NTNU
Innovation and Creativity

Optical Detection of Surface Acoustic Waves

Torbjørn Rostad

Master of Science in Electronics
Submission date: July 2006
Supervisor: Astrid Aksnes, IET
Co-supervisor: Helge Engan, IET
Arne Rønnekleiv, IET

Norwegian University of Science and Technology
Department of Electronics and Telecommunications

Problem Description

The main goal in this master thesis is to upgrade the laser probe setup used to study surface acoustic waves, in particular to develop a new LabView program. The laser probe is based on the knife-edge method. A modulated laser beam is focused onto the surface and the angle (tilting) of the reflected beam is dependent on the surface wave amplitude and phase. The tilting is converted into translation. By using a position sensitive detector the modulation is converted into an electrical signal.

The LabView program will initialize and control GPIB communication with different equipment (x-y translation table, lock-in amplifier, signal generators, oscilloscope, multimeter etc) in order to scan the device-under-test (DUT) and perform measurements of amplitude and phase. Low-level commands must be implemented to control the Newport x-y translation table as no driver is available for this old equipment. To verify that the program and setup function properly, measurements on SAW devices will be carried out. By combining phase and amplitude data, signal processing can be performed. Two dimensional Fourier transform techniques using MATLAB can be applied to separate the SAW propagation directions.

An additional goal in the thesis is to perform measurements using both the heterodyne interferometer - and the laser probe setups and to compare the results with respect to linearity, sensitivity and reproducibility. Absolute measurements cannot be performed using the laser probe setup. The aim is to calibrate the laser probe system by use of measurements from the heterodyne interferometer setup.

Assignment given: 27. January 2006
Supervisor: Astrid Aksnes, IET

Abstract

This project was worked on during the autumn 2005 at the Norwegian University of Science and Technology, Department of Electronics and Telecommunications. The assignment was to write a new LabVIEW programme that is to run the measurement procedure of a laser probe setup. The setup is used in characterization of surface acoustic waves(SAW). A programme was written that contained the necessary functionality and proved to operate satisfactorily. Several measurements were made on a SAW transducer, accurately picturing the wave. Fourier analysis were performed on the collected data in order to separate the propagation directions. An absolute amplitude measurement was made on a heterodyne interferometer, and the result was compared to a similar scan made using the laser probe. The work shows that the setup is ready for calibration against the heterodyne interferometer, in order to enable the laser probe to measure absolute amplitude by itself.

Preface

The work that has been done this semester has been very interesting and has brought me knowledge and expertise that will be valuable in future work. Especially programming in LabVIEW has become an asset, as it is in widespread use in the professional life. In addition, I have learned to work independently, actively seeking information on my own. I wish to thank the people who have helped me through this work. Astrid Aksnes is the teacher who has helped me more than she is obligated to do, for which I am thankful. Helge Engan has proved to be exceptionally knowledgeable, both when dealing with the setup and in general electronics and photonics. I would like to thank Hanne Martinussen, my tutor, for the numerous times she has been of assistance to me and for the use of the heterodyne interferometer she is currently working on.

Contents

List of Figures

1 Introduction

Surface acoustic waves has been the subject of massive interest through the years. From this interest, many characterization methods have appeared. One of which is called the knife-edge technique and is applied in a laser probe setup. In order to enable the laser probe setup to measure absolute amplitude, data acquired using the setup will be compared to measurements performed on a heterodyne interferometer. In order to use the setup in a satisfactorily manner, it is in need of an upgrade on the computer software side. Consequently, a new programme is written and tested using the programming language LabVIEW.

Chapter 2 gives some background information on the nature and history of surface acoustic waves and their application in modern electronics. Some characterization techniques will be reviewed as well in order to provide a base of comparison. In Chapter 3, the laser probe setup is reviewed, initially describing the knife-edge technique. The setup is viewed in context of its influence on the signal, both optical and electrical.

Further theory on the detection process is reviewed in Chapter 4, looking into the mathematics behind the technique. In addition, the lock-in amplifier will have its manner of operation described. Chapter 5 deals with LabVIEW programming. The programme is described according to its order of execution.

In chapter 6, the results of the measurements performed are presented. Both regular scans and data manipulated using Fourier transform are included here. In addition, there is a comparison between scans from the laser probe and absolute amplitude measurements performed on another setup. The LabVIEW programme and the data from the measurements are discussed in chapter 7, and the conclusions are presented in Chapter 8.

2 Background

This chapter will provide some background in the area of surface acoustic waves (SAW) and their applications. Methods used in characterization of these waves will be reviewed in order to exemplify the diversity in techniques and provide background for understanding the method applied in this work.

2.1 SAW

Surface acoustic waves are, as the name implies, waves that propagate on top of, or close to, the surface of a material. The typical penetration depth is no more than a wavelength, as opposed to bulk waves which propagate through the material. The reason for this confinement near the surface is the difference in phase velocity between bulk and surface waves. Since a surface wave has a lower velocity, its energy may not propagate into the material. Although bulk waves are more efficiently generated, SAW is preferred in many respects. A reason for this is the already mentioned confinement of energy close to the surface, which makes it easy to manipulate and pick up the waves. Also, the compatibility with planar technologies is an essential feature [1] [2]. The waves are excited by metallic fingers setting up an electric field on a piezoelectric material. The field causes the material to contract or expand, depending on the polarization, thus creating a perturbation which then moves along the surface.

Today, the field of SAW is the object of much attention, much due to the widespread use of SAW filters in mobile telecommunication systems and other applications. The reason for this extensive use is the low propagation velocity of acoustic waves compared to electromagnetic waves. It is about 10^5 times smaller for acoustic than for electromagnetic waves, and devices such as filters and resonators may therefore be manufactured substantially smaller. This feature became important after White and Voltmer in 1965 managed to create surface waves at RF frequencies, a feat which professed the possibility of using SAW transducers in real-time high-frequency signal processing[1]. The need to characterize these waves, their manner of propagation and other parameters became apparent and soon lead to the development of numerous diagnostic methods. Some of these methods were not developed to study the surface acoustic waves themselves, but rather to study the material parameters the waves could provide, and the information surface waves can give about bulk waves[3].

2.2 Heterodyne interferometry

One optical method showing great promise in the study of SAW and vibrating surfaces, is the heterodyne detection method. An example of such a setup is shown in Fig. 2.1. This particular setup is thoroughly described in [4] and [5], while the method itself is also reviewed in [6].

The heterodyne interferometer exploits the phenomenon of interference between

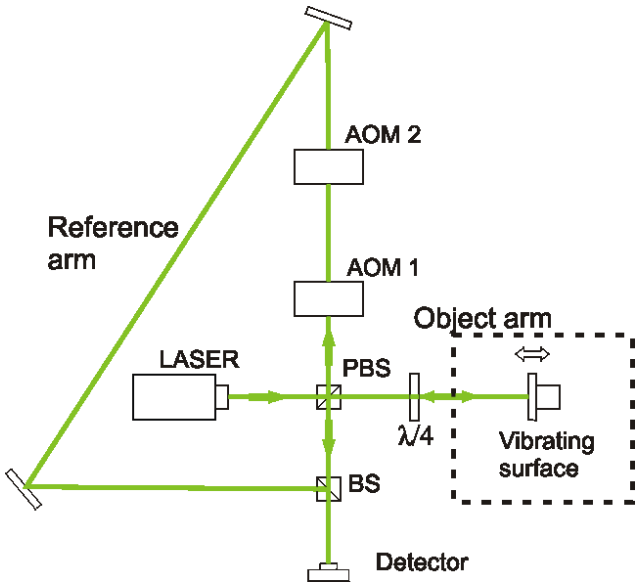

Figure 2.1: An example of a setup utilizing the heterodyne detection principle

two waves experiencing differences. As seen in Fig. 2.1, a laser beam is split into two arms. One, the object arm, hits a vibrating surface and is modulated with three frequency components due to a Doppler shift, one of which is equal to the frequency of the vibrating surface. The other arm, the reference arm, is also modulated, but this one by using two acousto-optic modulators(AOMs). Two AOMs are used in order to give the modulation a large dynamic range. The beam is modulated with almost the same frequency as that of the object arm, but with an additional $1kHz$. When these two beams of light interfere in front of the detector, they produce a frequency component equal to the difference between them. There is a visibility variation in the resulting beam, due to interference, which depends on the phase difference between the two arms. Due to the vibration, one arm experiences a different traveled length, which effects the spatial phase component. A lock-in amplifier locks on this frequency component and extracts amplitude and phase information.

Absolute amplitude measurements have been achieved in [4] and [5] by modulating the detected signal once more. This in order to isolate the part of the signal directly dependant on the amplitude of the SAW. Absolute amplitude means the exact height of the wave tops, and not just the height of the top compared to the bottom, which is often the case. The setup is presently able to detect and measure the amplitude of surface waves having an amplitude of about 7 pm.

2.3 The scanning tunneling microscope

Using the scanning tunneling microscope in SAW characterization provides a high spatial resolution. R. Koch et al shows in [3] that their setup is able to determine the topography of the surface wave, as well as the amplitude and phase. The setup

is depicted in Fig. 2.2.

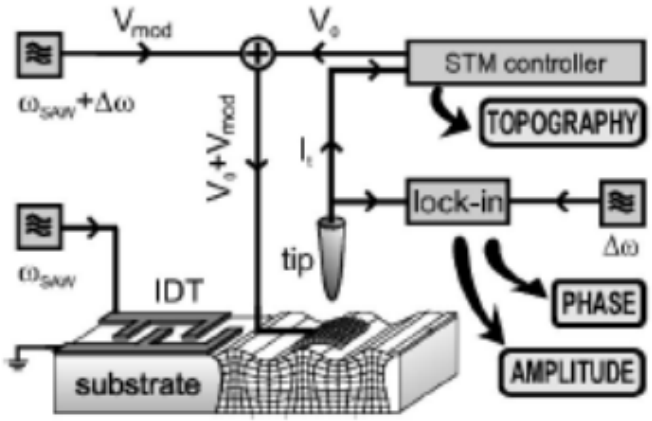

Figure 2.2: Characterization method using a scanning tunneling microscope.

The probe senses the variation in the tunneling current due to the surface perturbation from the surface acoustic wave. By modulating the tunneling voltage with a frequency slightly off the SAW frequency, the signal read from the probe will have a frequency equal to the difference between them, and may consequently be locked on by the lock-in amplifier. The drawback of this method is the rather extensive equipment needed when performing the measurements. The setup operates in very high degree of vacuum for best performance, a feature which is very impractical. In addition, the need to have a conductive film placed on the surface makes the technique somewhat volatile.

3 Setup and instrumentation

The laser probe setup which is used in this project was built in 1977 by Helge Engan and was first described in [7]. Several modifications have been made, but the basic principles remain the same. The setup has earlier been able to detect waves of frequencies as high as 1 GHz, with amplitudes as low as some tens of pm. In this chapter the setup of the laser probe will be reviewed and explained. There will be a description of the method applied in the setup as well as the instruments and components included.

3.1 The knife-edge method

This method is a way of visualizing the phase fronts of a travelling wave in a non-volatile manner. It was first used by Adler et al [8] in 1968 and setups using this technique have later been reported by e.g. Engan [7] and Kamizuma et al [9]. An overview of the method has also been presented in articles by Monchalin [6], Stegeman [1] and, to some degree, by De La Rue [10].

Rather simplistically, one may say that the knife-edge technique exploits the principle of reflection angles. Picture a reflecting surface. If a beam of light is incident on the surface plane at an angle θ, the angle of the reflected beam, θ_r, will be identical relative to the surface normal. As the surface is perturbed by a surface acoustic wave, the reflecting plane will no longer stay horizontal, but rather become the tangent plane of the wave at the point where the beam hits, as illustrated in Fig. 3.3. As the wave propagates, the angle of this tangent plane relative to the horizontal plane will rock back and forth at a frequency identical to that of the surface acoustic wave. Any changes in the angle of the reflective plane will incur a change in angle of the reflected beam. Thus, the transversal variation in the reflected beam depends on the inclination of the reflecting plane, and thereby the shape of the wave.

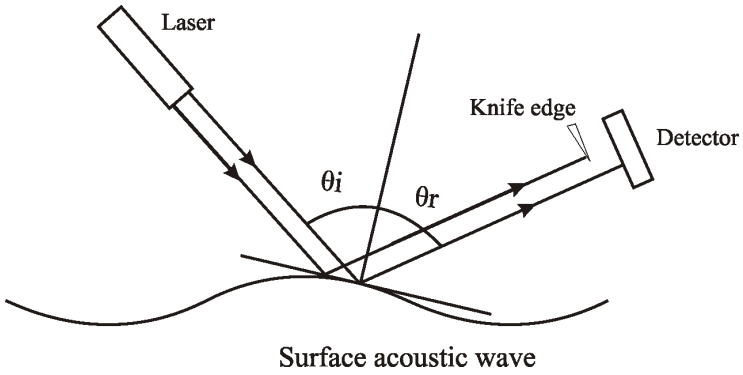

Figure 3.3: The original knife-edge method.

Fig. 3.3 shows the original knife-edge method, placing a sharp knife edge in front of the detector, partially covering it. As the reflected beam is scanned across

the detector, the detected signal is at a minimum when the beam is blocked and maximum when it is not, thus creating a harmonic signal of the same frequency as that of the surface wave. The beam spot hitting the detector is comparatively large, making a smooth transition between minimum and maximum. A small beam spot would either be either completely detected or completely blocked. Since the wave amplitude has an influence on the angle of the reflected beam, the detected signal also carries information on this parameter. It is apparent that for this method to function properly, the beam spot of the incident beam should be considerably smaller than the acoustic wavelength.

In his work, Monchalin[6] also describes a variation of the knife-edge method using a position dependant detector. The knife edge is removed and the detector replaced with one able to sense the position of the beam spot on its effective receiving area. One way of realizing such a detector is coupling two and two quadrants of a quadrant photo detector. A quadrant detector is a detector, most often circular, where each quadrant has its own output, as seen in Fig. 3.5. It is this variant that is applied in the setup[7] and the principle of its alignment is shown in Fig. 3.4.

After the beam is reflected, the objective lens brings the beam to its original inclination, pointing upwards. This realigning of the beam direction transforms the wave-dependant variation in angle of the reflected beam to a horizontal variation on the detector. As the laser beam is traversing the surface of the detector, perpendicular to the line dividing the two halves, the two halves experience different amounts of light. The sum of the resulting signals will remain the same, but by running them through a division circuit, subtracting one signal from the other, the result is a harmonically varying voltage similar to that of the original method. It is assumed that the beam hitting the detector does not have a beam spot smaller than the active area of the detector. At the same time, the spot should not be too small, as that would make the detector difficult to position.

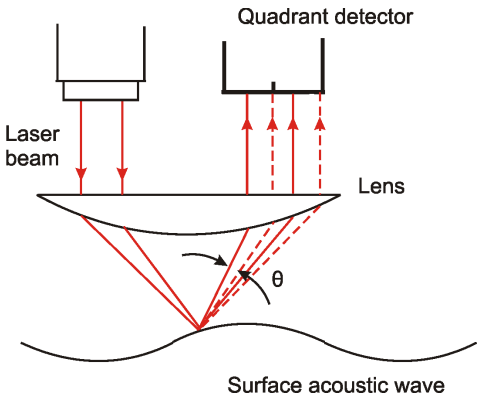

Figure 3.4: Principle of the modified knife-edge setup.

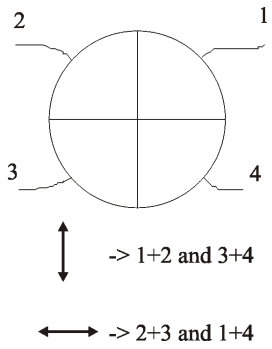

Figure 3.5: Geometry of quadrant detector.

Compared to the original knife-edge method, this variation in the setup makes use of the whole beam power, thereby doubling the signal strength[7]. Another ben-

efit from removing the knife edge is having one less optical component to align. All changes reducing the complexity of the system without deteriorating its functionality are welcome changes. Also, using the quadrant detector makes adjusting the detector to the beam easier, as it is only necessary to monitor the DC signal from the detector to place the beam in the centre. When the two halves of the detector are balanced, meaning the beam is centred on the detector, the output DC signal is zero and the sensitivity is at its best. This is because the beam then travels equally far onto each half, consequently making the detected signal revolve about zero with maximum values on each side.

3.2 Using the stroboscope effect

Surface acoustic waves are most often excited by signals of very high frequencies, up to several GHz. If the knife-edge method was to be used directly on waves of these frequencies, they would yield an output signal having the same frequency. Detectors capable to follow such rapid movements are quite expensive and often have much smaller active area. Therefore, it is of interest to mix the signal down to a lower frequency, while still keeping the necessary information. The information-carrying signal may then be extracted from the detected signal using a lock-in amplifier. In order to shift the frequency of the signal one may apply what one might call the stroboscopic effect, in lack of a better word.

The stroboscopic effect is apparent in the everyday life. An example of the phenomenon is the wheel rim of a car. When it spins at a certain speed, the wheel rim seemingly stops spinning, while the wheel itself clearly continues. This happens as the sample rate of the eye becomes identical to the frequency of the holes in the wheel rim spinning. The same phenomenon may be exploited in the setup. By chopping the laser beam at the same frequency as that of the surface acoustic wave, the wave seemingly stands still and the detector regards the surface wave as a standing wave. As there seems to be no propagation, the detected signal is DC and will only vary as a function of the beam spots movement across the device surface. This yields information on the wave, but will not be accurate enough [7]. Small maladjustments in alignment will have great consequences as only a small part of the light is effected by the wave, while most of the light merely sees the surface as a mirror. The area of interest is the wave, not necessary the surface itself, so finding a DC signal within the total detected light is difficult. In addition, information on the propagation of the actual wave would not be possible.

The difference in frequency is created by modulating the laser with a slightly higher frequency than the excitation of the SAW device, also called single-sideband modulation. The effect may be viewed as an oversampling of the surface wave, making the detector seemingly detect a wave of a frequency considerably lower than that of the actual wave. The lower frequency signal may then be extracted using a lock-in amplifier. Both details on this modulation and the lock-in amplifier may be found in Chapter 4.

3.2.1 Sample Movement and SAW Excitation

In order to characterize the device-under-test (DUT), the laser beam must be scanned across the surface. However, moving the optical setup is a cumbersome procedure due to its weight, size and need for stability. Moving the DUT instead is therefore a more practical choice. Fast and accurate movement is achieved by placing the DUT on a xy translation stage, which is connected to an external motion controller, all delivered by Micro Controle. The motion controller is connected to the computer using the GPIB interface, allowing control of the controller from there. This connection is presented in Fig. 3.6, where all connections in the setup are included. The GPIB interface and its uses are reviewed in the programming chapter5.

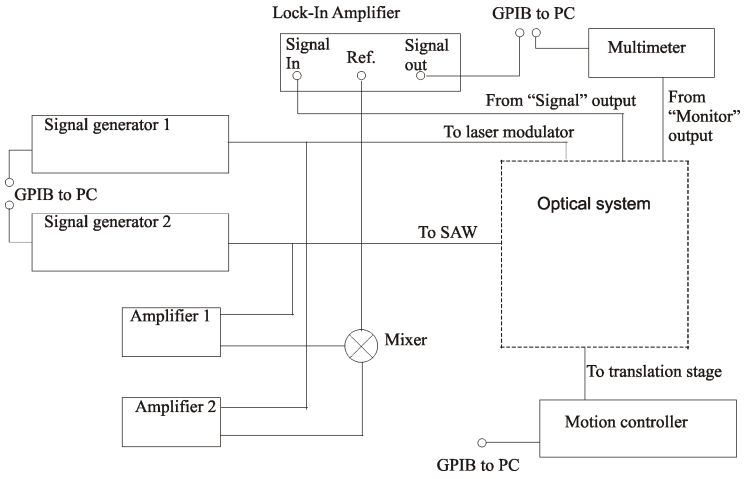

Figure 3.6: Setup of the electrical equipment.

By using the computer, the user may decide which part of the DUT he wishes to study and whether the scanning of the surface should be one-dimensional or two-dimensional. The stage has a resolution of 1 micrometer in each direction. A tilt table is fastened on top of the stage in order to ensure that the sample is oriented completely horizontal as this is very important. A slight tilt of the sample causes the surface to be in focus of the optical system at one end, and out of focus at the other end, consequently making the measurement data unreliable. On top of the tilt table is a metal cylinder made to elevate the sample to an appropriate height. The cylinder was made using specifications fitting the current DUT, which means that should one wish to measure upon another sample, the elevation block will have to be refitted or replaced. Upon deciding the specification the orientation of the sample in the horizontal plane was not stressed. This resulted in a slight misalignment of the sample relative to the directions of the stage. The consequences were not severe in this work, but will nonetheless be reviewed in the discussion of results.

As is explained in Chapter 2, the excitation of surface acoustic waves occurs when setting up an electric field over a piezoelectric material. The signal setting up

the field is most often modulated with a high frequency. In this setup, the signal originates in signal generator 2, as seen in Fig. 3.6. The signal is first wired to a amplifier, which amplifies the signal by 40 dB, before being split in two. One of the two signals are then wired to the SAW device mounted on the translation stage.

3.3 The optical setup

As the characterization technique applied is optical, the optical part of the setup is of great importance. Light leaves little room for errors, due to the very nature of it. Its short wavelength makes even small errors have large consequences, but therein lies also its great strength, its ability to detect small changes.

3.3.1 The microscope

Perhaps the most prominent part of the optical setup is the microscope presented in Fig. 3.7. Manufactured by Minolta, this instrument has several practical features which will be reviewed shortly. First, consider the optical path through the instrument. A laser is attached to the top left ocular directing the laser beam into it. The 5 mW laser emits light at wavelength 670 nm, placing it in the red part of the spectrum. Built into the laser emitting device is an electro optic modulator. The signal driving the modulator is wired from the signal generator marked "laser modulation", as seen in Fig. 3.6. After modulation, the laser beam enters the microscope and is directed, by mirrors and beam splitters, through the microscope objective at the bottom and onto the sample. The five available objectives are placed on a revolving wheel, making it easy to switch between them. As is explained earlier, the beam reflects off the surface and re-enters the microscope through the same objective as on its way down. There, mirrors and beam splitters directs the light up through the microscope and into the detector unit placed atop the right ocular.

The microscope allows the laser beam to exit and re-enter through the same objective. This is beneficial compared to using two objectives, as that would greatly increase the complexity of the system in form of aligning and adjustment. Having to apply two objectives also raises the question of having room for them both. When using an objective with 20X magnification, the focus length is not very long, and having two of them placed together would be close to impossible. Another great benefit is the practical placing of the camera on top of the microscope. It is a Philips solid state camera which displays the surface of the sample on a nearby screen, greatly simplifying the positioning of the beam on the SAW device.

The positioning benefit is the main reason behind including the microscope in the setup. Without it, one would have to rely on measurement scans to find the position, a procedure which takes more time and is less accurate. In addition, using the camera will simplify the adjustment of height by viewing the beam spot size. Using the monitoring feature while measuring, however, is not favourable, since a lot of the signal power then is directed to the camera, instead of to the detector. This problem is solved by the mode selection mechanism seen in Fig. 3.7, switching

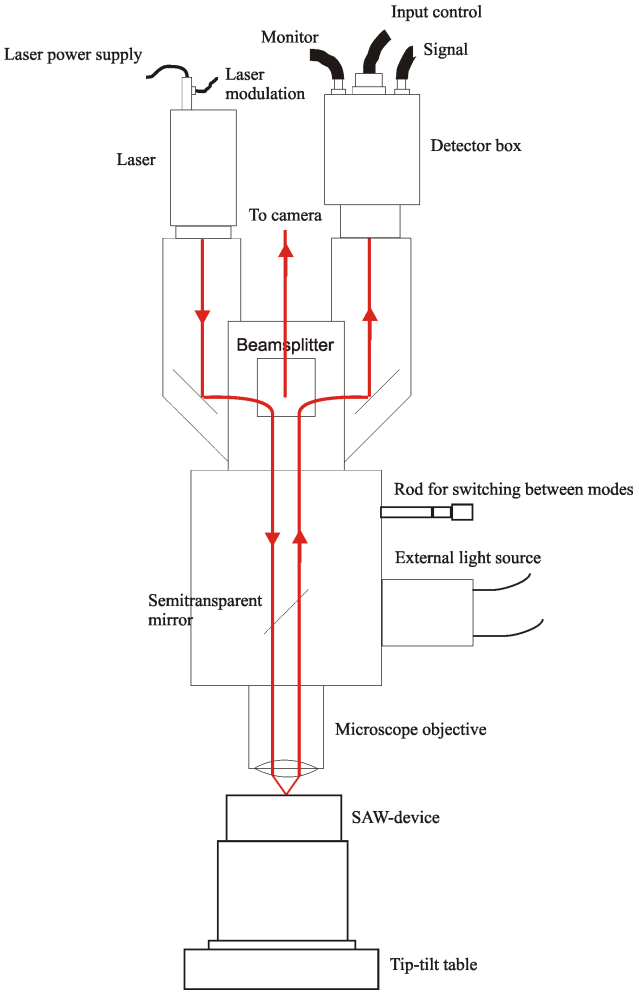

Figure 3.7: Simplistic view of the optical microscope used in the setup.

between positioning mode and measuring mode by pulling a handle/rod on the side of the microscope. The handle is connected to mirrors and beam splitters in such a way that when the rod is pushed all the way in, the reflected laser beam continues straight to the detector. Pulling the rod halfway out brings a beam splitter into position, dividing the reflected beam between detector and camera. Pulling it all the way out brings a mirror into the beam path, blocking out the laser beam completely, thereby displaying only the surface on the screen. The second position is the one used when positioning the beam.

Looking at the surface requires an additional external source of light besides the laser beam. Placing the microscope in positioning mode only displays the laser beam as a somewhat weak diffuse dot of light, leaving the surface features indistinguishable. For this purpose Olympus has included an opening in the side of the microscope, guiding a bright beam of white light into it. Using a beam splitter, this light is guided down towards the surface. The light is powered by a power source of its own, with possibility of adjusting the brightness. Keeping this light on during measurements has been found to deteriorate the data. This was not, however, tried during the measurements in this work.

Using a microscope in the setup is not entirely without drawbacks, but the positive aspects are found to outweigh the negative. The paramount concern is the loss of signal power. It is already explained how the laser beam is guided through the microscope using beam splitters and mirrors. Each time a beam splitter is applied, some fraction, normally fifty percent, of the light is lost. This is due to the semi transparent property of the beam splitter. In order to exemplify this, consider the reflected beam. It follows the same path as the incident beam and is therefore partially directed back towards the laser and partially towards the detector on its way back. Moreover, the rigid construction of the microscope hinders individual adjustments or changes, should that be desired, but this has not been an issue so far.

3.3.2 Detector and detector electronics

The detector used in this setup is, as described earlier, a quadrant detector with two and two quadrants connected. It is seen in Fig. 3.7, positioned atop the microscope facing the right ocular. Both detector and detector circuitry is built into a box, in order to prevent external noise from reaching it. Most frequency components have only small influence on the measurements, since they are not picked up by the lock-in amplifier. There have, however, been problems with noise radiating from the signal generator cables, as the signals they carry are in the desired frequency area. This type of noise is called synchronous or coherent noise and was found to almost disappear when enclosing the detector circuitry in a metal casing.

An interesting feature of the detector circuitry is the ability to change between the orthogonal dividing lines of the detector, meaning that it is possible to change which two quadrants of the detector that is connected. By measuring in both directions, one perpendicular to the other, one is able to discern between both lateral

and transversal wave components. This switching between directions is controlled by the control box shown in Fig. 3.8. The switches in the upper right corner will change the amplification the signal experiences as it passes through the box.

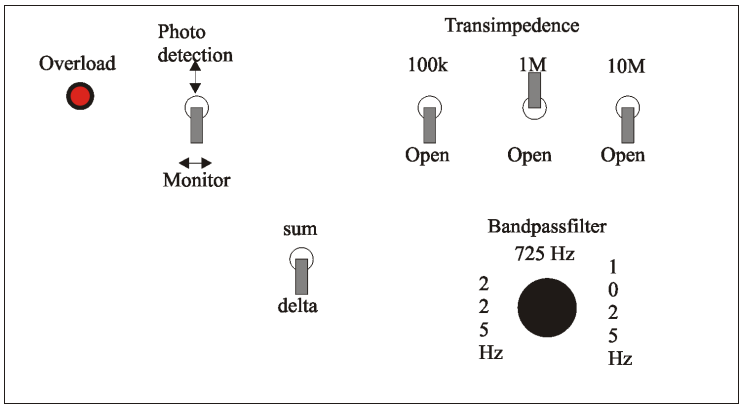

Figure 3.8: Control box for detector process.

In addition to being able to switch between lateral and transversal measurements, the box enables the user to decide whether to use the difference or sum signal of the two halves of the detector. Measurements always make use of the difference signal, but when adjusting the detector to balance the two detector halves, the sum signal is useful. When the sum signal is at a maximum, the difference signal should be zero, something which should be the case for both directions. If this is achieved, the position of the detector is optimal for measurement. Adjusting the detector is easily done by using the two millimeter screws on which the detector box is mounted. The signal used when positioning the detector is the output signal marked "Monitor" in Fig. 3.7, fed to a multimeter which then displays the signal.

3.3.3 Precautions

The microscope is a massive instrument, and with all the extra equipment like camera, detector box and laser attached, its great weight causes some instability. Mainly, this instability is a consequence of the instrument only being supported in the back, creating quite an overhang. Small vibrations deriving from the surroundings may escalate as the optical setup experiences a slight pendulum effect from that overhang. There have been taken measures against this. In order to hinder the vibrations from reaching the setup, the microscope is mounted on a massive iron T beam. This reduces some of the pendulum effect. Also, the optical setup is supported in the front by a rod fastened to the table beneath. This rod also remedies another problem of instability. It may be practical to elevate the microscope when finding the focus length. This is done by turning a wheel driving a slide on which the microscope is mounted. Due to the weighty nature of the microscope, this slide has a tendency to slide out of position when the equipment is handled,

thereby rendering the beam spot unfocused. The front support takes away some of the stress on this mechanism and hinders the misalignment.

3.4 Wiring and signal processing

As is seen in Fig. 3.6, the signal generator marked "Laser modulation" feeds the electro-optical modulator in the optical setup. Prior to that, the signal is split into two and the second wire is connected to an amplifier. Here the signal is amplified by 24 dB and is subsequently wired to a mixer. In the mixer, the SAW signal and the laser modulation signal are mixed together, thereby creating an output signal of a frequency equal to the difference between their individual frequencies. This operation mimics the frequency manipulation incurred by the stroboscopic effect in the optical system. The low frequency electrical signal is then fed to the reference input channel of the lock-in amplifier. This instrument constitutes a vital part of the setup by extracting the desired information from the detected signal. Its signal input channel is fed the signal from the detector via the control box. The detector signal experiences amplification both in the detector circuit and in the control box.

4 Signal processing

There is more to the detection method chosen for this setup than has been dealt with in chapter SETUP. In this chapter, the details of the detection will be explained, focusing on the extraction of information from the detected signal. The signal will be regarded from a mathematical point of view, thereby visualizing the information therein. In addition, the lock-in amplifier used to extract this information will have its manner of operation reviewed.

4.1 Information of the detected signal

The basics of the knife-edge method were explained in the previous chapter. A laser beam is reflected off the surface acoustic wave and scans the detector as a result of the tilting the wave incurs. It has also been stated that the frequency of the detector signal is the same as that of the wave it "sees". What it "sees" is a partially standing wave with frequency equal to that of the difference between laser modulation signal and SAW excitation signal. How does all this come together in the detected signal?

Assume the SAW excitation signal having a frequency ω_r. The signal modulating the laser beam, however, has a frequency $\omega_r + \omega_l$, where ω_l is a considerably lower frequency. As the light hits the detector, a photo diode current, I_d, is created. Using this single-sideband modulation [7], the diode current from the detector may be described by

$$I_d = C_1[1 + b(\cos\omega_r + \omega_l)t][C_2 + C_3 A_{SAW}\cos(\omega_r t - kx + \phi)]. \tag{1}$$

It is apparent that I_d depends on several recognizable parameters. C_1 is a scaling factor depending on the laser beam intensity and the detectors ability to convert light into electricity. The effect of the laser modulation on the detected signal is apparent in the first bracket. A part of the laser beam is not modulated and is therefore not contributing a frequency component, hence the 1 in the bracket. The b indicates how well the laser is modulated.

The second bracket in (1) deals with effect of the surface acoustic wave on the signal and other elements in this interaction. C_2 indicates how large the offset value of the detected signal will be due to inaccurate positioning of the detector relative to the beam. If the beam is not centred in the middle, one half of the quadrant will always provide a proportionally larger signal than the other. Ideally, this constant should be zero. Last of these constants is C_3, which depends on beam geometry. The A_{SAW} is the amplitude of the surface wave. The phase difference between the two waves influencing the detected signal is given by ϕ.

Using simple conversion rules, the frequency component of interest, that with frequency ω_l, may be found. This is the part that is locked by the lock-in amplifier. It was found to be

$$I_{d,l} = \frac{1}{2} b C_1 C_3 A_{SAW}\cos(\omega_l t + kx - \phi). \tag{2}$$

Notice the familiar components from (1). As seen in the equation, the current is a harmonic variation of frequency ω_l, the frequency of the apparent surface wave "seen" by the detector. The wave number, however, is that of the actual surface acoustic wave. Through this frequency manipulation one has successfully moved the amplitude and phase information of the high frequency device to a much more convenient frequency band. There is also a spatial component of the phase, which is found by sampling at multiple positions on the surface. Having information on this part of the phase allows the separation of propagation directions.

Often the surface acoustic wave is a superposition of two waves propagating in opposite directions. These wave components may be indicated by splitting the signal in (2) in two. The first factors of the equation are not influenced by this dividing, but A_{SAW} is split according to the amplitude of each wave component. 2 may then be written as

$$I_{d,l} = \frac{1}{2} bC_1C_3[A_{SAW,1}\cos(\omega_l t + kx - \phi_1) + A_{SAW,2}\cos(\omega_l t - kx + \phi_2)]. \quad (3)$$

This representation will be used further in this review, but in order to understand the information extraction process, the lock-in amplifier will be described first.

4.2 Lock-in amplifier

The lock-in amplifier constitutes a vital part of the post processing of the signal. The main feature of this instrument is to isolate the part of the signal where the useful information lies and then detect the phase of it[11][12]. It is also possible to measure the magnitude of this signal as well, but this is somewhat cumbersome using a single-phase lock-in amplifier. One might use a dual-phase lock-in, but to understand the basics the single-phase manner of operation will be reviewed first.

The instrument has two inputs, one is the input signal to be measured, the other a reference signal at the same frequency as the input signal. The input signal frequency is assumed to be 1025 Hz, as is the case in the present work. Using band-pass filters, this frequency band is isolated and is fed to the most important part of the instrument, the phase-sensitive detector, PSD. Most of the noise which lies outside the pass-band will then be filtered out. The input signal may be described by

$$V_{in} = A_{in}\cos(\omega t). \quad (4)$$

The PSD creates an internal sinusoidal reference signal from the reference signal

$$V_{ref} = A_{ref}\cos(\omega t + \theta), \quad (5)$$

which is then multiplied with the input signal. The resulting signal then becomes

$$V_{psd} = A_{in}\cos(\omega t)\cdot A_{ref}\cos(\omega t+\theta) = \frac{1}{2}A_{in}A_{ref}\cos\theta + \frac{1}{2}A_{in}A_{ref}\cos(2\omega t+\theta). \quad (6)$$

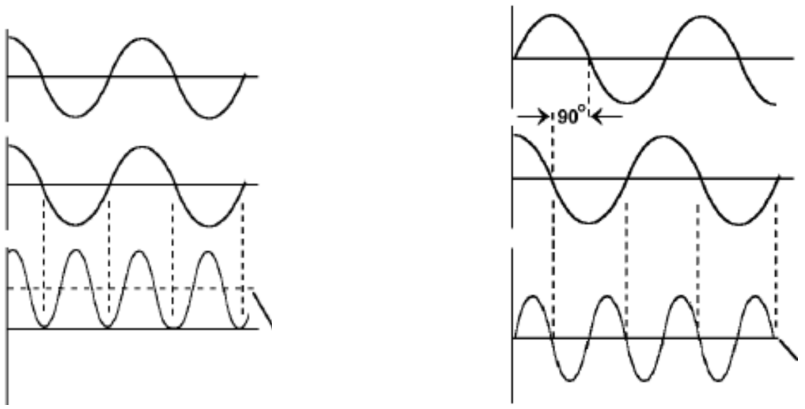

(a) In phase. *(b)* 90 degrees out of phase.

Figure 4.9: Relation between input and reference signal and the outcome.

When the signals are multiplied, the resulting signal is also sinusoidal, but with twice the frequency and a mean level dependant on the phase difference between the signals, θ.

Fig. 4.9a shows the effect of having no phase difference at all. In Fig. 4.9b the phase of the input signal is delayed by 90 degrees with respect to the reference signal and it is clear that by doing this the mean level is moved from a maximum level to a minimum. Consequently, by rectifying and smoothing the resulting signal the output of the amplifier is a DC signal that is a measure of the phase difference.

In addition to measuring the phase it is also desired to measure the amplitude of the 1025 Hz signal. This is possible using a single-phase amplifier, but then one would have to measure twice at each point, with the phase of the reference signal moved 90 degrees between the two measurements. A modified version of the instrument, called the dual-phased lock-in amplifier, will continuously display a constant amplitude. It has the same two inputs, but differs in that that it has two PSDs. Each of these components uses the reference signal fed to them to create internal reference signals, but these signals are 90 degrees out of phase, i.e. a sine and a cosine wave of the same frequency. They then operate as they would in the single-phase amplifier, but the output of them are denoted X and Y, respectively. The output may be presented in a system of coordinates. If the input and reference signal are in phase the magnitude of the input 1025 Hz signal is that of X. If they are 90 degrees out of phase the magnitude is given by Y. It is often practical to transform these Cartesian coordinates to polar coordinates. From the relation above it is easy to see that an arbitrary phase difference gives a magnitude of the signal given by V_{out} in

$$V_{out} = \sqrt{X^2 + Y^2}. \tag{7}$$

Thus, the magnitude will remain the same, although the difference in phase between the input signal and the reference signal should change. The phase of the

signal is given by

$$\theta = \arctan(\frac{X}{Y}). \tag{8}$$

The transformation may also be presented graphically as in Fig. C

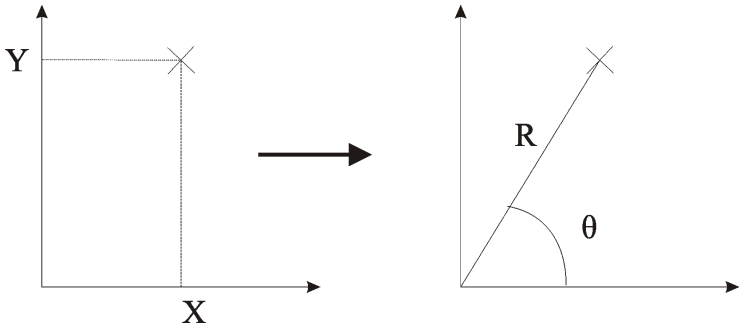

Figure 4.10: X and Y transformed into R and θ.

4.3 The signal being being stored

In the present work, the data is logged in the form of X and Y. Transforming the detected signal into a DC signal does not remove the spatial variation of the surface wave part of the signal. In (8), the spatial variation is included in the amplitude A_{in}. In this specific case, the resulting X and Y component of the processed signal, I_X and I_Y respectively, can be written as

$$I_X = \frac{1}{4} bC_1C_3[A_{SAW,1}\cos(kx + \phi_1 - \phi r) + A_{SAW,2}\cos(kx + \phi_2 - \phi r)], \tag{9}$$

$$I_Y = \frac{1}{4} bC_1C_3[A_{SAW,1}\sin(kx + \phi_1 - \phi r) - A_{SAW,2}\sin(kx + \phi_2 - \phi r)]. \tag{10}$$

These two signals may be seen as the real and imaginary part of the frequency independent part of the detected signal. It is apparent that these waves will be of equal shape and period as the actual surface wave, and may therefore be used to image it. However, the amplitude of the signal will only show the shape of the physical amplitude, as there are too many unknown constants to provide absolute amplitude measurements.

4.4 Absolute amplitude

The missing constants in (9) and (10) may potentially be found by comparing absolute amplitude measurements and measurements using the laser probe. Through identifying each parameters influence on the signal, one might find the relationship

between actual wave and detection signal. The measurements would, of course, have to be made on the same sample, using the same settings.

There is no certainty that such a feat will be successful. If the relation found was not general, but rather device-dependant, a part of the benefit would disappear. This because every device would require extensive absolute amplitude studies on another setup before absolute amplitudes could be measured using the laser probe. However, should the relation be found and it works satisfactorily on all devices, the benefits would be great. The laser probe setup is potentially very fast and very robust against vibrations, which makes it a very formidable diagnostic method, should absolute amplitude measurements be achieved

5 Programming

A large part of the work laid down this semester has been in programming the execution of the measurement procedure. In this chapter, the programming aspect of the work will be presented. In order to fully comprehend the work, a short description of the programme used, along with its main features, will be given. The programming segment itself will be organized according to the finished programme, rather than the order in which the different parts were made. This will make it easier to follow the course of the programme and comprehend the manner of operation.

The purpose of the programme written is to enable the user to measure the amplitude and phase of the surface acoustic waves traveling across a sample at predefined positions, either as a line scan or as a two-dimensional scan. The settings of the instruments in the setup should be available for adjustment from the programme and the measurement data should be stored in a file.

5.1 LabVIEW

Important in all measurements are the precise control of all settings and the logging of all results. This may be achieved by using a computer with appropriate software such as LabVIEW from National Instruments. LabVIEW was first developed for Apple Macintosh in 1986, but is today available for a series of operating systems. The programming language used in the programme is graphical, making it fairly easy to comprehend and work with. In addition, a long range of tutorials are available on the Internet and the programme itself offers an easy-to-follow help feature. The company also offers drivers for an impressively long list of instruments, which eases the communication issues between instrument and computer[13]. A driver is a collection of functions which allows the program to communicate with the instrument without having to resort to low level commands, thereby simplifying the programming process.

The LabVIEW equivalent of a regular programme is called a Virtual Instrument (VI). Three components constitute such a VI: a block diagram, a front panel and a connector pane. The block diagram is perhaps the most comprehensive part when comparing to other programming languages, since it is here the actual programming takes place. Functions are represented by blocks which are connected by wires to provide inputs and outputs. These blocks may be placed inside "for loops", cases or sequences which will decide the course of events in the programme. The front panel is an interface which allows the setting of the values fed to the functions in the block diagram. The front panel of a VI has much the same relation to the block diagram as the actual front panel of an instrument has to the circuitry inside. The connector pane is a graphical representation of the VI, allowing it to be called by putting them inside other VIs.

In order to communicate with an instrument, the computer needs to know which instrument it is currently talking to. This is achieved by giving each instrument its unique GPIB address number, which is normally easily set on the instrument. After

setting the address, LabVIEW always refer to the instruments by this number.

LabVIEW includes several helpful diagnostic tools which have proven valuable in the programming process. One particularly helpful feature is the highlighting option, accessed by clicking the button picturing a light bulb in the upper left corner. When this tool is activated the execution of the programme is slowed down considerably and each signals progress highlighted. This way, the order of execution for the different functions and structures in the programme may be found. Also, the signal values are shown as they are issued and any errors originating from incorrect values will be detected. Probing the connections also presents the signals sent, but the signal values are then available both in slow motion and real time execution. A third helpful feature in LabVIEW is the built-in error report output found in most of the original functions and also many of the functions in the instrument libraries. This is an output gate of the function which reports if any errors have occurred with the function. When connecting an indicator to the output, a window appears on the front panel of the programme. If an error occurs in the function, a message containing the error code will be displayed in the window. Sometimes the error within the code is too severe to run the programme. The "run" button picturing a white arrow then changes into a broken arrow in grey. Pressing the button in that case will bring up an error list describing the reason.

5.2 The programme

The previous programme used to run the setup was last updated in 1997. After so many years, most of the functionality of this programme became obsolete, and both computer and programme was in need of a update. The old programme was written in and run by LabVIEW 3.1.1. The computer on which the new programme is to run on presently uses LabVIEW version 7, while version 8 is available for download. The gap between version 3.1.1 and 7 proved too great, so simply transferring the previous programme to the newer computer and then update the code was impossible. Therefore, it was necessary to write the programme from the start. It is also worth mentioning that the previous version of the measurement programme contained functionality that was unnecessary in the present scope of the project. It was therefore not included.

There is more than one way to organize the contents of a programme. One way to do it in LabVIEW is creating several subVIs which may be called inside higher level VIs. This way, a lot of programme files are created and navigation may become more difficult. However, it may also lead to a more orderly main execution file, since most of the actual programming is then hid in subVIs. This programming method was only used initially in the programme. The programme and the execution are further described later. Since the new programme is only meant to perform one task, the measurement of the surface acoustic wave amplitude and phase, the need to put functionality into subVIs falls away. If additional tasks were to be performed, key functionality, like the movement of the translation stage, would be put in a subVI to prevent unnecessary duplication. Determining the wave velocity is an example

of such an additional task.

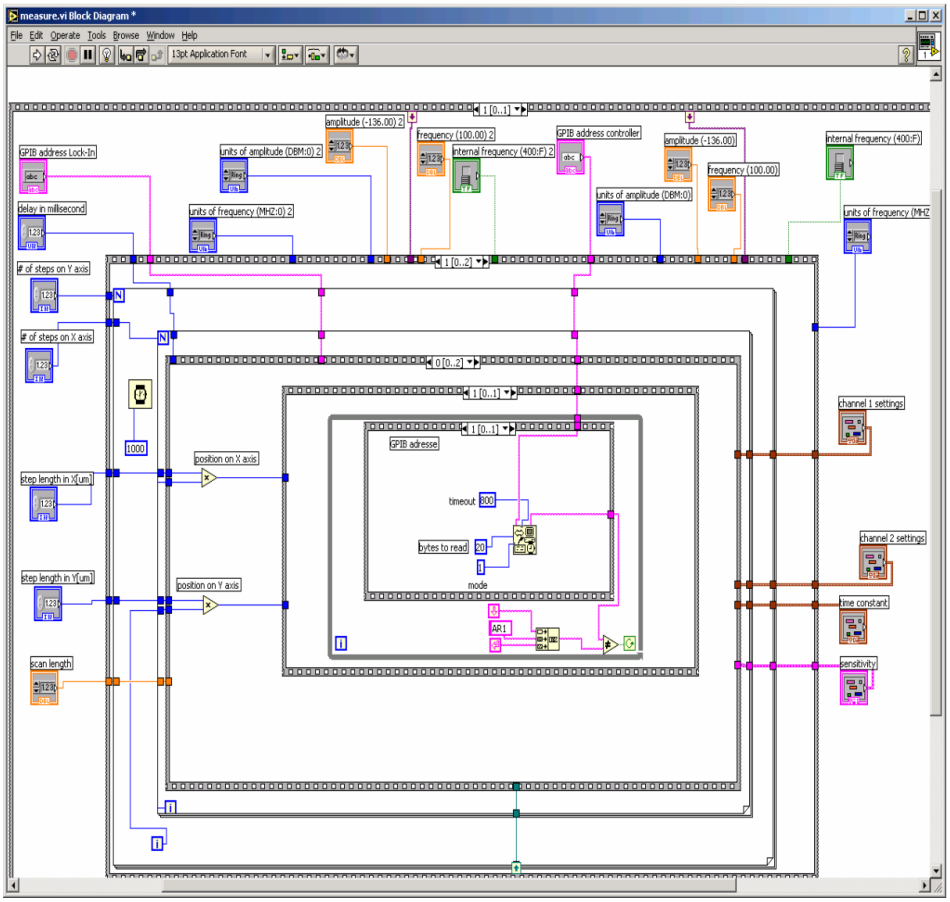

Figure 5.11: A screenshot of the block diagram showing the hierarchy of stacked sequence structures.

The programme created in this project has been organized in a sequential manner. The stacked sequence structure is a structure block which contains one or more frames. The frames are organized by numbers which determine their order of execution, meaning frame number 1 is executed only after frame number 0 has finished. A screenshot of the block diagram window is shown in Fig 5.11. This organization ensures that some elements of the programme are sure to be executed before others. An example may be that measurements are performed only after the ramp movement is finished, not during or prior to it. Sequence structures may be put inside other sequence structures, thus creating layers or levels in the code. In order to satisfactorily explain the execution order of the programme, a flow chart is shown in Fig. 5.12 using the sequence levels to represent a hierarchy. The execution order is then indicated by placing arrows between frames and between layers. A more substantial review of the execution is given later and is organized in the same way, by levels and frames. Relating the review to the flow chart will hopefully ease

the understanding of the execution.

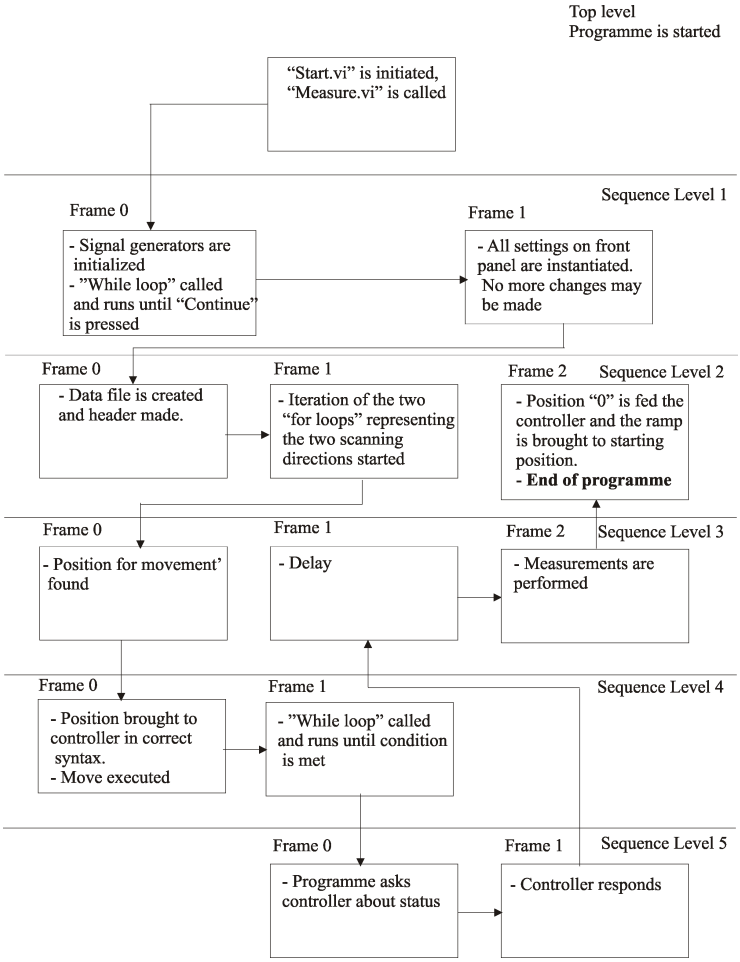

Figure 5.12: Flow chart indicating the order of execution.

5.3 Starting the programme

In order to run the programme one needs to initiate the file "start.vi", which is part of the library measure1. When initiating this file, LabVIEW automatically starts up and displays the front panel of "start.vi". This VI is a small programme constructed to present the user with a choice of operation. The choices which are available at the present time is looking at an older measurement result or creating a new one. The list of choices may be extended easily whenever more functionality is required.

The choice of operation is made by pressing one of the two buttons shown in Fig. 5.13a. They are linked to Boolean controls in the block diagram, placed inside a while loop which runs continuously until the proper signal is received by the red

(a) In front panel. *(b)* In block diagram.

Figure 5.13: Block diagram and front panel representation of choice buttons.

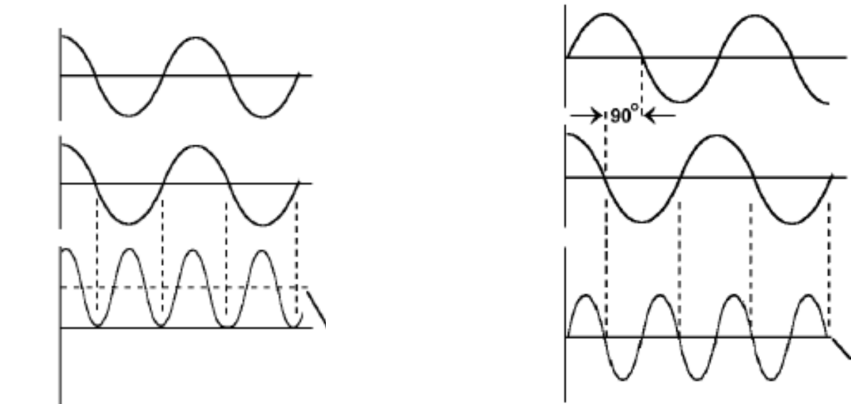

(a) Showing the entire initial front panel. *(b)* Decision structure for initial choices.

Figure 5.14: Image of the initial front panel, where the old measurement is shown, and its block diagram equivalent.

button seen in Fig. 5.13b. The signals from the Boolean controls are inputs on an "exclusive or" Boolean function. It continues to deliver the value "False" as long as none or both buttons are pushed. When a choice is made, the value "True" is sent from one of the buttons to the red button, thereby ending the continuous run of the "while loop". The programme then continues to the next sequence frame.

As illustrated in Fig. 5.13b, the values of the two buttons are wired out of the frame and into the next. There, the Boolean values are inputs on two case structures, which is shown in Fig.5.14b. Case structures have two alternative executions, depending on the Boolean input value. In the two case structures in the second frame of "start.vi", an action will only be performed if the input value is "True". If the "Look at old measurement" button is pressed, the programme initializes the subVI "Read From Spreadsheet File.vi", placed in the case structure wired to that button. Since no file path to a file is given, the function opens a dialogue window which allows the user to browse through the computer for the desired file.

After having found the desired file, the data is available from the output gate of the function as a single precision array. Using the function "Index Array.vi", the

X component of the measured data, which is placed in column five of the data file, is extracted. This data is chosen to represent the measured data. The values are shown in the waveform graph window seen in Fig. 5.14a. In addition to the actual data, also information on the measurement is shown. The header of the file is read using "Read Characters From File.vi" and is shown in a text window also found in Fig. 5.14a. Finally, the number of elements in the measurement is found using "Array Size.vi" and displayed in the front panel.

5.4 Making a new measurement

The other available choice is performing a new measurement. Pressing that button makes the programme call the subVI measure.vi, which lies in the same directory as start.vi. This is the main part of the programme and most of the functionality lies here. Observe that the hierarchy of the flow chart originates from this file, with the exception of the top level which represents the file "start.vi". When the subVI "measure.vi" is called, the subVIs front panel is opened. Fig. 5.15 shows a screenshot of this window.

5.4.1 Sequence level 1, frame 0

When called, the "VI measure.vi" automatically opens in running mode. At this point the user may make the necessary changes to the settings, before starting the actual measurement procedure. All instruments which are possible to control from the computer are available for adjustment from this window.

The execution is kept from moving to the next frame while the adjustments are made. This is achieved by placing a "while loop" in the first sequence frame. The "while loop" is continuously fed the value "False", until the button marked "Continue" is pressed, working in much the same way as the buttons on the start.vi front panel. In the same frame as the while loop, the two signal generators used in the setup are initialized. This is achieved by calling the function "Initialize.vi", which is a part of the signal generator library. These functions, one for each instrument, require the GPIB address for the two signal generators before returning an internal ID for them. These IDs are wired to a sequence frame further down the execution order, where the frequencies are set.

5.4.2 Sequence level 1, frame 1

As the button "Continue" is pressed the programme progresses to the next frame, which is the frame of most importance in "measure.vi". It is the outermost frame in Fig. 5.11. Around the outskirts of this frame, all settings which are set in the front panel are read from their respective controls. The values are placed on the input gates of the sequence within, where they will be read later.

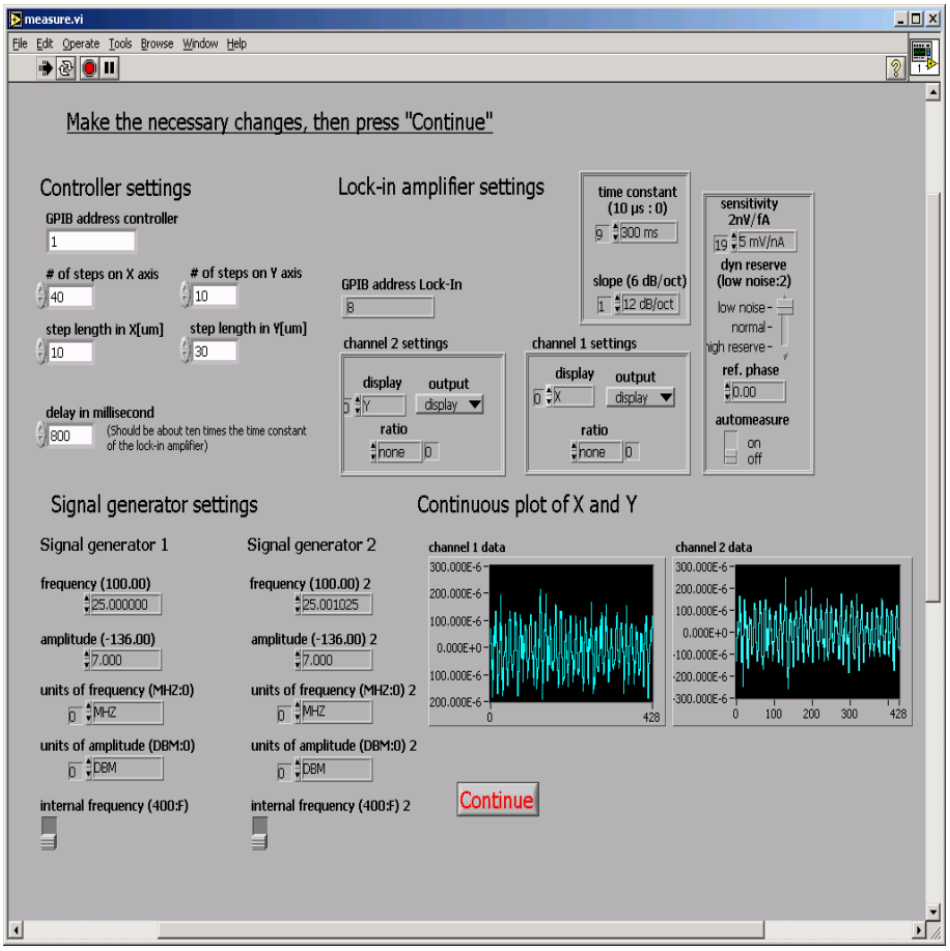

Figure 5.15: The front panel of the subVI "measure.vi" is opened in running mode, but paused so that changes may be made to the settings.

5.4.3 Sequence level 2, frame 0

Most of the settings described earlier take effect in this frame. The functions controlling the signal generators are found here. The function used is called "HP8648A Application Function.vi" and is also found in the library of the signal generators. In Fig. 5.15, the signal generator settings are placed in the lower left half of the window. Through the function, more settings are available, but these were not deemed necessary for now. Should the need arise, more settings are easily added.

In the same frame, the function "SR830.vi" is found. This function handles the operation of the Lock-In Amplifier, and is found in the library within the Lock-In driver. In the front panel shown in Fig. 5.15, the settings for the Lock-In Amplifier is found in the upper right corner. The Lock-In Amplifier function needs to know the GPIB address of the device in order to communicate with it. The multimeter settings are also set here. However, these settings are all constants, meaning they are not configurable from the front panel.

The last element executed in this frame is the creation of a file in which to log the measured values. This is done by calling the function "Format into file.vi". The function is able to take in a number of formats, transform them to text and write them into the text file. In this part of the program, the only writing being made in this file is making an informative heading. This heading will later provide the necessary information about the measurement. The information in the header includes, in addition to some of the settings from the front panel, the time of the measurement, obtained from the function "Get Date/Time String.vi". The function "Format into file.vi" has a file path input so that the function may know where to put the newly created file. This input is left open,thereby forcing the program to open a dialogue box, asking for the file path. When this is done, the programme moves to the next frame.

5.4.4 Sequence level 2, frame 1

The next point in the execution of the programme is moving the ramp on which the sample is placed. The objective of this part of the programme is to enable the ramp to move in a two-dimensional manner, according to predefined values. In order to achieve this, communication with the motion controller is vital. Since it is rather old, there was not found a driver for it, making it necessary to communicate directly with it, instead of relying on functions from the instruments own library. The direct method is far more difficult and time consuming than using the driver. Every instrument has its own set of instructions and commands, and writing these without knowledge of the syntax used is close to impossible. Most manuals, however, provide an overview of the specific commands and the syntax, as does the manual for the motion controller [14]. When speaking directly to the controller one may use the function "GPIB Write.vi", which is LabVIEWs own function for this purpose. Which device to speak to is defined by the address string input and refers to the GPIB address of the instrument.

The user can change most of the settings concerning the movement of the ramp. These are available in the upper left corner of the front panel, see figure (ref front panel). The size of the scan is decided by the number of steps in each direction, and the length of them. The control for number of steps in y-direction is wired to the count terminal of a "for loop". This makes the "for loop" run as many times as the control field indicates. Within this "for loop", there is placed another "for loop", this one connected to the control for number of steps in x-direction. In this way, the ramp first finishes a line scan in x-direction, before returning, moving one step in the y-direction and repeating the scan in x-direction. The controller does not depend on the step size directly, but rather moves according to coordinates defined from a certain starting point. These coordinates are calculated by multiplying the signal from the step size control and the signal from the iteration terminal of the "for loop". Consequently, for each new iteration, the coordinate value increases by one unit step size and the ramp moves accordingly.

5.4.5 Sequence level 3, frame 0.

The two position signals, one for each dimension, now enter a new stacked sequence structure. First the signals are run through a case structure to check for negative figures, and rectify them if necessary. Here, the numbers are formatted into a string to fit the "GPIB Write.vi" function.

5.4.6 Sequence level 4, frames 0 and 1

In frame 0, the syntax of the motion controller truly becomes essential. The different elements of the command are here combined using the function "Concatenate Strings.vi" and fed to the function "GPIB Write.vi", which sends the command to the controller. At this point the translation stage moves to the coordinates defined in the string. A "while loop" is placed inside the next sequence frame. It runs until the proper condition is met. The condition is decided by the stacked sequence structure inside the "while loop".

5.4.7 Sequence level 5, frames 0 and 1

In the first frame, the programme asks the controller what its status is, e.g. if it had finished its movement. The next frame introduces the function "GPIB Read.vi", used when listening to the response from the controller. The string arriving from the controller is checked with the desired response, stopping the "while loop" if the signal is identical. As the movement is now finished, the programme pointer moves back to sequence level 3.

5.4.8 Sequence level 3, frames 0 and 1

Just after the translation stage has made its move, there are fluctuations in the values measured by the lock-in amplifier. This is mostly due to the time constant of

the lock-in amplifier, but also due to vibrations introduced by the sharp movement of the stage. In order to let the Lock-In Amplifier values stabilize, there is a delay placed between the moving and the measuring part of the program, in frame 1 at Sequence level 3. The length of this delay may be set on the front panel, in the field labeled "delay in milliseconds", just below the controller settings.

The following frame is where the actual measurement is performed. The function "SR830 Data Storage Example.vi" receives the measured data from the lock-in amplifier and places it on the output gate in the format double. Since the lock-in is dual phase, there are two output channels. These may be represented in two ways, either as Cartesian coordinates X and Y, or as polar coordinates with amplitude R and phase θ. Their relation will be discussed in the Chapter 4. Switching from one format to the other is easily done by switching the channel setting marked "display" on the front panel. The output channels are connected to two waveform charts which are displayed in the lower right corner of the front panel. They will continuously display the measured values from each measuring point along the surface. The time the lock-in uses to acquire data may be set in the front panel in the field called "scan length". The multimeter used in the setup did not have a library provided in a driver. Therefore, the function "GPIB Write.vi" had to be used here as well. The reading of data is triggered by the function "GPIB Trigger.vi", which is called when the programme enters the frame.

The measured values are not shown continuously on the front panel, but are merely logged in the measurement file. The actual procedure of logging the results are managed by the function "Format Into File.vi", which is one of LabVIEWs own VIs. It is placed in the lower half of the frame. There are several inputs to the function. First, there is the "refnum" input which is a reference number created in sequence level 2, frame 0, along with the log file in order to identify the file for later use. Without this identification, the function would open another dialogue box and then create a new file. The other inputs are values put into the file, arranged in columns after their origin. In the first two columns, the coordinates of the measurement are placed, indicated by position in x- and y-direction respectively. The next two columns find the measured values from the lock-in amplifier. The inputs specified for these columns are wired to the logging function from the function "SR830 Data Storage Example.vi". The last column in the file displays the measured data from the multimeter. These values are fed to the function from "GPIB Read.vi", which reads the data from the multimeter. The measurement procedure here described will run once for each iteration of the "for loops", logging measurement data for each move made.

5.4.9 Sequence level 2, frame 2

After having finished the entire scan, the ramp will need to be repositioned to the starting position. This is achieved by copying the code related to the movement into this frame, but feeding the function the value 0 on both axes, effectively telling the controller to move the ramp to starting position. This concludes the execution

of the programme. The "measure.vi" front panel is still displayed, but is no longer in the running mode.

6 Results

Throughout the programming process, different elements in the programme were continuously tested. However, regular measurements were not performed until the programme was finished. Some of the results from these measurements will be presented here. There will also be a review of the measurement procedure, followed by some thoughts on the significance of the results. In addition, the setup itself will be discussed, with emphasis on sources of errors.

6.1 SAW device

The measurements were performed on a y-cut $LiNbO_3$ SAW transducer. A drawing of the device is presented in Fig. 6.16, which is not perfectly accurate, as the fingers on the device are so small as to be almost indistinguishable. The device has been studied earlier using the heterodyne interferometer that the laser probe is to be calibrated against. A frequency scan was made on the device, showing a resonance frequency at 25 MHz [4]. Consequently, this was set as the frequency of the SAW excitation signal throughout the measurement process.

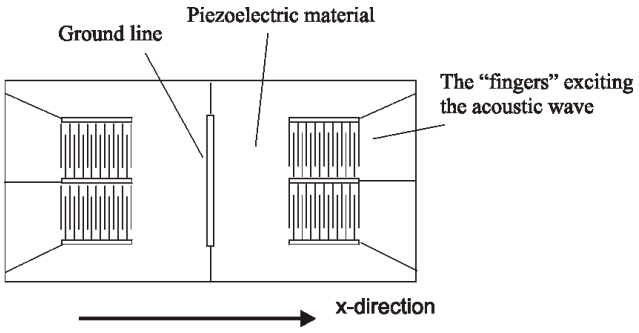

Figure 6.16: The SAW transducer under study.

The measurements were made on the lower left part of the SAW device, and only that part of the surface was excited. There were performed numerous scans on the device under test, but only the more significant of them will be presented and discussed here.

6.2 Initial measurements

At the start of the measurements, there were many elements that were not perfectly clear. The settings on the various instruments were adjusted on a trial-and-error basis, but with some prior knowledge from earlier measurements made with the setup. During the measurement period, most settings were kept unchanged as long as they worked, mainly in order to create equal premises for the individual scans. Also, as the process was trial-and-error, the risk of ruining components was present when making adjustments.

The ground line seen in the middle of the device provided a high reflectivity surface on which to perform the initial measurements. High reflectivity provides a high signal-to-noise ratio and the strip will therefore provide decent results despite the settings not being optimal. The logged data is plotted using MATLAB. Only the x-component of the data is plotted, as it is sufficient to display the shape of the wave. One of the initial scans is presented in Fig. 6.17. The wave shows that the measurement was good, with very little noise. Its period is about 140 μm, which corresponds with the period calculated from the frequency of the SAW. Observe the slight twist of the wave shape, not easily seen in the figure. This is an effect of the misalignment of the metal cylinder on which the device-under-test is fastened to, which was mentioned in Chapter 3.

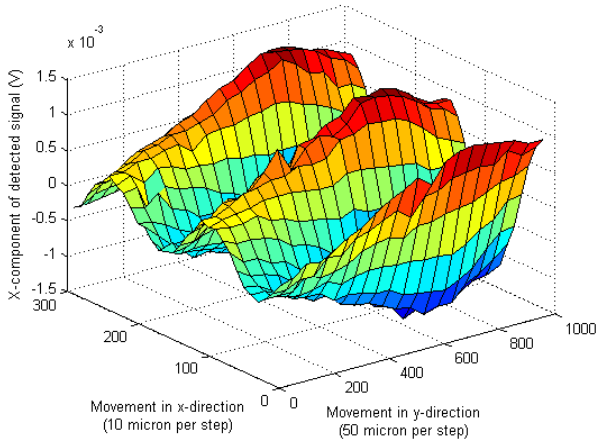

Figure 6.17: One of the initial measurements on the SAW transducer.

6.3 Further measurements and Fourier analysis

A long range of scans were made, on different parts of the transducer and with variable degree of success. Probing the ground line generally gave good results, while probing the $LiNbO_3$ crystal introduced a considerable amount of noise, as will be showed. In order to fully characterize the wave propagation, Fourier analysis was employed in order to separate the propagation directions.

6.3.1 Fourier analysis of ground line measurement

The data from the lock-in amplifier is logged as X and Y, which is the real and imaginary part of the signal, respectively. Having both these parts makes it possible to Fourier-transform the wave, thus regarding it in view of the wave number k. The process is indicated in Fig. 6.18, showing images of each step. It will be explained subsequently.

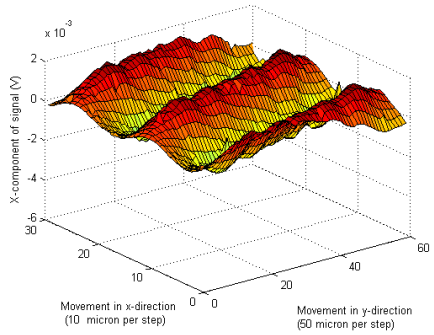

(a) Real component of detected signal.

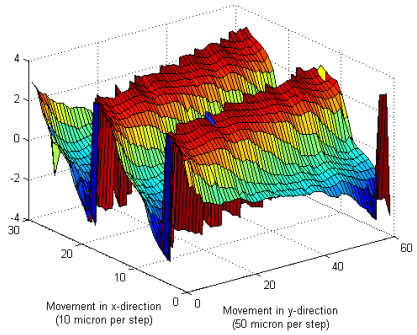

(b) Phase of detected signal.

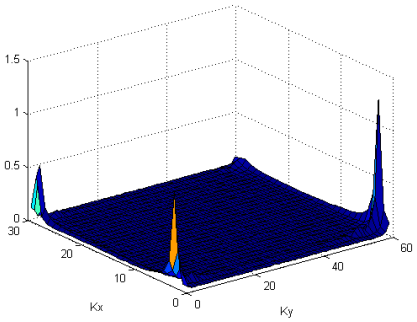

(c) Wave Fourier-transformed into k-space.

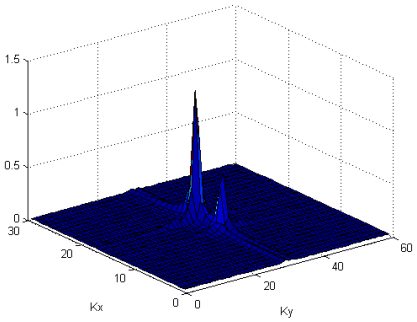

(d) Fourier-transformed wave shifted.

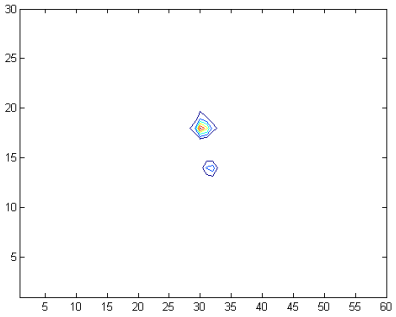

(e) Contour image of shifted wave.

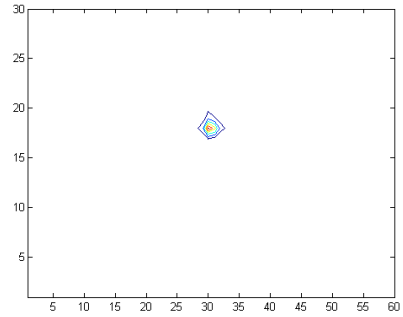

(f) Contour image after filtration.

Figure 6.18: Fourier manipulation on wave in order to split the propagation directions, images a-f

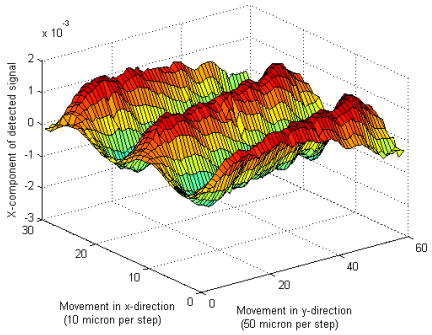

(g) Wave component in positive x-direction.

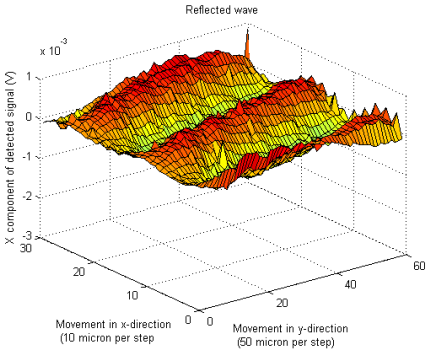

(h) Wave component in negative x-direction.

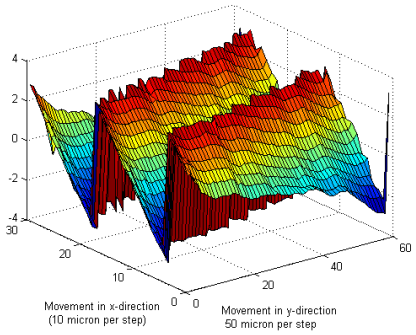

(i) Phase of wave in positive x-direction.

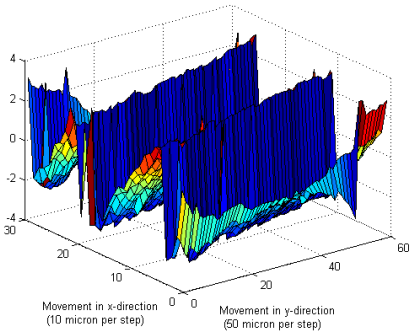

(j) Phase of wave in negative x-direction.

Figure 6.18: Fourier manipulation on signal from ground line measurement, images g-j.

Fig. 6.18a shows the wave as it was prior to Fourier manipulation. The scan was very wide, almost covering the entire wave transversally, that is, in y-direction. Observe the minor bumps on the amplitude tops. These are thought be caused by diffraction, a phenomenon often found on such devices. In Fig. 6.18b, the phase of the wave is presented. The phase indicates which propagation direction which is dominant in the surface wave. As is seen, the phase image is rather misshaped.

When the logged data is Fourier-transformed, the wave will appear as in Fig. 6.18c. The wave components will be in each corner, showing that the propagation directions are directly opposite. In order to show the propagation directions in a more informative way, the transformed data is shifted. Shifting the data means switching the quadrants of the k-plane. The resulting representation is presented in Fig. 6.18d. This relation may be viewed as contours, as in Fig. 6.18e. The upper half of kx signifies the part of the wave traveling in the positive x-direction. Setting all components below this line equal to zero, as shown in Fig. 6.18f, will leave only the the positive component. The same procedure may be followed in order to isolate the negative component.

After having isolated each propagation component they are inverse shifted and inverse Fourier-transformed. The forward propagating wave component is shown in Fig. 6.18g, with phase in Fig. 6.18i. Observe the improvement of these plots now that the other component is removed. The reflected wave with corresponding phase is shown in Fig. 6.18hand 6.18j. The signal is not as large as in the forward direction and the shape is not as pure. This component is probably the wave being reflected from the metal fingers situated on the opposite side of the transducer.

6.3.2 Fourier analysis of measurements on the crystal

SAW transducers may be unidirectional, meaning they only excite waves moving in one direction and suppressing the other. It was found interesting to determine if this was the case with the device-under-test. A scan was made on the crystal behind the fingers, on the left side as the DUT is depicted in Fig. 6.16. The scan showed a wave present, indicating that the transducer is bidirectional. The data from this measurement was also analyzed using Fourier transform.

The original wave is shown in Fig. 6.18a. As is seen, this wave is much more misshaped than the previous, showing much noise in general and one giant spike in particular. The spike might originate from an error at a measurement point or from some particle on the surface. If it is a particle, it is extremely small, no larger than 15 micron according to the step size. From Fig. 6.19d, it is apparent that the wave consists of only one wave component. The wave is propagating in the negative x-direction, as is indicated by the phase in Fig. 6.19b. Notice the strange image in Fig. 6.19f. In this plot, the wave component propagating in the positive x-direction is displayed. However, since this wave has no such wave component in that direction, only noise is displayed. The noise is also apparent in Fig. 6.19g, where there is no hint of any wave.

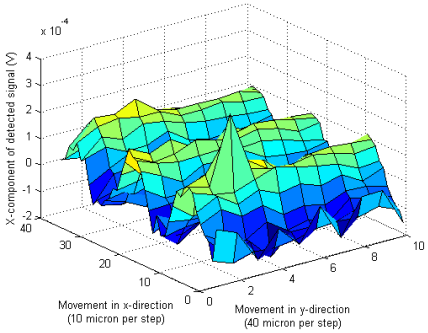

(a) Real component of detected signal.

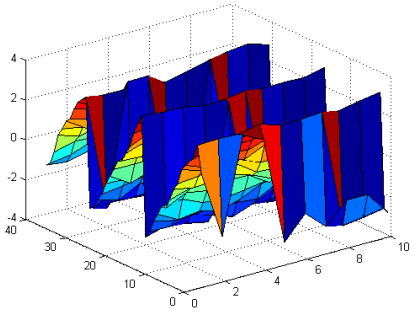

(b) Phase of detected signal.

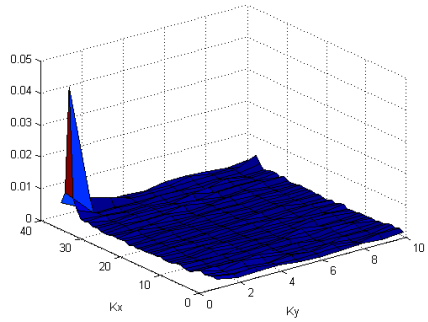

(c) Wave Fourier-transformed into k-space.

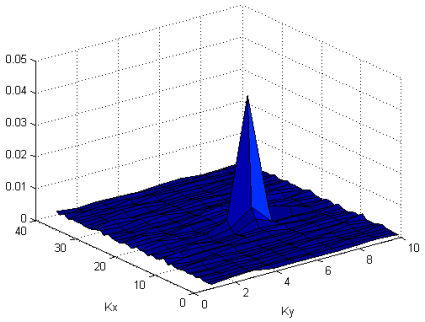

(d) Fourier-transformed wave shifted.

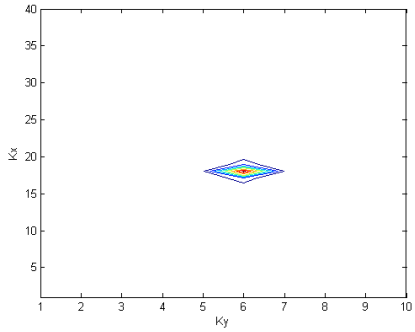

(e) Contour image of shifted wave.

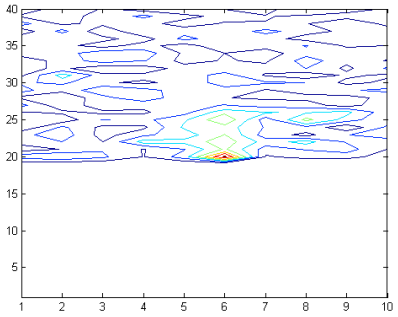

(f) Contour image after filtration.

Figure 6.19: Fourier manipulation on wave in order to split the propagation directions, images a-f.

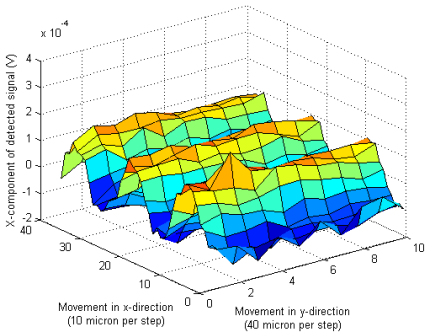

(g) Wave component in positive x-direction.

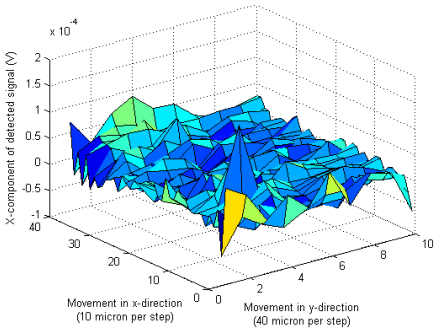

(h) Wave component in negative x-direction.

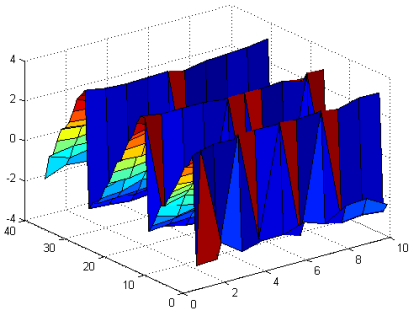

(i) Phase of wave in positive x-direction.

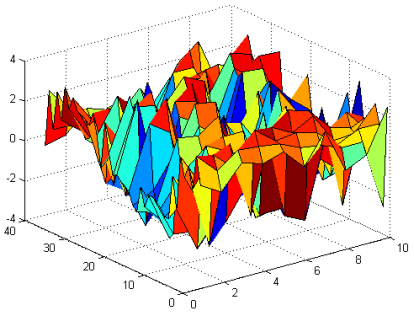

(j) Phase of wave in negative x-direction.

Figure 6.19: Fourier manipulation on signal from measurement on crystal, images g-j.

6.4 Normalizing the detected signal

The difference in reflectivity between crystal surface and metal surface is considerable. This becomes a problem as a scan involves both media. Fig. 6.20 shows a scan covering two sets of metal fingers with the ground line in between.

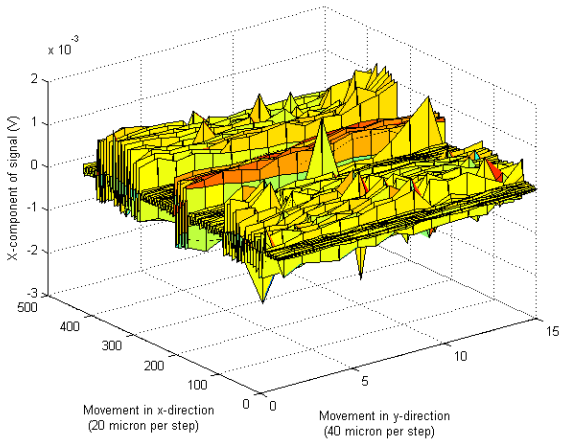

Figure 6.20: A scan covering both fingers and ground line, showing the problem of difference in reflectivity.

The idea was to see how the wave behaves across the entire surface, but it is impossible to read anything from the image due to the reflectivity problem. The detected signal is many times higher when probing a metal surface than the crystal, although the actual surface wave amplitude is the same . A solution to this problem is to normalize the detected signal by dividing it on the total amount of light, which is only dependant on the reflectivity and not the surface acoustic wave. In order to achieve this, a multimeter was included in the measurement process. It was connected to the "Monitor" output of the detector, hoping that it would receive the sum of the signals from all quadrants of the detector. Scans showed, however, that the multimeter received the same signal as the lock-in amplifier. The circuit drawings of the detector circuitry showed that it would be able to retrieve the sum signal. However, since this problem surfaced as late as it did, there was no time to start modifying the detector box. One chose instead to perform measurements scanning the surface twice. Once logging X and Y, and once using the summed signal to log total amount of light. Unfortunately, this did not work either. At arbitrary positions in the scan, the multimeter would log the value zero. The reason for this was not found and continues to be a problem. There were made multiple scans over the same area in order to rule out a physical reason for the faulty value. They proved that there was the surface itself was not responsible. It might be that the programming concerning the multimeter is faulty, as there was no drivers found and all communication and orders had to be relayed using basic commands.

6.5 Absolute amplitude calibration

The heterodyne interferometer reviewed in Chapter 2 is capable of determining the absolute amplitude of a surface acoustic wave. Performing measurements with both the heterodyne setup and the laser probe may produce a scaling factor for the laser probe, thus enabling independent absolute measurements. Both setups made one scan on the same SAW device, with identical settings. The measurements were performed on the ground line in order to avoid noise. In these measurements, the data displayed is in the form of polar coordinates R and θ. R is the amplitude of the wave. A plot of the amplitude from the laser probe measurement is depicted in Fig. 6.21. Observe that the period seems to be 70 μm, despite the frequency being 25 MHz, which should have made the period 140 μm. The reason for this is the rectification of the signal involved in the amplitude calculation, as is apparent in (7). The amplitude unit is Volts, as this is merely the voltage amplitude of the signal.

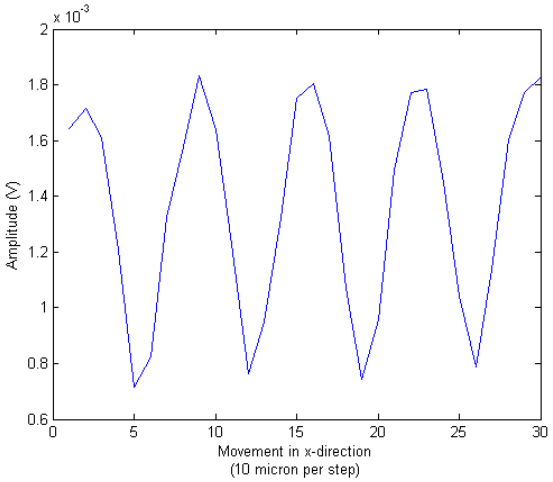

Figure 6.21: The SAW transducer measured upon.

The amplitude unit in the heterodyne scan is nanometer, indicating the exact amplitude of the surface acoustic wave. The scan length was fairly similar, as was the position on the ground line. The plot from the heterodyne measurement is presented in Fig. 6.22.

Finding a relation from these measurements alone is impossible. Simply stating that measuring 1.8 mV using the laser probe setup equals 3.5 nm in absolute amplitude is inaccurate, at best. In this particular case, with this device and these settings, the assumption is correct, but there is no way of knowing if this relation is general. Hence, an extensive study is required in order to map the different variables in the measurement process. The study will have to be made on several devices, with variations in excitation voltage and frequency, this in order to fully map potential influences. The nature of the relation will thus become apparent and

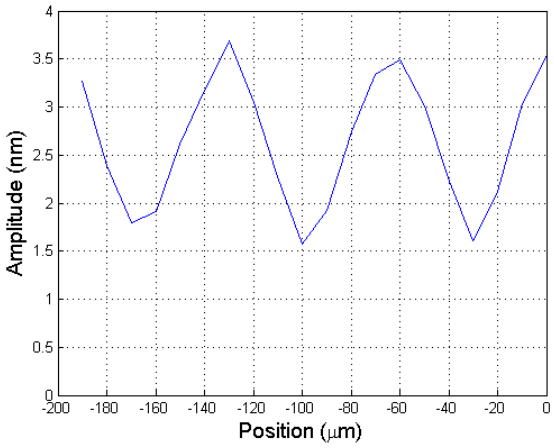

Figure 6.22: Absolute amplitude measurement made on SAW device using heterodyne interferometry.

one finds out if it universally applicable. Should the various dependencies prove nonlinear and strictly material-dependant, an universal relation between the two setups would be difficult, or indeed impossible, to calculate.

7 Reflections

There are many aspects of the setup, programming and results that need to be reviewed, both in view of sources of errors and functionality. These aspects will be dealt with shortly.

7.1 The setup

Sources of errors exist in most measurement setups, and the present setup is no exception. One familiar source is noise, both internal and external. Electrical wiring radiates electromagnetic noise which may influence other wires. This subject was mentioned briefly in Chapter 3. Although most of the problems with synchronous noise disappeared when casing the detection circuitry, some frequency components may still be left to influence slightly, especially when measuring on the crystal itself, where the signal is low. The chosen low frequency component, ω_l, was set to 1025 Hz in order to escape external noise originating from the power grid. This noise is therefore not taken into consideration.

The knife-edge method applied in the laser probe setup is very insensitive to vibrations, compared to other optical characterization tools. The precautions mentioned in Chapter 3 adds to this feature. However, at the time of the measurements, the air-supported table on which the sample is placed was not functional. This may have caused vibrations otherwise damped, to reach the sample, thereby deteriorating the quality of the measurements. There were no evidence, or even indications, of this in the present work, but it should be taken into consideration on further measurements. The surface acoustic waves here measured upon were excited with relatively high voltage, creating a surface wave with a high amplitude. High amplitude leads to a strong signal and a high signal-to-noise ratio, making it fairly impervious to minor disturbances.

The original setup was built in 1977, and it has been modified several times through the years. Both control box and detector box were designed and manufactured locally. The circuitry inside is a collection of relatively old soldered wires and components, which may introduce a source of error in the work. There was no evidence of this being a problem.

Several beam splitters are present in the microscope in the optical setup. They are very useful in diverting a certain amount of light, but they are also a great source of signal loss. Losing about fifty percent of the light at each beam splitter is very much, but as it is, the benefit of using the microscope is larger than the signal weakening it incurs. The microscope objectives were not entirely clean when starting the work, but wiping them proved difficult due to the hard-to-reach lens inside the objective. Possible dirt on the lens may account for some deteriorating effects on the measurements.

A 20X microscope objective was chosen for the measurements in this work, giving a 10-15 μm spot size. It worked very well in the measurements performed on the ground strip, but had some problems when scanning the fingers of the transducer.

First, the difference in reflectivity made it difficult to see a wave pattern in the plotted wave. The data was further deteriorated by the microscope objective picking up all border effects from the interface between metal and crystal. Lowering the magnification to 10X might clean up the signal some. In addition, there is need of normalization, but this was not achieved due to the faulty multimeter or programme.

The $LiNbO_3$ SAW transducer has been the subject of many measurements, and shows signs of this. Being left open to the surroundings has left the surface rather dirty. The spots of dirt influences the reflectivity at that position, producing unexpected measured values.

7.2 Programming

The programming produces few sources of error, as it exercises little physical influence on the signal. However, there are some elements in functionality that are worth mentioning. During some of the measurements, an error message concerning the function "GPIB Read.vi" was displayed. The occurrences of this message were random, and their cause was not found. As the error did not seem to influence the measurement, the programme was allowed to continue. Still, there should not be any errors in the programme, as they question the accuracy of the measurement.

In an attempt to normalize the detected signal, a multimeter was included in the setup. Like the other instruments, this one was connected to the computer through the GPIB interface. When logging the values from this instrument, the multimeter infrequently displays the value zero. This is thought to be a result of some glitch in the programme, although no such glitch has been determined. When communicating with the multimeter, the data is retrieved by using the function "GPIB Read.vi". There might be a problem with the function triggering the signal reading. Should the triggering be delayed for some reason, the measured value will not be stored, but instead write the value zero. The nature of the glitch speaks against such a possibility. The zeroes do not appear at the same place or with equal spacing. In addition, they arrive in pairs, which is found to be strangely organized.

When handling the measured data, MATLAB was found to be a very useful tool. Its range of signal processing features and plotting functions are extensive. As it is now, the analysis are made in MATLAB seperately, on data from the files generated by LabVIEW. There have been made scripts in order to simplify the operations, but they are still time-consuming and requires switching between the two programmes. In LabVIEW, however, exists the possibility of including MATLAB functionality in the measuring programme, but this feature was not explored in this work due to insufficient time. The possibility did not become apparent until the later part of the period, and was not prioritized since it was not strictly necessary for the purpose of the programme. It may prove useful in future variations of the programme.

The programming required introduction to the LabVIEW graphical programming language. Before starting, basic principles were studied using a instructional book written by Finn Haugen[15] and by browsing the National Instrument webpage. Principles were tested using programming examples included in LabVIEW.

This study took some time, much due to the unfamiliar graphical interface of the programme. The study of the programming language and optical detection method caused the initial period of time in the work to be unproductive, but this was not to be avoided.

7.3 The measurements

The measurement procedure includes several delays put in to avoid unnecessary fluctuations in the data and erroneous probing points due to timing difficulties. Using the current settings, measurements take too long, from one to one and a half second per probing point. As its relatively high measurement speed is a major feature of the laser probe, the time per point should be decreased. This may be achieved by experimenting with the settings and stabilizing the setup further. Activating the air-supported table under the sample will perhaps dampen the vibrations caused by the movement of the ramp, shaving off some of the stabilizing delay mentioned in Chapter 5.

Measurements performed in the ground line have proven excellent, with very little noise and other irregularities. The crystal produces far more noise, and the plotted surface is therefore not as smooth as that of the ground line. The wave shape is easily discernable, however, and the data serves it purpose in imaging the wave. In more accurate measurements, like when calibrating for absolute measurements, the toothed surface of the detected signal from the crystal may prove too inaccurate. Decreasing the magnification of the microscope objective may exclude some of the noise factors of the surface.

Fourier analysis proved to function as it was intended, separating and isolating the different wave components. The work also indicates the possibility of suppressing noise components and other undesired elements by handling the signal in k-plane.

8 Conclusions

The work described in this thesis has provided a thorough understanding of the optical method applied and all the elements it comprises. Programming the setup behaviour has been a challenge, finding a solution to missing drivers and other problems. The expertise in LabVIEW derived from this work, and the understanding of characterization principles, will be a valuable asset to bring into professional life.

Through this period, numerous successful measurements were performed. Ground line measurements provided a high signal-to-noise ratio, while scans made on the crystal introduced considerable amount of noise. During these measurements, both setup and programme functioned adequately, without major problems. Later measurements, which crossed the boundaries between metal and crystal were useless due to the large difference in reflectivity. Fixing the problem by normalizing the signal was fruitless due to a multimeter readout error.

The possibility of analyzing the stored data using Fourier analysis has been proven, successfully separating the propagation directions of two waves of different nature. The method professes that further manipulation is possible, like suppressing noise.

An amplitude measurement on a heterodyne interferometer provided absolute amplitude data. The plot was compared to a plot from a similar measurement on the laser probe setup, thus indicating the manner in which calibration of the laser probe will transpire. The goal is to be able to make absolute amplitude measurements on the laser probe alone.

Further work should include improving the measurement settings in order to improve the signal-to-noise ratio of the crystal measurements. Accurate settings will also help decreasing the time each probing point takes.

In order to calibrate the laser probe against the heterodyne interferometer, numerous scans must be made on each setup. The scans should be made on several SAW devices, with different excitation voltages and frequency.

MATLAB functionality should be included in the LavVIEW programme, in order to simplify the handling of the stored data.

References

[1] George I. Stegeman. Optical probing of surface waves and surface wave devices. *IEEE Transactions on Sonics ans Ultrasonics*, SU-23(1):33–63, january 1976.

[2] Richard M. White. Surface elastic waves. *Proceedings of the IEEE*, 58(8): 1238–1276, August 1970.

[3] R. Koch et al. Nanoscale imaging of surface acoustic waves by scanning electron microscopy. *Journal of Applied Physics*, 2005.

[4] Astrid Aksnes et al. Characterization of acoustic vibrations on micro- and nanostructures with picometer sensitivity. *SPIE Optics and Photonics, Interferometry XIII:Applications*, 2006.

[5] Hanne Martinussen et al. Heterodyne interferometry for high sensitivity absolute amplitude vibrational measurements. *Interferometry XIII: Techniques and analysis, Proc. SPIE 6293*, 2006.

[6] Jean-Pierre Monchalin. Optical detection of ultrasound. *IEEE Transactions of Ultrasonics, Ferroelectrics, and Frequency Control*, UFFC-33(5):485–499, september 1986.

[7] Helge Engan. Phase sensitive laser probe for high-frequency surface acoustic wave measurement. *IEEE Transactions on Sonics and Ultrasonics*, SU-25(6): 372–377, november 1978.

[8] Robert Adler et al. An instrument for making surface waves visible. *IEEE Transactions on Sonics and Ultrasonics*, SU-15(3):157–159, july 1968.

[9] Hiroshi Kamizuma et al. Development of fast-scanning laser probe system based on knife-edge method for diagnosis of rf surface acoustic wave devices. *IEEE Ultrasonics Symposium*, pages 1604–1609, 2005.

[10] Richard M. De La Rue. Heterodyne optical probing of surface acoustic waves in a partial standing wave situation. *IEEE Transactions on Sonics and Ultrasonics*, SU-24(6):407–411, November 1977.

[11] PerkinElmer Instruments. *What is a Lock-in Amplifier?*, 2000.

[12] *The Benefits of DSPLock-in Amplifiers.*

[13] National instruments webpage. www.ni.com.

[14] Solartron Instruments. *7150 Remote Control Operating Manual.*

[15] Finn Haugen. *Lær LabVIEW.* TehTeach, 1 edition, 2001. ISBN 82-91748-06-3.

A LabVIEW code from from "measure.vi"

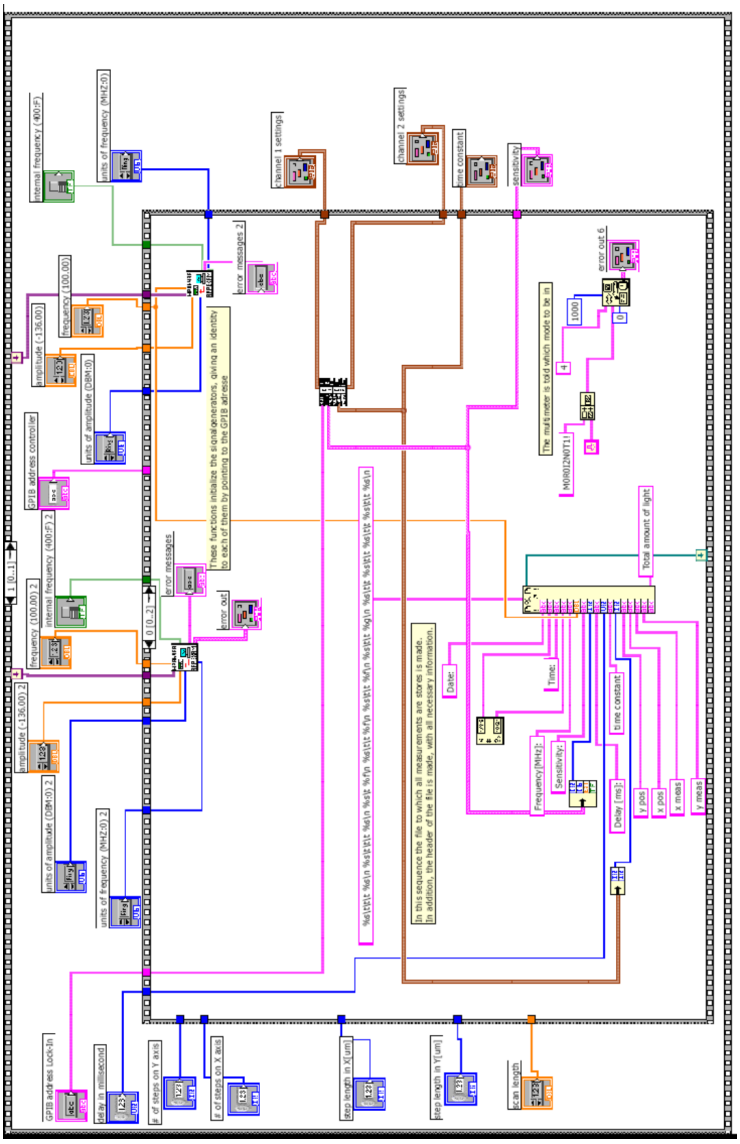

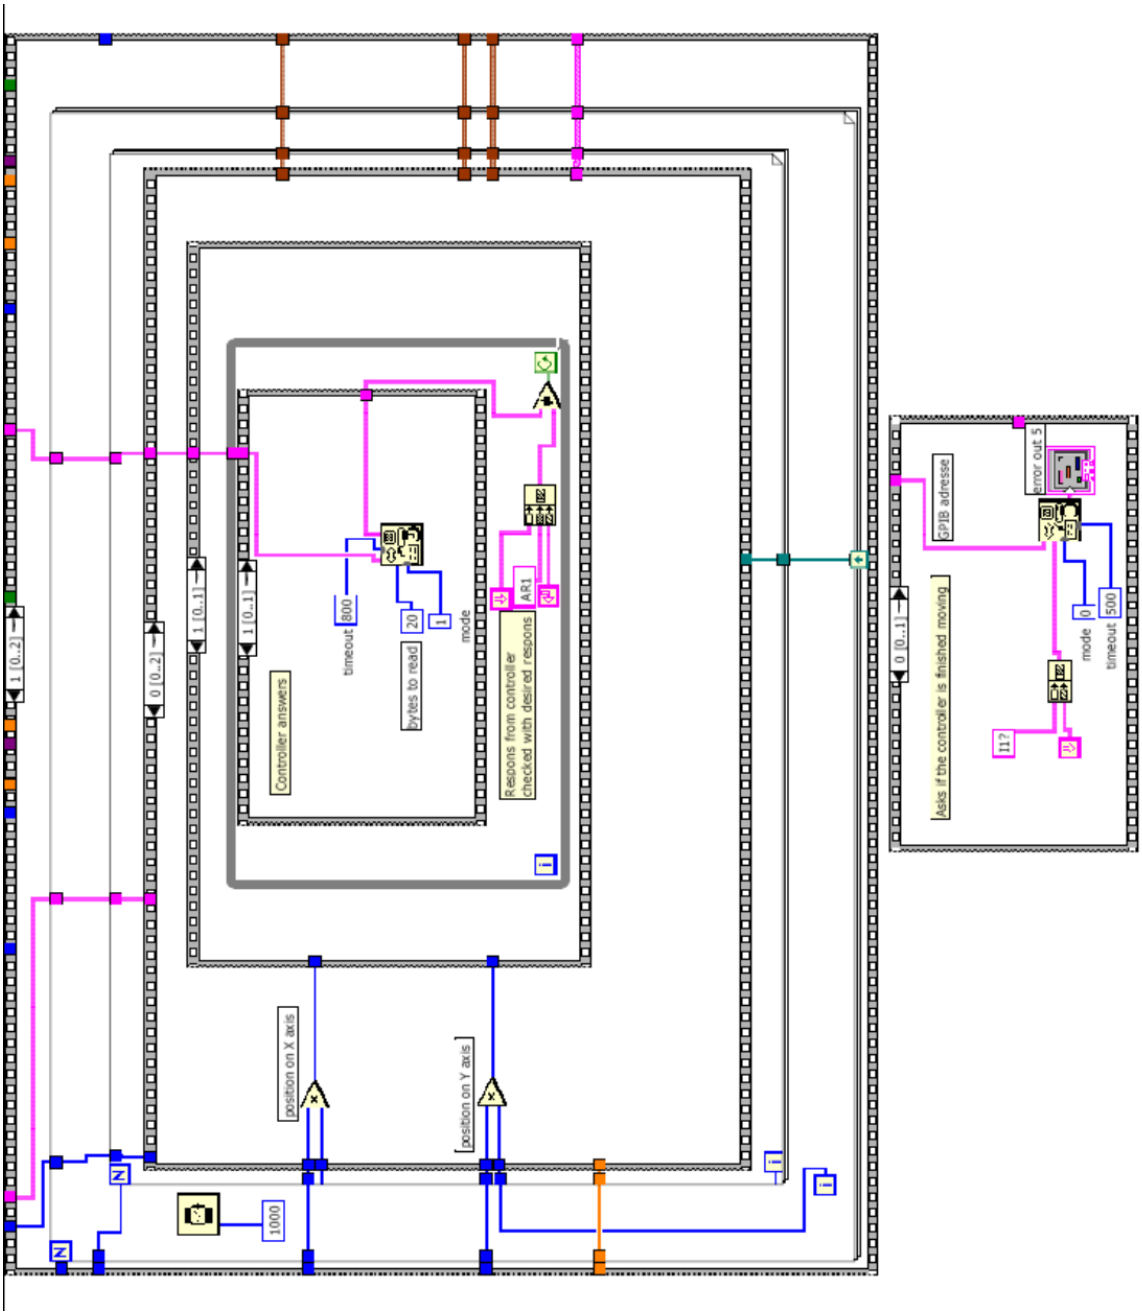
1 [0..2]
0 [0..2]
1 [0..1]
1 [0..1]
timeout
800
20
bytes to read
1
mode
AR1
Controller answers
Respons from controller checked with desired respons
position on X axis
position on Y axis
1000
0 [0..1]
GPIB adresse
error out 5
Asks if the controller is finished moving
mode
timeout
500
0
117

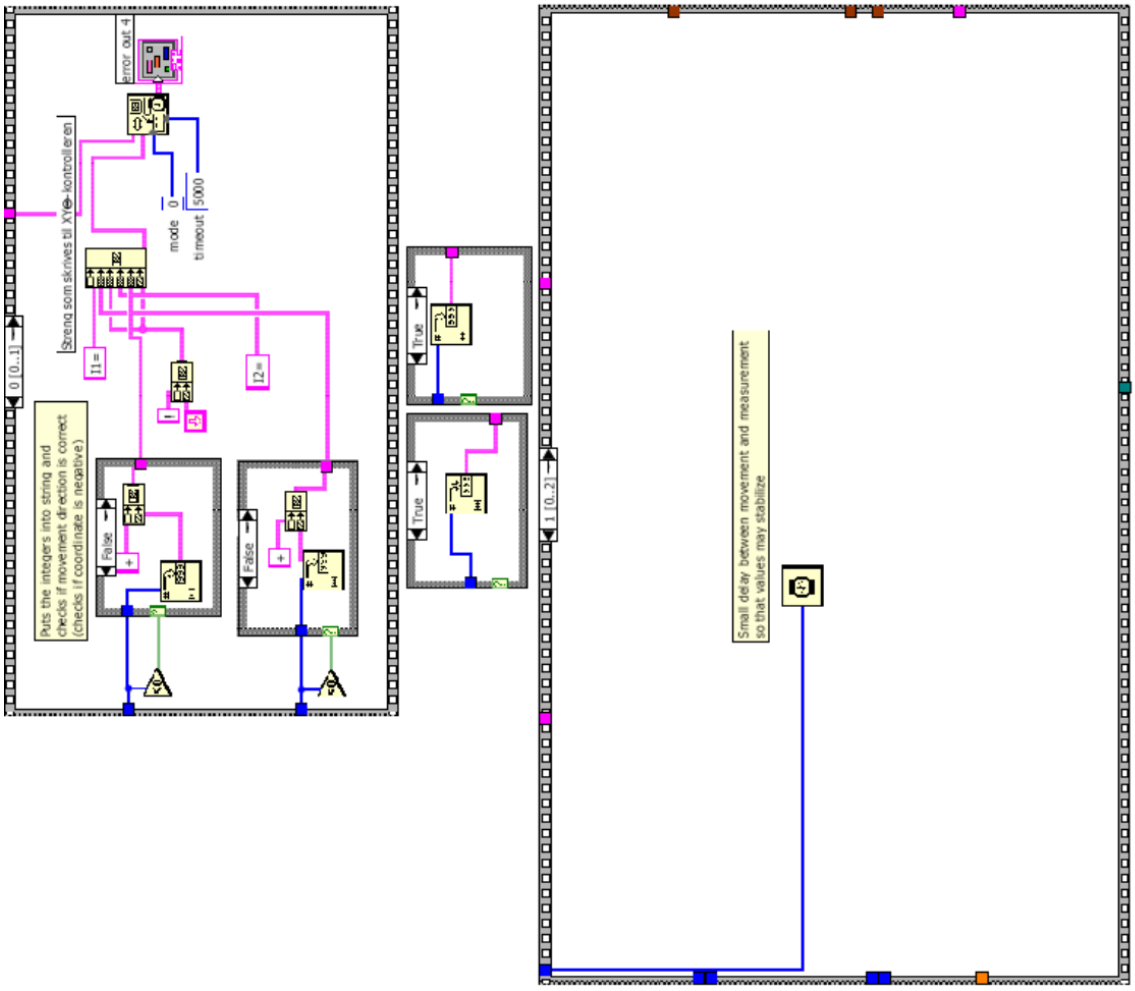
error out 4
Streng som skrives til XYE-kontrolleren
mode
timeout
5000
I1=
I2=
0 [0..1]
Puts the integers into string and
checks if movement direction is correct
(checks if coordinate is negative)
False
False
True
True
1 [0..2]
Small delay between movement and measurement
so that values may stabilize

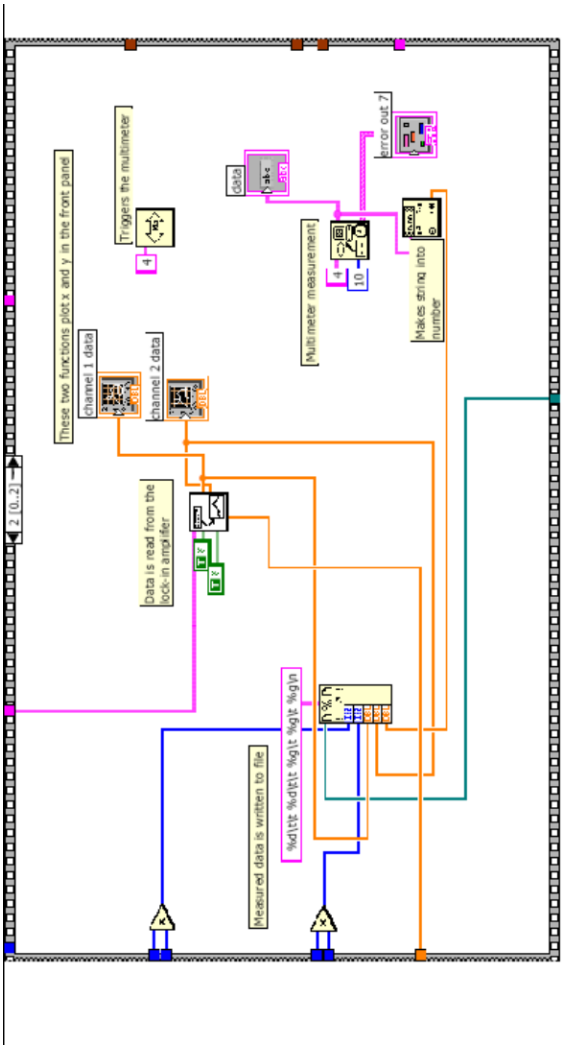
These two functions plot x and y in the front panel
Triggers the multimeter
data
error out ?
Multimeter measurement
Makes string into number
channel 1 data
channel 2 data
Data is read from the lock-in amplifier
Measured data is written to file
%d\t\t%d\t\t%g\t %g\n

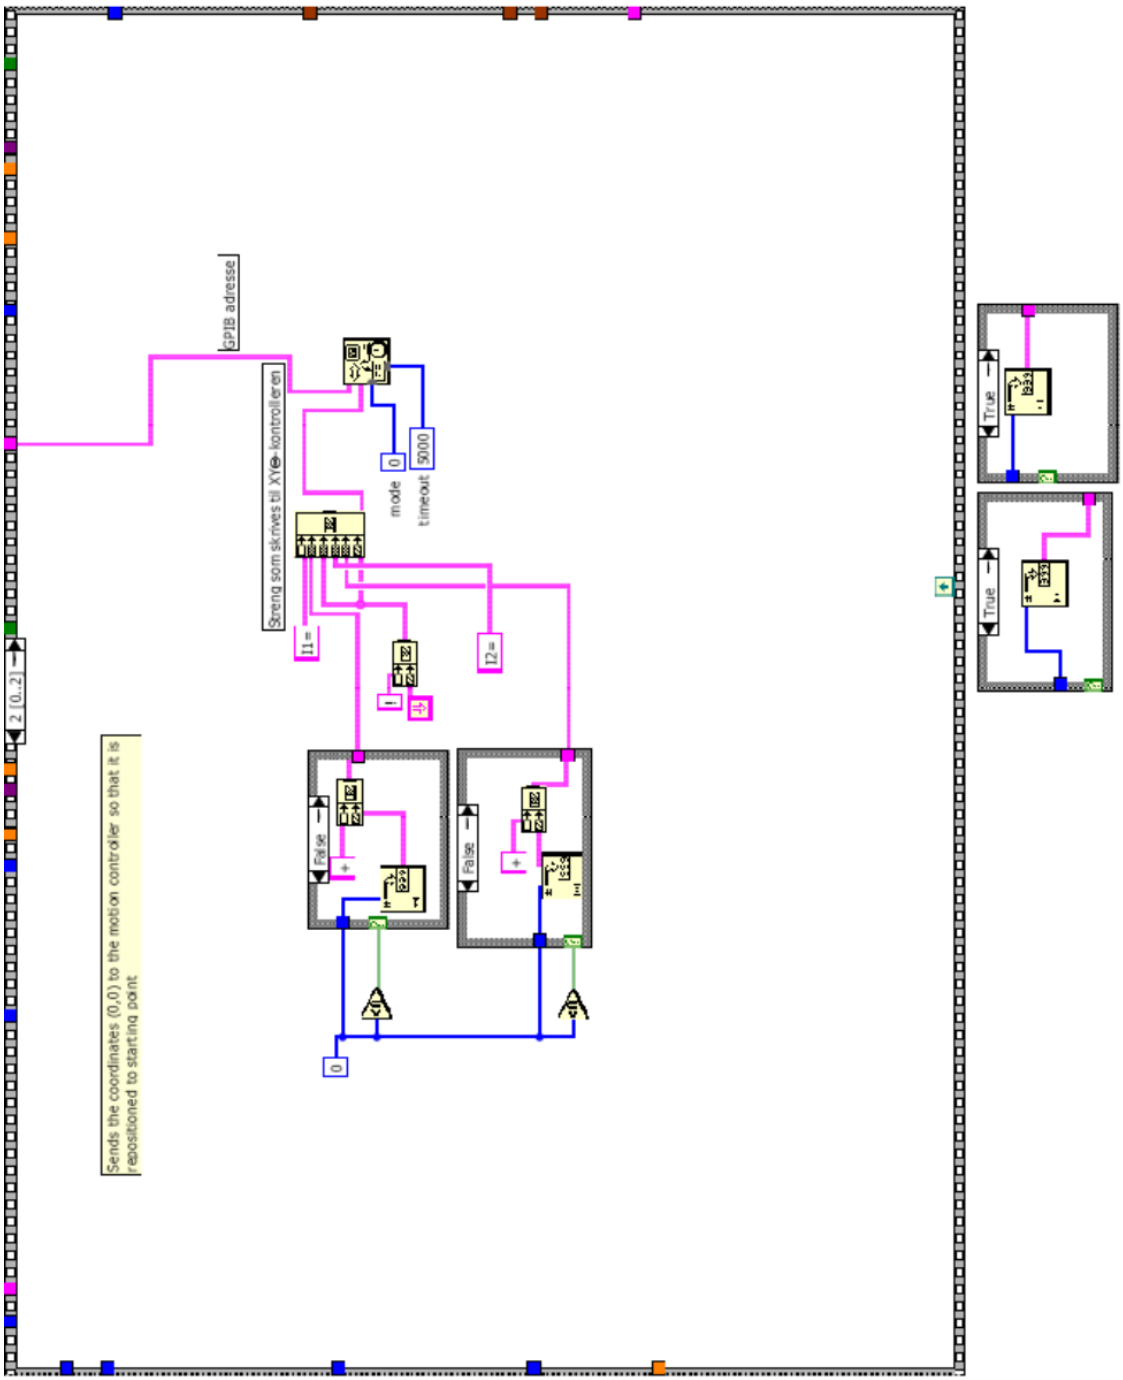
GPIB adresse
Sreng som skrives til XYe- kontrolleren
mode
timeout
0
5000
I1=
I2=
2 [0..2]
Sends the coordinates (0,0) to the motion controller so that it is repositioned to starting point
False
False
True
True

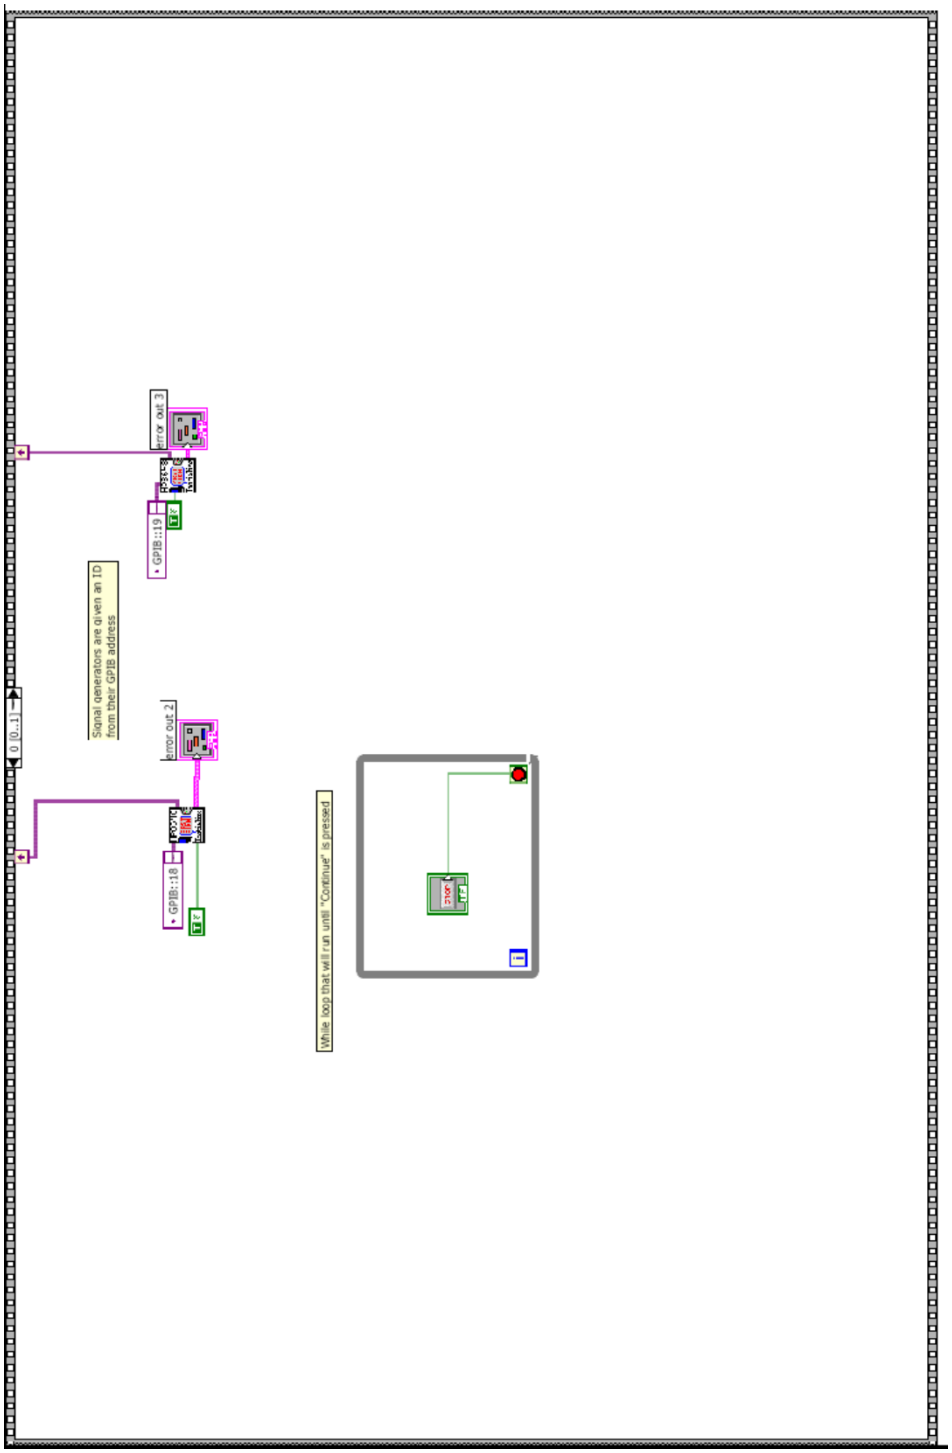
error out 3
GPIB::19
Signal generators are given an ID from their GPIB address
error out 2
GPIB::18
While loop that will run until "Continue" is pressed.
STOP

B LabVIEW code from "start.vi"

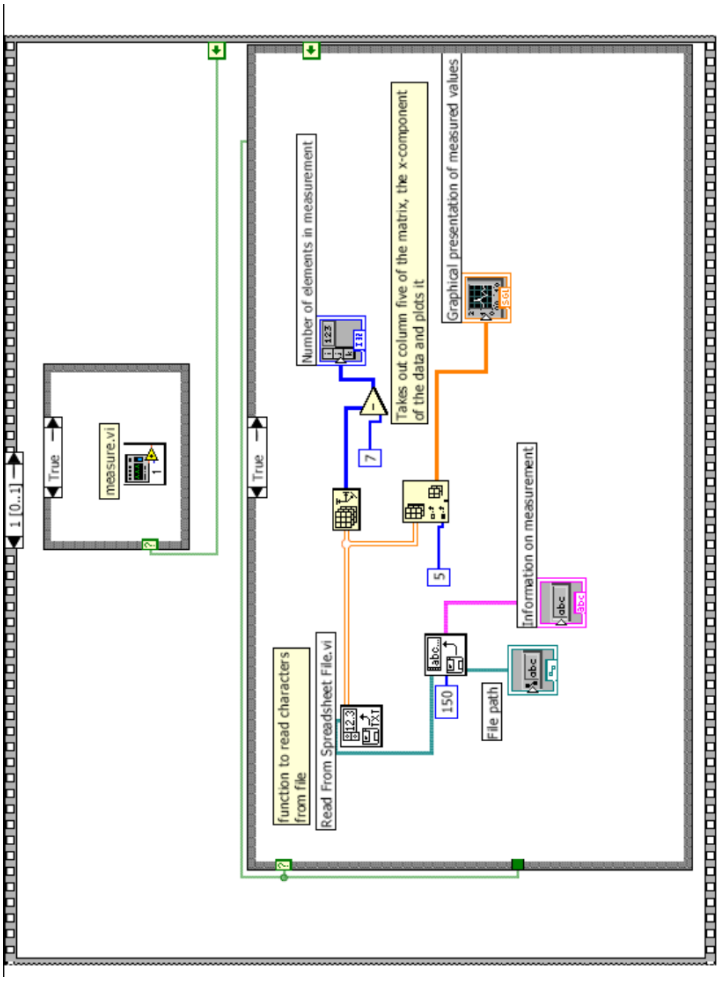

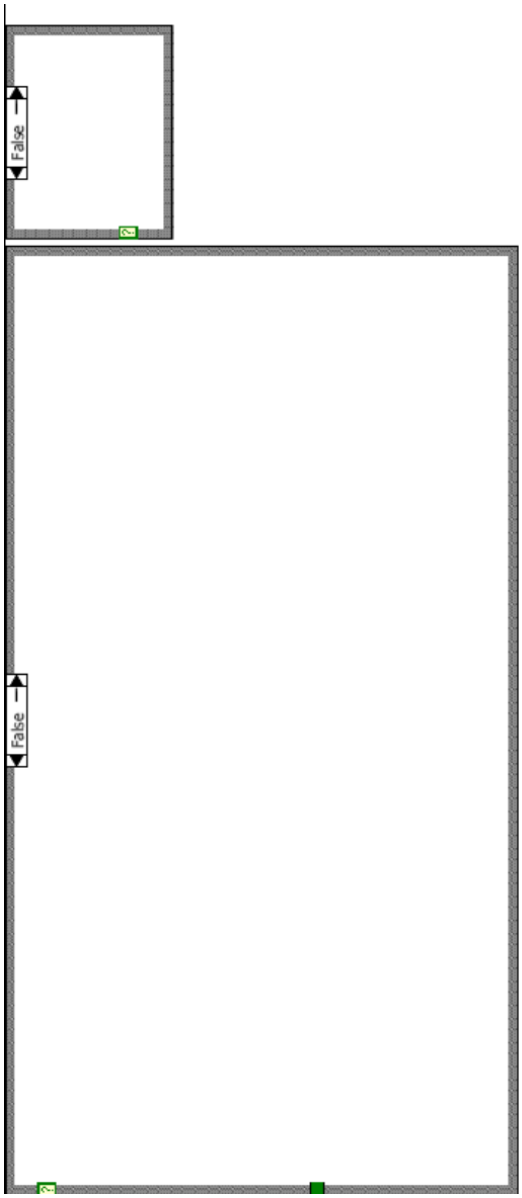
False
False

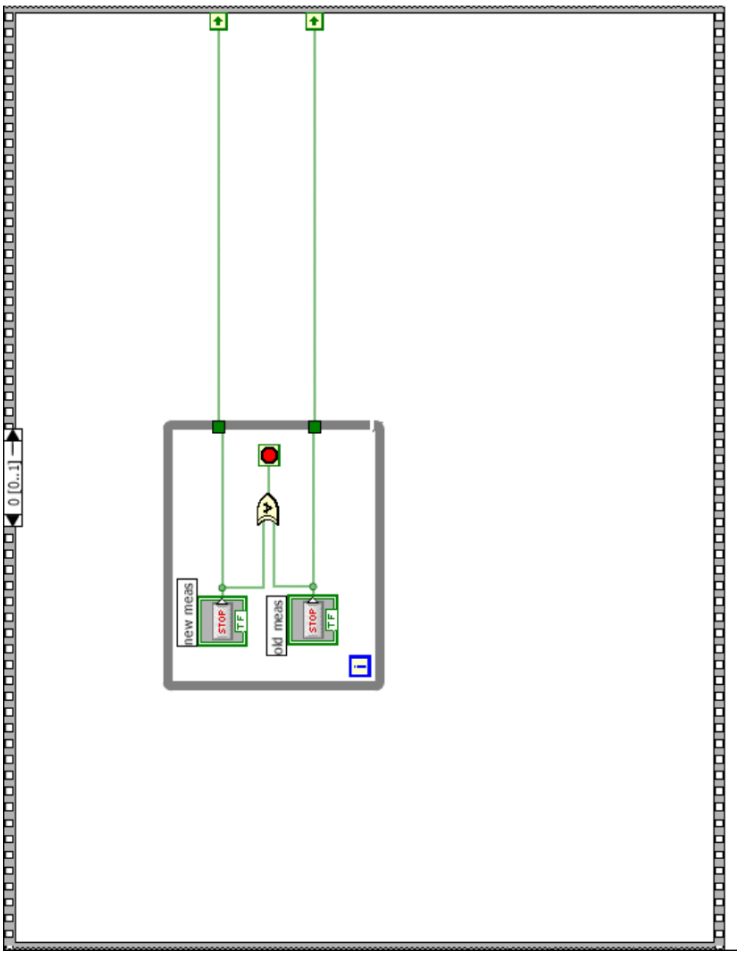
i
0
new meas
STOP
TF
old meas
STOP
TF

C MATLAB code

```
%This function takes in the two vectors vect1 and vect2 and transforms
%them into matrixes with dimensions specified in the input values n and m.

function [res1, res2] = skit(vect1, vect2,n, m)

matr1 = zeros(n,m);
for i = 1:1:m;
matr1(:,i) = vect1((i-1)*n+1:((i-1)*n)+n);
end

res1= matr1;


matr2 = zeros(n,m);
for i = 1:1:m;
matr2(:,i) = vect2((i-1)*n+1:((i-1)*n)+n);
end

res2= matr2;
```

```
%This function seperates the two propagation directions of
%the wave, displaying both X component and phase before and after
%


function varmt(xmat,ymat,n,m)

figure,
surf(xmat)%diplays matrix prior to manipulation

compmat = complex(xmat,ymat);%makes a complex matrix on the form A + iB
phase = angle(compmat);
figure,
surf(phase)

fourmat = fft2(compmat);%2D Fast Fourier Transform of complex matrix
figure,
surf(abs(fourmat))%displays wave in k-space

fourmat = fftshift(fourmat);%changes 1. and 3. quadrant, 2. and 4. for simp

figure,
surf(abs(fourmat))%displays wave in k-space shifted

figure,
contour(abs(fourmat))%displays contour plot of wave k-space prior to filter
fourmat2 = fourmat;%fourmat 2 is to become the wave component propagating b

%here follows the filtering process, values in upper half of x axis
%represents the forward propagating wave
for i = 1:1:n
    for u = 1:1:m

        if i < (n/2)
            fourmat(i,u) = 0;
        end
    end
end

for i = 1:1:n
    for u = 1:1:m

        if i > (n/2)
            fourmat2(i,u) = 0;
        end
    end
end

figure,
contour(abs(fourmat))%contour plot of positive k-component

fourmat = ifftshift(fourmat);%matices are inverse shifted
fourmat2 = ifftshift(fourmat2);

compmat = ifft2(fourmat);%matrices are inverse Fast Fourier Transformed
compmat2 = ifft2(fourmat2);
```

```
phase = angle(compmat);%phase of positive wave component
figure,
surf(phase)

phase2 = angle(compmat2);%phase of negative wave component
figure,
surf(phase2)
figure,
surf(real(compmat))
%displays real part of complex matrix (X)
%in positive propagation direction

figure,
surf(real(compmat2))%same for negative diretion
```