is heated and reduced by a factor 7.8 when the complete platform is heated.

8. Summary, Conclusions and Recommendations

In this final chapter a summary is given in which conclusions are drawn about the work described in this report. This summary is discussed in Section 8.1. Furthermore, recommendations are given for future research on biosensors in Section 8.2.

8.1. Summary and Conclusions

Biosensors that employ GMR elements, combined with suitable immunoassays and magnetic labels can be used to measure the concentration of specific molecules, e.g. proteins or drugs-of-abuse, in samples of bodily fluids. The use of superparamagnetic particles as labels for molecular detection has some distinct advantages. Firstly, biological samples contain hardly any magnetic material; therefore the detection background is intrinsically low. Secondly, magnetic sensors for label detection can be integrated on-chip, enabling extreme miniaturization of diagnostic systems. Thirdly, magnetic actuation can be employed to speed up biochemical assays that are otherwise diffusion limited, by actively concentrating targets within a sensing zone.

A platform was developed for concentration measurements (prior art). A modulation and demodulation scheme is employed to separate crosstalk from the signal of interest and obtain optimal signal to noise ratio. In order to determine the sensitivity, linearity and limit of detection of the platform, the response of the detection platform to various numbers of beads on a sensor surface was experimentally determined by making use of an inhibition assay. It was found that the platform output depends linearly on surface bead density and shows a sensitivity of 5 nV/bead. It was furthermore found that a single bead is detectable with a platform having a bandwidth of 0.1 Hz.

The dependency of the platform output on environmental variations in e.g. temperature and magnetic field should preferably be small, such that environmentally induced platform output variations are smaller than those caused by the thermal noise of the GMR. The current platform lacks this robustness, since the sensitivity of the GMR is a function of temperature and externally applied magnetic fields. Moreover, the transfer of the analog front-end is a function of temperature. Finally, inherent capacitive and inductive crosstalk introduces errors in the measured signal. All these effects combined result in output signal variations that are two orders of magnitude larger than those caused by GMR thermal noise. The aim of this research was therefore to realise a more robust measurement platform. This is done by minimising the temperature dependency of the analog front-end. Moreover, robust detection algorithms were developed to increase robustness. To facilitate fast implementation and validation of new and more robust detection algorithms, a dedicated analog front-end was combined with a flexible digital backend in the form of an FPGA.

The analog front-end consists of an excitation current source, a sense current source and a pre-amplifier. In order to analyse the performance of the analog front-end, the transfer of the pre-amplifier was measured to determine the exact gain and filter characteristics. The gain at the frequencies of interest, i.e. 10 to 50 kHz, was specified to be 80 dB, which corresponds well to the measured gain of 79.7 dB. Subsequently, an excitation-

and sense current were applied to the sensor to measure the signal and noise performance of the front-end. The input-referred noise voltage spectral density was measured to be $4.6 \text{ nV}/\sqrt{\text{Hz}}$, which is close to the theoretical limit imposed by the thermal noise of the GMR, i.e. $2.9 \text{ nV}/\sqrt{\text{Hz}}$. Finally, the temperature dependence of the transfer of the single-stage pre-amplifier was measured. The temperature coefficient of the pre-amplifier gain was found to be $0.71 \text{ } \%/^{\circ}\text{C}$. Although the temperature dependency has been reduced significantly compared to the previous platform, it is still considerable, even though the number of analog filter orders has been reduced to a minimum. Further improvement of the robustness should be provided by the robust detection algorithms.

A model of the combined sensor and analog front-end was made in Matlab Simulink that facilitates easy performance testing of the robust detection algorithms. The model parameters were obtained from measurements performed on real sensors. The same spectral components with similar amplitudes and the same noise voltage spectral density are obtained when exciting the model, compared to when exciting a real sensor, such that the validity of the model is concluded.

Signal processing is required to extract the actual information-carrying signal that is linearly related to the bead density on the sensor surface and independent of temperature and external magnetic fields from the sensor signal. Two algorithms were developed to this end and subsequently implemented in the FPGA. The reference sensor algorithm combines the signal of a measurement sensor and a reference sensor in order to extract the signal of interest. The vector decomposition algorithm combines the responses of the sensor upon excitation at different frequencies in order to extract this signal.

In order to test the vector decomposition algorithm, simulations were performed in Matlab Simulink at different temperatures. From these, it could be concluded that the estimated values for the signal of interest correspond well with the expected values over the complete temperature range. Furthermore, it could be concluded that the algorithm is indeed capable of extracting, and referring back to 20°C , the in-phase component of the magnetic signal, for all temperatures. It was found that the maximum deviation of the signal of interest from its expected value was $22\text{e-}9 \text{ V}$, which corresponds to about 5 surface bound beads.

A joint phase and amplitude estimator (ϕ, A -estimator) is used to track variations in amplitude and phase of the sensor signal during the measurement. The obtained estimates are used as input for the detection algorithms. The estimator uses gradient-based least squares optimisation. A theoretical treatment of the ϕ, A -estimator was given, supported by some simulation results of a ϕ, A -estimator that was implemented in Matlab. The simulations confirmed the relation found between input signal SNR and amplitude estimate SNR. The simulations furthermore confirmed the obtained relation between the time constants and the adaptation gains of the adaptation loops.

Three criteria were formulated to test the reliability of the estimates obtained with the joint phase and amplitude estimator. The convergence criterion that states that the loops have sufficiently converged if the residual parameter mismatch is smaller than 3 times the gradient noise standard deviation, was used to obtain a suitable value for the adaptation constant μ . This constant was chosen $0.95\text{e-}6$, such that convergence occurs within 1 second.

The digital back-end comprises a waveform generator to produce a sense- and excitation- waveform, the ϕ ,A-estimator and a digital pre-amplifier to filter and amplify the magnetic signal before it is applied to the estimator. Since specific spectral components need to be attenuated, the filter consists of two comb filter stages.

In order to test their performance prior to the actual hardware implementation, the waveform generator and joint amplitude and phase estimator were simulated in Matlab Simulink using Xilinx System Generator models. These models were subsequently converted into VHDL optimised for Xilinx devices. The generated VHDL was finally, along with other digital back-end VHDL code, used to generate a program file to be loaded in the FPGA.

By simulation, it was tested whether the components in the spectra of the waveforms generated by the waveform generator, other than those intended, are sufficiently low. Contrary to the sense frequency, the excitation frequency is not an integer dividend of the clock frequency of the FPGA, resulting in the generation of subharmonics. It was, however, shown that these subharmonics are at least 80 dB smaller than the fundamental harmonic, which suffices. This simulation result was confirmed by measurements. Furthermore, the amplitude of the generated waveforms was compared for different frequencies, in order to check that the amplitude is constant across the frequencies of interest. The amplitude of the waveforms generated varied about 100 ppm across the frequency range of interest, which results in amplitude variations in the magnetic signal that are smaller than those as a result of the GMR thermal noise.

The performance of the ϕ ,A-estimator was tested by simulation. The GMR referred noise standard deviation of the amplitude estimate was found to be 12 nV. Subsequently, the GMR referred standard deviation of the amplitude estimate of a sensor was measured. The standard deviation of the estimate was found to be 16 nV corresponding well to the expected and simulated value.

In order to test if the filter in the digital pre-amplifier sufficiently suppresses the sense signal, as well as the capacitive crosstalk signal, the spectra of the magnetic signal before and after filtering were measured and compared. From these measurements, it can be concluded that the digital pre-amplifier filter shown notches at the intended positions and is capable of sufficiently attenuating the sense and capacitive and inductive crosstalk signals.

Two types of experiments were conducted in order to test the reference sensor algorithm. In the first type of experiments only the sensor was heated. The maximum error, i.e. difference between the compensated amplitude and the expected amplitude was found to be 50 nV. This error corresponds to about 29 surface bound beads. The vector decomposition algorithm is subject to further research, since unexpected signals were obtained. In the second type of experiment the complete platform, including the sensor, was heated in an oven. The maximum error that was observed is 273 nV, corresponding to about 109 beads. A residual temperature dependency of 0.06 %/°C was observed in the compensated amplitude signal. From the experimental results obtained with the heating experiments, it can be concluded that the temperature dependency of the platform output signal is negligible in the case that only the sensor is heated. When the complete platform

is heated, the temperature dependency of the platform output signal is reduced by a factor 7.8 as compared to the uncompensated signal.

8.2. Recommendations

In general it can be said for any sensor system that the sensor itself should be designed such that the sensor interface electronics can be kept as simple as possible. The amount of signal processing required to obtain accurate and reliable signals should in principle be kept to a minimum, in order to guarantee robustness. In other words, signal processing should be seen as a last resort to deal with problems such as a large dynamic range, crosstalk and lack of robustness that cannot be solved by proper sensor design. In the case of the biosensor platform, a lot of effort is taken during the detection electronics design to cope with the large dynamic range of the sensor signal. This problem is two-fold. Firstly, the small resistance variation in the GMR per surface bound bead compared to the large R_0 , i.e. a factor of about 10^{-9} , means that large signals coexist on the detection PCB together with the extremely small signal of interest. This greatly complicates the design and lay-out of the PCB, since the large signals on PCB are bound to result in crosstalk and harmonic distortion, e.g. in the sense amplifier. It can therefore be recommended to use sensors that show a large magneto-resistive effect. Tunneling Magneto-resistances (TMR), which show a magneto-resistance effect that is much larger than that of GMRs might for instance be employed. Secondly, the dynamic range of the magnetic signal itself is large as a result of magnetic crosstalk. This crosstalk, which is typically a factor 10^4 larger than the signal per surface bound bead, is dependent on temperature and therefore reduces the robustness of the output signal against temperature variations. Even small variations in the GMR sensitivity result in large variations in the platform output signal, compared to those as a result of changes in surface bead density. It is therefore recommended to use sensors that have a reduced magnetic crosstalk signal. This can be achieved by centring the GMR and current conductors in the same horizontal plane. Alternatively, the magnetic crosstalk can be compensated by adding magnetic crosstalk with the same amplitude but of different sign. This can be achieved by applying a current through a third current conductor, specifically designed for this purpose. This so-called compensation conductor is shown in Figure 17.

A trade-off exists between the measurement speed of the platform, which is determined by the convergence time of the joint phase and amplitude estimator, and the noise variance in the platform output signal, which is determined by the bandwidth of the adaptation loops. Both the convergence time and the bandwidth depend on the adaptation constant used in the estimator loops. The aforementioned trade-off can, however, be circumvented by using an adaptive adaptation gain instead of the adaptation constant. By choosing the adaptation gain large at the starting point of the estimation, convergence will occur fast. By subsequently reducing the adaptation gain, the noise variance of the estimate is kept small.

It can be recommended to include a subsystem in the estimator that certifies that convergence has occurred prior to storing the estimates. This is especially important if the phase of the magnetic signal that is applied to the estimator is shifted exactly 180 degrees with respect to the reference signal used by the estimator. In this case the

estimator may theoretically take infinitely long to converge, such that inaccurate estimates are stored.

Since the convergence of the phase adaptation loop depends on the amplitude of the magnetic signal applied to its input, it is recommended to guarantee a certain amplitude on the input of the estimator. This could be done by controlling the amplitudes of the sense- and excitation currents. However, to obtain optimal signal to noise ratio of the magnetic signal, it is advisable to maintain the amplitude of the sense- and excitation current at their respective maximum values. Therefore, it is recommended to adapt the gain in the digital pre-amplifier, such that the signal at its output is of approximately constant amplitude. The obtained estimates should of course be scaled by the inverse of the applied gain adaptation.

Although not harmful, the spikes that occur in the amplitude estimate as a result of switching between GMRs that can be observed in Figure 78, can be prevented by starting the estimator directly after a switching event, using the stored amplitude estimate before switching as an initial value. This also results in faster convergence, such that the adaptation constant, and therefore noise variance, can be reduced. A smaller adaptation constant would, however, result in some initiation time after a reset, to allow the estimates to reach the correct values.

It is recommended to perform further measurements using the vector decomposition algorithm, in order to determine the cause of the reduction of the signal upon switching from a lower to a higher excitation frequency, where an increase is expected. In other words, further experiments should be conducted to find the cause of the improper functioning of the platform when employing the vector decomposition algorithm.

During the implementation of the analog front-end model, it has become clear that Simulink is not reliable and user-friendly when it comes to solving differential equations, e.g. as a result of the inclusion of a differentiator in the model. The results depend heavily on the solver selected. In-depth knowledge of the selection of specific solvers is thus required to obtain reliable results. Furthermore, the transfer of Simulinks analog filter models depend heavily on chosen solver. It is therefore recommended to take the effort to discretise any physical model manually and implement the model in Matlab code. This code can subsequently be imported in Simulink in order to perform simulations with the rest of the system.

System Generator has been found to be an excellent tool in the process of FPGA prototyping. It allows simulation of the an FPGA design in the rich simulation environment of Simulink. Moreover, the conversion from System Generator Simulink models to VHDL and finally hardware has been proven to be virtually effortless. The same applies for the incorporation of System Generator output in a larger VHDL design. It should, however, be noted that the decision whether to use System Generator or write VHDL should be made per subsystem designed. Some subsystems are easily implemented in a few lines of VHDL and would take considerably more effort to design as a System Generator model. On the other hand, subsystems that are more easily realised by connecting a few System Generator blocks, or subsystems that required extensive simulations, should be designed using System Generator. It can therefore be recommended to use System Generator for FPGA prototyping, when designing systems,

or subsystems, that require thorough simulation and/or when this is expected to reduce the design time.

References

- [1] D.R. Baselt, G.U. Lee, M. Natesan, S.W. Metzger, P.E. Sheehan, R.J. Colton, *Biosensors Bioelectr.* 13 (1998) 731.
- [2] R. Kotitz, H. Matz, L. Trahms, H. Koch, W. Weitschies, T. Rheinlander, W. Semmler, T. Bunte, *IEEE Trans. Appl. Superconductivity*, 7 (1997) 3678.
- [3] Measurement performed by Femke de Theije.
- [4] D. Wild, *The Immunoassay Handbook*, Elsevier, 2005.
- [5] D.L. Graham, H.A. Ferreira, N. Feliciano, P.P. Freitas, L.A. Clarke, M.D. Amaral, *Sensors Actuators* 107 (2005) 936.
- [6] L. Lagae, R. Wirix-Speetjens, C.-X. Liu, W. Laureyn, G. Borghs, S. Harvey, P. Galvin, H.A. Ferreira, D.L. Graham, P.P. Freitas, L.A. Clarke, M.D. Amaral, *IEE Circuits Devices Systems* 152 (2005) 393.
- [7] R. Wirix-Speetjens, W. Fyen, K. Xu, J. De Boeck, G. Borghs, *IEEE Trans. Magn.* 41 (2005) 4128.
- [8] H. Fukumoto, K. Takeguchi, M. Nomura, H. Endo, *IEEE Transducers '05* (2005) 1780.
- [9] M. Megens, M. Prins, *J. Magn. & Magn. Mat.* 293 (2005) 702.
- [10] R.L. Millen, T. Kawaguchi, M.C. Granger, M.D. Porter, *Anal. Chem.* 77 (2005) 6581.
- [11] K.H.J. Bushow, *Handbook of Magnetic Materials*, 15 (2003).
- [12] L. A. Currie, *Limits for Qualitative Detection and Quantitative Determination: Application to Radiochemistry*, *Anal. Chem.* 40, 586-593 (1968).
- [13] R.J.M. van de Veerdonk, P.J.L. Beliën, K.M. Schep, J.C.S. Kools, M.C. de Nooijer, M.A.M. Gijs, R. Coehoorn, E.J.M. de Jonge, *J. Appl. Phys.* 82 (1997) 6152.
- [14] B.M. de Boer, J.A.H.M. Kahlman, T.P.G.H. Jansen, H. Duric, J. Veen, *Biosens. Bioelectron.* 22 (2006) 2366
- [15] J. Kahlman, private communication.
- [16] H. van Zon, J. Nieuwenhuis, private communication.
- [17] Measurement performed by Theo Jansen.

- [18] J.W.M. Bergmans, Adaptive signal processing systems, course notes Adaptive Systems, 2005
- [19] L.W. Couch, Digital and analog communication systems, Prentice Hall, 1997.
- [20] J.G. Proakis, D.G. Manolakis, Digital Signal Processing, principles, algorithms and applications, Prentice Hall, 1996