Thomas Kunzfeld

Development and Implementation of a Real-Time Docking Guidance System for an Autonomous Electric Vehicle Charger

Master's Thesis

Graz University of Technology

Institute of Electronics
Head: Univ.-Prof. Dipl-Ing. Dr.techn. Bernd Deutschmann

Supervisor: Univ.-Prof. Dipl-Ing. Dr.techn. Bernd Deutschmann

Graz, Mai 2019

Statutory Declaration

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources.

Graz, _____ _____

Date Signature

Eidesstattliche Erklärung[1]

Ich erkläre an Eides statt, dass ich die vorliegende Arbeit selbstständig verfasst, andere als die angegebenen Quellen/Hilfsmittel nicht benutzt, und die den benutzten Quellen wörtlich und inhaltlich entnommenen Stellen als solche kenntlich gemacht habe.

Graz, am _____ _____

Datum Unterschrift

[1] Beschluss der Curricula-Kommission für Bachelor-, Master- und Diplomstudien vom 10.11.2008; Genehmigung des Senates am 1.12.2008

Abstract

This thesis describes the development of an ultrasonic position measurement system intended for an autonomous charging system for electric vehicles developed by the company Volterio. It explains the functionality of the critical components and explores various alternatives to stated challenges in order to provide the possibility to modify the resulting system to suite new requirements. This is because this thesis should also serve as a reference for the developed position measurement system and the development process for Volterio has shown that requirements often change and new approaches to already solved problems need to be found quickly when the need arises. Due to the highly fluctuating requirements a broadly usable system has been developed, which can be used with two different types of sensors. One of these types of sensors is clearly favourable over the other in terms of signal quality, but could never be used in the actual prototype of the charging system successfully due to mechanical constraints. These constraints could be overcome without too much difficulty, which is why a separate measurement setup has been created using the better suited type of sensor. The results show that while the developed system is usable, there is still room for improvement in terms of accuracy. At the end of this thesis, ideas for improving it are provided.

Keywords: Position measurement, ultrasound, robotics, piezoelectric transducers

Kurzfassung

Diese Arbeit beschreibt die Entwicklung eines ultraschallbasierten Positionsbestimmungssystems für das automatisierte Ladesystem für elektrische Autos der Firma Volterio. Es wird die Funktionalität aller Komponenten beschrieben und es werden alternative Ansätze zu diversen Herausforderungen präsentiert. Diese Alternativen sollen als Anhaltspunkte dienen falls das resultierende System in Zukunft modifiziert werden muss, um neuen Anforderungen zu genügen. Der Entwicklungsprozess für Volterio hat gezeigt, dass sich Anforderungen dort sehr schnell ändern können und oft neue Ansätze für bereits gelöste Probleme gefunden werden müssen. Bedingt durch diese stark fluktuierenden Anforderungen wurde ein System entwickelt, dass in diesem Kontext flexibel einsetzbar ist. Es können zwei verschiedene Arten von Sensoren verwendet werden, wobei eine davon von einem messtechnischen Standpunkt klar zu bevorzugen ist. Diese Sensoren konnten aber nie im Prototypen des Systems verbaut werden, bedingt durch mechanische Herausforderungen, welche aber leicht überkommen werden könnten. Deshalb wurde eine eigene Messung mit den besser geeigneten Sensoren durchgeführt. Die Ergebnisse beider Messungen zeigen, dass das System zwar den Anforderungen genügt, aber noch Raum für Verbesserungen bietet. Deshalb wird am Ende dieser Arbeit ein kurzer Überblick über mögliche Verbesserungen geboten.

Contents

List of Figures

1 Introduction

1.1 Motivation

One often overlooked advantage of electric vehicles over fuel based ones is the number of different charging possibilities. Filling up the tank of a fuel based vehicle is a rather delicate process where flammable fuel could be spilled, making automation of this process uninteresting. Electric vehicles on the other hand can be charged by a variety of different ways, providing a lot of freedom in the design of automated charging systems.

Most current automatic chargers for electric vehicles use electromagnetic induction as a way to transfer energy to the vehicle. This approach has several drawbacks, including the poor efficiency and the weight of the coils on the vehicle. A positive aspect of these systems is that the relative position of the receiving coils on the car to the transmitting coils in the charging station is not very critical. A sub optimal configuration simply reduces the efficiency even further, but charging is still possible.

This is not the case for conduction based systems. They provide a much higher efficiency and can reduce the weight of the receiving side drastically, but they rely on a high precision docking mechanism. At the time of writing, no automatic conductive charging system is available on the end user market. Systems that do exist in research facilities use industrial style robotic arms to mimic human motions and rely on complex signal processing and kinematic algorithms to perform their task. The advantage of the Volterio electric vehicle charging system lies in its use of a proprietary connector. This connector is conically shaped, which allows for rotational freedom and mechanical self alignment. This allows for a very simple robot which performs the docking procedure and for cheap and not too complex position measurement, which will be developed in the scope of this thesis.

Using commercially available products for position measurement was not in the interest of Volterio when designing the system, since they tend to

be costly and more importantly, they need to be fitted to the robot which performs the docking process. This robot and the docking process and environment impose some strict requirements onto the guidance system. Therefore, it has been decided that it would be cheaper and easier to develop a system from scratch, since the requirements can be already considered in the very first stages of the design.

The design of the electronic guidance and position measurement system is the focus of this thesis. At the beginning, a short outline of possible approaches will be presented. Then, the implementation of the most promising approach will be discussed in detail and measurement results will be presented.

1.2 Basics of the Volterio Charging System

As already described in the previous section, the Volterio charging system is based around the proprietary plug design. This plug allows for full rotation and a displacement of about 2.5cm in each direction during the plugging procedure. The system consists of a vehicle unit, which holds the female part of the plug, and the base unit or robot which moves the male part of the plug. This robot has three degrees of freedom, while the vehicle unit is static, with the exception of a sliding cover which protects the plug from dirt. The three degrees of movement consist of the rotation of the complete base, the lifting motion of the arm of the robot which holds the plug and the linear movement of this arm along the long side of the base. A rendering of the system in the latest available revision at the time of writing can be found below in figure 1.1

The position measurement system can be mounted anywhere on the robot, however, it is desirable to mount it on the base since there is enough space available. The plug can also be altered slightly to accommodate sensors if absolutely necessary. Mounting the sensors on the base has the additional advantage that the system could be used to give feedback to the driver while they are parking so they can steer in the right direction and come to a stop inside the range of the robot. However, the range of the robot is rather large, making this a comfort feature rather than an absolutely necessary one.

Figure 1.1: Rendering of the Volterio robot and vehicle unit. [12]

1.3 Requirements and Aims

Apart from precision and accuracy, the position measurement system needs to fulfil some additional requirements. Firstly, it needs to be robust against influences from the environment, such as moisture and dirt. One part of the system is likely to be mounted on the underside of the vehicle, and therefore can get at least partially covered in mud or other substances. Also, dirt, snow, leaves etc. can be carried on the underside of the vehicle into the garage or parking space where the charging system is located, dropping some of it onto the sensors of the guidance system. This will deteriorate the performance of almost every possible system without extensive counter measures. However, some systems will fail more easily than others and the aim of this thesis is to develop one which offers a good trade off between complexity and robustness.

Additionally, the whole charging system will need to be able to uniquely and securely identify vehicles, for example for billing. Therefore, some sort of active component will be needed on the vehicle. If the position measurement system also needs an active component on the vehicle, it would be highly beneficial if these two could be combined on a single

module. Such an active component would need to run on separate batteries for a very long time, even if the battery of the vehicle is completely drained in order to still allow charging. This requirement makes passive positioning systems or systems that receive their energy wirelessly from the base station favourable. RFID or induction based systems would fall into this category. Even though the standby time needs to be long, the energy consumption requirements are not quite as strict when it comes to the actual measurements, since it is very likely that the system will be supplied with energy in a short time. However, the power consumption should generally be as low as possible.

If the system is not entirely passive, there needs to be a way to trigger the system when the vehicle comes near the charging port. This could be done either by using some channel provided by the measurement system itself or by using external communication hardware. If additional hardware is used, it would again be favourable to combine this system and the system needed for identification. Some near field RF communication system such as Wifi or Bluetooth would be the obvious choices for this. An additional advantage of such a system would be that it could also be used to transmit the measured positions from the robot to the vehicle, in order to provide some feedback to the driver (or the vehicle itself in case of a self driving one). This feedback is not absolutely necessary, but it should be supported if there is any way to achieve it with the developed system.

In order to provide feedback while the vehicle is approaching, fast measurements are needed. However, these measurements do not need to be as exact as the measurement of the position of the plug while docking. Therefore it could be advantageous to for example use separate systems for docking and feedback, or use different algorithms and calculations for the two phases.

Sensors of the system can either be mounted on the base of the robot or directly on the plug. Both of these solutions have their respective advantages and disadvantages, which is why it should also be considered to mount sensors on both of these locations. Very shallow measurement angles are to be expected, since the whole system should be able to connect to vehicles with an underfloor height of just 14 cm.

For the feedback while the vehicle is approaching, an accuracy of smaller

than 5 cm should be achieved. The measurement for the connector position for docking should of course be more accurate, in the range of about 1 cm . Deviations in this range can be corrected by the mechanical design of the connector. Theoretically it should be able to correct for translations of up to 2.5 cm as stated before, but experiments showed that at the current state of the connector, this was not reliably possible.

1.4 Limitations

As already mentioned, some unique challenges arise when designing a localization system for this specific purpose. Both the environment and the requirements given by the design of the robot have to be considered. The biggest factor coming from the environment is soiling of sensors and all parts of the system. For every measuring method, there is some material or effect coming from the environment that will deteriorate the accuracy of the measurement or prohibit measurements all together.

Docking and therefore position measurement will be performed on the underside of the vehicle and the robot covers a large area where it is able to dock to the vehicle. Because of this, the sensors for the positioning system will encounter a wide range of measurement angles. The designed system needs to be able to handle this by either employing redundant sensors at different positions on the robot to reduce the impact of varying angles or by being resistant to these angle variations in the first place.

Additionally, there are a lot of metal parts on a vehicle, especially on the underside, which will have an influence on some electromagnetics based systems. In the case of electric vehicles there are also high currents and strong magnets in the powertrain, which could also have an effect.

The underside of the vehicle together with the concrete or asphalt ground on which the charging robot is mounted form two hard, parallel planes. This has to be considered in sonic systems, such as ultrasound based ones since these planes will provide excellent reflection surfaces. Sound waves will bounce in a non deterministic way because of the irregular structure of the vehicles underside and only be dampened slowly. Such a system would need to be aware of theses reflections and either have a way of

detecting and filtering them or work in a way that they do not influence the measurement, for example by only sending the next pulse after the reflections of the previous one have a negligible amplitude.

1.5 Structure of this Thesis

During the writing of this thesis, the requirements for the positioning system frequently changed, often dramatically. This is why throughout the following chapters, multiple solutions to the same problems will be presented. Most of them were developed for different requirements, but they are still noteworthy and could be used in the future, should the requirements change again.

The first chapter of this thesis will briefly present some possible solutions and describe problems with some of them, which make them unsuitable for these requirements. One suitable approach is then selected and presented in detail.

The next chapter will then present the basic components of the system, along with several variations and their respective advantages and drawbacks.

Following this is a chapter about the properties of the signal which is generated by the sensors and a mathematical approximation. This approximation is used to generate test data for evaluating signal processing algorithms, which are being presented in chapter 5.

In order to calculate a position, the geometry of the system needs to be described. This is done in the next chapter, which presents two basic approaches of calculating a position with the selected sensors, and argues for the advantages of one of them over the other in this application.

The second to last chapter will then tie everything together and present the actual implementation of the developed system, which has been integrated into the robot. The measurement results will be evaluated and in the final chapter and possible improvements will be presented.

2 Possible Solutions and Their Properties

There are multiple possibilities for measuring a three dimensional position in the required range. In order to identify the solution best suited for this application, an initial survey has been performed. Three basic approaches were found and are briefly discussed below.

2.1 Optical Approaches

Using an optics based solution seemed like the most straight forward choice on first glance. On the underside of the vehicle, optical markers could be attached which could be detected by a wide angle camera in the base of the robot. These markers could either be passive, like QR Codes for example, or active, like LEDs. This would allow for two dimensional measurements, which would probably suffice for this application. A second camera could be added for three dimensional measurements. However, this approach has some obvious drawbacks. As soon as the underside of the car gets dirty, the performance of the system would quickly deteriorate. Passive markers would be undetectable under even thin layers of dirt and even LEDs would only continue to be detectable under thin layers. Adding a lot of redundant markers would reduce this risk, but never make it small enough. The optics of the cameras themselves could also easily be covered by dirt, snow or even just moisture. Additionally, image processing requires substantial amounts of processing power, which would increase the size, cost and power requirements for such a system compared to other solutions. All of these drawbacks led to the conclusion that an optical system is not a good choice for this application.

2.2 Electromagnetic Approaches

Another possibility is to utilize coils or antennas to receive electromagnetic waves. There are different approaches of estimating the position of the transmitter when multiple receivers are in use. Measuring time of flight is the typical approach when measuring distances in the range of kilometres or tens of meters, and even there the accuracy of the measuring system needs to be very high. For calculating the position of the car in this context, distance measurements with millimetre resolution and high accuracy need to be performed, which would not be feasible with this approach. Expensive components and complex design would be necessary, which is not desirable for this application.

A slightly simpler approach would be to only measure time differences instead of times of flight. This could be reduced to measuring phase shifts, which relaxes the requirements slightly. Still, the system would be rather complex, which should be avoided if possible.

Another approach which was quickly discarded is using a transmitting coil and multiple receiving coils to get a direction indication via received signal strength. The receiving coil would create an alternating electric field which would induce a current in the receiving coils. The voltage drops created by this induced current in all receiving coils could then be compared and coils further away from the transmitter would have a lower voltage drop. Then, by steering the robot in a way to minimize the difference in voltages, the target would be found. However, the environment in which this system is to be deployed is one of many electromagnetic disturbances like fields created by strong currents in electric motors. Also, induced currents would vary depending on the position of the coils since a lot of metal is to be expected on the underside of a car, changing the characteristics of each coil depending on its location. This approach would also demand sensors directly on the connector, which should be avoided.

RFID based positioning would have also been a possibility. However, this technology is still in active development and commercially available receivers are expensive and do not offer the desired resolution.

2.3 Ultrasonic Approaches

As was already stated, positioning on bigger scales is almost always done by measuring the time of flight of electromagnetic waves. This is infeasible on this scale, but what can be measured rather easily is the time of flight of slower waves, like sound waves. Then, many principles of positioning using electromagnetic waves can be applied to this approach. Additionally, not a lot of disturbances are to be expected when using ultrasound. This approach has been chosen to be explored further and has proven to be adequate for this problem.

3 Components of the Ultrasound Based System

3.1 Revised Requirements

The requirements stated in the introduction where general requirements for the positioning system. This section provides more detailed requirements for the implementation with ultrasonic sensors.

Since very shallow angles are to be expected, also in combination with large distances while the vehicle is approaching it has been deemed acceptable to use separate sensor sets for the approaching phase and the docking phase. One of these sensor sets can be angled towards the vehicle, therefore improving the angle between transmitter and receiver for increased signal amplitude.

The system needs to be able to perform measurements at a reasonably high rate for the feedback. The feedback does not need to be as accurate as the measurement for the docking, it would therefore be acceptable to have two different algorithms, where one provides fast, but less accurate measurements and the other takes longer, but produces a very accurate measurement. “Reasonably high” in this case was defined to be more than ten updates per second, ideally twenty. This way, together with interpolation and filtering, the measurement appears smooth to the driver.

If possible all sensors should be mounted on the base of the robot and not on the connector, which would simplify the cable management and reduce the number of printed circuit boards in the robot.

3.2 Ultrasonic Transducers

Ultrasonic transducers are devices which use a piezoelectric material to emit and/or receive sound waves with frequencies in the ultrasonic band. A transducer can either be used as a transmitter or as a receiver. However, there are specialized transducers which are designed to work as high power transmitters or as sensitive receivers. The main component of an ultrasonic transducer is a disk of piezoelectric material, whose thickness d determines the frequency of the emitted wave f_0 with the dependency $f_0 = (2n+1)\frac{c}{2d}$, with $n = 0, 1, 2...$ giving the harmonic frequencies and where c denotes the speed of sound inside the material. [7, p. 9]. Most ultrasonic transducers are based on the piston-style model, where backing material, the piezoelectric crystal and matching layers are stacked in a cylindrical form. A simplified representation of the construction of a typical enclosed ultrasonic transducer can be found in figure 3.1.

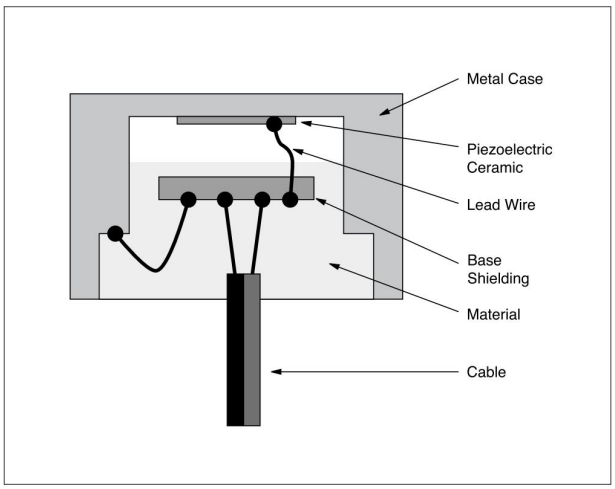

Figure 3.1: Construction of a typical enclosed ultrasonic transducer [8, p. 6]

The physical shape of the piezoelectric material and possible backing- and matching layers determine the angle dependent radiation and reception characteristics of a transducer, also called directivity. For circular piston-type transducers the pressure distribution is rotationally symmetric along the transmission axis. By altering the geometry, non-symmetrical distributions can be achieved. The angle of divergence of the main lope of pressure emitted is determined mostly by the factor $\frac{D}{\lambda}$, where D is the diameter

of the disc and where λ denotes the wavelength of the centre frequency of the transducer. Narrow beams are achieved by selecting $D >> \lambda$. [9, p. 247] A typical directivity pattern can be found in figure 3.2. Note that this transducer displays a non rotational symmetrical pattern, which is used in this thesis to limit disturbances to adjacent systems by directing most of the acoustic pressure along the underside of the car.

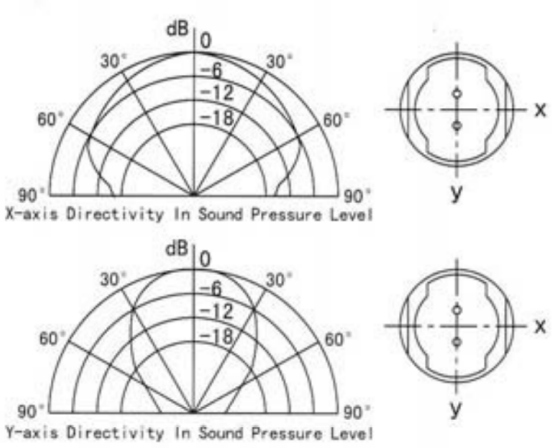

Figure 3.2: Directivity of the A-14P20 ultrasonic transducer. [1]

Also note that directivity patterns as shown in figure 3.2 are only true for the far field of the transmitter. The near field consists of many local minima and maxima in pressure. This effect happens because of the interference between all emitting elements dS of the surface of the transducer. For an approximation of the position of the boundary between near and far field a simplified plane piston model of a transducer can be considered. The place of the last maximum in pressure marks the beginning of the far field and it lies at $x_M = \frac{D^2-\lambda^2}{4\lambda}$ along the main axis. Of course real transducers do not follow this model exactly, however, for distances $x >> x_M$, only the far field needs to be considered. [9, p. 246-247]

In order to analyse piezoelectric transducers in the electronic domain, an equivalent circuit model is needed. A transducer can be represented as a series resonator with a parallel capacitance, which models the coupling of the electromechanical properties of the piezoelectric material [7, p. 9], see figure 3.3. [11] offers a detailed description of the components in the equivalent circuit and how they depend on the geometry of the sensor.

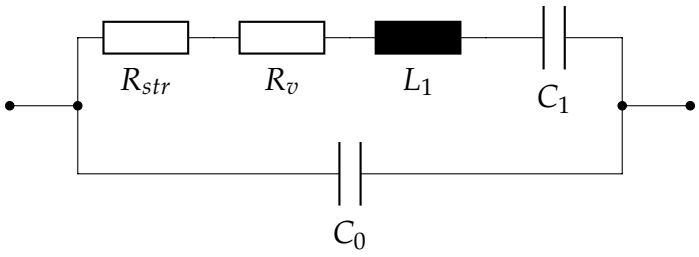

Figure 3.3: Equivalent circuit of a piezoelectric resonator

Typically, the resonator that is the piezoelectric system has a very narrow bandwidth. Figure 3.4 shows a typical impedance curve of an ultrasonic transducer. As expected, the plot shows that the impedance is smallest at the resonance frequency and that it is purely resistive in this point. The plot was generated by simulating the circuit shown in figure 3.3 with the following values: $R_{str} + R_v = 320\,\Omega, L_1 = 53mH, C_1 = 300pF, C_0 = 2000pF$ using the model shown in figure 3.3. The values were taken from [8]. Most datasheets of ultrasonic transducers only show values for C_0, the other values can rarely be found.

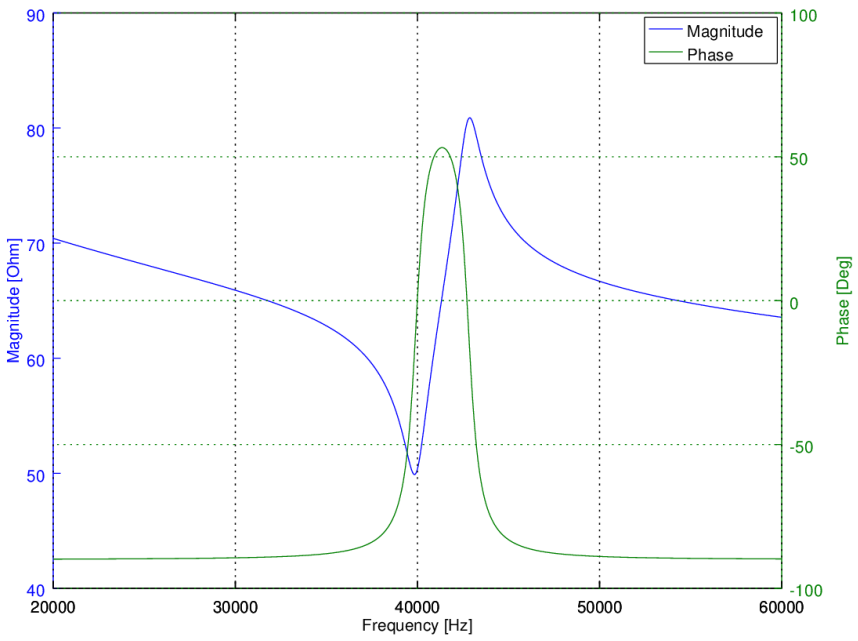

Figure 3.4: Typical impedance curve of an ultrasonic transducer

The next chapter will show how ultrasonic transducers can be used as transmitters.

3.3 Driving Transmitters

In order to excite an ultrasonic transducer, a voltage signal with the centre frequency of the transducer needs to be applied to its terminals. As noted in the previous chapter, these devices tend to have narrow bandwidths, meaning that the frequency is given very precisely by the chosen transmitter and can not be altered other than by a few 100 Hz up or down without sacrificing acoustic pressure. This also implies that matching transmitters and receivers need to be chosen for a system if a transducer should also be used as ultrasonic receiver. Chapter 3.4 will present an alternative receiver, which can be used for larger frequency ranges.

The deformation of the piezoelectric material, also called the displacement D, is proportional to the strength of the electric field across it according to $D = \varepsilon \cdot E$ [9, p. 223]. In order to achieve a high acoustic pressure, the displacement needs to be as high as possible, which in turn requires high driving voltages. For the A-14P20 transducer, which is used in this thesis, voltages of up to 140V can be used. However, it can be advantageous to use lower voltages, most notably to save space, cost and complexity by not requiring boost converters and high voltage switches. Thus, both possibilities will be explored in this chapter.

When looking at the electrical equivalent circuit of an ultrasonic transducer in figure 3.3 it can be observed that it does not consume any static current apart from leakage. However, dynamic currents need to be considered when designing a driver circuit since the transducer acts as a purely resistive load at the centre frequency, as shown in figure 3.4. This figure also hints towards typical impedance values at resonance, which will be important for designing drivers.

Low Voltage Transmitter Drivers

The simplest way of driving ultrasound transmitters is to apply a low voltage square wave with the centre frequency of the transmitter. Voltages of 3.3V, 5V or 12V are most likely already available in the system where the transmitter will be used since they are needed for supplying the control logic. Then, the most basic circuit could be an NPN transistor configured as a switch with a current limiting resistor. Figure 3.5 shows how such a simple driver could look.

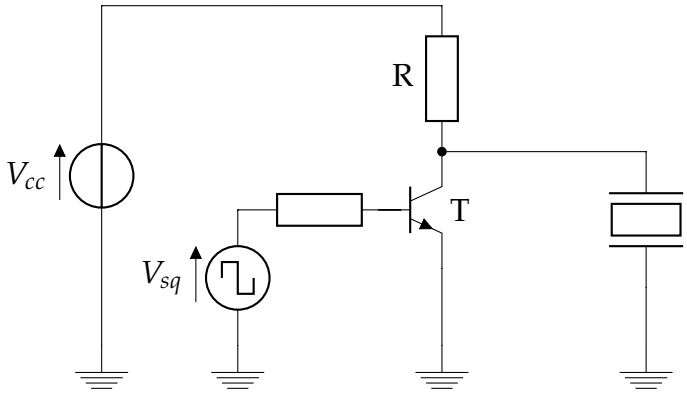

Figure 3.5: Simple transmitter driver

This circuit has several obvious drawbacks, the biggest being the resistor which is needed to limit the current when the switch is closed. This is limiting the current supplied to the transducer when the switch is off, deteriorating performance. Of course the circuit could be improved for example by adding an emitter follower between the output of the switch and the transducer, but the remainder of this section will show that there are better methods which do not rely on discrete transistors, while still being simple and small.

In figure 3.5, the voltage difference across the terminals of the transmitter is equal to the supply voltage V_{cc} in the best case where the voltage drop across R can be neglected. Since only the voltage across the terminals is important and no reference to ground is needed for the transmitter, a differential signal can be applied to boost the amplitude to roughly $2 \cdot V_{cc}$.

This is a typical application for an H-bridge. There are a number of available integrated bridges which can be used for this application. Since these devices are typically used to drive DC motors it is not difficult to find one which can supply the necessary current. A typical transducer will have an impedance of about 50 Ω, which translates to a current of about 250 mA at a driving voltage of 12 V.

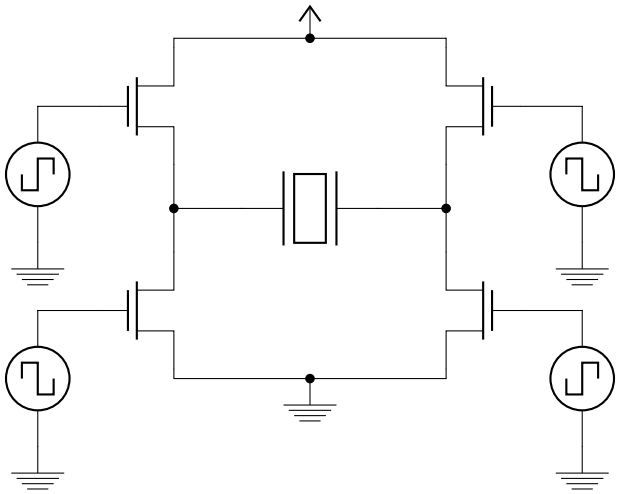

Figure 3.6: Transmitter driver using an H-bridge

For the sake of completeness an alternative is also presented at this point which can sometimes be found in literature about driving low voltage ultrasonic transducers. This approach uses a number of inverters (typically a hex inverter is used as these devices are abundantly available) to create a differential signal from a single ended input. Figure 3.7 shows how this can be achieved. In this circuit, two outputs are connected to each other respectively in order to increase the available current. This approach could be used for very low power transmitters where the voltage doubling has an advantage but there are strict space or cost constraints, but is is not considered a valid solution for the problem discussed in this thesis.

If the application only demands low distances between transmitter and receiver and the noise level is low enough, such low voltage transmitter drivers are a valid option, since they use very little space and power. This is of great advantage especially if such a system needs to be supplied via battery and/or has to be integrated into a tight space.

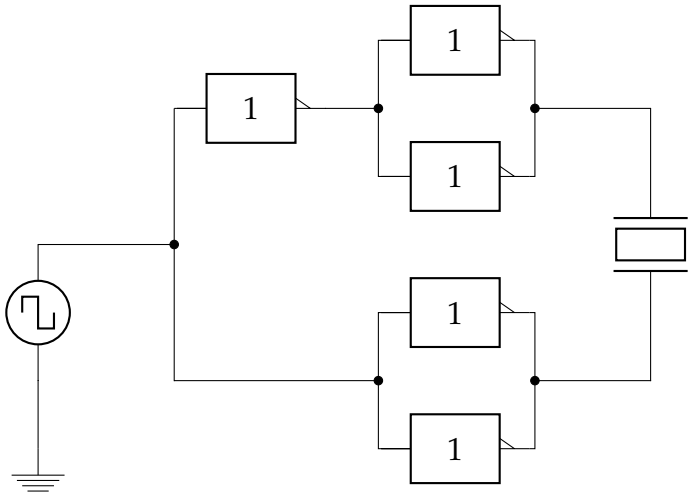

Figure 3.7: Differential transmitter driver

High Voltage Transmitter Drivers

In noisy environments or for measurements of larger distances, higher acoustic pressure and therefore higher driving voltages are needed. As mentioned before, the transducers used in this thesis allow for up to 140V across the inputs. [13] demonstrate a circuit which could deliver up to $\pm 300V$ at the outputs. It uses two supply voltage generators, one for positive and one for negative voltages. However, such a design uses a lot of components and therefore a lot of space and adds complexity to the design. A simpler approach would be to generate one positive voltage and use the principle introduced in the previous section to drive the transmitter at twice this voltage. This is exactly what is done by commercially available drivers for piezoelectric elements like the LT3572 by Linear Technology ([6]). It offers two channels, which can be driven independently by a square wave with arbitrary frequency of up to 80kHz. In this application, the LT3572 is set to generate a voltage of 40V, which translates to a voltage swing of 80V at the output of the H-Bridges. This voltage is still a lot lower than the maximum voltage of the transmitter, however early experiments showed that the acoustic pressure generated by the A-14P20 transducer driven with this voltage is adequate. Additionally, the size of the circuit is still reasonably small. Higher voltages could be achieved by either using different driver ICs or designing a non integrated solution consisting of a separate boost converter and a full bridge driver, or some form of step up

circuitry which boosts the amplitude of an input signal

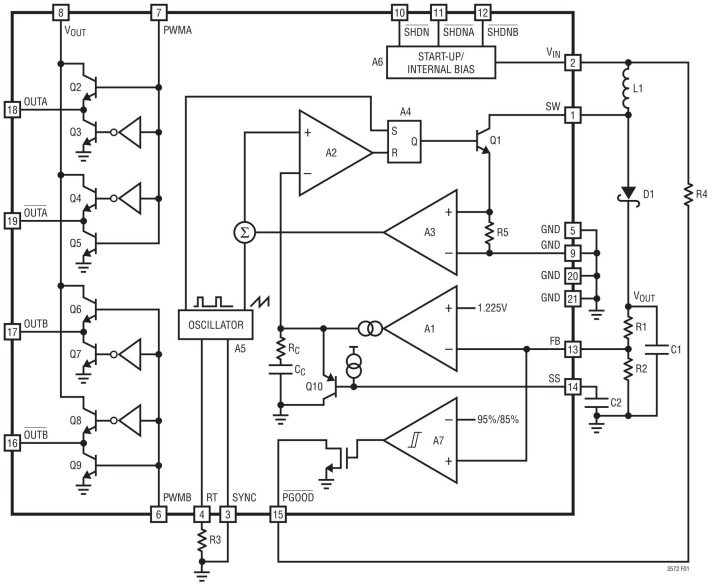

Figure 3.8: LT3572 Block Diagram ([6, S. 6])

3.3.1 Selected Transmitter Driver

For the prototype developed in this thesis, the transmitter driver is based on the LT3572 since it provides a good trade off between size and voltage swing. However it should be noted that this driver IC is not available with an automotive specification, which could be needed in the final product. At the time of writing this thesis it was not clear whether or not the product had to comply to automotive standards and it was decided to use non automotive parts for the prototype and transition to another solution should it be required at a later time. Since there are not other integrated drivers for piezoelectric elements available with automotive specifications, this solution would probably have to be a combination of a boost converter and an H-Bridge.

3.4 Ultrasonic Receivers

For receiving the ultrasonic signals there are two basic approaches. The first one is to use ultrasonic transducers as sensors, the second one is to use microphones which are capable of receiving signals up to ultrasonic frequencies. Both of theses solutions have their respective benefits and drawbacks over the other and the requirements for the system were never clearly enough specified to favour one of them over the other. This is why both types are considered in this thesis. The signal shape is slightly different for both types, but the signal processing algorithm which extracts trigger points from the received signals should be designed to handle these differences and to work with both receivers with only slight modifications or no modifications at all.

3.4.1 Transducers as Receivers

When using ultrasonic transducers, they of course need to be chosen with the same centre frequency as the transmitter. They could even be the same part as the transmitter, which would drive down the cost. A significant drawback of transducers as receivers is that they posse a very narrow bandwidth, which deteriorates the edges of the received burst signal, making signal processing more difficult and ultimately triggering less reliable. Additionally, commonly available transducers are rather big with diameters of over one centimetre, which limits the possible placements in the system. A huge benefit of this type of receiver is that many of them come encapsulated in aluminium. This gives them inherent robustness against water and dirt in the environment without additional measures. Of course the signal quality will be worse or even unusable if the sensor is covered to a certain degree, but this is true for all types of sensors, not only ultrasonic ones.

Since ultrasonic transducers are mostly capacitive sources with a high internal resistance (see chapter 3.2), a buffer or even a pre amplifier should be used directly at their output. Otherwise the cables or traces leading from the output of the amplifier to the input of the ADC would need to be shielded carefully. For this purpose, a printed circuit board has been designed which can be mounted directly on the transducer. The

main component of this board is a simple non inverting amplifier which offsets the signal from the transducer to half the power supply voltage for simple signal processing without the need for symmetrical power supplies. Additionally, it includes protection circuitry complimenting the already ESD hardened operational amplifier to ensure robustness. See figure 3.9 for the schematic of this board and figure 3.10 for a picture of how the board is mounted on the sensor.

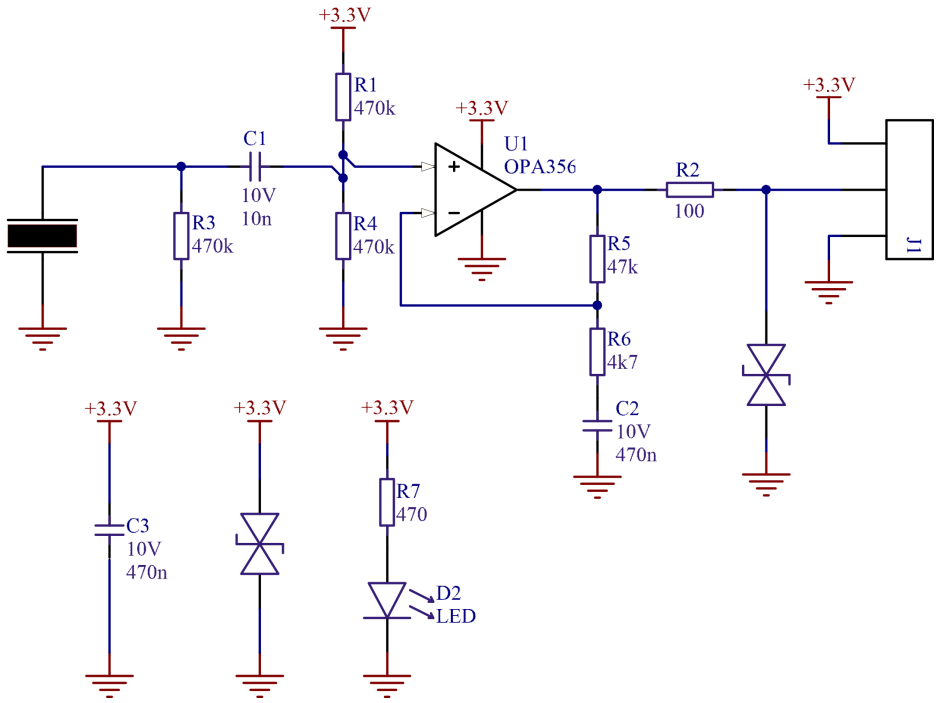

Figure 3.9: The schematic of the pre amplifier for ultrasonic transducers.

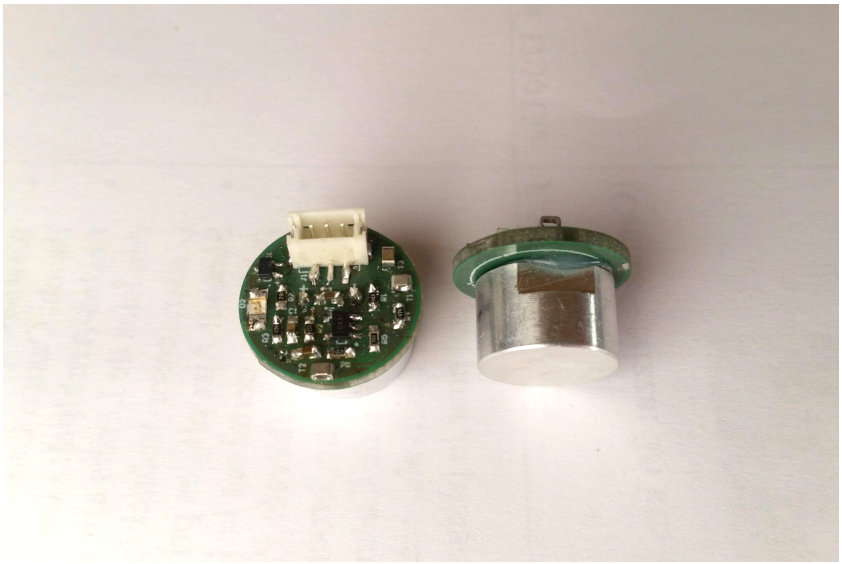

Figure 3.10: The pre amplifier board mounted on an A-14P20 ultrasonic transducer.

3.4.2 Microphones as Receivers

There are some microphones on the market which are sensitive to frequencies up to $80kHz$, mostly constructed as micro electromechanical systems (MEMS). For this application, the SPU0410LR5H-QB from Knowles, which has a sensitivity of $-40\frac{dBV}{Pa}$ at a frequency of $40kHz$ is used [14]. This microphone is also very small with a footprint of three by four millimetres and a height of only one millimetre. It is mounted upside down one the underside of the PCB, with a small hole for the sound input.

The major benefit of this type of receiver is the high bandwidth which allows for steep edges of the burst signal, simplifying trigger algorithms and making them more robust. One drawback of these microphones is their unprotected opening, which is not suitable for non clean environments. They need to be protected from water and dirt without influencing signal quality too much. This can be achieved via special acoustic materials which posses about the same acoustic impedance as air. In the same way as signals are reflected at mismatched impedances in electronics, acoustic waves are reflected at mismatched acoustic impedances. [9, pp. 244-245] Additionally, these materials are coated so that they repel water and dirt. However, using this material to cover the sensors increases the manufacturing complexity

and cost.

A small mounting board has been designed for these MEMS microphones. The microphones include an internal amplifier, so there is no need for an external pre amplifier. However, a buffer was included which uses the same ESD hardened operational amplifier as the board previously described to further protect the microphone. If not needed this can be bridged and simply not populated. The board was designed to have the same outline as the pre amplifier board for the transducers to allow for easy interchangeability.

4 Received Signal Properties

In order to be able to identify a trigger point in the received signal, burst signals are used instead of continuous signals. The shape of the received signal depends on the type of receiver used and on the number of periods of the carrier frequency applied to the transmitter, i.e. the length of the burst. Another factor influencing the shape of the received signal is the interference from echoes and reflections. This cannot be influenced by the design of either the transmitter or the receiver and is solely depending on the environment.

As mentioned in chapter 3.2, piezoelectric ultrasonic transducers have a rather small bandwidth compared to MEMS receivers. This results in a signal edge which is not as steep as the ones generated by MEMS receivers, which can make trigger algorithms less reliable and more difficult. Figure 4.1 shows the rising signal edges produced by a piezoelectric sensor and by a MEMS microphone side by side. It can be observed that the signal from the microphone reaches its maximum amplitude in less periods than the signal from the piezoelectric receiver and therefore displays a steeper edge.

The signals shown in figure 4.1 were filtered with a FIR bandpass filter with cut off frequencies of $30kHz$ and $60kHz$ with a Kaiser window and $N = 132$.

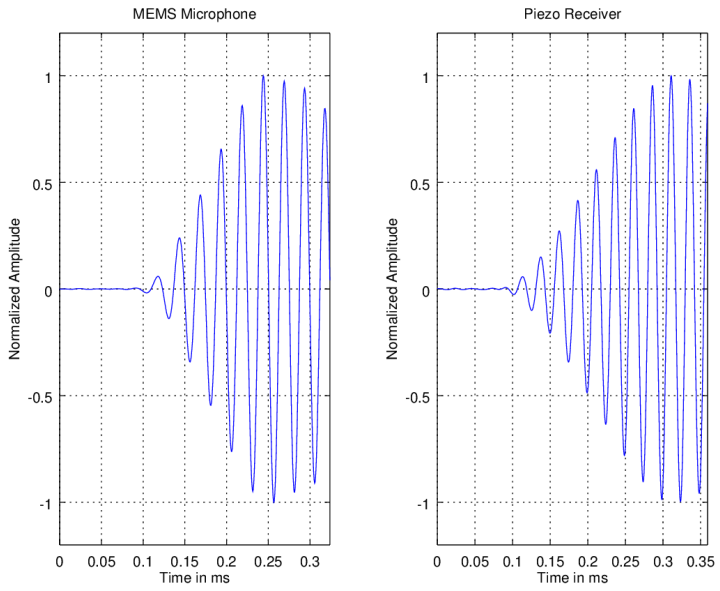

Figure 4.1: Signal shape comparison

4.1 Mathematical Model

In order to better understand the waveform and to provide a means of testing algorithms without the need for hundreds of precise measurements for reference data, a mathematical model of the burst signal has been created. This model mimics the real world signals very closely and provides parameters which can be adjusted to provide an accurate representation for different transmitter and receiver types and environments. These parameters have been selected empirically by comparing the output of the model with measured waveforms.

One received burst signal consists of multiple echoes and reflections of the real burst sent out by the transmitter which interfere with one another at the receiver. These echoes are modelled as well, since they influence the shape of the received signal to a great extend and are different for each positional configuration between transmitter, receiver and objects in the environment.

The simple burst without echoes can be modelled as follows:

$$P(t, t_0, f_0, s) = (t - t_0)^2 \cdot e^{-s \cdot (t - t_0)} \cdot sin(2\pi f_0 (t - t_0)) \cdot \Theta(t - t_0) \qquad (4.1)$$

This equation only models the shape of the signal, not its amplitude. In the tests for this thesis the output of the model is normalized and transformed into discrete values to mimic the analogue to digital conversion of the real system.

In equation 4.1, the overall shape of the burst is comprised of a polynomial and an exponential part, which mostly model the rising and falling edge respectively. t_0 is the starting point of the signal and can be used in the test fixtures for the trigger algorithms in a later section to measure the error of the tested algorithm. f_0 is the frequency of the carrier, which is $40kHz$ in the system discussed in this thesis. The shape of the burst can be adjusted via the factor s. Figure 4.2 shows two different values for s. In the real world signal, the shape is mostly influenced by the type of receiver and the length of the burst signal.

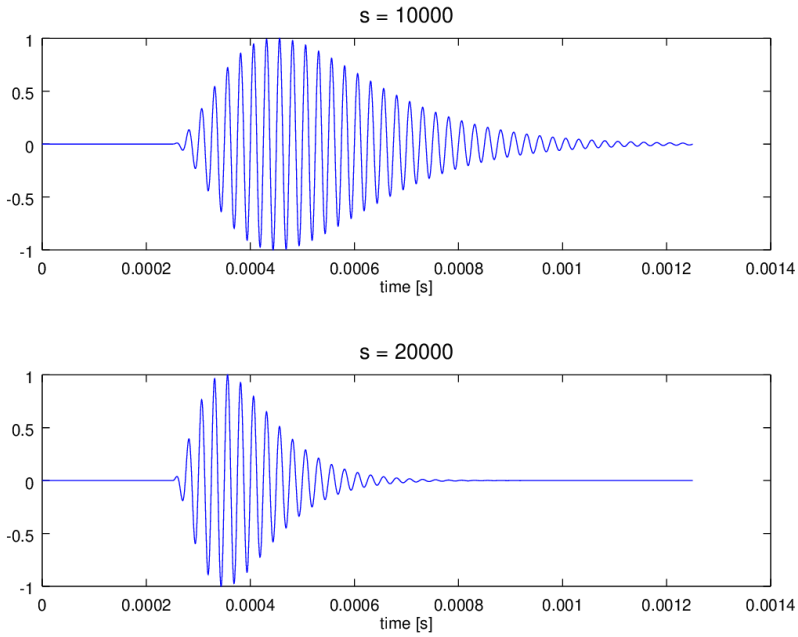

Figure 4.2: Influence of the shape factor s

In order to model echoes as well, multiple time offset $P(t)$ are weighted by a factor inversely proportional to the difference in t_0 to the first burst, which has weight 1, and summed. The number of echoes and the actual value of the weights again strongly depend on the environment and have again been selected empirically.

$$P_{echo} = \sum_{n=0}^{N} w_n \cdot P(t, t_{0,n}, f_0, s) \tag{4.2}$$

Figure 4.3 shows a selection of different bursts which were created with equation 4.2 and different weights, time delays and number of echoes. They reflect real world signals very well, especially when noise is added.

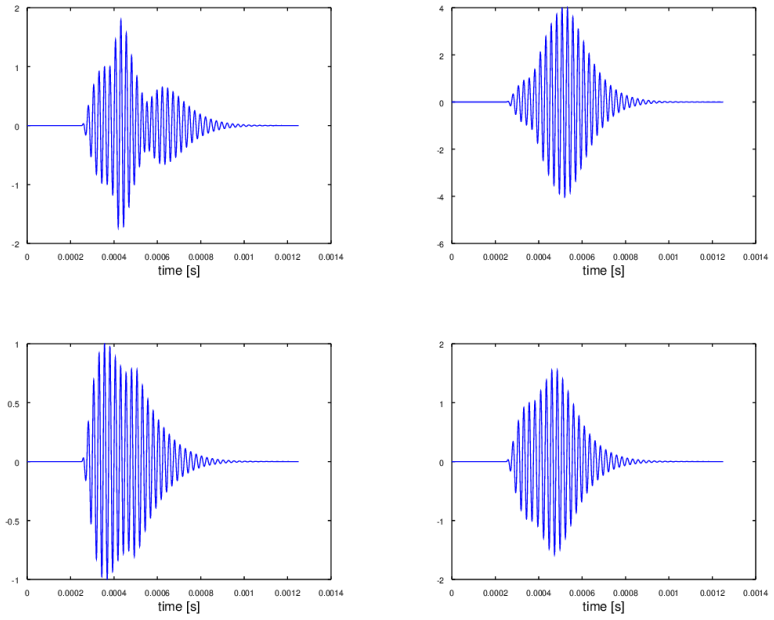

Figure 4.3: Examples of generated burst signals with echoes.

4.2 Angle (In-)Dependence

Since greatly varying angles are to be expected in the system the dependence of the signal shape on the angle of arrival needs to be investigated.

For receivers similar to the MEMS microphone introduced in section 3.4.2 where the opening of the sensor itself is very small and can be simplified to a perfect point, no strong dependence on the angle is to be expected. Shallower angles result in a reduction of signal amplitude, but the overall shape will stay the same.

Bigger sensors on the other hand could introduce a shift in time where the signal is received. For shallow angles the leading edge of a burst signal hits the perimeter of the sensor before the centre, whereas for steep angles, the whole surface of the sensor receives the signal at roughly the same time. It has been decided to investigate the existence or impact of this effect empirically.

For this experiment, a simple setup has been created, on which two ultrasonic transducers are mounted at the same height 10 cm apart. One of them acts as transmitter and is fixed in its position, the other acts as receiver and can be rotated around the centre of the receiving surface. See figure 4.4 for a picture of the measurement setup. The transmitter is on the right, the receiver on the left. On the bottom of the receiver mount there are markings for reading the current angle.

Figure 4.4: The setup for measuring the angle dependence of the received signal with two ultrasonic transducers.

The measurement was performed in the following way. A function generator was used to generate a 40 kHz square wave burst signal with four periods and an amplitude of 10 V . The trigger output of the function generator was used to start a capture on a USB oscilloscope, whose data can easily be exported for later analysis. Three measurements were performed, one at 0 degrees, meaning the transmitter and receiver faced each other (the green trace in the figure), one at 45 degrees (red trace) and one at 80 degrees (purple trace).

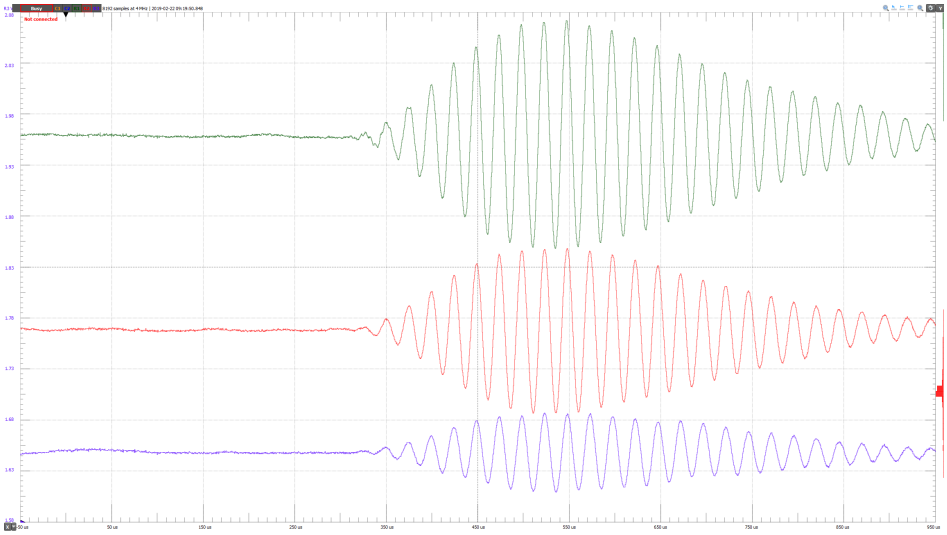

Figure 4.5: Results of the angle dependence measurement

Figure 4.5 shows no significant shift in the time of flight for varying angles, which is also why more in depth measurements have been deemed unnecessary. For a speed of sound of about $343\frac{m}{s}$ and a diameter of the sensor of about 1 cm , a shift of about half a period would be expected if the receiver is excited as soon as the pulse hits the edge when it is tilted relative to the transmitter. Since this is not the case, the effect stated at the beginning of this section is therefore either non existent or negligible.

5 Trigger Algorithm

5.1 Algorithm Requirements

The goal of the trigger algorithm is to robustly find a point in the received waveform which corresponds to the start of the received burst. This point needs to stay the same for all received bursts, regardless of their shape and amplitude. As already observed in chapter 4, the signal shape varies considerably, depending on the situation. This can make determining the start of the signal difficult even for humans.

Additionally, the algorithm should be able to work with limited resources as it will run on an embedded system alongside other software. Resources in this context are:

- **Clock cycles.**
  The execution time of the algorithm is important since the positioning system will also be used to give feedback while the car is approaching, therefore a high refresh rate is needed. However, there is a lower limit for this, since bursts of signals need a certain delay in between them to allow for echoes of previous signals to reach an amplitude low enough not to interfere. An execution time lower than this will not provide a lot of benefits, but should be pursued as an optional goal.
- **Flash memory.**
  The algorithm should not be overly complex, since this will in most cases directly translate into using a lot of code and therefore flash memory. The whole robot software should be able to fit into the internal flash memory of a standard microcontroller.
- **RAM.**
  RAM is among the most valuable resources in an embedded system and since the software of the robot is based on a real time operating system a lot of it is used for tasks. This implies that the algorithm should not use big arrays, recursive function calls etc.

The algorithms considered here work by sampling the signal over some time during which a pulse is expected and then applying filters and the

algorithm itself to the signal. In this case, the sampling frequency plays an important role, since it defines the number of samples which need to be stored. Of course it would be beneficial if the algorithm could also be used online. This is possible with some approaches to at least some extend, but tightens the requirement for the maximum execution time heavily by imposing a hard limit. Additionally, it limits possible signal processing options.

The most important property of a trigger algorithm is that it has minimal jitter. It needs to be able to identify the same point within the signal each time, even if the shape of the signal varies significantly. Such a point is always a zero crossing for the algorithms tested for this thesis, since it is the most robust characteristic with time information of a sine shaped signal. Since the identified zero crossing will not correspond directly to the real arrival time of the signal, the time offset needs to be calculated once by calibrating the system with a known distance between transmitter and receiver.

The next chapter will provide a brief overview of some characteristics of the received signal waveform which can be used to identify a unique point in time, i.e. a trigger point. Be advised that theses trigger points are only used to identify a period within the burst, the real trigger point which is then returned is calculated by using one of the zero crossings within this period.

5.2 Possible Triggers in the Received Signal

Before discussing the signal in more detail, the naming of the important parts of the signal should be stated. The transmitter driver excites the piezoelectric transmitter with a number of periods of a square wave to form a burst signal. In this thesis, the individual periods of the excitation signal are referred to as pulses, the whole of the excitation signal as burst. This also holds true for the counterparts in the received waveform. There, the term pulse refers to a single period in the carrier sine wave, while burst again refers to all pulses that make up the received signal. Note that the burst from the transmitter is not clearly visible in the receiver, since echoes and slow dampening increase the number of pulses in the

received waveform. Rising edge denotes the first few pulses which increase in amplitude until a maximum is reached. The number of pulses in this rising edge loosely correlate with the number of pulses in the burst in the transmitter.

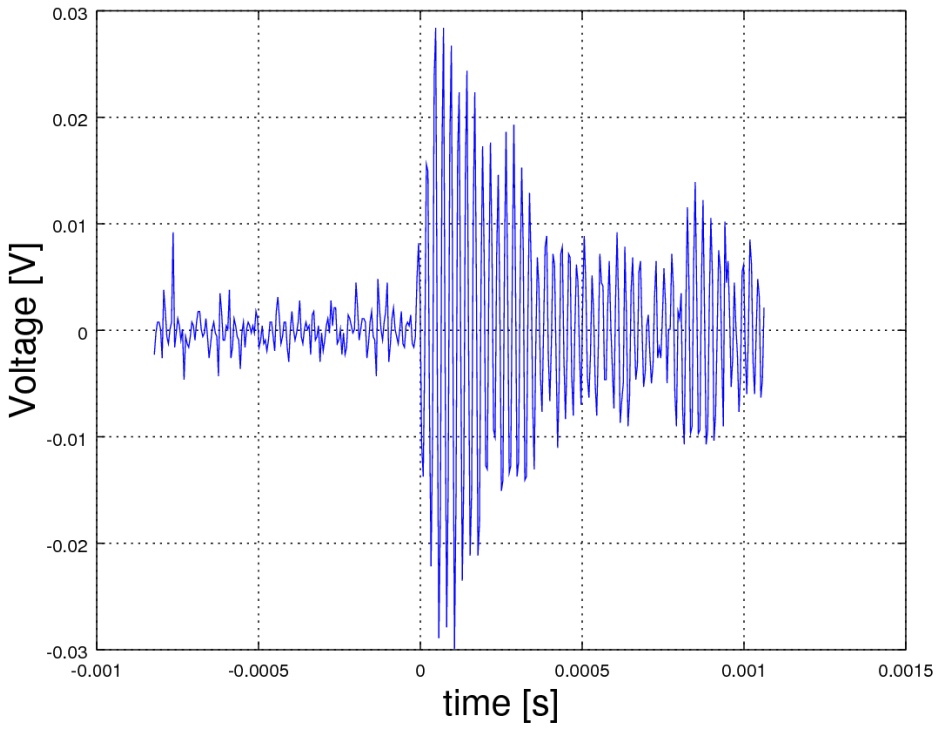

Figure 5.1: Unfiltered pulse as received by a MEMS microphone.

Figure 5.1 shows a typical waveform as seen by a MEMS based receiver. It shows the initial noise, the rising edge and interference of various echoes. The noise in this case is coming directly from the microphone, therefore it has to be considered in the algorithm. The electrical component of the noise could potentially be reduced by using different microphones or additional filtering of the power supply, but the acoustic component will always be present and potentially be worse in an outdoor environment. Additionally, the acoustic noise will mostly consist of audible frequencies with minimal amplitudes in the ultrasonic range. Because of this, the unfiltered signal from a piezoelectric receiver looks a lot cleaner. There, the acoustic noise is filtered by the inherent band pass filter which is the receiver itself.

The burst in figure 5.1 rises to its maximum in three periods, but this is not always the case. For example, this number increases when the number of periods used to excite the transmitter is increased, which is done to increase the acoustic pressure for longer distance transmissions. This number also changes if interference from reflections of the burst itself come into play. It also heavily depends on the type of receiver as stated before.

In standard ultrasonic distance measurement there is a fixed threshold which is used for triggering on the received signal. When inspecting the shape of the signal for a MEMS receiver with high bandwidth as shown in figure 5.1, this might seem like a valid approach also in this case. However, this is not considered suitable for this application. Firstly, the signal amplitude undergoes large swings for different situations, e.g. shallow reception angles, large distances etc. Secondly, especially for large distances, low SNR is to be expected. But the most important factor for why this approach is not suitable here is that it does not perform well if the rising edge of the signal is not steep, as is the case with small bandwidth piezoelectric receivers. The algorithm should be applicable to both MEMS and piezoelectric receivers.

The simple threshold trigger should however be considered as a starting point for other algorithms, which could improve its performance by employing additional knowledge about the signal.

The position of the maximum amplitude is prone to high fluctuations, even for successive measurements, it should therefore not be considered as a candidate for a trigger point. However, at least as a fallback for high noise environments the general shape of the signal could be used. This has been considered and investigated by [4]. As the shape of the signal can vary significantly depending on the environment this should only be considered as a fallback or for the feedback during the vehicle approaching phase, where high noise is expected, but the requirements for accuracy are not quite as high.

One property of the edge of the signal that could be leveraged for triggering is the growth of the signal amplitude, either within one pulse from one half wave to the second, or from one pulse to the next. A typical example for a system which uses an algorithm similar to this is Loran-C, a radionavigation system. There, the shape of the received RF burst is similar

to the shape of the ultrasonic bursts encountered in the system developed for this thesis. [2] outlines how the so called half cycle peak ratio can be used to uniquely identify a particular zero crossing in the burst to allow for accurate triggering.

The region after the rising edge should not be considered for algorithms since this is where most of the echoes of the signal cause constructive or destructive interference, which alters the shape tremendously. Additionally, depending on the type of noise and acoustic interference in the environment, the received signal can contain what looks like bursts before the real burst. This phenomenon is mostly encountered when narrow band band pass filtering is used and short, high amplitude pulses are coupled into the signal line as noise from a switch, clock signal or similar. If the algorithm makes use of such a filter and this kind of noise is to be expected, this needs to be considered in the algorithm and a suitable approach needs to be found to differentiate between the real burst and interference. If the signal is sampled before the algorithm is applied the simplest way to achieve this would be to search for the maximum in the sampled data and only consider a small region before this for the actual pulse detection.

As a rather brute force approach would be to perform a cross correlation between the received signal and one or multiple stored signals. As one would expect this is computationally very expensive and requires a lot of memory, first for storing the sample bursts and for the large number of samples which result from the necessarily high sampling frequency. However, this approach could be used for detection where high accuracy, but low speed is required.

The following section describes the process of developing an algorithm suitable for triggering on a pre sampled signal by combining the ideas presented in this chapter.

5.3 Development of a Trigger Algorithm

Rather than presenting a finished algorithm and analysing it, this section presents the actual development process in a bottom up style in order to show why certain choices have been made. Additionally, the finished algorithm has some parameters which can be tweaked depending on the situation and environment, whose purpose and impacts on the performance become much clearer this way.

Algorithms are evaluated using an octave framework which generates test data by employing the formula presented in chapter 4.1. This approach was favoured over using measured data since otherwise the error made by measuring the distances and positions when recording test data would add to the error made by the algorithm, falsifying the results. Additionally, in order to get meaningful results about the performance of the algorithms, hundreds of runs need to be performed and evaluated, which is infeasible when using measured data. The final algorithm will of course be evaluated further using real world data.

Relative Amplitude Triggering

As a starting point a simple approach which has been given the name "relative amplitude triggering" has been chosen. This approach is based on the observation that the rising edge and, specifically, the growth of the pulses leading up to the first maximum of the burst look the same for bursts which are generated with the same number of excitation pulses. This observation can be used to define a threshold based on the signal itself, to which the amplitudes of the individual pulses are compared. The first pulse whose amplitude crosses this threshold gives the trigger point. Be aware that the above observation only holds true if the angle of reception is about the same for each burst and this shape can change considerably as soon as echoes occur. These problems will be approached later in this chapter.

The first step in developing this simple algorithm is finding a way to approximate the amplitude of a pulse given the sampled points. Since the

sampling rate will most probably not be high enough to get a measurement of the amplitude directly, curve fitting can be employed to recover the amplitude with a reasonably small error. For this application, fitting a quadratic function to the sampled data has proven to provide a good trade off between computational complexity and accuracy of the approximation of the amplitude. When imposing the restriction that the sampling frequency must be an even multiple of the carrier frequency, the error made this way is about the same for every pulse, which diminishes its influence.

The next step is to define a metric on which the threshold calculation will be based. This could be the maximum amplitude of the burst, the total energy in the signal, the RMS value or any other metric which is tied to the burst itself. Since simplicity is favoured in the design of this algorithm, the maximum amplitude has been chosen as a starting point. If this proofs to not be reliable enough the other approaches will be considered. The actual threshold is defined to be a certain percentage of the maximum amplitude, with the percentage being dependent on the signal shape.

When a pulse whose amplitude crosses the threshold is found, the algorithm needs a way to calculate a certain point which is then used as a trigger point. If the same pulse is found for each signal, the start of the signal can be retrieved via a simple subtraction of the fixed offset. Zero crossings are the best candidate for an accurate, precise and robust trigger point. They can be easily calculated by drawing a line through adjacent sample points above and below the x axis and finding the zero crossing of this line.

The basic algorithm is as follows:

1. Search for the maximum amplitude in the sampled data and calculate a representing value
2. Search for individual pulses before the maximum and calculate their amplitude
3. For the first pulse whose amplitude crosses the threshold calculate the position of the zero crossing

In order to identify a suitable trigger point, the figure 5.2 has been created. It shows the amplitudes of the pulses on the rising edge of the burst relative to the maximum amplitude of the burst. It can be observed that

the variation of the relative amplitude varies significantly and increases with increasing pulse number. This can be explained by the fact that the influence of echoes increases with time, so the first one or two pulses are mostly unaffected, but later ones experience constructive or destructive interference.

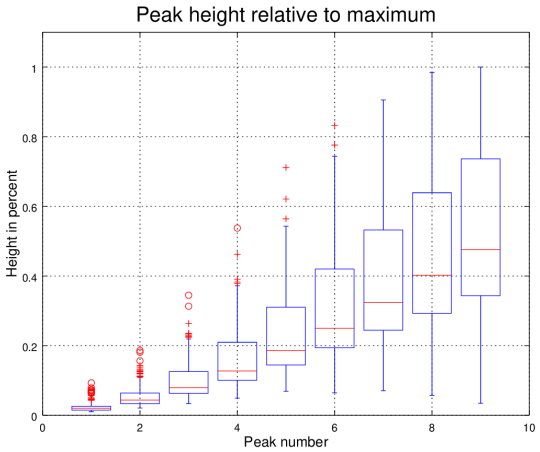

Figure 5.2: Amplitudes relative to the maximum.

As a preliminary test, an algorithm which simply triggers on a given percentage of the maximum signal amplitude has been developed using octave. There are no optimizations and only basic bandpass filtering. It is used to establish a baseline with which improved algorithms can be compared. Tests were performed using generated test data. Two runs were performed, one without any noise and one with added white noise at an amplitude of ten percent of the maximum signal amplitude. The triggering threshold was swept from 0 percent to 100 percent and for each threshold, 500 runs where performed. The results can be found in figure 5.3

Figure 5.3 shows that noise only influences the results at very low thresholds and that the mean error and the variance of the error increases with increased threshold, as one would expect when echoes are present. It is difficult to remove the influence of echoes on the trigger algorithm since they modify the shape of the signal drastically. It is therefore generally preferable to use a very low threshold and to try to trigger on the first pulse. This requires finding a way to reduce the influence of noise on the algorithm. The most obvious way to achieve this would of course be

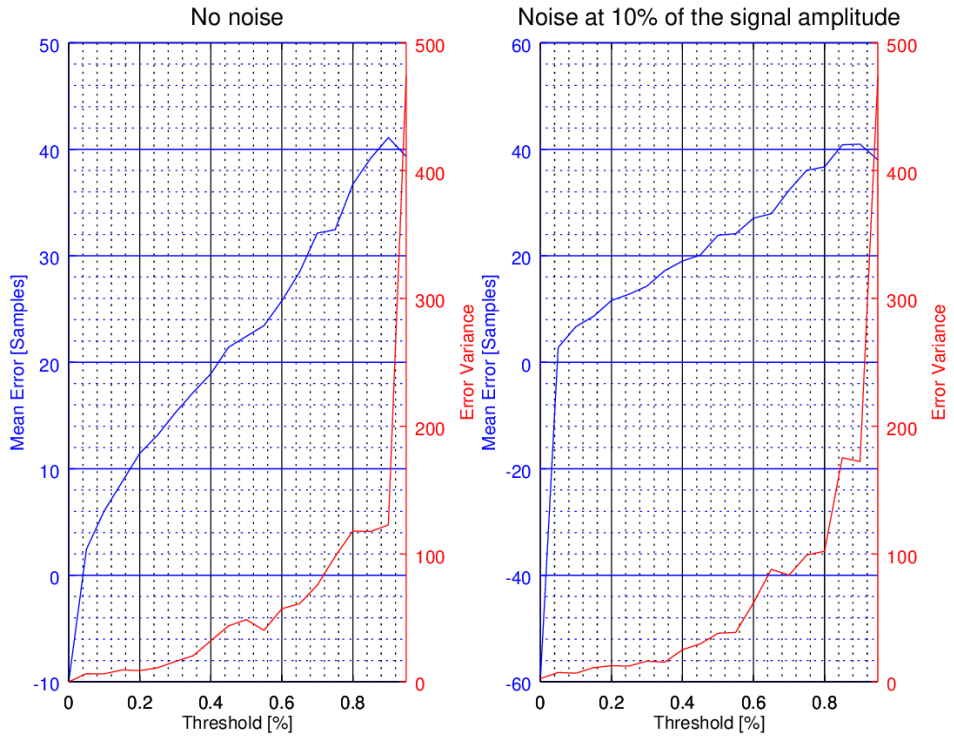

Figure 5.3: Results of the basic relative amplitude trigger algorithm

filtering the signal heavily, however, this does not always produce the desired outcome. As soon as noise contains the carrier frequency of the pulse it will not be removed by a filter and cause incorrect results when low thresholds are used. It is therefore desirable to use a priory knowledge of the signal shape to formulate requirements for the trigger points. These criteria will be checked by the algorithm for each potential trigger point and are presented below:

- The start of the pulse must lie within a certain range of the maximum amplitude
- The number of data points above and below the x axis is given by the ratio of sampling frequency to carrier frequency, plus or minus one
- The distance between zero crossings of the quadratic function fitted into the data points will correspond to roughly half the period of the carrier frequency
- The distance to the next pulse is given by the carrier frequency

- The amplitude of the next pulse has to be within certain bounds (as implied by figure 5.2 and 5.4). This criterion can be generalized to also include more than only the next pulse.

It should be noted that checking these requirements will generally work better with unfiltered or slightly filtered data, since a steep and narrow bandpass would only leave the components of the carrier frequency, which would make most of the criteria above ineffective.

The bound for the amplitude of the following pulse can be defined more precisely by measuring the amplitude increase from one pulse to the next, which results in figure 5.4. The numbers on the x axis denote the nth increase, which is defined as $\frac{A_{n+1}}{A_n}$, with A_n designating the amplitude of pulse n.

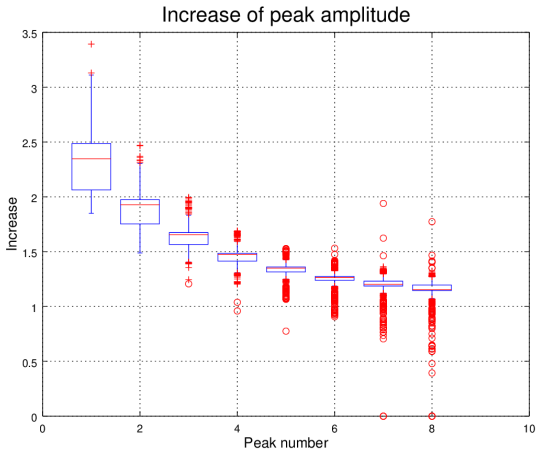

Figure 5.4: Increase of pulse amplitudes.

It can be observed that the mean increase follows a specific pattern and the variance decreases for increasing n. Most importantly the figure implies that not only the amplitude of the following pulse should be considered, but at least also the one following that because of the high variance of the first increase.

The basic algorithm from before has been extended to employ the conditions named above in order to reject false trigger points. The best results were achieved by bandpass filtering the input signal ($f_{low} = 38kHz, f_{high} = 45kHz, N = 26$) and a threshold of two percent of the maximum amplitude.

The algorithm can also reject measurements if the pulse which would meet all the conditions is above a certain threshold, which decreases the likelihood of a wrong measurement. This can happen if one or more of the conditions is not met for the real first pulse, but all are met for one of the following pulses. In this case, rejecting measurements and trying again is preferable over a wrong measurement. The algorithm could also easily be extended so that instead of rejecting measurements it returns a confidence value which can then be interpreted by the calling program.

The octave implementation of this algorithm has been tested using generated test data. At this stage, the performance of the algorithm under good conditions was tested, i.e. in a low noise environment. Since echoes can always be present in the real signal this test also includes them. As stated before reflections of the burst signal can drastically alter the shape of the signal, which the algorithm should be able to handle. 1000 runs have been performed with random number and location of these echoes and for a single shape factor $s = 11000$. This shape factor corresponds to a real world burst signal at a piezoelectric receiver which was generated by applying a burst with a duration of four periods to the transmitter. The real time of flight of the burst in this test corresponds to 50 samples. The results are summarized here:

- Mean error in samples: 64.601
- Variance of the error: 5.65
- Percentage of rejected measurements: 0%

Note that while these results look promising, the distribution of the error produced by the algorithm does not, as shown in figure 5.5. It shows that the error of one period appears almost 50 percent of the time. It can only take on discrete multiples of the carrier frequency period plus or minus a small, mostly negligible jitter. If the frequency of this error can be reduced, this knowledge can be used to improve the performance of the system.

Also note that the performance of this algorithm heavily depends on the signal shape factor and the signal shape in general. The influence of echoes is reduced by triggering on the very first detectable periods of the burst, but as soon as the length of the burst changes or other sensors with different bandwidth are used, it needs to be carefully tuned again.

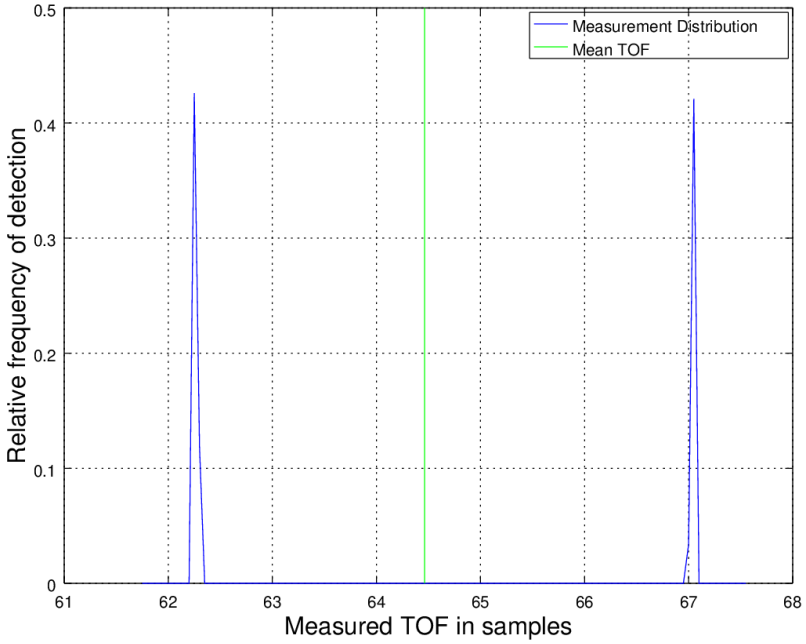

Figure 5.5: Distribution of the measured TOF.

The influence of noise on this algorithm is not shown here as the performance of the algorithm even without noise is not satisfactory. The next chapter will present a possibility to improve the performance.

5.3.1 Using Linear Regression to Improve Performance

The relative amplitude trigger algorithm performs quite well when it is trimmed to the expected signal shape. However, high noise and echoes lead to higher error counts. It would be beneficial to relax the criteria the signal has to meet in order to provide a valid trigger point, while maintaining low variance in the produced error.

One approach to achieve this is to use the rising edge of the burst to get information about its starting point, rather than using individual periods within the signal. This can be achieved by fitting a straight line into the rising edge and calculating where this line intersects the x axis. This point

can then be used to calculate a more precise value for the trigger point by employing the same approach as presented earlier.

In order to identify the rising edge of the signal, the relative amplitude triggering algorithm is used, but with less strict requirements and checks since in this case, it is not as important to identify the same points in the signal each time. It is just used to detect the rising edge. It would also be possible to use a fixed threshold for triggering which is guaranteed to be above the noise floor, however it is beneficial to reuse the already existing functions from the previous algorithm. The performance of the algorithm in terms of speed could be slightly increased by using the different approach, but this can be implemented at a later time.

Once the rising edge has been identified, the actual fitting is performed by calculating the approximate position of the maximum of a number of pulses and performing a linear regression algorithm. The number of used pulses can be varied depending on the shape of the signal, but using three pulses has been found to work best in most situations. Fitting is done on the rectified signal.

An example of the result of the linear regression fit is shown in figure 5.6. It shows an example pulse recorded from the sensors and the operations performed by the trigger algorithm. The quadratic functions fitted into the data to recreate the amplitudes can be seen, as well as the linear function fitted into the amplitudes of the rectified signal, approximating the rising edge. Finally, the zero crossing of this linear function is represented as a vertical line.

The result of the linear regression is then used to calculate a zero crossing. This point acts as a reference for calculating the actual trigger point. It is not returned directly since this would decrease the precision of the algorithm drastically. Instead, two steps are performed to increase precision. First, move a fixed offset forward in time from the calculated point to get a point within the rising edge of the signal instead of at the start of the signal, where the noise is relatively high. Next, search for a zero crossing in the actual signal before and after this point. The one closer to the reference point is then returned.

The same evaluation of the algorithm was performed, providing the following results:

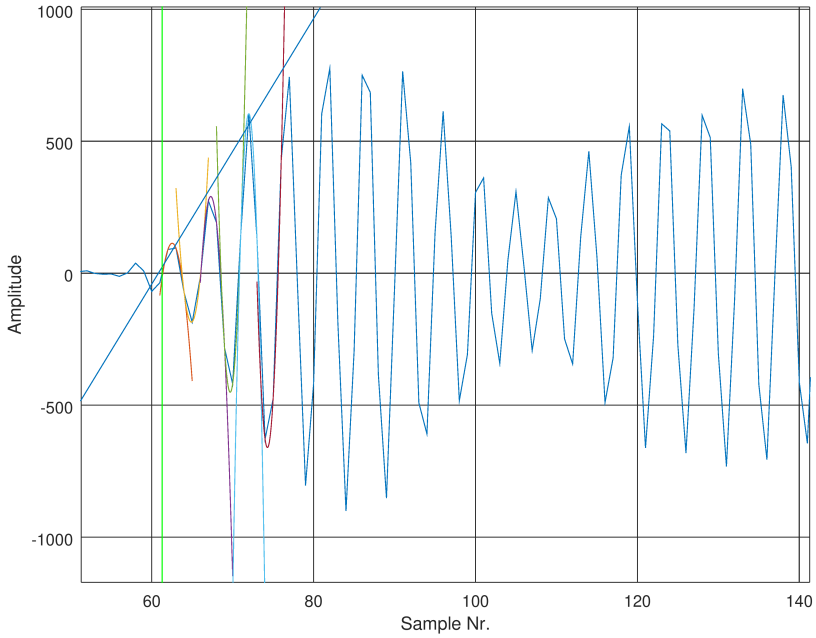

Figure 5.6: Representation of the linear regression trigger algorithm.

- Mean error in samples: 66.49
- Variance of the error: 2.33
- Percentage of rejected measurements: 0%

The performance of this algorithm is significantly improved over the previous algorithm, which is also shown in the distribution of the measured TOF in figure 5.7. There are again some errors with values of multiples of the period of the carrier frequency, but they are much more rare than before.

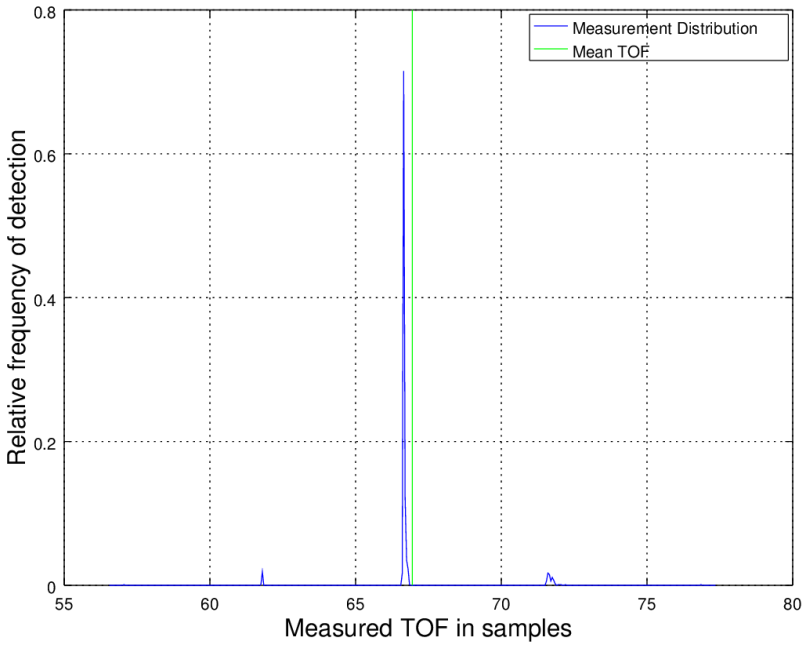

Figure 5.7: Distribution of the measured TOF.

6 Geometry

6.1 Position Measurement Approaches

There are two main ways of measuring a position when using wave propagation based methods, namely Time Of Flight (TOF) and Time Difference Of Arrival (TDOA) [3, p. 27-32]. The first and most intuitive is to directly measure time of flight. Measuring the time of flight of one signal yields a sphere of possible positions of the transmitter relative to the receiver with a radius given by the measured time multiplied by the propagation speed of the signal. By intersecting the resulting spheres, positions can easily be calculated. However, there needs to be some synchronisation mechanism since the exact time when the signal has been emitted needs to be known to calculate the TOF. This can be achieved by either using synchronised clocks and knowledge of the exact transmission time or by making use of a second channel with much higher propagation speed. In the case of using ultrasound as a measurement channel, a RF channel could be used to provide a synchronization signal between transmitter and receiver, eliminating the need for synchronized clocks.

To calculate a three dimensional position, at least three distinct lines of propagation of the signal need to be measured. This can be achieved by utilizing one of two ways. Using three transmitters at different locations and a single receiver, where the transmitters transmit at predefined, non overlapping time slots, or using one transmitter and three receivers. The second way enables faster and more efficient measurements, since only one burst needs to be generated and transmitted.

Using three measurements yields two solutions on opposite sides of the plane created by the three receivers or transmitters respectively. Depending on the configuration, one of these results can be discarded if a solution on a specific side of the plane is not realistic. If this is not possible, an additional transmitter or receiver is necessary to make the position calculation unambiguous.

An alternative to TOF based measurements is to measure the time difference of arrival. This approach has the advantage that no synchronisation between the transmitter and receiver is needed, but it requires an additional receiver or transmitter, depending on the configuration. A single measurement creates a hyperboloid of possible positions in contrast to the sphere created by a TOF based measurements. Calculating the true position therefore becomes a matter of intersecting multiple hyperboloids. A simple method of calculating the position given multiple time differences can be found in [5].

6.2 Selected Approach

For this thesis, both TOF and TDOA have been investigated as possible solutions. Ultimately, a TOF based approach has been chosen over a TDOA or hybrid one. The system needs an RF channel for communication either way, so it can also be used for synchronization. This means that the biggest disadvantage of TOF over TDOA does not play a huge role in this scenario. Additionally it has been found that TDOA based approaches suffer from severe error propagation when the receivers or transmitters are not placed far enough apart, which is not necessarily possible. The error propagation of the TOF based system is shown below, after presenting the formulas for calculating a position based on the measured TOF.

6.2.1 Calculating The Position Using Three Receivers

Using the minimum number of receivers for a TOF based system, the position can be calculated by intersecting the three spheres generated by the individual measurements of the TOF and knowledge about the speed of sound. This section will briefly describe the formulas involved, which are then needed to show how errors in the TOF measurement propagate and influence the result. It should be noted that since three receivers are used, the solution consists of two points with differently signed z coordinates. The receivers need to be mounted in a way which only allows one of these solutions.

Since three receivers will always lie in a plane, the normal vector of this plane is defined as the z axis of the coordinate system, which can be used to simplify the calculations. The resulting coordinates will need to be transformed into the coordinate system of the robot. This transformation is most likely quicker than the calculation of the position directly in the correct coordinate system. With this definition of the z axis, the formula for the sphere with radius d and centre $(x_m|y_m|0)$ is:

$$(x - x_m)^2 + (y - y_m)^2 + z^2 = d^2 \tag{6.1}$$

The coordinates $(x|y|z)$ should be calculated. In order to remove the squares of the solution and therefore simplifying the equations, the differences between the equations for the receivers (1), (2) and (3) are calculated. Another simplification introduced in this step is to set the centre of the sensor coordinate system to one of the receivers, in this case receiver (1). The following equation shows the difference between the equations for receiver (1) and (2), the equation for the difference between (1) and (3) follows the same scheme.

$$2 \cdot x \cdot x_{m_2} - x_{m_2}^2 + 2 \cdot y \cdot y_{m_2} - y_{m_2}^2 = d_1^2 - d_2^2 \tag{6.2}$$

The following substitutions will be used in the subsequent formulas to improve readability:

$$\begin{aligned} a_1 &= d_1^2 - d_2^2 + x_{m_2}^2 + y_{m_2}^2 \\ a_2 &= d_1^2 - d_3^2 + x_{m_3}^2 + y_{m_3}^2 \end{aligned} \tag{6.3}$$

Next, the variable x is eliminated from the equations, which results in

$$\begin{aligned} 2 \cdot y \cdot (y_{m_2} - \frac{x_{m_2}}{x_{m_3}} \cdot y_{m_2}) &= a_1 - \frac{x_{m_2}}{x_{m_3}} \cdot a_2 \\ y &= \frac{a_1 - a_2 \cdot \frac{x_{m_2}}{x_{m_3}}}{y_{m_2} - y_{m_3} \cdot \frac{x_{m_2}}{x_{m_3}}} \end{aligned} \tag{6.4}$$

The solutions for x and z are equally trivial:

$$\begin{aligned} x &= \frac{a_1 - 2 \cdot y \cdot y_{m_2}}{2 \cdot x_{m_2}} \\ z &= \pm\sqrt{d_1^2 - x^2 - y^2} \end{aligned} \tag{6.5}$$

6.2.2 Error Propagation

Since the measurement of the time of flight and the knowledge of the speed of sound are both subject to errors, the influence of these errors as a uncertainty in the distances $d_{1,2,3}$ on the position needs to be investigated. For simplicity, the error in measurement is assumed to be uncorrelated. This simplification can only be made for the error from the detection algorithm, not for the error made by an incorrect speed of sound or jitter in the starting time of sampling.

In order to calculate the uncertainty of y, the uncertainty of a_1 and a_2 is needed. This can be expressed as

$$\begin{aligned} \delta a_1 &= \sqrt{(\delta d_1^2)^2 + (\delta d_2^2)^2)} \\ \delta a_1 &= \sqrt{(2 \cdot d_1 \cdot \delta d_1)^2 + (2 \cdot d_2 \cdot \delta d_2)^2} \\ \delta a_2 &= \sqrt{(2 \cdot d_1 \cdot \delta d_1)^2 + (2 \cdot d_3 \cdot \delta d_2)^2} \end{aligned} \tag{6.6}$$

This uncertainty of y can then be calculated as

$$\delta y = \frac{1}{2 \cdot y_{m_2} - 2 \cdot y_{m_3} \cdot \frac{x_{m_2}}{x_{m_3}}} \sqrt{(\delta a_1)^2 + (\frac{x_{m_2}}{x_{m_3}} \cdot \delta a_2)^2} \tag{6.7}$$

This uncertainty then directly influences x, and both δx and δy influence z:

$$\begin{aligned}\delta x &= \frac{1}{x_{m_2}}\sqrt{(\delta a_1)^2 + (2 \cdot y_{m_2} \cdot \delta y)^2} \\ \delta z &= \frac{1}{2}\frac{1}{\sqrt{d_1^2 - x^2 - y^2}}\sqrt{(2 \cdot d_1 \cdot \delta d_1)^2 + (2 \cdot x \cdot \delta x)^2 + (2 \cdot y \cdot \delta y)^2}\end{aligned} \tag{6.8}$$

These equations give a rough idea of the influence of the random error in the detection time on the calculated position. Additionally, they can be used to find the optimal arrangement of sensors to minimize the influence of the error. This will be needed in the next chapter.

6.2.3 Further Considerations

The previous section showed that the propagation of the measurement error should in most cases not drastically influence the result. However, the system will use more than the minimum number of needed receivers or transmitters respectively to improve the accuracy and to enable sanity checks for the calculated positions. This also improves the reliability for cases where sensors are covered in dirt and will be discussed later.

Chapter 5 showed that errors tend to be multiples of the carrier frequency period. This stems from the fact that most algorithms will search for a half wave in the signal which satisfies some conditions. When a correct half wave has been identified an exact trigger time can be calculated via a zero crossing. The error introduced via this zero crossing calculation is in the range of single percentages of the carrier frequency period even for low sampling frequencies. However, if the algorithm identifies an incorrect half wave, the error will increase by a multiple of the period. This observation can be used to perform error detection and even error correction, when more than the minimum number of measurements can be performed.

The prototype consists of a system using two transmitters and four receivers, where the receivers are mounted on the robot and the transmitters near the plug on the vehicle. More receivers could be added to improve redundancy and system performance overall if this is possible in the future. But for the prototype, this configuration satisfies the minimum requirements for the TOF based system with error correction. The second

transmitter is necessary to calculate the rotation of the vehicle. The design of the plug on the car does not allow for a transmitter to be mounted exactly at its centre, meaning that the rotation of the vehicle relative to the robot needs to be known in order to calculate the true position of the centre of the plug.

As mentioned in the previous section, in order to measure TOF the exact time at which the pulse is transmitted needs to be known to the receiver. Synchronized clocks were ruled out as a possible solution for this problem at the beginning of the design, since this would require low drift clocks as external components or a way to synchronize the internal clocks of the microprocessors every so often. Such a synchronization requires some sort of communication channel, most likely an RF module. If such a module is used, it can also be used to transmit a start signal, provided the module itself supports such a feature. Radio waves have a propagation speed which can be assumed to be infinite for this application. The important requirement for the selection of an RF module for this system is that it provides a means of generating a transmission/reception complete interrupt without too much jitter. Most commercial RF modules do provide interrupts, but implement and abstract low level packet handling with automatic retransmissions which can cause high timing jitter between the transmission and reception interrupt when a packet was received, but the acknowledge packet was lost.

6.3 Sensor Placement and Mechanical Integration

In order to determine the optimal sensor placement, the equations from section 6.2.2 where used. The uncertainty for x and y should be minimized, the uncertainty of z is then automatically at its minimum. Since the sensors should lie in a plane, there are four parameters which can be modified, namely x_{m_2}, y_{m_2}, x_{m_3} and y_{m_3}. To determine the optimal placement, the set of parameters which produce the minimum mean squares error, $error = \sqrt{(\delta x)^2 + (\delta y)^2}$ can be calculated. There are limits imposed on the values of theses parameters by the mechanical construction of the robot, which is why the search was constrained to a $10cmx10cm$ grid.

The result of the parameter optimization shows that x_{m_3} and y_{m_2} should be as big as possible, while the values for x_{m_2} and y_{m_3} hardly influence the uncertainty. They should be more or less equal and not necessarily zero or at the maximum, but the uncertainty of x and y only increases slightly if these conditions are not met. This can be explained by the fact that the distance between the receivers should be as big as possible, which is achieved by setting x_{m_3} and y_{m_2} to the maximum possible values. However, when these parameters are at their maximum, changing the remaining parameters only slightly increases the distance between receivers (1) and (2) and (1) and (3), while simultaneously decreasing the distance between receivers (2) and (3). This placement could look like shown in figure 6.1

The receivers will need to be integrated into the aluminium cover of the robot. Since the sensors used are encapsulated in aluminium, they need to be mounted in a way to minimize coupling between the cover and the receiver. Because of the size of the cover and the shallow angles, the burst from a transmitter can get the metal to oscillate at the centre frequency of the transducers before it reaches all receivers. The speed of sound in metal is significantly higher than in air, resulting in a premature excitement of one or more receivers when improper decoupling is present.

Figure 6.2 shows the acoustic impedances present when the sensors are mounted in the aluminium cover. Z_A denotes the acoustic impedance of aluminium, which is roughly the same for the receiver and the cover material, and Z_B denotes the acoustic impedance of some other, dampening

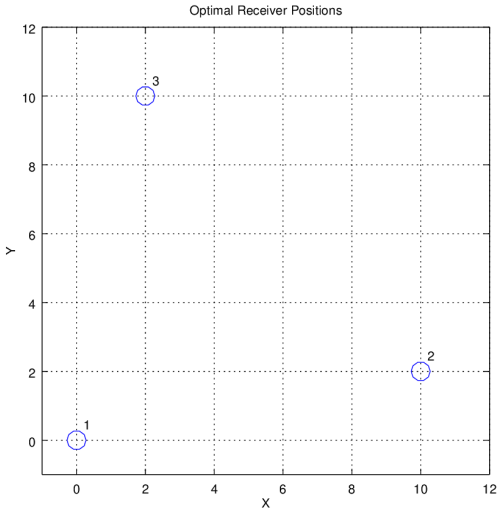

Figure 6.1: One possibility of the placement of three receivers for minimal uncertainty.

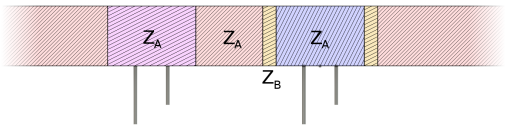

Figure 6.2: Acoustic impedances for the mounted receivers.

material. In order to achieve low coupling between receiver and robot, the acoustic impedances should be mismatched. This is not the case when the receiver is mounted directly in the cover, as can be seen on the left side of figure 6.2. When an additional layer of a different material with a highly different acoustic impedance is inserted between the two aluminium parts, two planes of mismatched impedances are created, as shown on the right side of figure 6.2. For the prototype, different materials where tried out, and a flexible foam like material was chosen, as it showed the highest dampening when inserted. For the future, a material will need to be found which also meets all the mechanical requirements.

6.4 Integration of a Fourth Receiver

As mentioned in chapter 6.2.3, a fourth receiver will be added to the prototype along the necessary three in order to improve performance. With a fourth receiver present, the system is over determined, which can either be used to calculate the three dimensional position more accurately, or using the additional information as a sanity check to test if the trigger algorithm produced a result off by one period or more. The details of the implementation will be presented in chapter 7.2.

In order to determine the optimal location for this fourth receiver, a similar approach to before was chosen, with the positions of the other receivers fixed and only optimizing the position of the fourth receiver. It has been found that as long as the distance between all receivers is kept reasonably high, the influence of errors does not increase drastically. This information together with the requirements of the mechanics resulted in a layout shown in figure 6.3

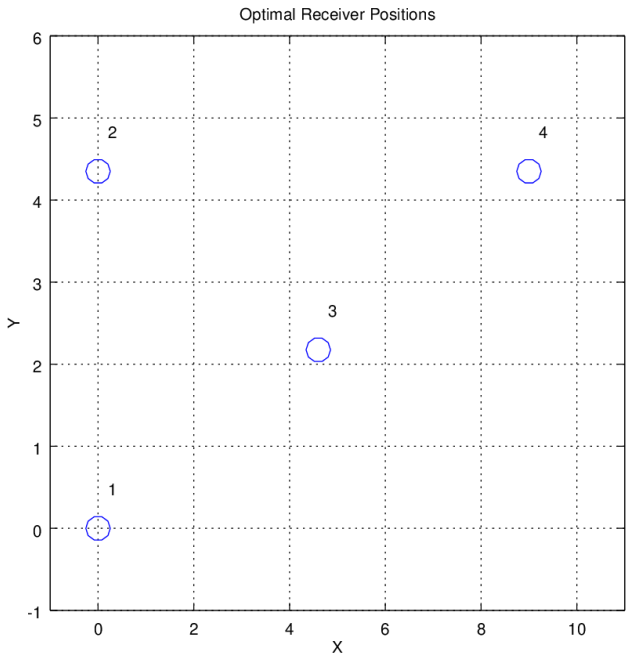

Figure 6.3: The final placement of the four receivers

7 Implementation of the Prototype

7.1 Hardware

The prototype of the Volterio system is composed of two main components, a robot controller board and a vehicle unit controller board. Both of them have the relevant components and software for the ultrasound position measurement system integrated. They communicate via a 2.4GHz radio module, which is used to exchange status information and which acts as a synchronization mechanism for the time of flight measurement. The following sections describe the components necessary for the ultrasound positioning system in more detail.

7.1.1 Sensors

Over different iterations of the robot, different sensors have been used. In the first version, the sensors where embedded in the connector, which required small sensors. Due to the better availability of MEMS microphones over small surface mount ultrasonic transducers these sensors where used. However, space constraints forced very small and mechanically unstable solutions, which caused many problems. Because of this the sensors where moved to the base of the robot, where more space is available. At the base of the robot more mechanical stress and dirt is to be expected, which is why it has been decided to use piezoelectric receivers. As chapter 8 will show this decision proofed to be a bad choice for a prototype, especially since ways of protecting the better suited MEMS receivers were provided. For this thesis, a separate receiver module has been created which uses MEMS sensors in the same configuration as the piezoelectric receivers on the robot.

To be able to provide feedback to the driver or the car respectively while the vehicle is approaching, a separate set of sensors has been mounted angled

in the direction from which it will be approaching. Figure 7.1 shows the layout. The microprocessor which handles the ultrasound measurements can dynamically switch between the two sensor sets, based on signal quality or position of the car.

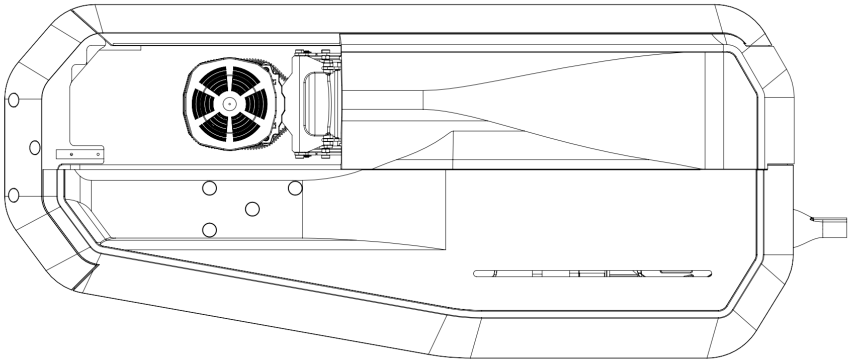

Figure 7.1: Locations of the sensors on the base of the robot (circles on the left half)

Figure 7.1 shows the locations of the seven sensors on the robot as circles on the left half of the robot. The sensor set pointing straight up is composed of four individual receivers, while the angled set consists of only three. This is again due to the need for exact measurements when the vehicle is parked and the fast measurements when the vehicle is approaching. Additionally, there is very little space available in the front of the robot to accommodate more than these three sensors.

Due to the presence of mechanical coupling the sensors as described in chapter 6.3 where surrounded by a foam like material before being mounted in the aluminium case of the robot. However, this material degraded the mechanical stability of the hole sensor system, increasing the potential advantage of MEMS receivers even further.

7.1.2 Receiver Module

The receiver circuit is integrated into the main robot control board and shares the controller with the rest of the robot software. The choice of

this controller was therefore heavily influenced by the requirements of the ultrasonic system. A 32 bit mixed signal ST microprocessor has been selected, specifically an STM32L443RC. It provides a fast 12 bit ADC, enough memory and some digital signal processing functions, which make it ideal for this application. Filtering is done almost completely in software, with the exception of a simple analogue high pass filter at the sensor inputs to bias the signal to half the supply voltage in order to simplify the circuit design by only using a single supply. The board supports two sets of four sensors each, with an analogue switch selecting the desired set before a measurement. The switch is followed by a pair of 2x programmable gain amplifiers, which are needed to guarantee a reasonable resolution of the sampled signal while not running into clipping issues since the signal voltage can vary greatly with distance and angles.

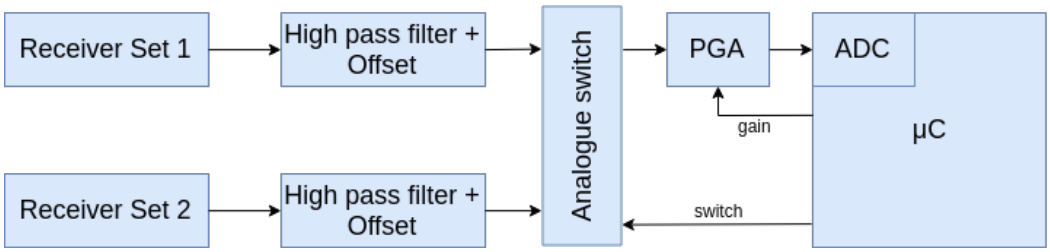

Figure 7.2: Block diagram of the receiver module.

Figure 7.2 shows a block diagram of the receiver module with its most important components.

As briefly described in section 6.2, a RF channel channel can be used to provide a synchronization signal. In the case of this prototype, a Nordic Semiconductor NRF24L01 2.4GHz module is used for this. It can be easily interfaced, provides complete integrated packet handling and an interrupt signal which can be programmed to indicate successful transmission or reception of a packet. This signal is used to start the transmission in the transmitter and to start the sampling in the receiver. Timing jitter introduced by this approach is in the scale of single microseconds and therefore not relevant for this application. The time between the "transmit complete" interrupt in the transmitter and the "receive complete" in the receiver can therefore be assumed constant and subtracted from the measured time of flight.

7.1.3 Transmitter Module

The ultrasonic transmitter module is integrated into the vehicle unit control board and provides support for two transmitters. It uses a LT3572, which was described in chapter 3.3 to drive the transmitters with a high voltage signal. However, this IC can only be used for the prototype, since there is no automotive version of it available.

The module supports two transmitters, which can be selected independently. It does not matter where the transmitters are mounted on the actual mechanical module, as long as they lie in a plane parallel to the ground. Small deviations from this plane which can arise purely from the cars suspension, uneven ground etc. have an impact on the accuracy of the calculated position of the plug on the car, but this should be handled by the mechanical design of the plug itself. As long as the positions of the two transmitters relative to the centre of the plug is known to the receiver it can calculate the target position for the plug using the measured positions of the transmitters.

The vehicle unit has a cover which protects the plug and the transmitters from dirt while the vehicle is driving. This cover can be opened with a DC motor, the driver for which is integrated into the transmitter module. Of course there is also a wireless module for the communication with the robot. One challenge here is the placement of the antenna, since the transmitter module with all the components sits in a metal box. For the prototype, the antenna is simply mounted to the underside of the car, away from the vehicle unit. This is of course no viable solution for the future. However, a more thought through design of the mechanical components of the vehicle unit could allow for a plastic cover, behind which the antenna would be mounted.

7.2 Software

7.2.1 Overview

Almost all of the complexity of the program lies in the firmware of the robot controller board. It handles the complete control of the robots motors and sensors, implements the high level state machine for handling interactions with vehicles and performs position measurements. The software for the transmitter module in the car is intentionally kept as simple as possible. It simply reacts to requests from the robot and has no internal state, so that temporary power loss, system errors etc. do not have a disrupting impact on the overall functionality.

In order to simplify the system design for the robot firmware the main components were split into individual tasks, which run on a real time operating system, namely FreeRTOS. The most important task for this thesis is the task which handles the measurement and calculation of the position of the plug on the vehicle. It sleeps as long as there is no request from another task.

7.2.2 Sampling and Triggering

The four supported receivers are always sampled in parallel so that only one pulse needs to be sent for one measurement. Since the internal ADC of the STM32L4 is used, which does not support true parallel sampling of multiple channels, the internal multiplexer is used. The ADC is capable of performing one $MSample/s$, which translates to a maximum sampling frequency of 250kHz per channel when four channels are used. For this application, the sampling frequency was set to 200kHz, which provides five points per period for a 40kHz signal, which proved to be enough.

The basic procedure for performing a measurement is as follows:

1. Send a request for an ultrasound pulse to the transmitter. The transmitter to be used and the number of excitation pulses is transmitted along with the request.

2. Wait for a response from the transmitter. The interrupt generated by the wireless module is used to trigger transmission as well as sampling.
3. After the sampling is complete the data needs to be reordered due to a limitation in the direct memory access controller of the STM32L4.
4. Then, the data is band pass filtered and the signal processing algorithm is applied.

If enough RAM is available a ping pong buffering scheme can be applied, where two separate data buffers are used. This allows for parallel sampling and processing, since no active data is overwritten. However, since the update frequency was high enough and the RAM was rather full, this feature remained unused. It can be activated in the future should the need arise.

7.2.3 Implementation of the Trigger Algorithm

The original algorithm was developed in octave and therefore needed to be ported to C in order to be usable in the system. Since testing and verification of the C implementation should not be performed on the embedded system, but rather on a PC with the necessary debugging tools a test bench was set up. This test bench supplies an interface with which the implementation can be called from the octave program which was used to test the initial implementation. It also implements the DSP functionality which is present on the target device so that the implementation of the algorithm does not need to be modified before being used on the target. This simplifies testing and debugging.

Filtering

To improve measurement accuracy and to remove outliers, the measurement results are being filtered. There are two basic options in this case, either filtering the calculated positions or filtering the measured distances before calculating the position. Since the result of the position calculation changes dramatically for even small errors in the distance calculation, it

has been decided to employ two filters, one before calculation and one after.

Which filters are used depends on the type of measurement to be performed. When performing a position measurement while the vehicle is approaching, a high update rate is required. Additionally, the real position of the car changes continuously. However, the position does not need to be exact, it is only used for feedback to the driver. This is why the measured distances as well as the calculated positions are filtered using a moving average filter of length $n = 5$ for both filters.

For the measurement of the position of the parked car filtering can be improved using knowledge about the trigger algorithm. Since this algorithm produces roughly discrete errors of multiples of the carrier frequency period a filter algorithm can be designed to take advantage of this characteristic. The errors are referred to as being roughly discrete since the calculation of the zero crossing introduces a non discrete error on top of the discrete error made by selecting the wrong period in the signal. The basic concept of the algorithm will be described in the following paragraph.

After multiple distance measurements have been performed and their results are stored in an array, the algorithm is applied for the results of each transmitter separately. It goes through the distances and groups them together based on the difference between them. All distances with a difference of less than half a carrier frequency period are grouped together. Then the group with the most members is selected and the mean over these values is returned as a result.

This returns a single set of distances which are then used to calculate the position of the transmitter. Multiple such measurements are then performed to get a set of positions. This set is then filtered using a median filter, whose result is then the final result of the measurement.

Correcting The Time Offset Between ADC Channels

Since the microcontroller used for this system only has one ADC with an analogue multiplexer, there is a slight offset in the actual sampling time between each channel. If not corrected, this offset translates directly into

a small, but not insignificant error in the calculated TOF, which in turn causes an error in the calculated position.

The offset between two channels is the same as the sampling or conversion time T_{conv} for one channel, which can be calculated as $T_{conf} = T_{sample} + T_{SAR}$ [10]. T_{sample} is the time required for the actual sample step of the conversion, which was set to be 12.5 ADC cycles in this application. T_{SAR} is the time taken by the successive approximation ADC to perform the measurement and is also 12.5 ADC cycles, because of the resolution of 12 bit. The ADC is clocked at $32MHz$, which gives a total conversion time of $T_{conf} = \frac{12.5+12.5}{32 \cdot 10^9} = 781.25\,\text{ns}$. Uncorrected, this time would lead to an error of about 0.23 mm from one channel to the next, or almost 1 mm from the first channel to the last.

This calculation was also verified by a measurement which used a small voltage drop at the input of the ADC at the start of a conversion to measure the delay between the sampling of two consecutive channels. Figure 7.3 shows the result of this measurement. The two traces shown correspond to two consecutively sampled channels. This measurement was also performed to verify the correct operation of the analogue circuitry. The voltage drop of about 10 mV seen in figure 7.3 is within acceptable limits and does not influence the result, since it is gone by the time the sampling step is complete.

Error Detection and Correction

Because of the drastically varying shape of the received burst signal which is depending on the structure of the environment, it was not possible to design a trigger algorithm with negligible error rate. However, as shown in chapter 5.3, the error produced by the best algorithm is in most cases one period, and almost always with the same sign. This information together with the additional measurement from the fourth receiver is used to implement tests to check if such an error has occurred in one or more processed signals and to try to correct this error.

The most trivial check which can always be performed, even with only three receivers, is to check if the term under the square root in equation 6.2.1 is positive and if the resulting coordinates are within the expected

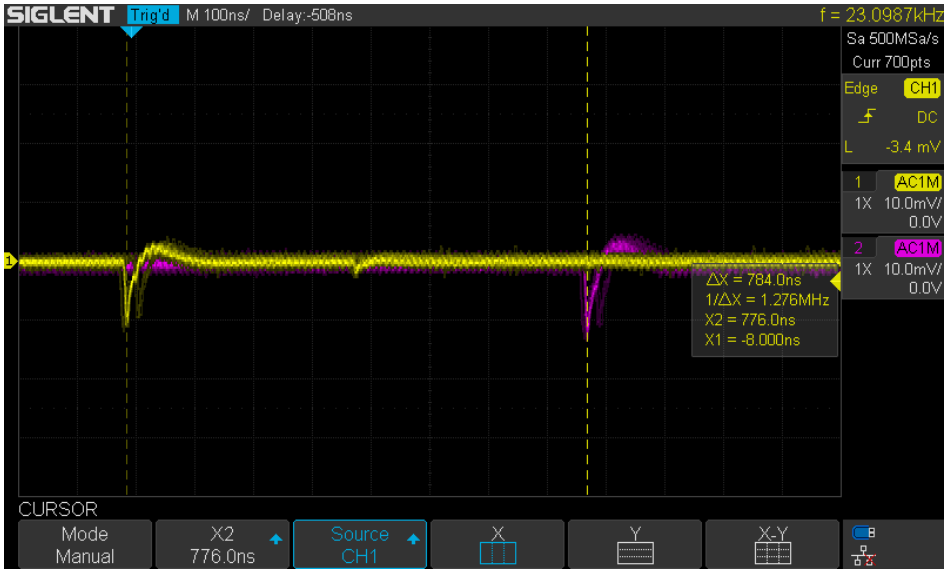

Figure 7.3: Verification of the time offset between two ADC channels.

range. Also, the last calculated positions can be stored if measurements are performed while the vehicle is approaching, so that this information can be used to further constraint the coordinates.

In order to perform further error detection and error correction, the speed of sound is calculated using the calculated times of flight from all four receivers. An error in each of them will influence the resulting speed of sound, most likely altering it to something unrealistic. The equation for the change in speed of sound given an error in one of the times of flight is rather tedious and therefore not shown here, but follows exactly the same principle as presented in chapter 6.2.2. From this formula it can be estimated that the speed of sound changes by some value between about one meter per second and tens of meters per second when a single time of flight is off by one period of the carrier frequency. Two errors can result in errors in the hundreds of meters per second. The real speed of sound can be estimated via an on board temperature sensor. If the speed of sound calculated from the times of flight differs too much from the estimated speed of sound, an error in one of the times of flight is assumed.

The algorithm then tries to correct this error by repeating the previous calculation of the speed of sound while modifying one TOF by the expected error. If the result is within the tolerated limits, this set of times of flight is

assumed to be correct. However, the calling function is notified that this result was achieved via error correction. If the result is still not valid, the error correction algorithm tries to modify two times of flight at once and tests again. After all pairs of measurements have been tested and still no valid result has been found, the algorithm returns with an error, since the probability of producing an invalid result is too high at this point.

8 Results

In order to evaluate the performance of the final implementation, a test setup has been created which allows for a single transmitter to be moved in all three directions relative to the receivers. The receivers were mounted in the same configuration as in the robot on a 3D printed part. Measurements were performed along the X and Y axis as well as on a line parallel to the X axis with an offset in the Y direction. These measurements were performed on two different heights. This way, the performance of the position measurement system could be evaluated in the range necessary for docking. Measurements in the range needed for the approaching phase were not performed. This decision was made due to the fact that the results for these measurements depend highly on the configuration of the transmitters on the underside of the car in combination with the parallel and highly uneven planes of ground and underfloor creating reflections. No vehicle was available for measurements and the results would not be meaningful as they could not be generalized. However, the results of the presented measurements can give an insight into the possible performance of the system in this case.

The measurements were performed with both types of sensors, as it became clear early on that the piezoelectric receivers yielded unsatisfactory results. The information that piezoelectric receivers do not provide the intended performance was made clear long before the end of this project and before this final measurement was performed.

8.1 Results Using piezoelectric Receivers

For measurements with the piezoelectric receivers a 3D printed part was designed which allows for the sensors to be mounted exactly like in the robot, with the exception of the coupling characteristics due to different materials. The coupling was minimized by mounting the sensors with the same acoustic dampening material as in the real robot enclosure, minimizing the effect. The transmitter was mounted on a rack which allowed it to be placed at precise positions relative to the receiving unit and to be moved in all three dimensions. The way the transmitter is mounted ignores the potential reflections from the underfloor of the vehicle. These reflections can safely be ignored due to the fact that the distances travelled by reflected signals is much higher than the distance between transmitter and receiver, which places echoes in the received signal far behind the actual signal, thus removing interference and influence on the measurement.

Only the results for the best performing algorithm are presented, which was the linear regression based algorithm. Its parameters were trimmed to the characteristics of the receivers empirically.

First, the measurements along the X axis are presented. For this, the transmitter was mounted at a height of 15 cm above the receivers, centred in both X and Y direction. Then, it was moved along the X axis to the positions ± 2.5 cm, ± 5 cm, ± 7.5 cm, ± 10 cm, ± 15 cm and ± 20 cm at $y = 0cm$. To test the behaviour at bigger distances, the same sweep along the X axis was repeated at a Y offset of 20 cm. The same measurements were repeated for a Z value of 7 cm. These measurements span the area of intended operation of the system. For each position, three measurements were performed, where each position measurement consisted of ten individual TOF measurements, which were filtered as described in chapter 7.2.3 before calculating the position. The positions themselves were not filtered by the system.

Figure 8.1 shows the results of the first measurement, where the transmitter was moved along the X axis, at a height of 15 cm. It shows the mean errors of the X, Y and Z components of the measured positions, as well as the 2D distance between the real and measured positions. The 2D distance was chosen over the 3D distance since for this application the accuracy of the X and Y components is far more important than that of Z because the

robot always moves straight up from a given 2D position when docking and does not rely on the data from the positioning system to know when to stop. The following figures also show the variance of the measurements. This does not show the real statistical variance due to the low number of measurements for each point, but can be used to show were jumps in the calculated position occur due to invalid error correction or wrong detections which still lead to a valid result.

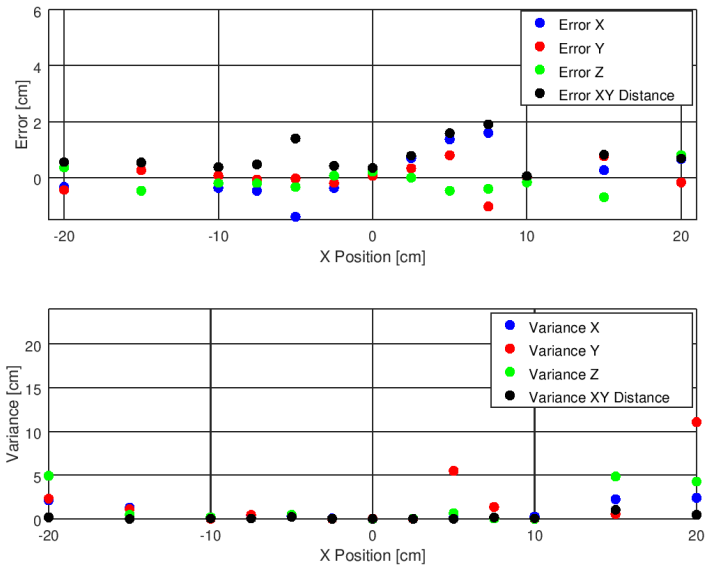

Figure 8.1: Results of the measurement along the Y axis at a height of 15 cm

Figure 8.2 shows the results of the same measurement repeated at a height of 7 cm. As expected, the performance decreased slightly due to the lower signal amplitude, which in turn is an effect of the shallower angles.

The results of the measurements performed by performing an offset sweep of X can be found in figures 8.3 and 8.4. There the performance decreases further due to the higher distances and shallow angles.

From the presented results it is clear that the performance of the piezoelectric receiver based system is not adequate for the requirements. The mean error is often bigger than the required 1 cm and is not distributed randomly enough to simply perform oversampling and filter the resulting distances. This is due to the fact that the error in detection does not

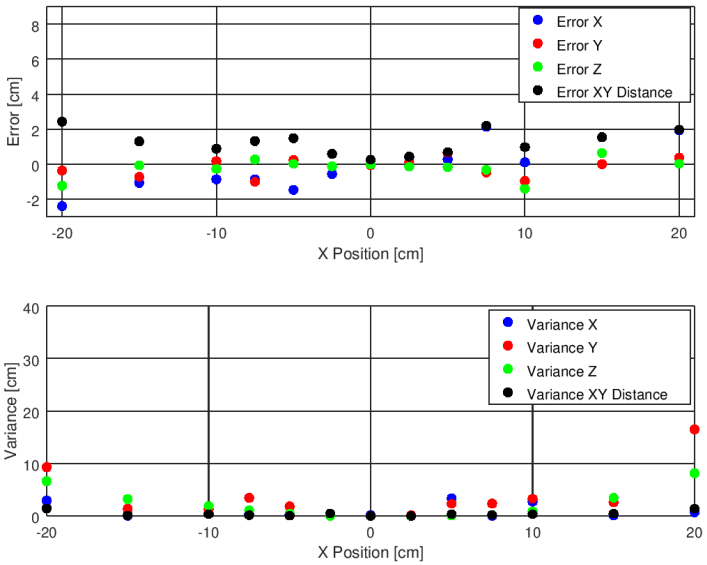

Figure 8.2: Results of the measurement along the Y axis at a height of 7 cm

only stem from random noise, but from the varying signal shape which stays consistent between multiple measurements for a single position. The true origin of this varying shape has not been identified, it seems to be a combination of angle dependency, coupling via the material where the sensors are mounted and the fact that the transmitter is also not perfectly decoupled, leading to excitation of the surrounding material, which also produces ultrasonic waves. These factors in combination with the slow rise time of the received signal due to the small bandwidth make it very hard to design an algorithm which can trigger reliably.

To reduce the influence of the biggest factor in algorithm performance, the above measurements have been repeated using MEMS receivers with a much higher bandwidth.

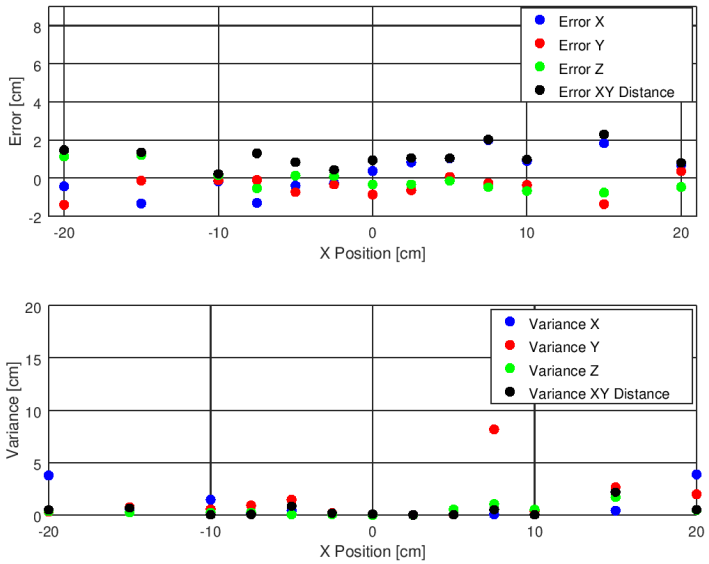

Figure 8.3: Results of the measurement parallel to the Y axis at an offset of 20 cm and at a height of 15 cm

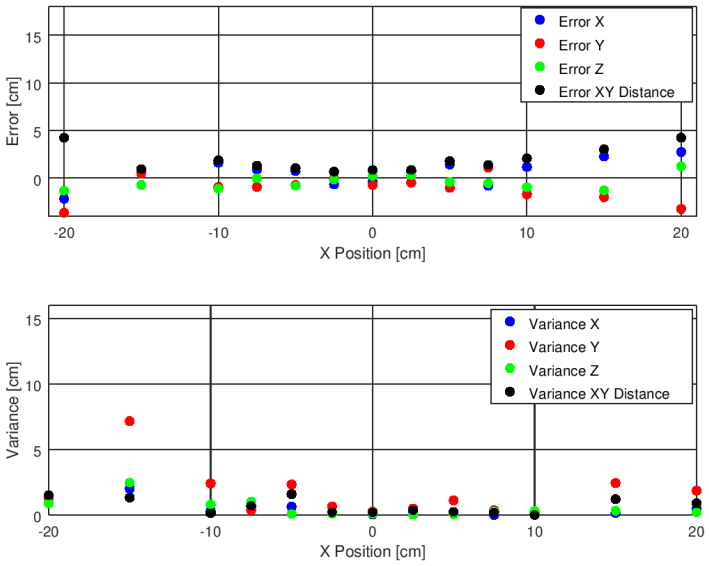

Figure 8.4: Results of the measurement parallel to the Y axis at an offset of 20 cm and at a height of 7 cm

8.2 Results Using MEMS Receivers

The same measurements from the previous chapter were repeated with MEMS receivers. The parameters of the algorithm were not changed, it worked out of the box for these types of sensors as well. The results for the sweep along the X axis are presented in figures 8.5 and 8.6. The measurements at a Y offset of 20 cm can be found in figures 8.7 and 8.8.

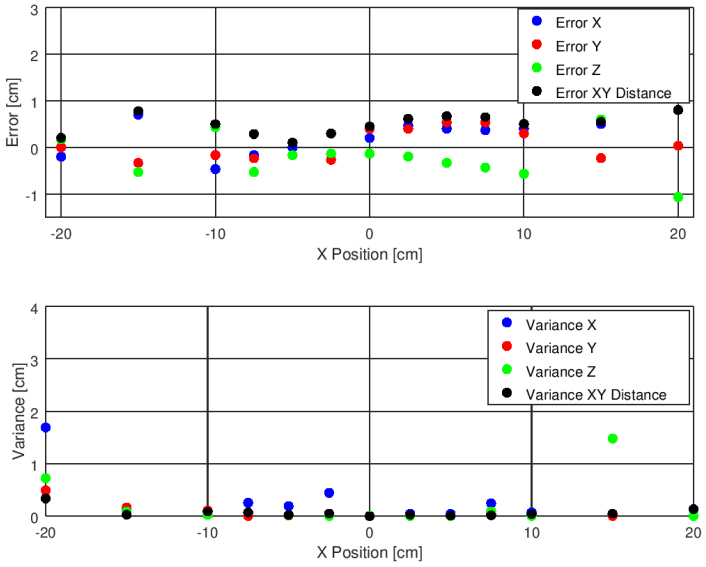

Figure 8.5: Results of the measurement along the Y axis at a height of 15 cm using MEMS receivers

As expected, the performance is exceeding the one of the piezoelectric based receivers, mostly due to the increased bandwidth. The signal edge provides unambiguous trigger points, which are almost always detected correctly by the algorithm.

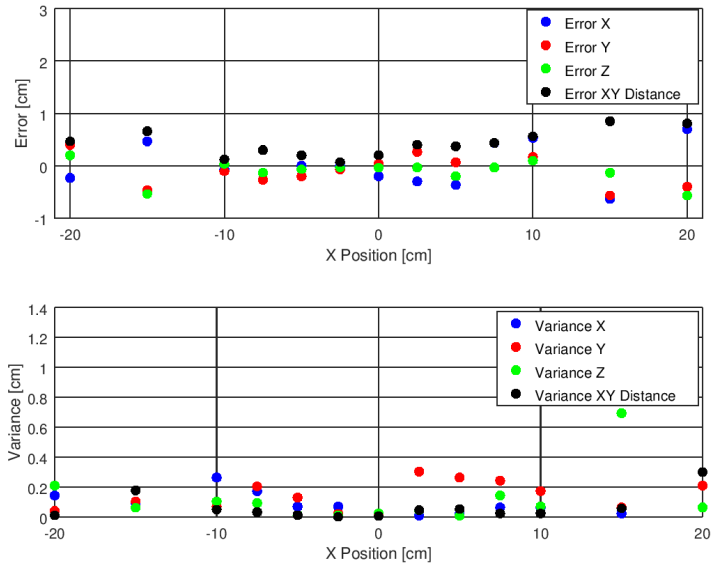

Figure 8.6: Results of the measurement along the Y axis at a height of 7 cm using MEMS receivers

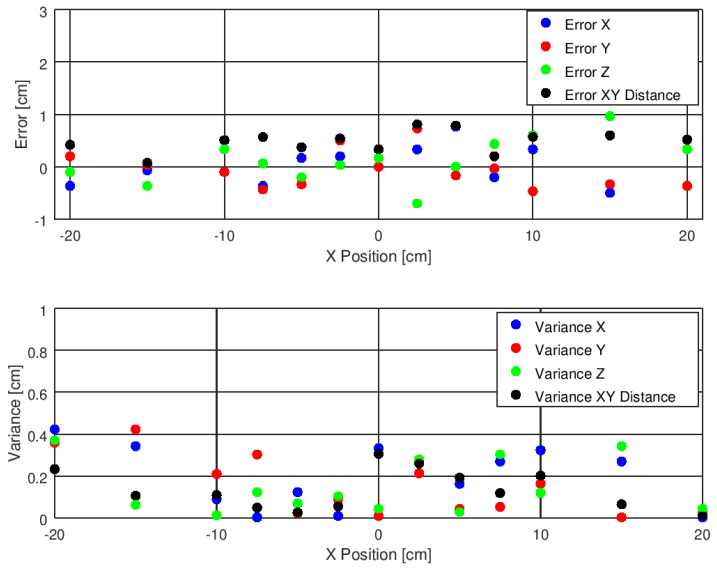

Figure 8.7: Results of the measurement parallel to the Y axis at an offset of 20 cm and at a height of 15 cm using MEMS receivers

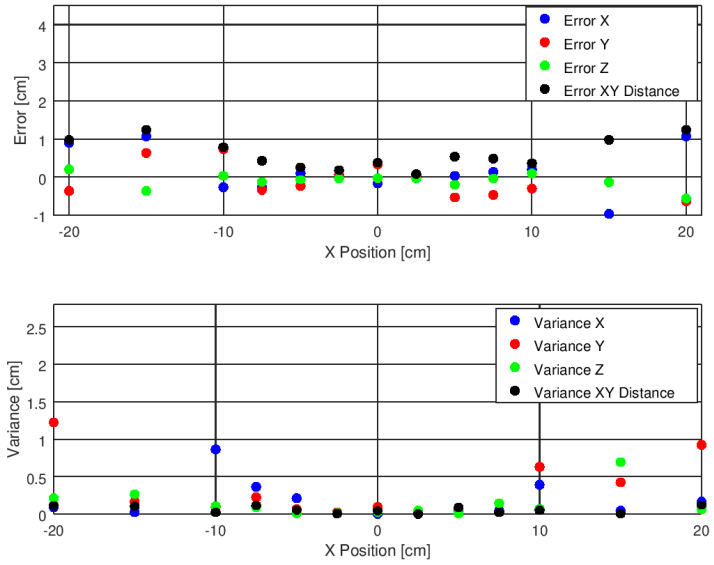

Figure 8.8: Results of the measurement parallel to the Y axis at an offset of 20 cm and at a height of 7 cm using MEMS receivers

9 Future Improvements

In this thesis, a prototype for a positioning system using ultrasound has been developed. As such, the resulting system does not provide the performance necessary to be directly integrated into the robot. For this integration, more development is necessary. However, feedback from other teams developing solutions for the same problem also indicated that using ultrasound is the best approach. The improvement with the highest impact would probably be to switch from piezoelectric receivers to MEMS receivers as already mentioned before in this thesis and repeatedly indicated to Volterio. The higher bandwidth of MEMS receivers compared to piezoelectric ones improves the shape of the received signal by increasing the steepness of the rising edge, ultimately improving trigger algorithm reliability.

Other improvements could be made in the electronic signal conditioning. It was discovered near the end of the project that due to the design of the programmable gain amplifiers which have a small input resistance the phase shift of the signals can vary between individual paths solely because of part tolerances. This can easily be fixed by using a different programmable gain amplifier and did not influence the result in this case.

Regarding the software there are a lot of potential improvements to be made for the trigger algorithm. Its performance heavily depends on the shape of the received signal and the environmental conditions like noise and reflections. There are many options available, and all of them need to be tested in the real system to get clear answers on their performance. The algorithm used in this thesis performed well for the task and was the result of a lengthy trial and error period with wildly different approaches. However, performance could potentially be improved significantly by employing a different algorithm.

At the time of writing it was not clear whether or not the vehicle unit needed to fulfil the automotive standard or what standard exactly applies to the robot. Should the vehicle unit need to be automotive, which is very likely, it would impact the transmitter unit as well as the robot. This is because the wireless module used for the prototype is not available as an automotive part and both the vehicle unit and the robot need to use compatible or the same module. Additionally, the LT3572 ultrasonic transmitter driver IC cannot be used. It would need to be replaced by a discrete boost converter and H-bridge.

Bibliography

[1] *A-14P20 Waterproof type ultrasonic sensor datasheet*. KM.

[2] A.J. Fischer. *The Loran-C Cycle Identification Problem*.

[3] Alan Bensky. *Wireless Positioning Technologies and Applications, 2nd Edition*. 2nd ed. Norwood: Artech House, 2016. ISBN: 1-60807-951-1. URL: https://ebookcentral.proquest.com/lib/gbv/detail.action?docID=4821275.

[4] Billur Barshan. "Fast processing techniques for accurate ultrasonic range measurements." In: ().

[5] Bertrand T. Fang. "Simple Solutions for Hyperbolic and Related Position Fixes." In: *IEEE Transactions on Aerospace and Electronic Systems* 26.5 (1990).

[6] *LT3572 - Dual Full-Bridge Piezo Driver with 900mA Boost Converter Datasheet*. Linear Technology. 2007.

[7] Marek Parusel. "Entwurf und Realisierung von Ultraschallwandlern für die Abbildung von spiegelnden Konturen." PhD thesis. Ruhr-Universität Bochum, 2005.

[8] Ltd. Murata Manufacturing Co. *Ultrasonic Sensor Application Manual*.

[9] P. P. L. Regtien. *Sensors for mechatronics*. 1st ed. Elsevier insights. Amsterdam and New York: Elsevier, 2012. ISBN: 978-0-12-391497-2.

[10] *RM0394 Reference Manual STM32L41xxx/42xxx/43xxx/44xxx/45xxx/46xxx advanced Arm-based 32-bit MCUs*. Rev. 4. STMicroelectronics. Oct. 2018.

[11] S. Sherrit, H.D. Wiederick, and B.K. Mukherjee. *Accurate Equivalent Circuits for Unloaded Piezoelectric Resonators*. 1997.

[12] Volterio. *Volterio (Volterio GmbH) Website*. 2019. URL: http://www.volterio.com/ (visited on 03/21/2019).

[13] Dingguo Xiao, Jiameng Shao, Huiling Ren, and Chunguang Xu. "Design of a high voltage pulse circuit for exciting ultrasonic trans-

ducers." In: *2013 Far East Forum on Nondestructive Evaluation/Testing: New Technology and Application*. IEEE, 2013, pp. 224–230. ISBN: 978-1-4673-6020-3. DOI: 10.1109/FENDT.2013.6635562.

[14] *Zero-Height SiSonic Microphone*. SPU0410LR5H-QB. Revision: H. Knowles. Mar. 2013.