

Beamforming for novel matrix arrays

Fresnel Zone Beamforming and Stolt Migration

Fabian Fool

Delft University of Technology

BEAMFORMING FOR NOVEL MATRIX ARRAYS

FRESNEL ZONE BEAMFORMING AND STOLT MIGRATION

by

Fabian Fool

in partial fulfillment of the requirements for the degree of

Master of Science

in Applied Physics

at the Delft University of Technology,

to be defended publicly on Thursday March 23, 2017 at 13:30 PM.

Supervisor:	Dr. ir. Martin Verweij	TU Delft / Erasmus MC
Thesis committee:	Prof. dr. ir. Nico de Jong	TU Delft / Erasmus MC
	Dr. ir. Michiel Pertijs	TU Delft
	Dr. ir. Rik Vos	Erasmus MC / TU Delft

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

ABSTRACT

Healthcare is a major societal expense and the costs rise each year. Further development of ultrasound could help reduce the costs. Currently a lot of effort is put into developing matrix arrays which allow for volumetric imaging and new applications. There are however multiple problems. Compared to arrays currently in use which have in the order of 128 elements, matrix arrays can easily contain 1000 to 10000 elements. If all elements would be connected independently, the cable would become very thick. Therefore, beamforming methods are required that can operate with fewer transmit and receive channels. Furthermore, the room for electronics on the chip is limited. So, the required electronics for the beamforming methods should be kept simple.

Currently, to focus in transmit all elements have to be wired out. Transmitting an unfocused beam would be an option to reduce the channel count, but this severely limited penetration. So, a new method is required to focus in transmit that requires as few as possible channels. The amount of receive channels also has to be reduced. This might be done by combining the data from each element in some way. However, this results in problems with image formation as the current golden standard, Dynamic Receive Focusing (DRF), uses the data from all individual elements. Thus, we need a new imaging algorithm which can work with fewer channels, but on the other hand have a similar or better image quality.

In this thesis we will propose solutions to the above mentioned problems with matrix arrays. This will be done separately for the transmit and receive part, but they do in no way exclude each other. To be able to focus pulsed waves in transmit we propose a method based on Fresnel zone plates which are used in optics to focus continuous wave light. Our method only requires a single continuous-wave excitation signal to be present, which is connected and disconnected on demand to each element. We have evaluated our method with measurements and simulations. As compared to the conventional focusing method, the spatial resolution is not affected by our method, but the Contrast-to-Noise ratio is 5% lower for shallow depths and up to 20% lower deeper into the medium. Overall though, the differences were relatively small and so it is clear that our new focusing method works very well. If needed, better results can be obtained by trading in frame rate. In this case the results are almost indistinguishable from the conventional focusing method.

To solve the problem with image formation, we have developed a frequency domain two stage beamforming method for use with matrix arrays, which does not require all element data to be present. This has been done for two matrix types. For the first method we have confirmed with simulations that it performs similar to the respective results obtained with two 2D frequency domain two stage beamforming method that have already been experimentally verified. For the second method we have evaluated the performance with simulations and measurements. Our method was able to obtain a 25% better spatial resolution as compared to DRF, without additional artefacts. As an alternative to the last method, we have also developed a frequency domain beamforming method that does require all element data, but only requires a single insonification by a spherical wave. This method did perform worse than the method discussed before in both simulations and measurements, but it does outperform DRF applied to spherical wave data.

CONTENTS

	Page
1 Introduction	1
2 Background knowledge	3
2.1 Wave propagation and extrapolation	3
2.2 Ultrasound imaging	5
2.3 Virtual points sources and receivers	7
2.4 Synthetic Aperture Sequential Beamforming	8
2.5 Stolt migration	9
2.6 Imaging performance	12
2.7 Fast Hankel Transform	13
2.8 Spherical Harmonic Transform	14
I Fresnel Zone Beamforming	17
I.1 Theory	19
I.1.1 Fresnel zone plates and lenses	19
I.1.2 Fresnel zone plates in acoustics	21
I.1.3 Improved Fresnel zone focusing	22
I.2 Simulation and experimental Setup	25
I.2.1 Beam profile	26
I.2.2 Image quality	27
I.3 Results	29
I.3.1 Beam profile	29
I.3.2 Image quality	36
I.4 Discussion and conclusions	45
I.4.1 Discussion	45
I.4.2 Conclusions	46
II Cylindrical Stolt Migration	49
II.1 Method development	51
II.1.1 Derivation	51
II.1.2 Comparison to polar and linear Stolt Migration	53
II.1.3 Implementation	53
II.2 Simulation setup	55
II.3 Results	57
II.4 Conclusions	61

III Spherical Stolt Migration	63
III.1 Method development	65
III.1.1 Derivation	65
III.1.2 Comparison to polar Stolt migration	67
III.1.3 Implementation	67
III.1.4 Computation time	68
III.1.5 Extension to spherical waves	70
III.2 Simulation and experimental setup	71
III.2.1 Setup 1	71
III.2.2 Setup 2	72
III.3 Results	75
III.3.1 Comparison	75
III.3.2 Least squares transformation	77
III.3.3 Computation time	78
III.3.4 Spherical Wave Stolt Migration	79
III.3.5 Verification	82
III.3.6 PSASB vs Stolt	87
III.4 Discussion and conclusions	89
III.4.1 Discussion	89
III.4.2 Conclusion	92
Acknowledgments	93
Bibliography	95
A Hankel function approximation	99
A.1 Approximations	99
A.2 Comparison	100
A.3 Effect on results	105
A.4 Conclusion	105

1

INTRODUCTION

Healthcare is a major societal expense and the costs rise each year. Since 2000 the costs in the Netherlands have more than doubled to 95.5 billion Euro in 2015, or about 5.628 Euro per person.[1] Because of these high costs it is important to find new ways to prevent, diagnose and cure diseases that will have a positive impact on the costs, but also of course on quality of life.

One of the imaging modalities used in healthcare is ultrasound, which uses acoustical waves with a high frequency. An ultrasound system is portable and much cheaper than both an MRI and CT scanner.[2] Also, there is no radiation burden. However, it is not applicable everywhere. Bones reflect most of the ultrasound, which makes imaging the organs behind the ribs difficult. Also, the penetration depth is limited, which becomes especially problematic if there is a thick fat layer, and the resolution is not always sufficient.

Further development of ultrasound could help reduce the costs. Currently a lot of effort is put into developing matrix arrays which allow for volumetric imaging and new applications. There are however multiple problems. Compared to arrays currently in use which have in the order of 128 elements, matrix arrays can easily contain 1000 to 10000 elements. If all elements would be connected independently, the cable would become very thick. This is for example especially a problem for a matrix array to be used in transesophageal echocardiography, where a probe is put into the oesophagus through the throat to image the heart.

A simple way of reducing the required amount of channels that have to be used in transmit is to excite all elements at the same time. In other words, use unfocused beams instead of focused ones. This however severely limits penetration depth in highly attenuating media, especially if high frequencies have to be used to allow for a higher resolution. This is for example the case in the brain of pre-term babies, where one would like to image the brain and measure the perfusion. A lack of perfusion is a common health problem which can lead to cognitive deficits, motor disabilities or even death. So we need a way to focus a beam in transmit with as few as possible channels. There is however another restriction. For small probes, as well as arrays using high frequencies, the electronics should be kept simple. This makes it for example impossible to integrate individual pulsers on the ASIC for each element.

The amount of receive channels also has to be reduced. This might be done by combining the data from each element in some way. However, this results in problems with image formation as the current golden standard, Dynamic Receive Focusing (DRF), uses the data from all individual elements. Thus, we need a new imaging algorithm which can work with fewer channels, but on the other hand have a similar or better image quality.

In this thesis we will propose solutions to the above mentioned problems with matrix arrays. This will be done separately for the transmit and receive part, but they do in no way exclude each other. To be able to focus pulsed waves in transmit we propose a method based on Fresnel zone plates which are used in optics to focus continuous wave light. Our method only requires a single continuous-wave excitation signal to be present, which is connected and disconnected on demand to each element.

To solve the problem with image formation, we will develop a frequency domain two stage beamforming method for use with matrix arrays, which does not require all element data to be present. In literature, two stage beamforming methods in both time and frequency domain have already been developed for linear[3][4] and phased arrays[5][6]. All methods have to shown to outperform DRE. Furthermore, one of these methods has already been extended for use with one kind of matrix arrays[7]. This method is however not applicable to two other kinds of matrix arrays, for which we will develop the methods here. As an alternative, we will also modify one of the developed beamforming methods to work with spherical waves. This does require all element data, but the frame rate can also be much higher.

THESIS OUTLINE

We will first introduce background knowledge in the next chapter that will be helpful in understanding the rest of the thesis. This will include wave propagation, wavefield extrapolation and the basics of ultrasound imaging. Furthermore, two two-stage beamforming methods that are already developed in literature, Sequential Aperture Sequential Beamforming and Stolt Migration, will be introduced.

After this, the thesis branches into three parts. These parts can be read independently from each other. The first part is about Fresnel Zone Beamforming that only requires a single high-voltage signal for focusing. In chapter I.1 of this part the theory behind Fresnel zone focusing and our method are introduced. This will be followed up by the simulation and experimental setup used for the investigation of the new method. Chapter I.3 contains the simulation and experimental results and in the next chapter some open ends are discussed and the conclusions are presented.

Part II is all about the development of a migration method in cylindrical coordinates. In the first chapter of this part the mathematical derivation is given. Also, a comparison with Linear and Polar Stolt is made and an implementation scheme is given. In chapter II.2 the simulation setup is described and the results obtained with this setup are shown in the next chapter. The final chapter of this part contains the conclusions.

In the final part Stolt Migration in a spherical coordinate system is developed and extended to allow for insonification with spherical waves. Like the previous part, the first chapter contains the mathematical derivation and a comparison with Polar Stolt. Also, we derive a quasi fast Hankel transform like algorithm for a possible faster migration and give the implementation scheme. At the end of this chapter, we modify the algorithm to work with spherical waves. Chapter III.2 contains the simulation and experimental setup that is used to evaluate the performance of the newly developed methods in comparison with the current golden standard. In the next chapter the results are presented and discussed. Finally, in chapter III.4 a few open ends are discussed and the conclusions are given.

The appendix contains an investigation in the approximation that we use for the Hankel function in this thesis. We will show how good this approximation is and how it might have affected the results. Furthermore, we introduce two alternative approximations and compare them with the approximation that we have used.

2

BACKGROUND KNOWLEDGE

In this chapter we will introduce some background knowledge that will be helpful for understanding the rest of this thesis. First wave propagation and extrapolation will be discussed. After that, we will look at the basics of image formation in ultrasound. Next various image formation techniques will be introduced, including Sequential Aperture Sequential Beamforming(SASB) and Stolt migration. Then we look at ways to quantify the performance of the image formation techniques. Finally, we will look at the Fast Hankel Transform and the Spherical Harmonic Transform, which are used in the third part of this thesis.

2.1. WAVE PROPAGATION AND EXTRAPOLATION

In short, ultrasound is just a pressure wave with a frequency in the order of MHz. The sound that we can hear are also pressure waves, but the frequency is much lower. The propagation of these pressure waves can be described by the linear wave equation[8]:

$$\nabla^2 p(\underline{r}, t) = \frac{1}{c^2} \frac{\partial^2 p(\underline{r}, t)}{\partial t^2} \quad (2.1)$$

where p is the pressure, c the speed of sound in the medium and ∇^2 the Laplace operator. This wave equation does neglect higher order effects, but for relatively low pressures these effects are negligible. The time domain version of the wave equation as given above is usually very cumbersome to work with. It is much more convenient to Fourier Transform equation 2.1 and work in the frequency domain. In this thesis we will use the following definition of the Fourier Transform

$$\mathfrak{F}(f(t)) = \int_{-\infty}^{\infty} f(t) e^{-2\pi i f t} dt. \quad (2.2)$$

By applying the above definition of the Fourier Transform on the wave equation we get

$$\nabla^2 \hat{p}(\underline{r}, f) + k^2 \hat{p}(\underline{r}, f) = 0 \quad (2.3)$$

where $k = 2\pi f/c$ is the wavenumber. This equation is known as the Helmholtz equation.

The wave equation can be solved in various coordinate system by substituting the corresponding Laplace operator. Solving this equation in a Cartesian coordinate system leads to plane waves that are of the form

$$\Re\left(e^{i(k \cdot r - 2\pi f t)}\right). \quad (2.4)$$

Taking the real part of the above equation is required because pressure is a real quantity, but it is more convenient to do the derivations in the complex domain. The results are different in a cylindrical or spherical coordinate system. In parts II and III the wave equation will be solved in these coordinate systems.

2.1.1. ATTENUATION

Equation 2.4 seems to describe a wave that can travel infinitely far. However, there are always losses due to absorption and scattering. This can be described by making k complex: $k = k' + ik''$. The plane wave equation then reads

$$\Re \left(e^{i(k' \cdot r - 2\pi f t)} \right) e^{-k'' r}. \quad (2.5)$$

We see that there is now a term which attenuates the wave exponentially with distance.

It turns out that k'' , which describes the losses, can be described using a power law[9, p.74]:

$$k''(f) = \alpha_0 + \alpha_1 |f|^\gamma. \quad (2.6)$$

For most media α_0 is zero and γ is usually between 1 and 2. This does mean that higher frequencies are attenuated more and results in a changed pulse shape.

Beside the attenuation, the phase velocity c is also a function of frequency. However, it is often considered to be a small effect and therefore neglected.[9, p.75].

2.1.2. WAVEFIELD EXTRAPOLATION

A commonly used technique in seismology is wavefield extrapolation.[8] Wavefield extrapolation allows us to place a virtual receiver at any position inside the medium.

The basis of wavefield extrapolation is the Kirchhoff integral, which reads:

$$P(\underline{r}_a) = - \oint_S (P \nabla G - G \nabla P) \cdot \underline{n} \, dS, \quad (2.7)$$

where $P(\underline{r}_a)$ is the pressure at point \underline{r}_a inside the volume V with boundary S , G is the Green's function and \underline{n} the outward pointing normal of the boundary S . This equation can be derived from the Helmholtz equation using Green's theorem.

The Kirchhoff integral is almost never used for extrapolation as it requires measurements along a closed surface of both the pressure P and $\nabla P \cdot \underline{n}$, which is related to the velocity perpendicular to the boundary. It is however possible to circumvent these requirements and use the Rayleigh I or II integral instead. The Rayleigh II integral reads

$$P(\underline{r}_a, f) = \frac{z_a}{2\pi} \iint P(r, f) (1 - ik\Delta r) \frac{e^{ik\Delta r}}{\Delta r^3} dx dy, \quad (2.8)$$

where $\Delta r = r - r_a$ and the integral runs over the receiver positions. The Rayleigh I integral looks similar, but requires measurement of $\nabla P \cdot \underline{n}$ instead. Using the above integral we can transform the measured field at the receiver locations to the field that would have been measured if a receiver had been at position \underline{r}_a inside the measured medium. By repeating this for different positions at a certain depth we effectively placed a virtual receiving transducer inside the medium.

In certain cases it might be more efficient to use a different formulation of the above integral. This can be obtained by applying a Fourier Transform in both x and y to obtain

$$P(k_x, k_y, \Delta z, f) = P(k_x, k_y, \Delta z = 0, f) e^{ik_z \Delta z}, \quad (2.9)$$

where $k_z = \sqrt{k^2 - k_x^2 - k_y^2}$. This equation produces similar results as equation 2.8, but extrapolates the whole wavefield to a new depth Δz with one multiplication.

2.1.3. IMAGING CONDITION

Wavefield extrapolation can be used in a three step procedure to construct an image. The first step is to extrapolate the measurement back to a point r_a inside the medium, which is sometimes also referred to as back-propagation. Next the incident wavefield is forward propagated to the same point. Finally, the extrapolated measured data is deconvolved with the incident wavefield and the sample at zero-time is taken. This last step is referred to as the deconvolution imaging condition and can be mathematically described by[10]

$$b(r_a) = \int \frac{P^-(r_a, f)}{P^+(r_a, f)} df, \quad (2.10)$$

where P^- is the back-propagated wavefield, which could be obtained with equation 2.8 and P^+ the forward propagated incident field.

The above equation is almost never used in practice as there is a risk of division by zero. A possible way of solving this is to multiply the top and bottom by the complex conjugate of P^+ and introducing a small stabilization factor ϵ . This results in

$$b(r_a) = \int \frac{P^-(r_a, f) (P^+(r_a, f))^*}{|P^+(r_a, f)|^2 + \epsilon} df. \quad (2.11)$$

There are other ways to stabilize equation 2.10, but all do result in an imperfect deconvolution.

In various occasions we are not interested in the exact amplitude. Furthermore, we don't always know the exact bandwidth of the source. In this case the incident wavefield P^+ can be described by the phase only and is of the form $e^{if(kr)}$. Equation 2.11 then reduces to a cross-correlation integral:

$$b(r_a) = \int P^-(r_a, f) (P^+(r_a, f))^* df. \quad (2.12)$$

We should note that the resolution obtained using the above equation is worse in comparison to the deconvolution imaging condition.

2.2. ULTRASOUND IMAGING

Imaging in medical ultrasound is based on the concept of pulse-echo imaging. Here a pulse is transmitted by a transducer and the reflections of the pulse on the inhomogeneities in the medium are measured by the same transducer. The image is thus a map of the reflectivity scaled with the amplitude of the incoming pulse. This reflectivity is derived from the differences in speed of sound and density of the medium, but we do not image those quantities themselves.

To determine the position of the reflection it is assumed that the speed of sound is constant in the human body. This assumption works well for most soft tissues[9, p.5] and makes it straightforward to relate travel time to depth. Still, the depth in the image might be different from the real depth in the medium due to this assumption. Furthermore, refraction, which is related to the speed of sound differences, is not taken into account and therefore the lateral position might be wrong too. What also might result in an offset is the pulse itself as it does not have its maximum amplitude at zero time. Not taking this into account results in a small offset. This offset does not have to be constant as attenuation could change the pulse shape and thus result in a different position of

the maximum amplitude. Finally, multiple scattering is not taken into account, so this might also introduce artefacts.

A single ultrasound measurement is called an A-mode scan, in which we only get information along a single line. To get a 2 or 3 dimensional image multiple A-mode scans along different scan lines are combined into what is called a B-mode scan.

In the rest of this section we will first describe the different scans that are possible and then we introduce two methods to form each scan line for the final image. Two other imaging methods that do not construct each scan line independently, SASB and Stolt Migration, are discussed after this section.

2.2.1. SCAN MODES

In 2D imaging we can distinguish two kinds of B-mode scans: a linear and a sector scan. Both kinds are displayed in figure 2.1. In a linear scan, displayed in (a), the scan lines are perpendicular to the transducer and parallel to each other. The scan area can therefore only be directly below the transducer. For each scan line a subset of the elements is used. The type of array used for this kind of scan is a linear array that is relatively wide to obtain a large image width and has elements in the order of 1 wavelength or larger. This does result in grating lobes, but because the beam is usually not steered the grating lobes have a minor influence on the image quality.[9, p.183] Also, because the elements are wide, less of them are needed to obtain a large image width.

In a sector scan, displayed in figure 2.1 (b), the scan lines originate from the same point and are steered under different angles. The image width is here determined by the maximum steering angle and therefore not directly on the aperture size. For this kind of scan phased arrays are used, which are smaller compared to linear arrays and have an element size in the order of $1/2$ wavelength. This size is needed to avoid grating lobes for all steering angles, because grating lobes have destructive effects on the image quality in a sector scan.

Volumetric imaging is done in a similar way. There are effectively three kind of scan modes that are obtained by the possible combinations of the 2D scan modes. A linear scan can be done in both directions, a sector scan in both direction, but also a linear scan in one direction and a sector scan in the other. The element and aperture size requirements in each direction depend on how the matrix array is operated.

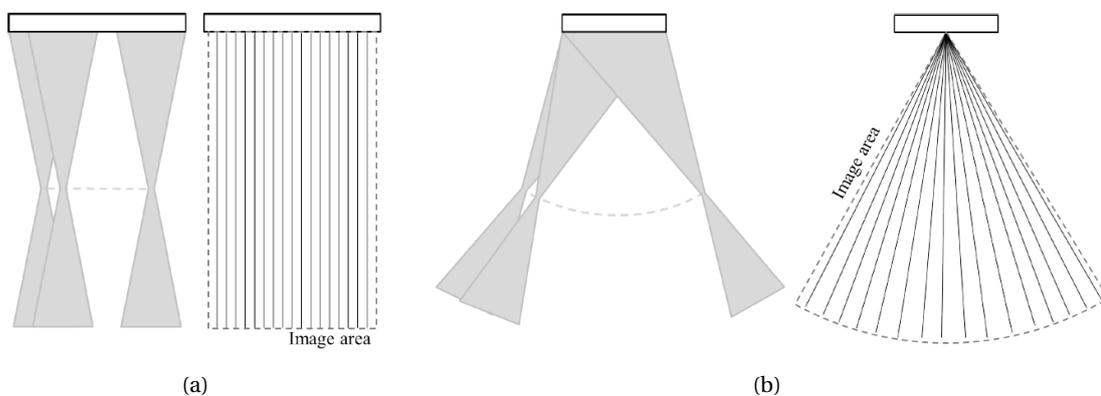


Figure 2.1: In (a) a linear scan is illustrated and in (b) a sector scan. The left part of the subfigure shows the subsequent scan beams that form the scan lines on the right inside. [6]

2.2.2. IMAGE FORMATION

To form each line of the B-mode scan we can use different techniques. The most simple technique is to focus at a single point in transmit and receive by applying the appropriate delays to each element. The transmit and receive focus are not required coincide. After this is done for each scan line, all scan lines are combined to form a 2D or 3D image. The image formed by this technique has the best resolution around the focal points in transmit and receive and deteriorates quickly at other depths. However, because there is only a single focus in receive, it is possible to do the focussing directly on the transducer. In this case only the already combined data has to be send to the computer.

To improve the image quality one would like to focus at multiple points. This can most easily be done in receive if all element data is separately available. In this case, the data can be dynamically focused at every depth in the computer. This technique called Dynamic Receive Focusing(DRF). The resolution is much better away from the transmit focus compared to technique described in the above paragraph. Currently DRF is the golden standard in medical ultrasound image formation.

2.3. VIRTUAL POINTS SOURCES AND RECEIVERS

An important concept for both Synthetic Aperture Sequential Beamforming(SASB) and Stolt Migration, which will be introduced in the next two sections, are virtual point sources and receivers.

To explain this concept we look at an A-mode scan made by using a single focus in transmit and receive, which is illustrated in figure 2.2 for a linear and a sector scan. In this setting the focal point can be seen as a virtual point source and receiver with a limited opening angle. A redatuming of the measured data to the focus transforms it into data from a pulse-echo measurement done by a single point. For a line made with a single focus in transmit and receive the redatuming can be applied by applying a time shift. Events that happen at positive time in the shifted data are pulse-echo reflections from scatters somewhere below the focal points. The events at negative time can after time reversal also be seen as pulse-echo reflections, but now from somewhere between the focus and the transducer. All focal points together from the different scan lines form an array of virtual point sources and receivers that do pulse-echo measurements with a limited opening angle. In case of a linear scan, the virtual array can be found at a depth of z_f , which is the focal distance. For a sector scan, the array lies on a circle with radius r_f .

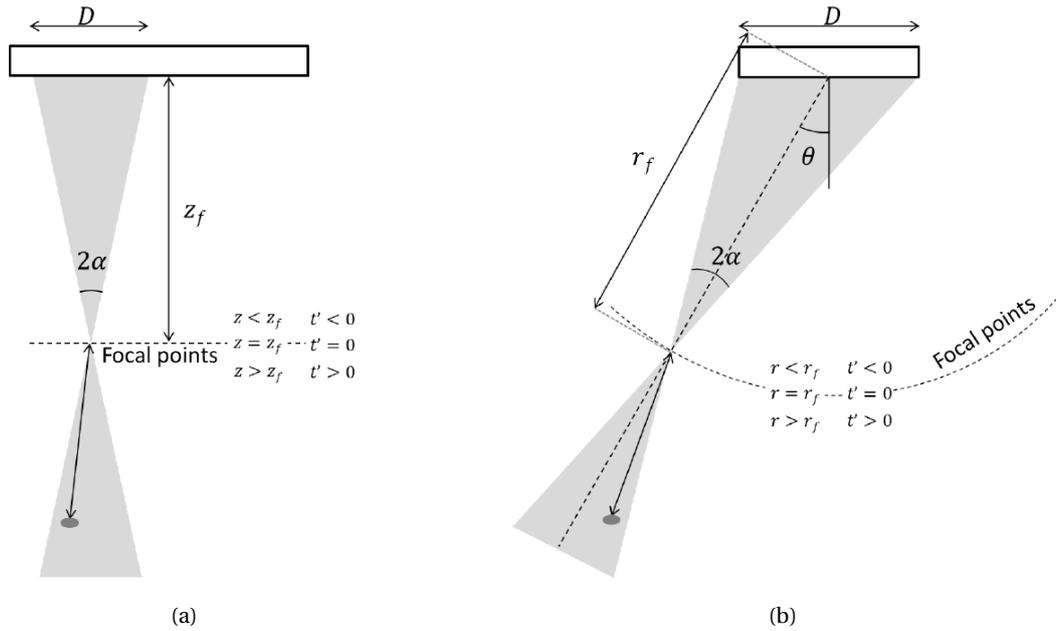


Figure 2.2: In (a) a single A-mode scan for a linear scan and in (b) for a sector scan. [6]

2.4. SYNTHETIC APERTURE SEQUENTIAL BEAMFORMING

Synthetic Aperture Sequential Beamforming (SASB) is a two stage beamforming technique that has been developed by Kortbek et al [3] for linear scans. A version for sector scans has been developed by Bera et al [5] and is named Phased SASB (PSASB). Currently a version for volumetric scans does not exist yet, but the extension should be possible. In this section we will shortly discuss the basic principles of SASB.

In the first stage of SASB low resolution lines (LRL) are created by focusing the transducer in transmit and receive at the same point. This does allow element data combination as only a single output is required per scan line. The focal points of each LRL are considered as an array of virtual point sources and receivers which do pulse-echo measurements. A single generation of such a LRL is shown in figure 2.3(a). In this figure we see an echo projected at the red dot on the LRL. This does however not have to be the real location of the scatterer that generates the echo. The reflection could have come from anywhere on the arc A. This fact is used in the construction of the final high resolution lines that are formed by weighted summing of points from multiple low resolution lines. This is illustrated in figure 2.3(b) where the points of the low resolution lines on arc B are summed to create the blue dot on the high resolution line. If there is really a scatterer at the blue dot, then the contributions from the various lines add up and the scatterer will be visible in the final image.

For a linear array SASB improves the lateral resolution by at least a factor 2 compared to DRF. Furthermore, the PSF width is almost constant with depth, while for DRF it increases linearly with depth. Interestingly, this large improvement is not observed in case of a phased array as the lateral resolution improved by only 20% and the PSF width did increase with depth. Still, in both situations SASB outperforms DRF while the required data throughput is lower as well.

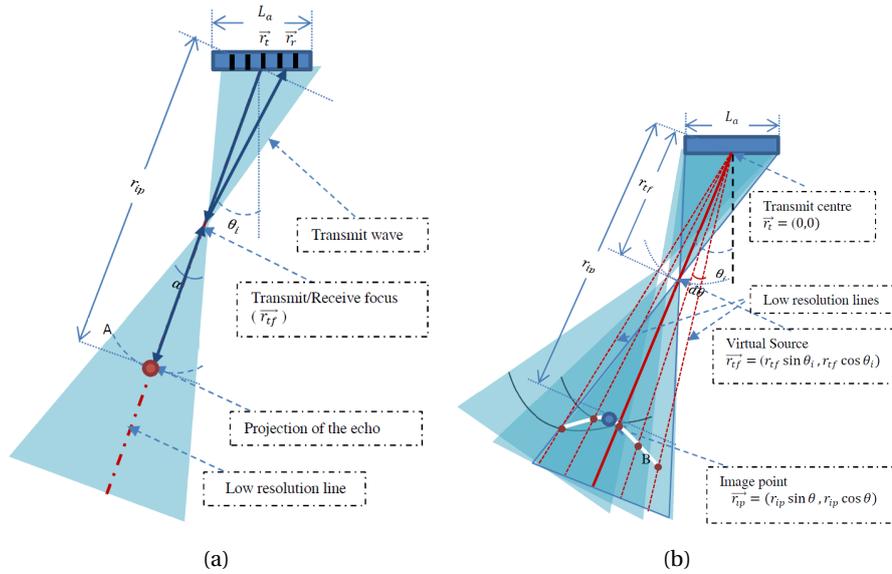


Figure 2.3: In (a) the generation of a single low resolution line and in (b) the construction of a high resolution line by a weighted summation of multiple low resolution lines. [5]

2.5. STOLT MIGRATION

Stolt migration is a technique originally developed for seismology in 1978[11]. Unlike other migration techniques at that time that use finite difference methods, the newly proposed method used the linear wave equation. However, the method is difficult to use with a non-constant velocity. In seismology this is usually the case. This is in contrast to medical imaging, where the propagation velocity differs only by a few percent in most tissues and is therefore assumed constant. [9, p.5] Stolt migration was applied to linear scans by Vos et al[4] and to sector scans by De Wit[6].

The first stage of Stolt Migration is equal to SASB. Therefore, the required data throughput is here also less compared to DRE. The second stage is quite different. Stolt Migration takes into account the whole wave, while SASB only uses time delays based on the geometrical path length. Effectively Stolt Migration thus uses more information than SASB to form the final image and hence one could expect a better resolution compared to SASB, which is what both Vos and De Wit observed.

In the derivations for the linear and phased array version of Stolt Migration the Exploding Reflector Model(ERM) is used. In a pulse-echo measurement the transmitted wave first travels to the scatterer and is then reflected back towards the transducer. In the ERM one assumes that instead the scatterer "explodes" at $t = 0$ and transmits a wave that propagates at an effective wave velocity of $\hat{c} = c/2$ towards the transducer. So instead of a two-way propagation along the same path we are left with a one-way propagation. To obtain an image using this model one has to extrapolate the wavefield back to each point and take the sample at $t = 0$, because at that time the energy was still located at each scatterer.

While this model is very helpful for understanding the situation, it does not take the amplitude correctly into account. Furthermore, it is also very limiting as it is impossible to use it for plane wave imaging without making approximations, because there is no two-way propagation along the same path. In this case the imaging condition described in section 2.1.3 has to be used. This is exactly what Volker [12] for example has done.

In the rest of this section we will first introduce linear and polar Stolt Migration and then look how both methods could be expanded forinsonification with respectively a plane or cylindrical wave.

2.5.1. LINEAR STOLT MIGRATION

In the first stage of linear Stolt migration a set scan lines are created that each have a single focus in transmit and receive. The focus points are used as virtual sources and receivers and the measurement data is redatumed to these points. The following substitutions effectuate this:

$$\begin{aligned} z' &= |z - z_f| \\ t' &= \left| t - 2\frac{z_f}{c} \right| \end{aligned} \quad (2.13)$$

The measurement data now corresponds to a pitch-catch configuration originating from the focus points. The near-field ($z < z_f$) is processed separately from the far-field ($z > z_f$), but the algorithm is similar.

We will skip most of the derivation of this algorithm, which can be found in the paper of Vos[4], and start at the equation which describes how to extrapolate the measured wavefield to a point:

$$p(x, z', t') = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \hat{p}(k_x, f) e^{i\hat{k}_z z'} e^{-i(k_x x - 2\pi f t')} dk_x df. \quad (2.14)$$

In the above equation $\hat{p}(k_x, f)$ is the Fourier transform of recorded field $p(x, z' = 0, t')$ and $\hat{k}_z = \sqrt{\left(\frac{2\pi f}{\hat{c}}\right)^2 - k_x^2}$. The last exponential in the equation is related to the 2D Inverse Fourier Transform, while the first exponential extrapolates the wavefield to a depth z' with velocity $\hat{c} = c/2$. The ERM requires that we evaluate the extrapolated wavefield obtained by equation 2.14 at $t' = 0$, which is the time at which the scatterers "explode". However, due to this the term $e^{2\pi i f t'}$ disappears and we do no longer have an 2D inverse Fourier Transform. Stolt solved this by noting that f is a function of \hat{k}_z and employing a change of integration variables. This yields

$$p(x, z', t' = 0) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{\hat{c} \hat{k}_z}{2\pi \sqrt{k_x^2 + \hat{k}_z^2}} \hat{p}(k_x, f(\hat{k}_z)) e^{-i(k_x x - \hat{k}_z z')} dk_x d\hat{k}_z. \quad (2.15)$$

The integral now describes a 2D inverse Fourier Transform again and the near and far field can be reconstructed using the above equation. Before the above equation can be evaluated, an interpolation step in the frequency domain is required to employ the change of integration variables from f to \hat{k}_z .

It is important in this algorithm to prevent spatial aliasing, which shows up as grating lobes in the final image if the first stage scan lines have a spacing larger than half the effective wavelength. The maximum spacing of the scan lines is given by[4]

$$\Delta x < \frac{\hat{c}}{2f_{max} \sin(\alpha)}, \quad (2.16)$$

where α is the half-width opening angle of the ultrasound field and f_{max} is the maximum frequency contained in the pulse.

Compared to SASB the lateral resolution is similar, but both methods do provide a much better resolution than DRF away from the focus. The advantage over SASB is that Linear Stolt makes use of an FFT, which is very efficient. Furthermore, it can be easily extended to 3D by including an extra FFT for the extra dimension and using a slightly different change of variables.[7]

2.5.2. POLAR STOLT MIGRATION

Linear Stolt migration as described above is incompatible with sector scans. The main problem is that a sector scan can not very well be described in Cartesian coordinates, as the scan lines are all

under a different angle and the virtual sources and receivers are located on a circle. This situation can easier be described in a polar coordinate system. So, the algorithm has been re-derived from the linear wave equation in polar coordinates.

To transform the first stage scan lines, made with a single focus in transmit and receive, into a pitch-catch configuration the following time shift has to be used

$$t' = t - 2\frac{r_f}{c}, \quad (2.17)$$

where r_f is the radius at which the transducer is focused at each scan line.

After a derivation from the Helmholtz equation in cylindrical coordinates, it turns out that the migration is described by[6]

$$p(r, \theta, t' = 0) = \sum_{m=-\infty}^{\infty} \int_{-\infty}^{\infty} \hat{p}(r_f, m, f) \frac{H_m^{(1)}(\hat{k}r)}{H_m^{(1)}(\hat{k}r_f)} e^{im\theta} df, \quad (2.18)$$

where $H_m^{(1)}$ is the cylindrical Hankel function of the first kind, which describes an inward propagating cylindrical wave, $\hat{p}(r_f, m, f)$ is the 2D Fourier Transform in θ and f of the time shifted data and $\hat{k} = 2\pi f / \hat{c}$ the ERM wavenumber, which contains the ERM velocity $\hat{c} = c/2$.

We see that as opposed to the linear case there is no change of coordinates necessary any more, but we now do have Hankel functions instead of the complex exponential. In this form it is impossible to do a 2D IFFT, which significantly decreases computational efficiency. Besides, calculating the Hankel functions itself is computationally expensive. It turns out that both arguments, kr_f and kr , are much larger than 1 for most of the imaging depths, so it is possible to use the asymptotic expansion for large arguments. These expansions can be found for both the modulus and the phase in Abramowitz and Stegun[13, eq.9.2.28-29]. The first few terms in the expansion of the square modulus and phase are:

$$|H_m^{(1)}(\rho)|^2 \approx \frac{2}{\pi\rho} + \frac{4m^2 - 1}{4\pi\rho^3} + \dots, \quad (2.19)$$

$$\arg(H_m^{(1)}(\rho)) \approx \rho - \left(\frac{m}{2} + \frac{1}{4}\right)\pi + \frac{4m^2 - 1}{8\pi\rho} + \dots \quad (2.20)$$

The phase is very important to get right and therefore De Wit took the first three terms for the phase. The third term turned out to be of high importance, but including more terms was unnecessary. Amplitude is of less importance, so only the first term was taken. Using this approximation we can describe the ratio of two Hankel functions as

$$\frac{H_m^{(1)}(\hat{k}r)}{H_m^{(1)}(\hat{k}r_f)} = \sqrt{\frac{r_f}{r}} e^{i\hat{k}(r-r_f)\left(1 - \frac{4m^2-1}{8\hat{k}^2 r_f}\right)}. \quad (2.21)$$

Another difference with the linear case is that the all depths can be reconstructed within a single calculation. Normally for processing the near-field ($r < r_f$) the time axis has to be flipped and $H_m^{(2)}$ has to be used instead of $H_m^{(1)}$. However, instead of reversing the time axis of the image lines, we can make use of the time-reversal property of the Fourier Transform that results in

$$\hat{p}(r_f, m, -f) \frac{H_m^{(2)}(\hat{k}r)}{H_m^{(2)}(\hat{k}r_f)} \rightarrow \hat{p}(r_f, m, f) \frac{H_m^{(2)}(-\hat{k}r)}{H_m^{(2)}(-\hat{k}r_f)} \rightarrow \hat{p}(r_f, m, f) \frac{H_m^{(1)}(\hat{k}r)}{H_m^{(1)}(\hat{k}r_f)}, \quad (2.22)$$

where we have used that $H_m^{(2)}(-\hat{k}r)$ is equal to $H_m^{(1)}(\hat{k}r)$, which can be derived from equation 2.20 together with the fact that $H_m^{(2)}$ is the complex conjugate of $H_m^{(1)}$ for real arguments.[13, eq. 9.1.40] So,

we can process the near-field with exactly the same algorithm and therefore near and far field can be reconstructed in the same calculation.

To prevent spatial aliasing, it is again important that the line density is dense enough. This can be expressed by rewriting equation 2.16 for a sector scan as

$$\Delta\theta \leq \frac{\hat{c}}{2r_f f_{max} \sin(\alpha)}. \quad (2.23)$$

For a small half width opening angle α the above equation reduces to [6]

$$\Delta\theta \leq \frac{\hat{c}}{Df_{max}}, \quad (2.24)$$

where D is the aperture width. The above equation can also be derived from the Nyquist frequency, indicating that it just meets the requirement to reconstruct the smallest details possible.

In comparison to PSASB with a hamming window, Stolt migration improves the lateral resolution by 25% and compared to DRF even up to 50%. However, if an equal weighting is used when forming the HRL in PSASB, the same resolution can be achieved as Stolt at the expense of extra artefacts in the final image.

2.5.3. PLANE WAVE STOLT MIGRATION

The ERM only works for pulse-echo type measurements where there is two-way propagation along the same line. This is the case when the concept of virtual sources and receivers are used in Linear Stolt Migration, but not in case of plane wave imaging(PWI). Garcia et al[14] still tried to make ERM compatible with PWI by approximating the travel times, however it is also possible to use a different but exact approach presented by Volker[12] to make Linear Stolt compatible with plane waves. This will be discussed in this section.

The approach by Volker is based on the imaging condition of equation 2.12 where the extrapolated measurement is deconvolved with the forward extrapolated incident wavefield. Here, the incident wavefield is a plane wave that can be described by $e^{-ik_z \Delta z}$. The measured data can be extrapolated using the Rayleigh II integral of equation 2.9.

To incorporate this into the framework of Linear Stolt a few changes have to be made. First of all no virtual sources and receivers are necessary any more, so the space and time coordinates do not have to be shifted. Second, we have to replace one of the exponentials in equation 2.14. The $e^{\hat{k}_z z'}$ both extrapolates the measured data and describes the incoming wave. For PWI this exponential has to be replaced by $e^{ik_z \Delta z} e^{ik_z \Delta z}$. We note that \hat{k}_z contains the ERM velocity \hat{c} , while k_z does not. The last change is the different change of variables to get a 2D-IFFT. After changing these 3 things, Linear Stolt is compatible with PWI.

2.6. IMAGING PERFORMANCE

Objectively measuring imaging performance is difficult as the interpretation of the image is highly dependent on the physician and equipment settings. In this section we will present objective metrics that are commonly used in literature and try to quantify the image quality. [15]

2.6.1. SPATIAL RESOLUTION

A point object can never be imaged as a point and will always end up as a blob. This blob is called the point spread function (PSF). The final image in an imaging system can be seen as the object, which

consists of many infinitesimally small points, convolved with the PSF. Therefore, the fine details of the object are blurred out. A small PSF is therefore an important part of a good imaging system.

For simplicity, it is assumed that the PSF of an ultrasound system can be split up along three axes. The axial PSF, which spans along the depth or radial axis, is mainly determined by the transmitted pulse. The shorter the pulse, and thus the larger the bandwidth, the lower the axial PSF. There is a limit to how short the pulse can be, as all transducer have a finite bandwidth. The two lateral PSF's are perpendicular to each other. The width is mainly determined by the beam profile and can vary with depth.

The measure that is mainly used to quantify the PSF is the full width at half maximum(FWHM), which is equal to the width at -6 dB. A smaller FWHM is usually better, but it also important that the PSF the does not have a long tail. This can be quantified using the full width at tenth maximum(FWTM), which is the width of the PSF at -20 dB.

For continuous waves there exists a theoretical lateral PSF width for a transducer without apodization in the far field that is given by[16]

$$PSF_{-6dB} = 0.88 \frac{\lambda z}{D}, \quad (2.25)$$

where λ is the wavelength, z the depth and D the aperture width. This measure does not work well with ultrasound as usually pulsed waves are used that contain many frequencies. Furthermore, apodization is also common. Therefore, De Wit[6] has developed a tool that calculates the directivity pattern while taking into account the different frequencies and the apodization. From the directivity pattern we can derive a more accurate lateral PSF width.

2.6.2. CONTRAST RESOLUTION

It is important in medical ultrasound to be able to distinguish different kind of tissues. Each kind of tissue backscatters a different amount of the incoming ultrasound. In literature there are various ways to quantify this. Here we will use the contrast-to-noise-ratio which is defined as[17]

$$CNR = \frac{\mu_b - \mu_c}{\sqrt{\sigma_b^2 + \sigma_c^2}}, \quad (2.26)$$

where μ_b and μ_c are the mean amplitudes before log-compression in the background and a hypoechoic cyst region and σ_b^2 and σ_c^2 the variance in those regions. The background and cyst region are equal in size and found on the same depth.

2.7. FAST HANKEL TRANSFORM

The Hankel Transform is defined as[18]

$$F(k) = \int_0^{\infty} f(r) J_n(kr) r dr, \quad (2.27)$$

where J_n is the n-th order Bessel function of the first kind. The inverse transform reads

$$g(r) = \int_0^{\infty} F(k) J_n(kr) k dk. \quad (2.28)$$

Evaluating one of these integrals numerically for a grid of N frequencies and N radii would require $O(N^2)$ multiplications, which results in lengthy calculations for a dense grid. Therefore, various algorithms have been explored to reduce the computation time. Here we present the Quasi fast Hankel transform(QFHT), introduced by Siegman[19].

The QFHT introduces the following change of variables

$$r = r_0 e^{\alpha x} \text{ and } k = k_0 e^{\alpha y}, \quad (2.29)$$

where k_0 , r_0 and α could be determined by the relations found in the paper by Siegman. The above change transforms equation 2.27 into an equation of the form

$$\hat{F}(y) = \int_{-\infty}^{\infty} \hat{f}(x) \hat{J}_n(x+y) dx. \quad (2.30)$$

This new form looks like a correlation integral that can be evaluated by three Fourier Transforms:

$$\hat{F}(y) = \mathcal{F}^{-1} [\mathcal{F}^{-1} (\hat{f}(x)) \mathcal{F} (\hat{J}_n(x))]. \quad (2.31)$$

This is faster as it scales as $O(N \log(N))$ instead of $O(N^2)$. However, it requires an exponentially spaced grid in both k and r , which do not appear in many problems. It is possible to get around this problem by using a Non Uniform Fast Fourier Transform (NUFFT) as was done by Liu and Zhang. [20] Various implementations of the NUFFT, which go from an equispaced grid to a non-uniform one, exist which mainly consist of an oversampled FFT and a smart interpolation. [21]

2.8. SPHERICAL HARMONIC TRANSFORM

The forward Spherical Harmonic transform, also called analysis, is defined as [22]

$$\tilde{f}(l, m) = \int_0^{2\pi} \int_0^\pi f(\varphi, \theta) \overline{P_l^m(\cos \varphi)} e^{-im\theta} \sin \varphi d\varphi d\theta, \quad (2.32)$$

while the inverse transform, sometimes called synthesis, is given by

$$f(\varphi, \theta) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \tilde{f}(l, m) \overline{P_l^m(\cos \varphi)} e^{im\theta}. \quad (2.33)$$

In both equations $\overline{P_l^m}$ are the normalized associated Legendre polynomial which are related to the unnormalized polynomials by

$$\overline{P_l^m}(\cos \varphi) = \sqrt{\frac{2l+1}{4\pi}} \sqrt{\frac{(l-|m|)!}{(l+|m|)!}} P_l^m(\cos \varphi). \quad (2.34)$$

The normalized polynomials together with the complex exponentials are called the Spherical Harmonics and form an orthonormal basis on a sphere.

2.8.1. DISCRETIZATION

To evaluate the transforms numerically they have to be discretized. From Nyquist's criterion we can determine the maximum l and m . For a spacing of $\Delta\theta$ and $\Delta\varphi$ between the angles this is given by

$$|m| \leq \frac{\pi}{\Delta\theta} \text{ and } l \leq \frac{\pi}{2\Delta\varphi}. \quad (2.35)$$

This means that the number of angles should at least be $N_\varphi = l_{\max}$ and $N_\theta = \min(N_\varphi, 2m_{\max})$. The last comes from the fact that the maximum m also has to be smaller than l , otherwise the associated Legendre polynomials are zero.

For θ we can use a FFT if the angles are equispaced. However, it turns out that equally spacing the inclination angles φ is not the best choice. Choosing the angles corresponding to the zeros of the Legendre polynomial of order l_{\max} is more convenient as it is then possible to use Gauss-Legendre quadrature. This turns the integral into the following summation:

$$\int_0^\pi f_m(\varphi) \overline{P_l^m}(\cos \varphi) \sin \varphi d\varphi = \sum_{i=1}^{N_\varphi} f_m(\varphi_i) \overline{P_l^m}(\cos \varphi_i) w_i. \quad (2.36)$$

In the above equation f_m are the Fourier coefficients for θ , φ_i the i -th zero of $P_{l_{\max}}$, which are also called Gauss nodes, and w_i the weights. A Gaussian quadrature rule, like the one above, is exact if f_m is a polynomial of order less than l_{\max} .

By using the above we can now write the forward and inverse SHT as

$$\tilde{f}(l, m) = \sum_{i=1}^{N_\varphi} \sum_{j=1}^{N_\theta} f(\varphi_i, \theta_j) \overline{P_l^m}(\cos \varphi_i) e^{im\theta_j}, \quad (2.37)$$

$$f(\varphi_i, \theta_j) = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l \tilde{f}(l, m) \overline{P_l^m}(\cos \varphi_i) e^{-im\theta_j}, \quad (2.38)$$

which can be evaluated numerically.

2.8.2. GAUSS NODES

Finding the zeros of a Legendre Polynomial is a non-trivial case. However, there exist an expression for the approximate position of the i -th zero of the polynomial of order l which reads[23]

$$\varphi_i = \left(1 - \frac{1}{8l^2} + \frac{1}{8l^3} - \frac{1}{384l^4} \left[39 - \frac{28}{\sin^2 \phi_i} \right] + \dots \right) \cos \phi_i, \quad (2.39)$$

where ϕ_i is given by

$$\phi_i = \frac{l - i + 3/4}{l + 1/2} \pi. \quad (2.40)$$

If required, the approximation can be improved by using Halley's method[24] or a similar algorithm.

2.8.3. LEAST SQUARES TRANSFORMATION

The Gauss-Legendre quadrature requires a very specific set of inclination angles. It is also possible to use equispaced inclination angles, but then twice the amount of angles is required. Another possibility is to use a least squares transformation with an arbitrary set of angles, which gives coefficients that best fit the data.

The least squares Spherical Harmonic transformation can be derived from equation 2.38 that in matrix form reads

$$\underline{F} = \underline{Y} \underline{C}, \quad (2.41)$$

where \underline{Y} is a matrix containing the spherical harmonics, \underline{F} is a vector which contains the function values at the grid points and \underline{C} is a vector with the coefficients $\tilde{f}(l, m)$. To obtain the coefficients one has to invert the equation, which can be done by multiplying with the inverse of \underline{Y} :

$$\underline{C} = \underline{Y}^{-1} \underline{F}. \quad (2.42)$$

This equation describes the least squares forward Spherical Harmonic transformation, although it also possible to define it differently. Note that the matrix \underline{Y} is usually not square and therefore a pseudo-inverse has to be used.

I

FRESNEL ZONE BEAMFORMING

I.1

THEORY

The current way of focusing used in medical ultrasound is by sending signals to each element with a different delay. This delay is such that the pulse from each element arrives at exact the same moment at the focal point. The required delay for each element as compared to the centre of the probe is mathematically described by

$$t_d(x, y) = \frac{1}{c} \left(\sqrt{(x_c - x_f)^2 + (y_c - y_f)^2 + z_f^2} - \sqrt{(x - x_f)^2 + (y - y_f)^2 + z_f^2} \right) \quad (\text{I.1.1})$$

where $(x_c, y_c, 0)$ are the coordinates corresponding to the centre of the transducer, $(x, y, 0)$ the coordinates corresponding to the centre of an element and (x_f, y_f, z_f) the coordinates of the focal point.

As long as the number of elements stays small we can use this kind of focusing, but for transducers used in 3D image formation the amount of elements becomes so large that it is impossible to send a unique signal to each element. Therefore, we will need a different method to form the transmit beam.

In this part we will introduce a focusing method that does not require all elements to be wired out independently. The idea comes from Fresnel zone plates that are used in optics, which provide a simple alternative to lenses for focusing continuous waves. The Fresnel zone plate and its varieties will be introduced in the next section. Then, we will consider a current application in acoustics. At the end of this chapter our new method is introduced.

I.1.1. FRESNEL ZONE PLATES AND LENSES

The normal way of focusing continuous-wave light in optics is with a lens that works on the concept of refraction. The most simple focusing lens consists of a single convex surface at which the light is refracted towards the focal point. This lens is shown in the right of figure I.1.1. While the concept is simple, the lens always introduces some degree of distortion or aberration which will cause imperfect focusing. [25, ch.5]

A different kind of lens is the Fresnel Zone Plate. The simplest one consists of rings which are either opaque or transparent and is shown in figure I.1.1 on the left. Instead of refraction, they use diffraction to focus the light. Only the light that constructively interferes at the focus is transmitted. Crudely, these are the parts that have a path length difference of less than half the wavelength compared to the centre, but also the parts that have a path length difference between 1 and 1.5 wavelength and so

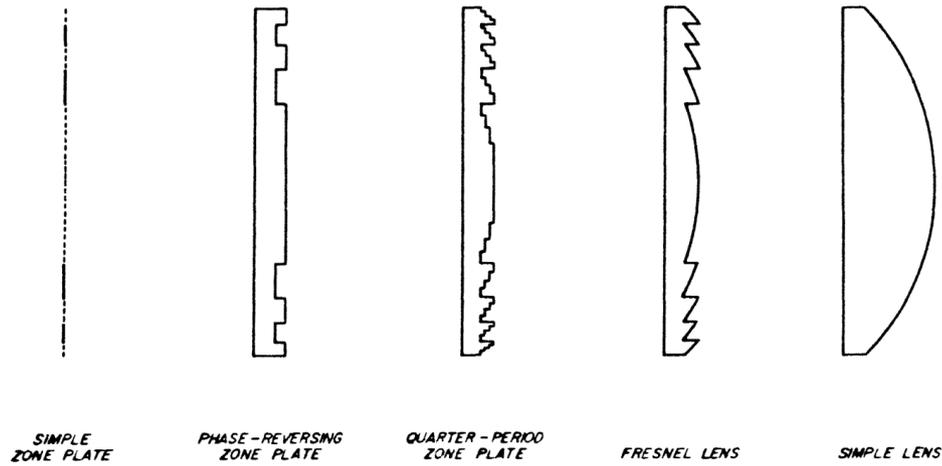


Figure I.1.1: An overview of the different kind of zone plates and lenses.[26]

on. This means that we must have a zone transition from transparent to opaque or vice versa at radii where the path length difference is $\frac{n\lambda}{2}$, where λ is the wavelength and n an integer. We can describe this by the following equation[27]

$$\frac{n\lambda}{2} = \sqrt{r^2 + z_f^2} - z_f, \quad (\text{I.1.2})$$

where r is the radius from the centre and z_f the distance of the focal spot. This equation can be rewritten to get the radii at which the transitions occur

$$r = \sqrt{\frac{n\lambda}{2} \left(\frac{n\lambda}{2} + 2z_f \right)}. \quad (\text{I.1.3})$$

While this zone plate is simple, it also blocks almost half the light and thus has a low efficiency. This can be solved by replacing the opaque zone by a phase correcting zone, which gives the light a shift of 1/2 periods compared to the transparent zones, which corresponds to a phase shift of π . For continuous waves this means that the maximum error in the focus will be brought to within 1/2 periods and therefore all light will constructively interfere at the focal point. This will significantly increase the efficiency, but the light transmitted at the edges of each zone still has a low contribution to the total intensity. This can be improved by introducing subzones. For example by using four zones the maximum difference at the focal point is less than 1/4 periods. The amount of subzones can be increased until we effectively have a continuous distribution, which will bring the light to a perfect focus. This kind of zone plate is called a Fresnel Lens.

It is however unnecessary to introduce a large amount of subzones to get a good intensity at the focus. We can see this when comparing the main-lobe amplitudes of the PSF in the focal plane for an array with quantized delays, as calculated by Peterson et al[28]. The relative main lobe amplitude between a Fresnel lens and a zone plate with N zones is

$$A(N) = \frac{\sin(\pi/N)}{\pi/N}. \quad (\text{I.1.4})$$

By using this equation we can get the relative amplitude and intensity for a zone-plate with 2, 4 and 8 zones, which is shown in table I.1.1. While the step from 2 zones to 4 zones is quite large, but the step from 4 to 8 is already significantly smaller. Thus, further increasing the number of zones seems unnecessary in case of continuous waves. Most likely side-lobes that appear due to the quantization will also contribute a bit to the amplitude at the focus, which makes the given numbers the worst case scenario.

Table I.1.1: Amplitude and intensity for a zone plate with 2, 4 and 8 zones relative to a Fresnel lens.

Number of zones	Relative amplitude	Relative intensity	Intensity (dB)
2	0.64	0.41	-3.9
4	0.90	0.81	-0.91
8	0.97	0.95	-0.22

I.1.2. FRESNEL ZONE PLATES IN ACOUSTICS

The concept of Fresnel zones has also been applied in acoustics. In this section we will look at a recent implementation by Nguyen et al.[29] They build up an ultrasound system which uses a limited set in delays for both the transmit and receive part. Here we will only look at the transmit part.

The proposed transmit part by Nguyen consists of 4 channels. Before every transmit, each element is connected to one of the channels using a set of switches. This means that it is only possible to send 4 unique time delays to all elements. The delays are chosen using the following method: First the time delays given by equation I.1.1 are quantized to multiples of $1/4$ periods and then the modulo operator is applied to bring all delays within 1 period. The result is shown in figure I.1.2(a).

Effectively this is an application of the quarter-period zone plate that is used for focusing continuous-waves on the pulsed waves in ultrasound. For continuous waves the modulo operator has no effect as the wave is periodic, but this is not the case for pulsed waves. For the centre elements in pulsed ultrasound this is not a problem yet, as the delay is still less than 1 period. The problem is with the outer elements, which have a timing error larger than 1 period. The timing errors becomes even worse for lower f-numbers as the delays for the outer elements increase. This is also the case for steered beams. Also, a higher centre frequency worsens the error as one period becomes shorter. This method is therefore certainly not applicable in every situation, but for specific settings the authors claim that the whole system can achieve an image quality similar to an image made without restrictions on the delays.

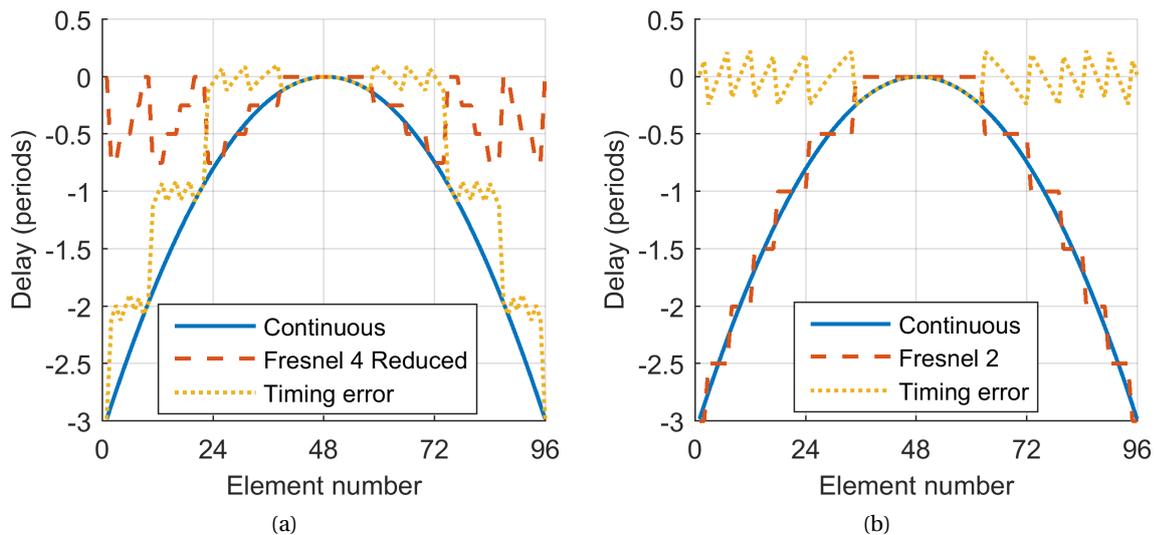


Figure I.1.2: Delay profiles and timing errors for (a) the method of Nguyen (denoted as Fresnel 4 Reduced), and (b) the improved version (named Fresnel 2).

I.1.3. IMPROVED FRESNEL ZONE FOCUSING

The approach presented above only requires 4 unique transmit signals, but this can still be too much to include on a small matrix arrays. Furthermore, applying a concept used for continuous waves on pulsed waves brings various problems. That is why we have developed a different approach that can excite the elements by connecting them to a single high voltage signal. The switching is accomplished by local transistors and timing hardware.

We will not dive into the exact electronic implementation here and will jump straight to the effect on the delays. Effectively what this method does is quantize the delays to multiples of $1/2$ periods, without applying the modulo operator as is done in the method by Nguyen. This is illustrated in figure I.1.2(b). The maximum timing error always stays within $1/4$ periods. This is clearly lower than the method of Nguyen could achieve for the elements at the edge of the array. We therefore expect better results from this method, especially at low f -numbers, high frequencies and steering angles.

If we could use multiple high voltage signals, we can extend the approach to effectively quantize the delays to $1/(2n)$, where n is the number of available high voltage signals, and further reduce the timing errors. For continuous waves we saw that the first few steps of finer quantization were the most important. We also expect that this is the case for pulsed waves. However, as compared to continuous waves, the sidelobes appearing due to the quantization will be of lower intensity with pulsed waves and thus the main lobe intensity increases. This is because pulsed waves contain many frequencies and the position of the sidelobes depends on the frequency.[28] Therefore, the sidelobes are diffused out. This is exactly what Magnin et al[30] observed, although for a quantization to multiples of $1/2$ periods the first pair of sidelobes, which have the largest intensity of all pairs, is effectively unchanged.

I.1.3.1. FURTHER IMPROVEMENT

Quantizing the delays to multiples of $1/2$ periods is quite coarse. It might be required to use a finer quantization, even though it might not possible to use extra channels. There is a way to use a finer quantization at the cost of the number of frames per second. The idea is to emulate a quantization of $1/(2n)$ by using n different pulses as illustrated in figure I.1.3(a) for a sequence where 2 shots are used per scan line. In the first shot we only use the elements which are best approximated by a delay of $N/2$ periods, where N is an integer, and in the second shot only the elements are best approximated by a delay of $1/4 + N/2$ periods.

While this might improve image quality, it also has major downsides. First of all this method reduces the number of frames per second. This can result in 'jerkiness' when only an image is formed after taking all the frames required for a single compounded image. To avoid this an interlacing technique can be applied, after each frame where one of the two shots is used for a scan line the image is updated. This approach is also taken for compound imaging and is illustrated in figure I.1.4. While this does prevent the 'jerkiness', it does cause persistence or image blurring artefacts if the transducer or target moves too rapidly.[31] The second downside is that this approach only works well if the system can be conceived as linear. In this case the two frames can be added without any problem to form the final image. In case of large non-linearities artefacts may show up.

Another problem that may occur in attenuating media is that the pressure becomes too low. Each shot uses about half of the elements and thus less pressure is generated per shot. This might obscure deeper lying reflectors and lower signal-to-noise levels. A possible way to solve this is by using partly overlapping shots. By this we mean that each pulse also includes elements that are normally only included in the other pulses. This overlap illustrated in figure I.1.3(b). If we compare this with figure I.1.3(a), we can clearly see that extra elements are included in each shot. This does increase the

pressure, but the extra elements might also decrease the imaging performance.

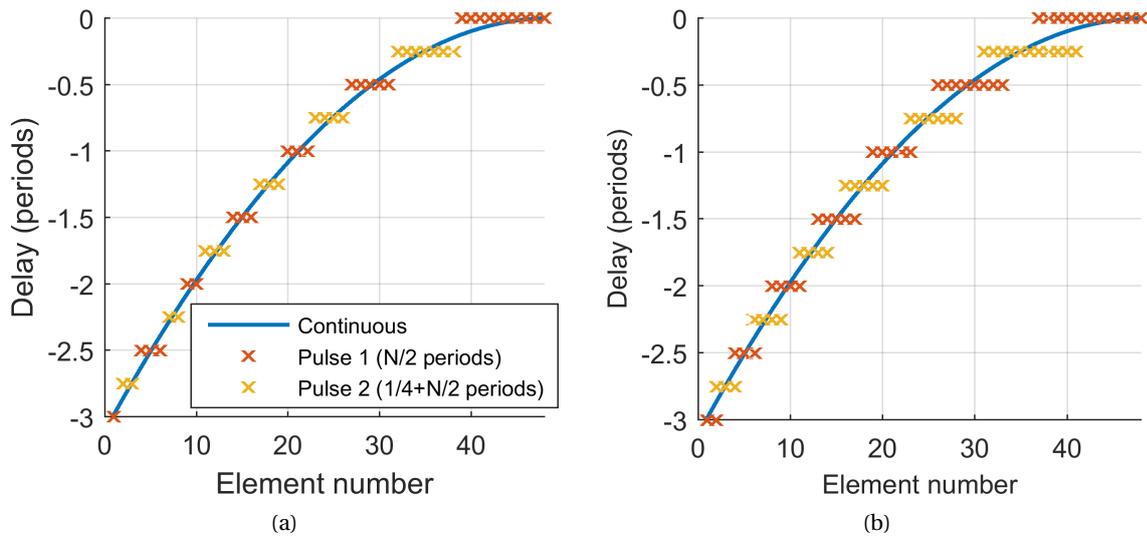


Figure I.1.3: In (a) a sequence used for emulating a delay quantization to multiples of 1/4 periods and in (b) a sequence that has some overlap between the two shots to compensate for the lower pressure per pulse.

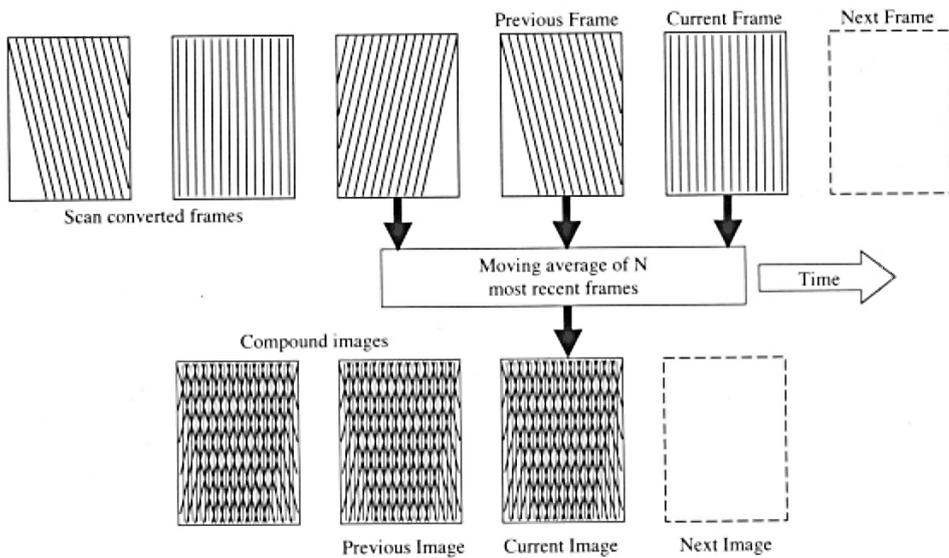


Figure I.1.4: A real time spatial compounding sequence. For each formed image the N most recent frames are used.[31]

I.2

SIMULATION AND EXPERIMENTAL SETUP

To assess the performance of the improved Fresnel zone focusing we will do simulations and measurements. Two versions of our method will be used: one with the delays quantized to $1/2$ periods and one quantized to $1/4$ periods. We will refer to these methods as Fresnel 2 and Fresnel 4. These two versions are compared to continuous delays and the 4 channel version of the implementation by Nguyen, which is referred to as Fresnel 4 Reduced. The described methods will be compared in two ways: the transmit beam profile and the image quality. The transmit beam profile will only be simulated, while images will be created by both simulations and experiments.

All simulations are done with FieldII[32][33]. The simulated transducer is a conventional phased array probe (ATL P4-1, Philips Ultrasound, Bothel, Washington, USA) with a centre frequency of 2.5 MHz and 96 elements. This probe was also used in the measurements, during which it was connected to a Verasonics Vantage System (Verasonics Inc., Kirkland, Washington, USA). The specific settings used during the simulations and measurements are summarized in table I.2.1.

Table I.2.1: Overview of the settings used during the simulations and measurements

	Parameter	Value
Transducer	Centre frequency	2.5 MHz
	Number of elements	96
	Pitch	295 μm
	Element width	245 μm
	Apodization	None
Scan settings	Number of scan lines	256
	Scan angle	$\pm 45^\circ$
	Focal Depth	30 mm
Verasonics	Excitation	2 half cycles
	Excitation type	Bipolar
	Sampling Frequency	10 MHz
FieldII	Excitation	4 cycle Gaussian windowed sinusoid
	Sampling Frequency	100 MHz
Wire phantom	Wire material	Copper
	Wire diameter	120 μm
	Speed of sound of medium	1480 m s^{-1}
	Attenuation of medium	Negligible
Tissue phantom	Wire target material	Nylon
	Wire diameter	100 μm
	Speed of sound of medium	1540 m s^{-1}
	Attenuation of medium	0.5 $\text{dB MHz}^{-1} \text{cm}^{-1}$

I.2.1. BEAM PROFILE

We will look at the beam profile in two ways. First the effect of the 4 methods on the pressure distribution is investigated. The difference should be as small as possible between the Fresnel methods and continuous delays. We expect that the more they differ, the worse the final image will be. However, most important is to get high pressure along the scan line, which is the area of interest, and at deeper depths, which is mainly important if the attenuation of the medium is high. Furthermore, sidelobes should be kept to a minimum. The methods will be compared in two situations: a focused beam at 0 degrees and one at +30 degrees. The focal distance from the centre of the transducer for both situations is 30 mm, which corresponds to an f-number of 1.1.

To get the beam profile, the pressure as function of time is simulated over an equispaced grid in the x-z plane with a spacing of $\lambda/10$. After the simulation, the absolute value of the Hilbert transform is taken over the time data at each point and the maximum value was taken, which represents the peak pressure. Before display the data is log-compressed. The medium in these simulations will be water and the attenuation is taken to be negligible.

The second way we look at the beam profile is by investigating the pulse arrival time from each element on the scan line. This is important because it is assumed that travel time is proportional to the depth. If this assumption is invalid, then scatterers will be displayed at wrong positions in the final image. Furthermore, if the difference between the first and last pulse to arrive is large, the points will appear stretched in the radial direction.

I.2.2. IMAGE QUALITY

The beam profile itself does only tell half of the story. The final image is what counts. This will be investigated with simulations and measurements. A sector scan will be done with 256 scan lines within a scan angle of 90 degrees and the probe will be focused at a depth of 30 mm. No apodization will be used in transmit and receive.

In transmission the four methods have different delays, but in receive DRF is used for all methods. To prevent discretization errors the measurement data was upsampled from 10 to 60 MHz. The simulation data was already sampled at 100 MHz and upsampling was therefore not required.

Two kinds of phantoms will be imaged. The first one is a wire phantom that consists of copper wires that are strung within two plates as shown in figure I.2.1(a). During the measurement the phantom was placed in a water tank with anti-reflection material at the sides as shown in figure I.2.1(b). The goal of these measurements was to see how the various methods affect the PSF. The FWHM and FWHM in the lateral direction are used as a measure to investigate this.

The second phantom is a tissue mimicking phantom (040-GSE, CIRS, Norfolk, Virginia, USA) that consists of nylon wire and cylindrical cysts. The layout of and the setup for measuring this phantom is shown in I.2.2. With this phantom we have investigated how the CNR is affected by the different methods. For this investigation 5 of the 6 anechoic stepped cylinders were used, which can be found at a depth of 4.5, 7, 10, 13 and 16 cm. The shallowest measured cylinder has a diameter of 6.6 mm and the other ones have a diameter of 10 mm. For the calculation of the CNR only data within a diameter of respectively 4.6 mm and 8 mm is used.

Before display, first a depth correction is applied to compensate for the overall trend. This is done by first fitting a low order polynomial through a set of scatterers and then dividing the scan lines through the fit. In the first phantom the copper wires were used for this, and in the tissue phantom the vertical distance indicators. After this correction, the absolute value of the Hilbert transform is taken over the scan lines and finally the image is log compressed.

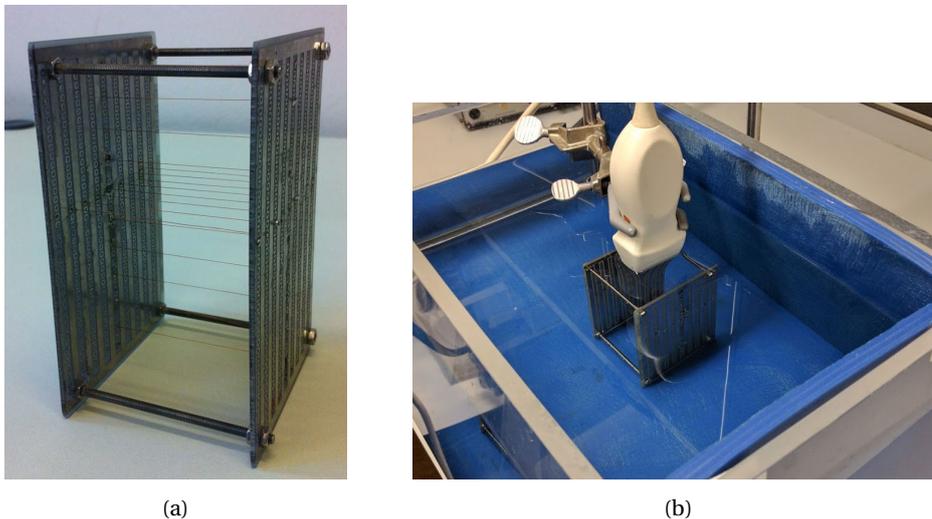


Figure I.2.1: The wire phantom used during the measurements. In (a) a view of the wires spun between the two plates and in (b) the measurement setup.

I.3

RESULTS

In this chapter the simulated beam profiles and the images constructed using continuous delays, Fresnel 2, Fresnel 4 and Fresnel 4 Reduced will be shown and compared.

I.3.1. BEAM PROFILE

In this section we will look at the simulated beam profile and look how the different methods affect the pressure distribution. After that we will investigate what effect the methods have on the arrival times and finally we look at the interlace method to emulate Fresnel 4.

I.3.1.1. PRESSURE DISTRIBUTION

In figure I.3.1 the beam profiles are shown for the 4 methods in the situation that the probe is focused at 30 mm and not steered. Comparing the four subplots makes it clear that continuous delays produce the smoothest beam profile. Below the focus almost all pressure is kept within a triangle and there is a clear focus visible.

The profile produced by Fresnel 2 looks different. It shows large variations, especially close to the transducer where the zones destructively interfere at the edges. Especially the centre zone has clear edges that run almost all the way towards the focus. Furthermore, the pressure is about 6 dB lower at this edge. Another interesting effect is also visible near the transducer. Instead of just a focusing beam, we also see a part of a diverging beam. This would especially have been visible with continuous waves, as the delay profile with just 2 zones for a diverging wave is then exactly equivalent to a focusing wave. In case of pulsed waves the delay profiles are mirrored. Still, close to the transducer only a part of the zones interfere so the diverging beam is still somewhat visible. Further in, more zones start to interfere, cancelling out the diverging beam. This interference also causes the field to become smoother further from the transducer.

Compared to continuous delays, the pressure distribution with Fresnel 2 is more diffuse. More pressure appears outside of the main beam and thus less inside it. This can also be seen in figure I.3.2(a) and (b). In (a) we can see that the pressure at the focus is about 3 dB lower and it drops off more after the focus. So it is more difficult to observe deeper lying reflectors. The lateral FWHM, as can be seen in (b) is almost unaffected by Fresnel 2, but the pressure does not drop further than about -20 dB. This might cause some extra artefacts in the image, because scatterers outside the main beam still receive a relatively high pressure. This is unwanted as it might obscure weaker signals.

Choosing a finer quantization improves the results, as is clear if we look Fresnel 4. The beam profile is more comparable to continuous delays than Fresnel 2. More pressure is contained within the main beam, resulting in an almost equal pressure at the focus. This, of course, also leads to less pressure outside of the main beam and thus fewer artefacts. Furthermore, the variations are reduced especially further away from the transducer. Close to the transducer the edges of each zone can still be observed, but less clear because the period difference between subsequent zones is now only $1/4$ periods instead of $1/2$ periods. Finally, the diverging beam has disappeared as the delay profiles needed, even in the continuous wave case, differ from the delays of a focusing beam.

Fresnel 4 reduced, which is shown in the rightmost subplot in figure I.3.1, produces a profile which differs a lot from the other 3 methods. Close to the transducer the field is similar to Fresnel 4, but it looks much worse further away. The focus is not clearly visible any more. In figure I.3.2(a) only a wide bump is visible around the focal depth of 30 mm with a much lower pressure. The lateral FWHM is a bit wider as compared to the other methods as can be seen in figure I.3.2(b), while the FWTM is even more than 5 times as wide.

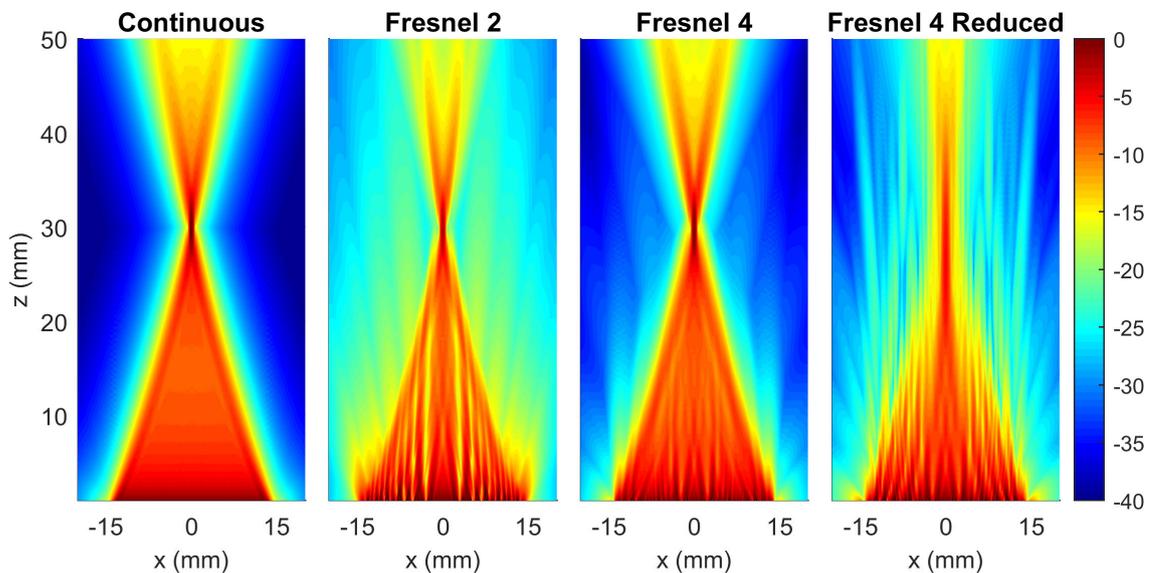


Figure I.3.1: Pressure distribution of a transducer focused at 30 mm using each of the 4 methods.

The beam profiles obtained with the 4 methods when the beam is steered are shown in figure I.3.3. Continuous delays, Fresnel 2 and Fresnel 4 show a similar behaviour as was visible in the unsteered beams, but Fresnel 4 Reduced looks much worse if the beam is steered. A large part of the beam seems to behave like a plane wave and just a part of it somewhat focuses. Due to this, the pressure at the focal point is more than 4 times lower in comparison to the other methods. Furthermore, the FWHM is significantly larger, as can be seen in figure I.3.4, and the FWTM is almost as wide as the simulated field. These things are the result of even larger focusing errors compared to an unsteered beam, which result in that even fewer elements constructively interfere at the focus.

From these simulations it seems that with Fresnel 4 one could expect very similar results as with continuous delays, except very close to the transducer. With Fresnel 2 some extra artefacts can be expected due to higher pressure outside of the main beam. The results with Fresnel 4 Reduced are expected to be the worst of all, especially at high steering angles.

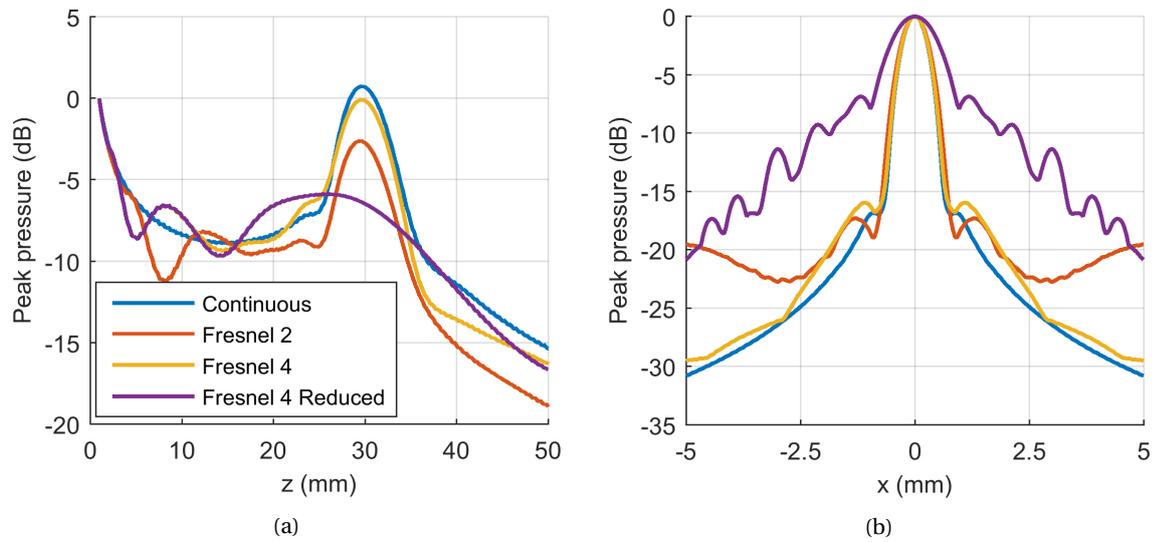


Figure I.3.2: In (a) the pressure distribution along the centre line and in (b) at a depth of 30 mm

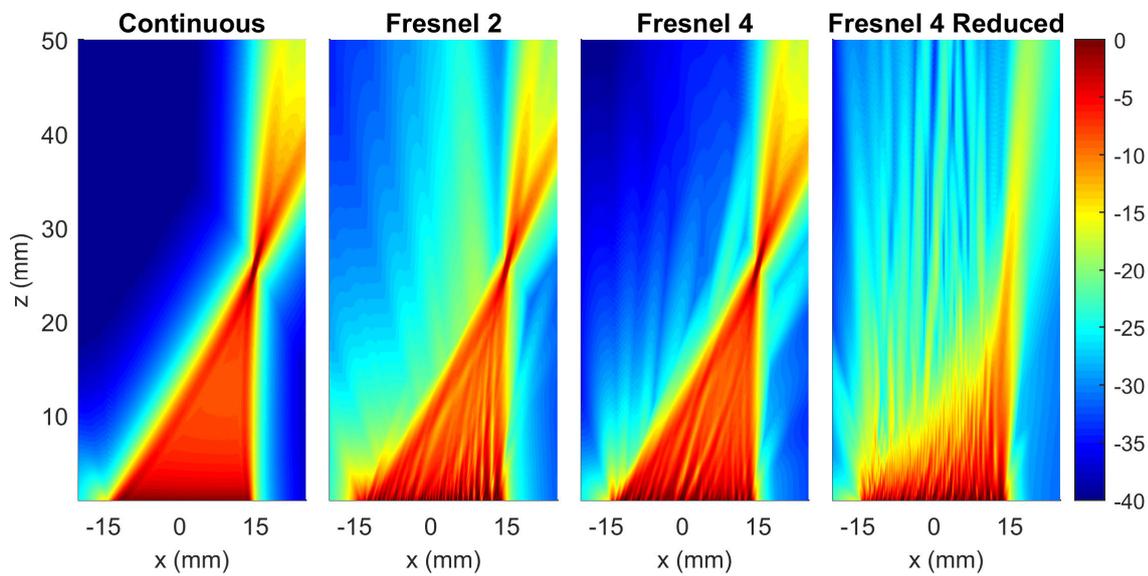


Figure I.3.3: Pressure distribution of a transducer focused at a radius of 30 mm and steered at 30° using each of the 4 methods.

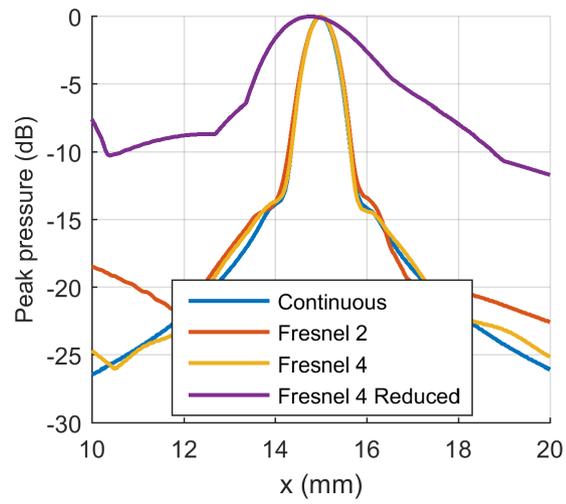


Figure I.3.4: Pressure distribution at the focal depth. For the beam focused at a radius of 30 mm and steered at 30° this is at about 25 mm.

I.3.1.2. ARRIVAL TIMES

To form an image the assumption is made that time can be directly convert to depth by use of the speed of sound. If this assumption is invalid, then the points will appear shifted from their original position. We have checked this assumption for the four methods by plotting the time it takes for the first pulse to arrive at a point on the scan line. This is shown in figure I.3.5 for the continuous delays and for Fresnel 4 Reduced. We see that Fresnel 4 Reduced clearly differs from continuous delays. For higher steering angles the first pulse arrives earlier than with continuous delays. This means that scatterers will appear closer to the transducer than they really are. Luckily, it seems like all times are just shifted a bit compared to the continuous case, which means that it could be corrected.

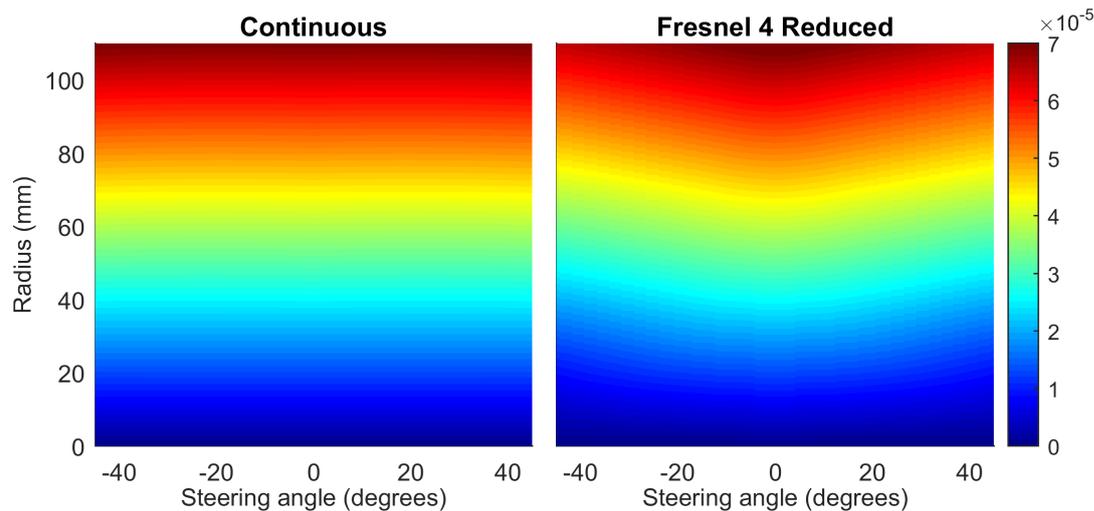


Figure I.3.5: The arrival time of the first pulse each point on the scan line for continuous delays and Fresnel 4 Reduced.

The first arrival times of Fresnel 2 and Fresnel 4 are also slightly different from continuous delays, but this would have been hardly observable in figure I.3.5. Therefore, in figure I.3.6 we have plotted the difference between first arrival times obtained with continuous delays and our Fresnel methods. First of all, we can see that the maximum difference is always less than $0.1 \mu\text{s}$ for Fresnel 2. This corresponds to the maximum timing error that is expected with this method. For Fresnel 4 the maximum timing error is twice as low and this is exactly what can be observed. The next thing that stand out is that we can distinguish two regions for both methods. Before and at the focus the difference is always negative, meaning that the pulse from the Fresnel methods arrives earlier. The second region can be found after the focus, where it looks quite erratic. For some scan lines the first pulse arrives strictly earlier with the Fresnel methods and for other scan lines exactly otherwise. If we would image a flat plate with Fresnel 2, it might not appear flat any more in the image, but it seems that it has some bumps. The height of these bumps for Fresnel 2 can maximally be twice the maximum timing error times the speed of sound, which corresponds to about $300 \mu\text{m}$, if we have two scan lines next to each other with exactly opposite behaviour.

The first arrival time is not the only thing that is different. In figure I.3.7 the difference between the first and last arriving pulse is plotted for continuous delays and Fresnel 4 Reduced. For continuous delays the difference is only very large in the near field and quickly becomes less until the difference is 0 at the focus. This is comes from the fact that it takes relatively long for the outer elements to arrive at the start of the scan line, while the centre elements arrive almost instantly. Fresnel 4 Reduced shows totally different behaviour, especially when the beam is steered. The difference between the first and last pulse becomes very large. This will result in elongated points in the radial direction. The reason is simple, if the beam is steered more, the outer elements will also have a larger difference

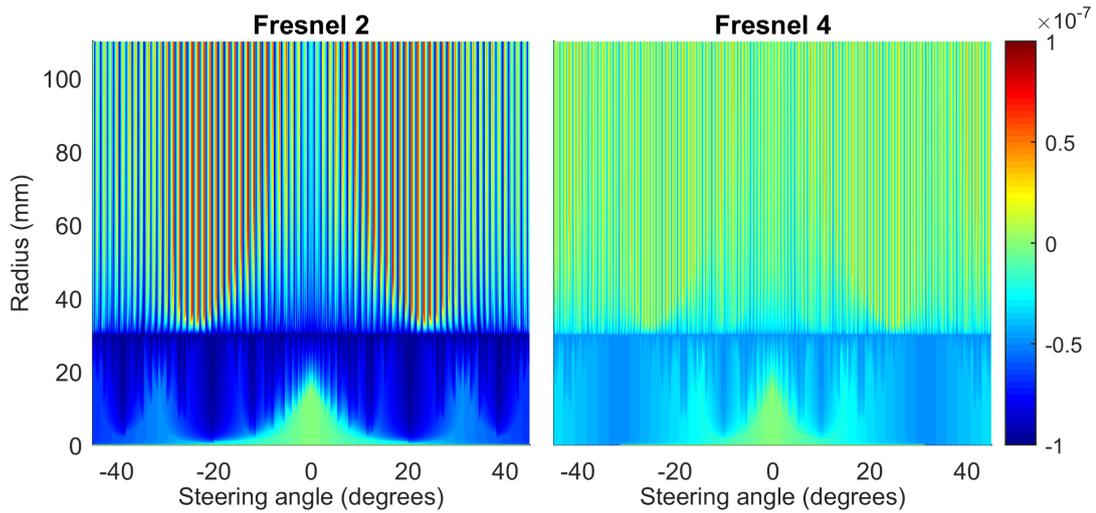


Figure I.3.6: The difference between the arrival time of the first pulse with continuous delays and with Fresnel 2 or Fresnel 4. A negative value means that the first pulse arrives earlier as compared to continuous delays.

in delay. The modulo operator applied to the delays of Fresnel 4 prevents this, keeping the delays within one period and making the mismatch in arrival time very large.

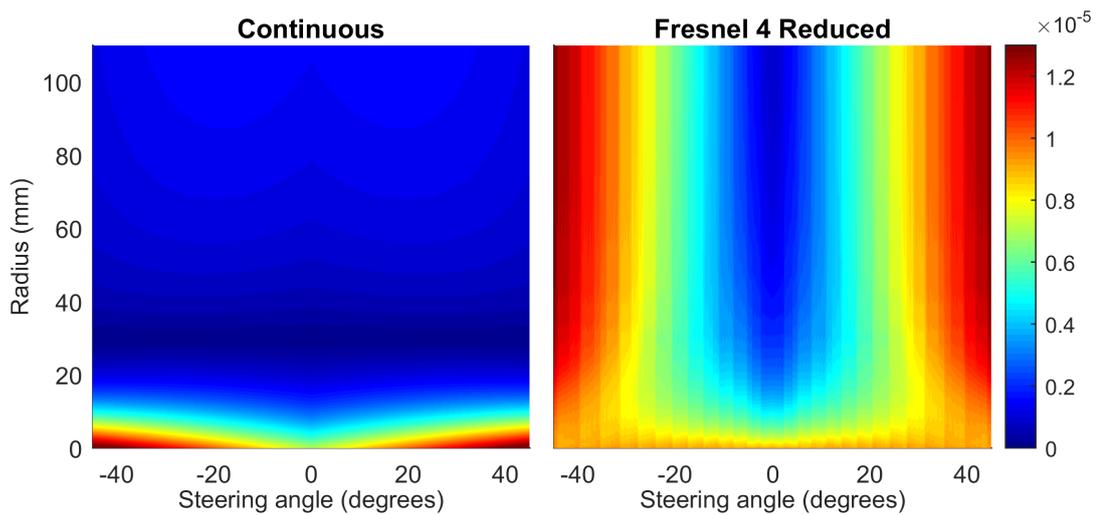


Figure I.3.7: The difference in arrival time between the first and last to arrive pulse at each point on the scan line for continuous delays and Fresnel 4 Reduced.

Again it is difficult to observe the difference between continuous delays and our Fresnel methods. So for Fresnel 2 and 4 we plot how much larger the difference in arrival time between the first and last pulse is as compared to continuous delays. This is shown in figure I.3.8. Again we see that Fresnel 4 is much more similar to continuous delays as Fresnel 2, which is of course as expected. Also, again the difference appears to be very erratic. Sometimes the difference between the first and last pulse is equal for the Fresnel methods and continuous delays, but there are scan lines where the difference jumps up to $0.2\ \mu\text{s}$. This could mean that at these scan lines the points appear more elongated in the radial direction and therefore also as less bright. Therefore, there might be some kind of modulation visible in the final images due to this, especially in Fresnel 2.

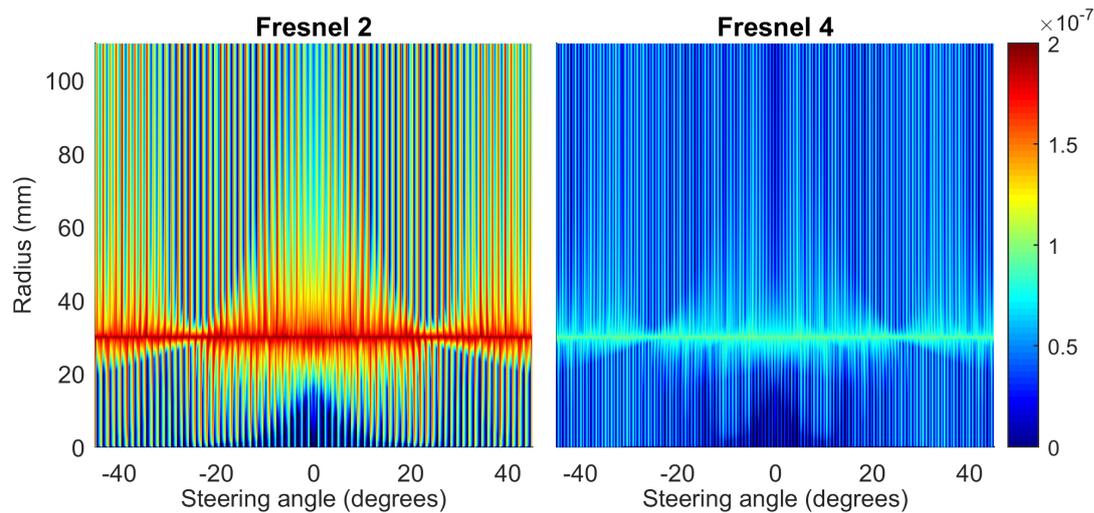


Figure I.3.8: The difference between continuous delays and our Fresnel methods for the time between the first and last pulse at each point on the scan line. A higher values means that there is more time in between the first and last pulse for the Fresnel methods.

I.3.1.3. EMULATION OF FRESNEL 4

There is no difference visible between the field of Fresnel 4 and the two-shot emulation because in the simulation we have a linear system with no attenuation and no noise. With attenuation and noise it is likely that deeper regions will not be correctly imaged by the 2 shot sequence as the weaker signal will sooner be overshadowed by noise. An overlap between the two shots, as illustrated in figure I.1.3(b), might help increase the pressure in the deeper regions, possibly at the cost of decreased performance. In the investigation we will use 3 different amounts of overlap: 4/3, 3/2 and full overlap. With a 4/3 overlap, for example, we mean that the maximum allowable timing error for each shot has been increased by 4/3 in comparison to when no overlap is used.

The pressure distribution for different amounts of overlap is shown in figure I.3.9. The field in the subplot where no overlap is used is exactly equal to the field by Fresnel 4. When including overlap, some interesting things happen. For the 4/3 and 3/2 overlap we see that the field inside the centre region appears smoother. Furthermore, we see that the field is more concentrated as at the outsides the pressure is slightly reduced. These two effects disappear when all elements are used in both shots. Also, at the scan line in the centre the pressure is slightly lower than the surroundings. This is better visible in figure I.3.10(a), where the pressure along the centre is clearly the lowest with full overlap.

The FWHM in the lateral direction is not affected by the different overlaps as visible in figure I.3.10(b). The FWTM is slightly lower for the 4/3 and 3/2 overlap, but the difference is rather small.

A full overlap thus seems to perform worst, but there is not really a clear distinguishing feature between the other two. Therefore, we have just chosen to only form an image with the 3/2 overlap later in this chapter.

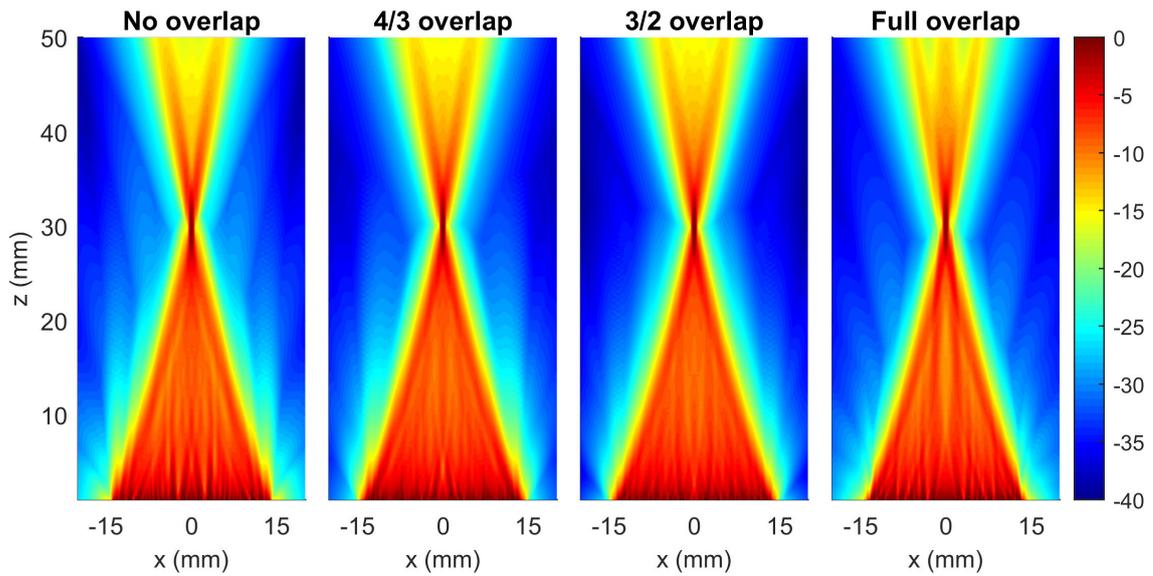


Figure I.3.9: Pressure distribution of a transducer focused at 30 mm using different amounts of overlap between the two shots.

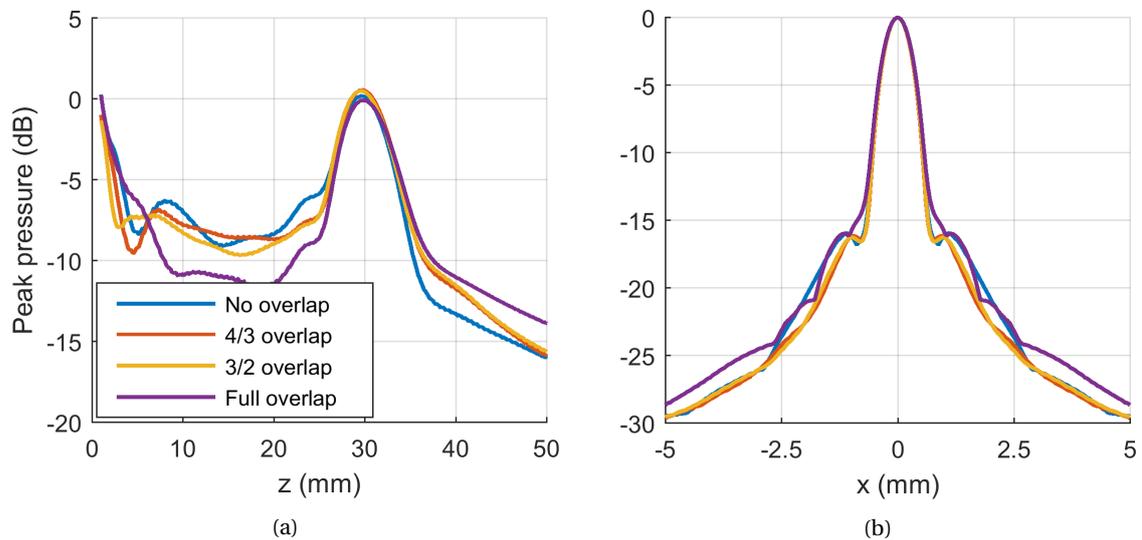


Figure I.3.10: In (a) the pressure distribution along the centre line and in (b) at a depth of 30 mm in (b)

I.3.2. IMAGE QUALITY

In this section we will investigate how the image quality is affected by the 4 methods. First we will check how the PSF is affected by the changed delays in the image of the wire phantom. Afterwards, we will look at the image of the tissue phantom and look how the contrast-to-noise-ratio(CNR) is affected. At the end of this section we will investigate the interlace method to emulate Fresnel 4.

I.3.2.1. PSF WIDTH

In figure I.3.11 the images formed from the measured data are shown for the 4 methods. All wires are visible in each image and at a depth of 20 mm the large screws are visible. Because we have used no apodization both in transmit and receive, at each point we can clearly see the effect of the limited aperture, which causes the 2 parabolas. Also, we see an extra reflection just above and below each

point. The extra points above are caused by the edge waves that are not suppressed as no apodization is used. The lower ones are most likely due to the internal reflections in the wire.

If we compare the image that is formed with continuous delays with Fresnel 2, there are two differences visible. First of all we see that the image of Fresnel 2 contains more artefacts. This is most likely caused by the increased pressure outside of the main beam as could be seen in figure I.3.1. However, the artefacts are only visible at about -40 dB and will most likely not be visible medical ultrasound images which include speckle. The second thing is that the parabolas caused by the finite aperture appear to be modulated. On some scan lines they appear brighter than on other scan lines, which is not the case with continuous delays. In figure I.3.8, we saw that the difference between the first and last pulse on each point on the scan line was for some lines much larger than others. This might have caused the modulation. It might cause some problems, as some details could be lost due to this, but the image of the tissue phantom has to verify this. Most importantly however, is that the points appear at the same positions and look similar.

A finer quantization of the delays is expected to improve the image and make it even more equal to the image formed by continuous delays. This is exactly what is visible when using Fresnel 4. The extra artefacts that were visible and the modulation in the image with Fresnel 2 have almost disappeared. Therefore, an even finer quantization for imaging purposes seems unnecessary.

Fresnel 4 Reduced is again very different from the other methods, as was also the case when comparing the beam profile. The points appear stretched in the radial direction, especially close to the transducer, which will certainly lower resolution, but on the other hand they seem similar in the lateral direction. Another difference is that the points are slightly shifted. This can be seen when looking at the points at a depth of 60 mm. The more the points are to the side, the more they are shifted upwards. This effect is also seen at the parabolas caused by the finite aperture. These parabolas are not continuous, but appear stepped. This is all caused by the delay profile that this method uses. As we have seen in section I.3.1.2 the arrival times differ from the continuous delay profile. At higher steering angles the pulses arrive at the scan line earlier than expected, resulting in the points being shifted upwards. Furthermore, the difference between the first and last pulse at each point also becomes very large for higher steering angles and close to the transducer, resulting in the points to appear stretched.

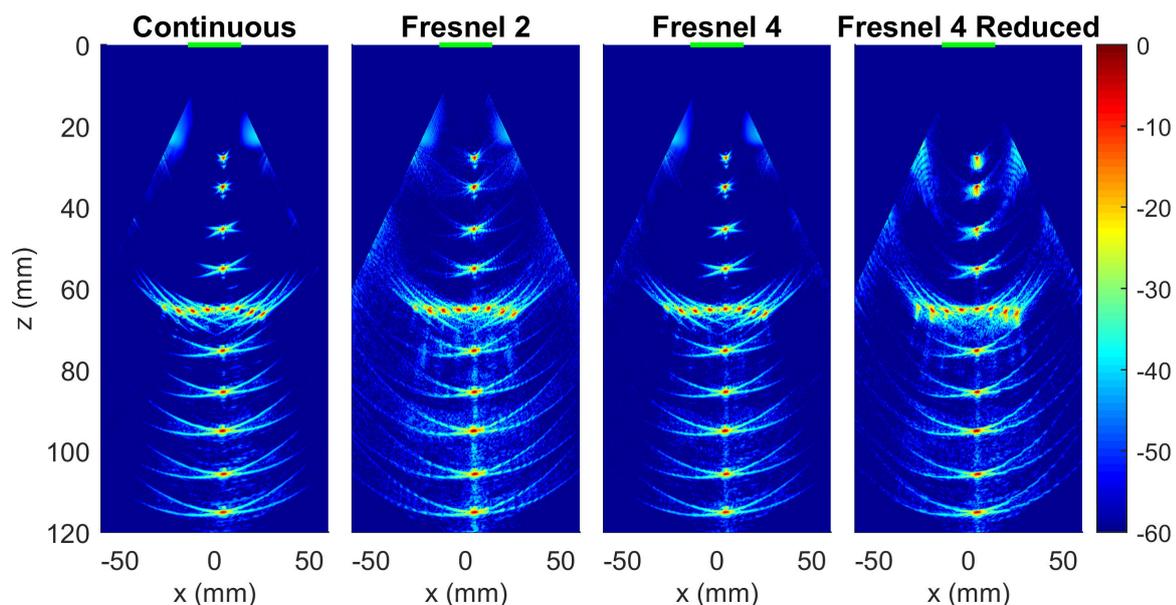


Figure I.3.11: Image formed from the measured data of the wire phantom using each of the four methods.

In all 4 images, the points appear very similar in the lateral direction. This is confirmed by looking at the FWHM and FWTM in figure I.3.12(a) and (b). There are only a few outliers and for most of the points all methods result in the same width.

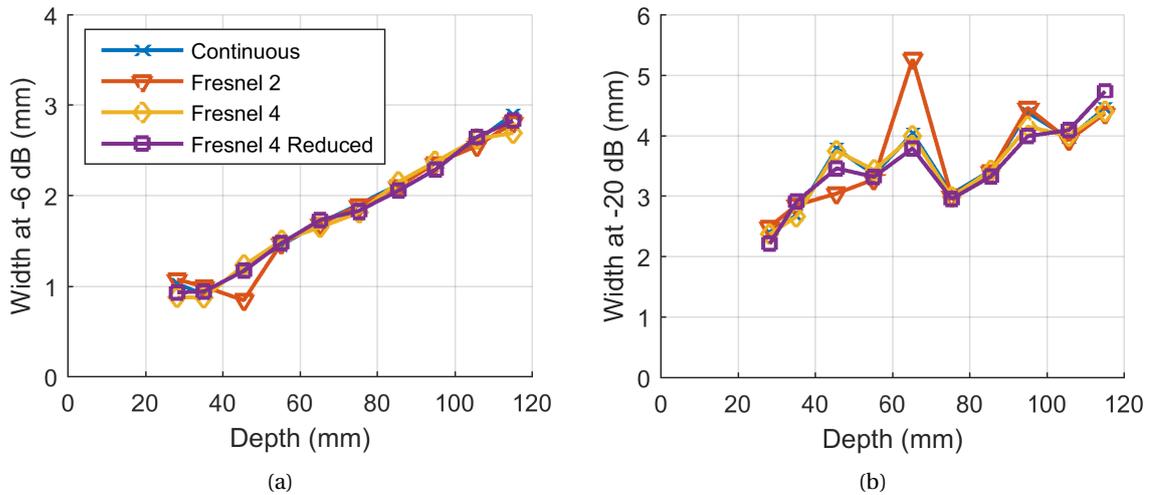


Figure I.3.12: In (a) the FWHM and in (b) the FWTM of the points visible in figure I.3.11.

I.3.2.2. COMPARISON WITH SIMULATIONS

We have also done simulations of the measured situations to confirm the results. The simulation results are shown in figure I.3.13. We see that the measurements and simulations are very similar. The screws have been omitted in the simulations and therefore do not appear. Furthermore, there are no internal reflections so we do not see extra points appearing just below the wires. Also, no noise has been included in the simulations, which reduces the amount of artefacts. The final difference is that the parabolas are a bit more visible. This is most likely caused by the inherent apodization of the transducer used in the measurements, which causes the elements at the side to be attenuated a bit, or due to the different directivity of the elements.

The FWHM, as visible in figure I.3.14(a), of the points in the simulations follow the same trend as in the measurements. The FWTM, shown in figure (b), does show some differences, especially for the points shallower than 70 mm, but the trend is still similar.

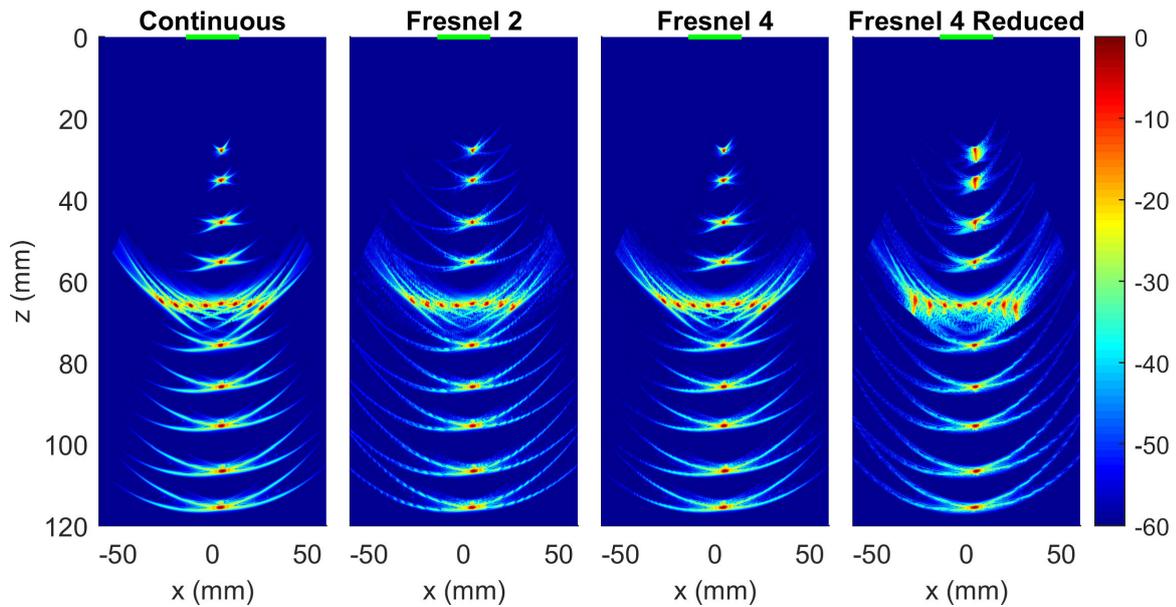


Figure I.3.13: Image formed from the simulated data of the wire phantom using each of the four methods.

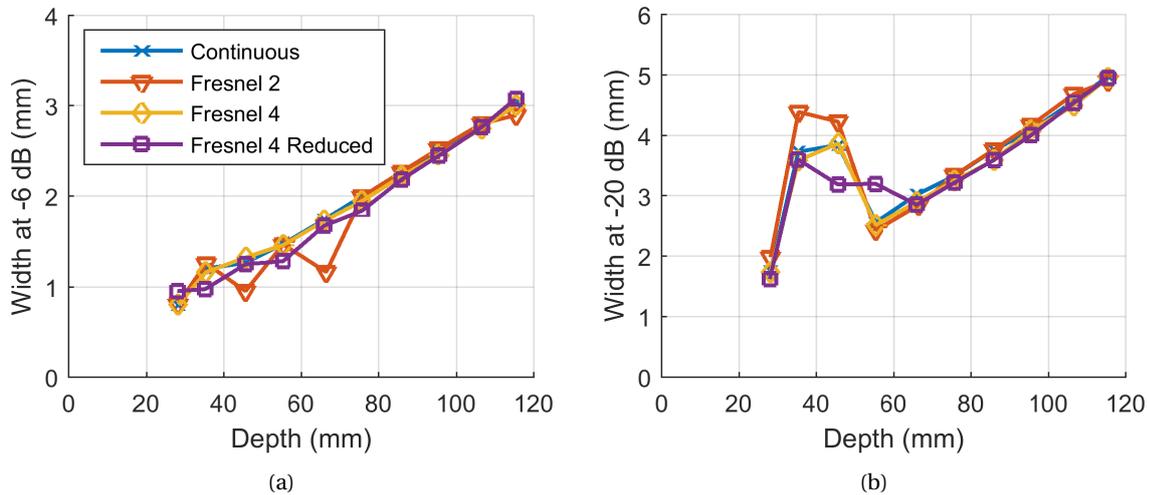


Figure I.3.14: In (a) the FWHM and in (b) the FWTM of the points visible in figure I.3.13.

I.3.2.3. CONTRAST-TO-NOISE-RATIO

The images of the tissue phantom created with each method are showed in figure I.3.15. It is very difficult to observe any difference between continuous delays, Fresnel 2 and Fresnel 4 in these images. The differences that were visible in figure I.3.11 are now less apparent. The modulation in the Fresnel 2 image can still be observed, but it is certainly less clear than it was in the image of the wire phantom. The image of Fresnel 4 Reduced is again clearly different from the other images. The rows of points at about 40 mm and 85 mm appear slightly bent upwards. Also, especially in the upper half, the image looks smeared out and details are much less clear. This results from the large difference in arrival time of the pulses from each element as could be seen in figure I.3.7.

To quantify the CNR we look at the 5 hypoechoic cysts that are more clearly visible in figure I.3.16 at depths of about 44, 67, 96, 123 and 153 mm and a lateral position of 32 mm. The calculated CNR for these cysts is shown in figure I.3.17. Continuous delays, Fresnel 2 and Fresnel 4 show the same trend with a decreasing CNR further away from the transducer. Fresnel 4 is also not distinguishable from

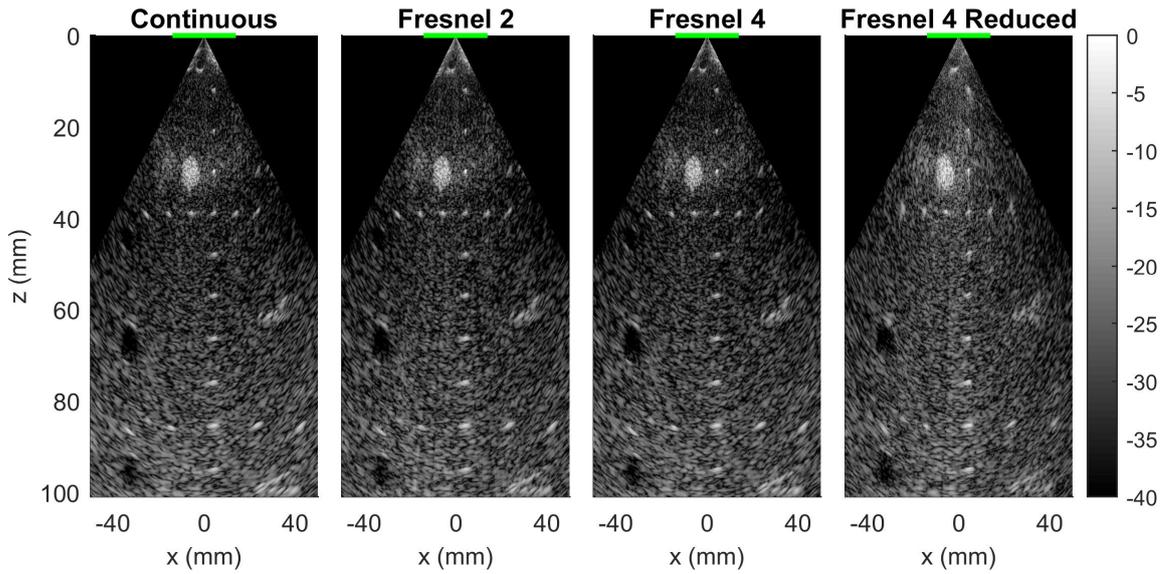


Figure I.3.15: Image formed from the measured data of the tissue phantom using each of the four methods.

continuous delays, indicating that for these applications a finer quantization is unnecessary. Fresnel 2 is consistently lower, but for the upper three cysts the CNR this is only by about 5% as compared to continuous delays. The difference does increase to about 20% for the lower two cysts. This might be because Fresnel 2 has a lower pressure at deeper depths as we already saw that the pressure was more diffuse in the beam profiles. Fresnel 4 Reduced shows a different trend. Especially the CNR of the first cyst is much lower by almost a half. This is not strange as this cyst cannot be distinguished in the image. For the deepest cysts, Fresnel 4 Reduces gives similar results as Fresnel 2.

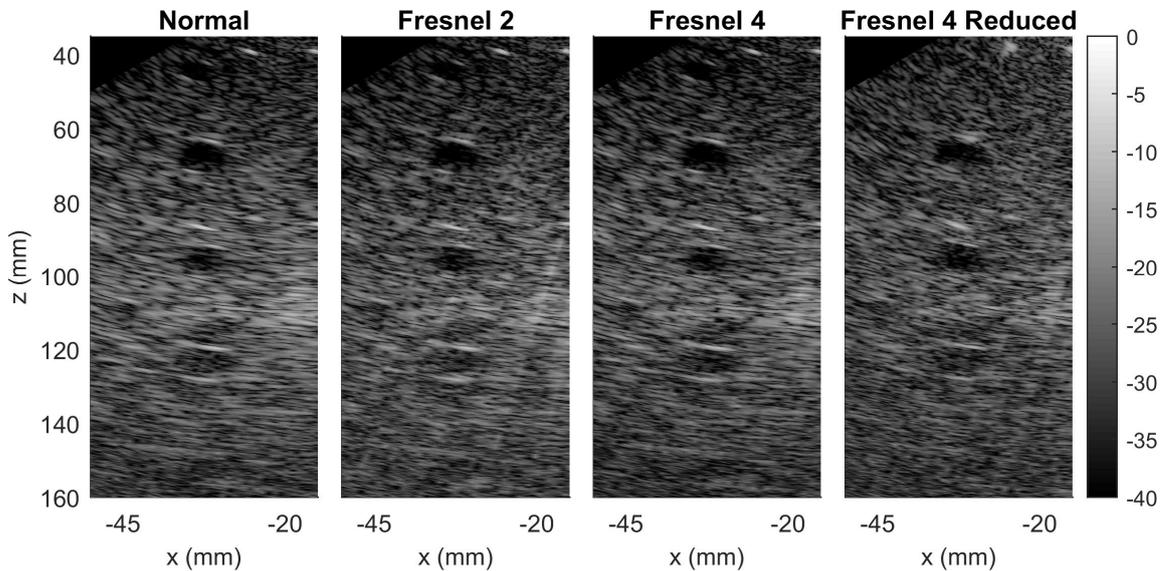


Figure I.3.16: Zoom in on the cysts used for determining the CNR.

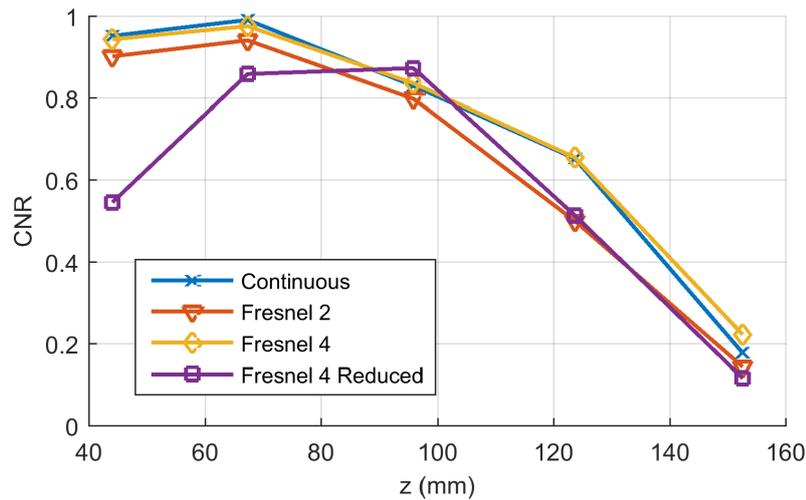


Figure I.3.17: Image formed from the measured data of the tissue phantom using DRF.

I.3.2.4. EMULATION OF FRESNEL 4

In figures I.3.18 and I.3.20 the images of the wire and tissue phantom using Fresnel 4, the two-shot emulation and a sequence with overlap are shown. There are no visible differences between Fresnel 4 and the two-shot emulation, indicating that the system behaves linearly. The FWHM and the FWTM of these two methods, visible in figure I.3.19, are also not distinguishable. However, there are some small differences in the CNR for the cysts as visible in figure I.3.21. For shallow depths the CNR's are equal, but the CNR of the deepest two cysts differ. This happens due to lower pressure of the separate shots, which becomes a problem deeper into the attenuating medium as the SNR becomes too low. In this case, with an attenuation of $0.5 \text{ dB MHz}^{-1} \text{ cm}^{-1}$ and a centre frequency of 2.5 MHz, we see that the CNR is still higher than achieved with Fresnel 2, of which the CNR is visible in figure I.3.15, but this might change with a higher attenuation or centre frequency.

As a possible solution for this low SNR at deeper depths, we proposed a small overlap between the two shots. We have investigated the results obtained with a $3/2$ overlap. With the $3/2$ overlap we mean that the maximum allowable timing error for each shot has been increased by $3/2$ in comparison to when no overlap is used. The overall behaviour using this overlap is similar to when no overlap is used. However, the modulation that was observed before is more clearly visible again. With Fresnel 4, most of it has disappeared. The FWHM and FWTM might also be affected by this modulation as can be seen in figure I.3.19. The overlap also affects the CNR, which has become lower over the whole depth range, while we had introduced this to increase it. Therefore, an overlap does not seem a good way to increase the pressure at deeper depths.

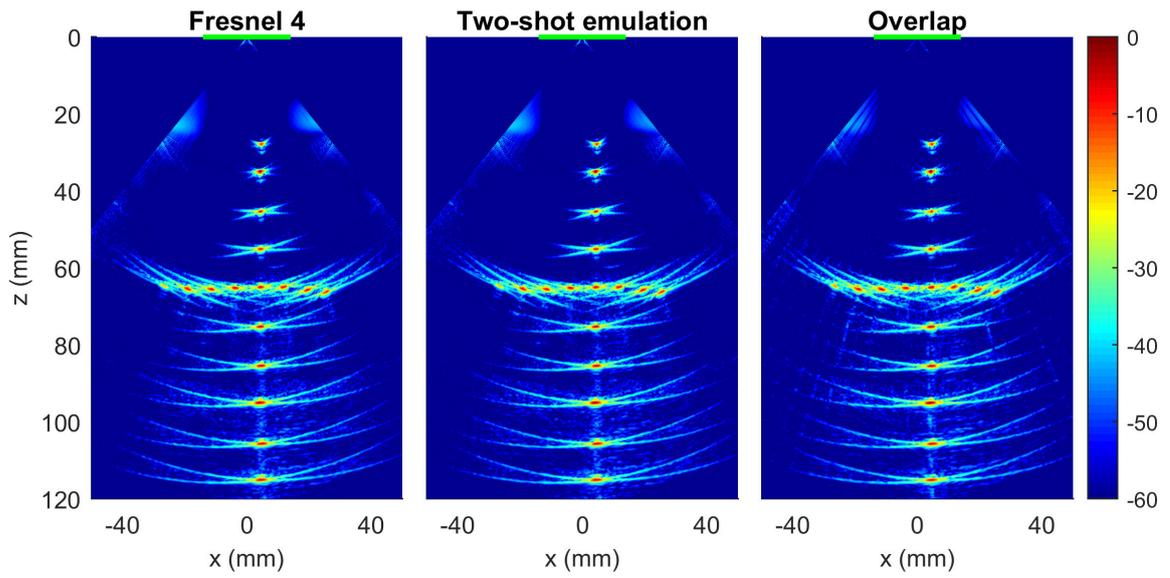


Figure I.3.18: A comparison between the images created of the wire phantom with Fresnel 4, the two-shot emulation and a sequence with overlap between the shots.

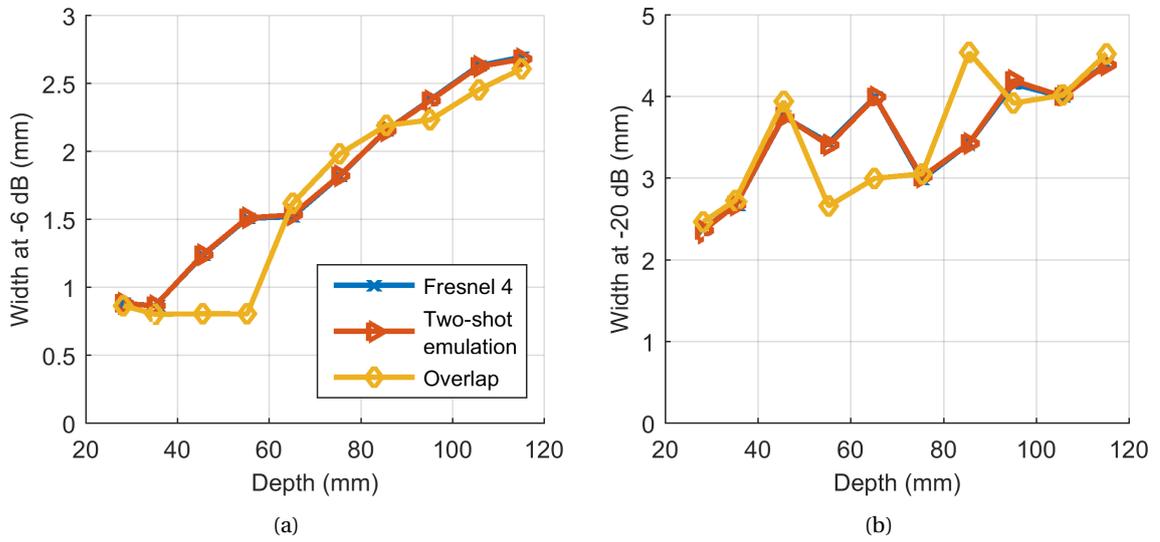


Figure I.3.19: In (a) the FWHM and in (b) the FWTM of of the points visible in I.3.18.

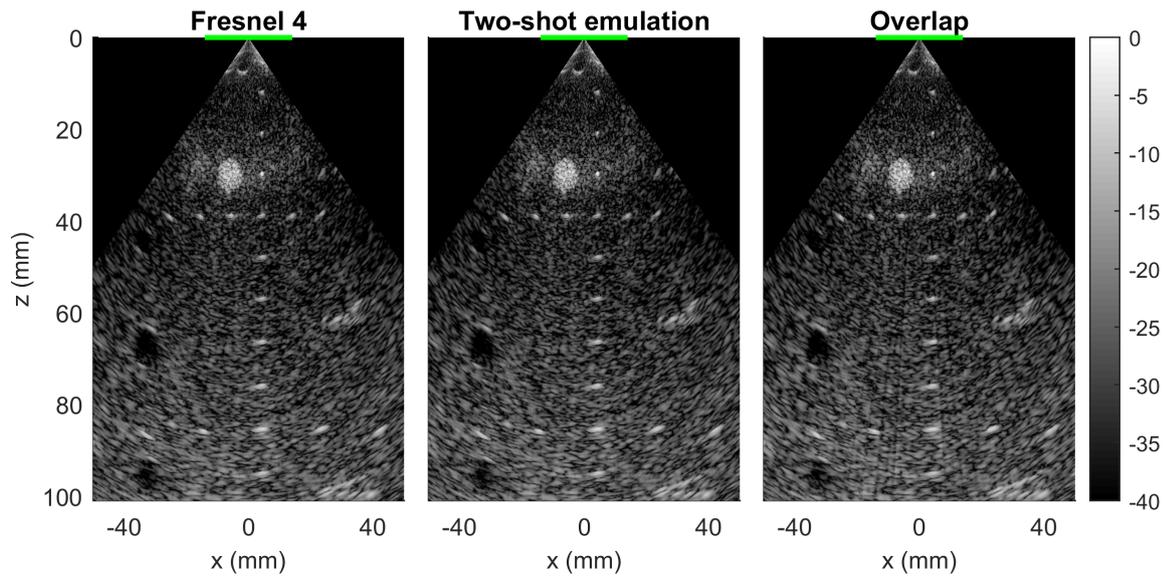


Figure I.3.20: A comparison between the images created of the tissue phantom with Fresnel 4, the two-shot emulation and a sequence with overlap between the shots.

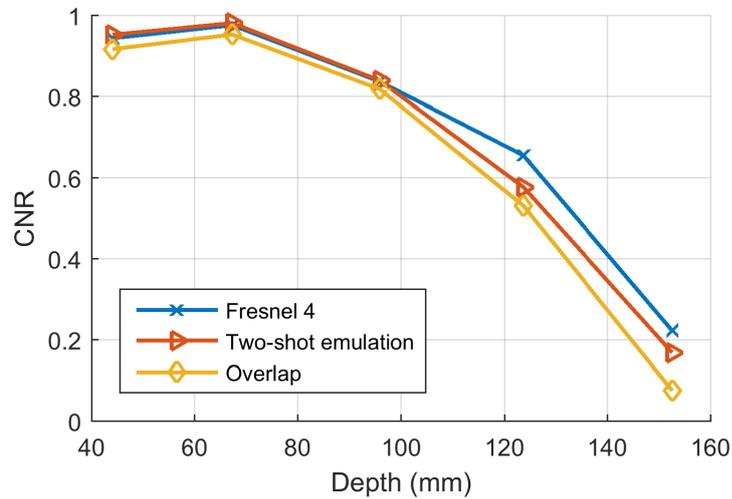


Figure I.3.21: A comparison between the CNR for the cysts in the tissue phantom. Three cysts are visible in I.3.20.

I.4

DISCUSSION AND CONCLUSIONS

In this chapter we will first discuss a few open ends and after that we give our conclusions.

I.4.1. DISCUSSION

There are still a few points of discussion left about our focusing method that were not addressed in this part. For example, we have used bipolar excitations, but the electronic implementation might only allow for unipolar excitations. How this exactly affects the pulse shape and results is unknown. Two other discussion points are apodization and how this focusing method works with other imaging methods. Specifically methods to measure the perfusion. This will be discussed next.

I.4.1.1. APODIZATION

In transducers where all elements are wired out, apodization can be applied in transmit by sending signals with different voltages. This is not compatible with the focusing method developed here as there is only a single high voltage signal available to all elements. In case of the tissue phantom apodization would not be necessary due to high speckle level, but it might improve image quality in other cases.

A simple method to achieve apodization is to just disable a fraction of the elements in an area so that it approximates the window function in that area. This can of course be done in many ways, but the results are most likely not very good. In the far-field we can use the Fraunhofer approximation to get the pressure. This effectively states that the pressure field is related to the Fourier Transform of the aperture.[34, p.74] Due to the sharp transitions from on to off, high spatial frequencies have to be involved, which will show up as side lobes with very similar amplitude as the main lobe. The near field will of course also be affected. One of the effects will be similar to the sharp zone transitions as were visible in the pressure fields of the Fresnel methods. Thus, this is not a good method to achieve apodization.

A multiple shots per line sequence is another method. This is similar to the two-shot emulation of Fresnel 4. In this case with 2 shots we could approximate the apodization function with 2 levels. This will most likely give better results than the binary apodization as discussed in the paragraph above, but it will most likely not be enough. Furthermore, this method relies again on the linearity of the system and it reduces the frame rate. Most likely, applying this is not worth it. There are no clear other possibilities to achieve apodization with this method. So applying a variable apodization in

transmit is currently impossible.

I.4.1.2. PERFUSION MEASUREMENTS

As mentioned in the introduction, this focusing method is developed to achieve larger penetration compared to a plane wave in highly attenuating media. This is for example the case in the brain of a pre-term baby. As shown in this part imaging seems to go well, but for this application measurement of perfusion in the brain is also wanted.

Doppler imaging is widely used in ultrasound for measurement of blood velocity. Continuous wave Doppler can give a indication of the velocities in a certain part, but cannot determine the exact locations. We have not investigated continuous wave ultrasound in this thesis, but it turns out that the difference between Fresnel zone focusing and normal focusing is much larger for continuous waves.[35] One of the problems with 2 unique delays is that the delays for focusing are exactly equal to defocusing. Steering is also problematic as the delays for steering at, for example, 30° are equal for steering at -30° . Furthermore, the field looks a lot more erratic with just 2 unique delays than in the pulsed wave case investigated here. The received frequencies are most likely not different, but the erratic pressure distribution might cause various problems with CW Doppler.

Pulsed wave Doppler methods, like Power Doppler for example, are expected to work better. The performance might be a bit lower due to the lower pressure in the region of interest and the more erratic field compared to the continuous delays case, the difference is expected to be small as is also the case with the final image.

Doppler techniques might however not be capable of measuring the velocity in the small vessels present in the neonatal brain. This is related to the very low backscattering strength of blood.[36] Furthermore, the blood flow is wanted, so the volume of blood that flows through the brain per unit time, and not just velocity. This might mean that different methods are needed, for example speckle tracking[37], or Doppler imaging has to be combined with additional information.[38] As long as these methods use pulsed waves, the results with this new focusing method should be comparable to focusing with continuous delays. Methods using continuous waves might suffer from an enormous decrease in performance with Fresnel zone focusing.

I.4.2. CONCLUSIONS

In this part we investigated a new focusing method that only requires a single high-voltage channel and local transistors and timing hardware. The focusing is required to achieve larger penetration compared to an unfocused beam in highly attenuating media, which is an alternative to achieve a cable count reduction.

We have tested the performance of our method by comparing it with continuous delay focusing. The pressure distribution of our method was quite different from continuous delay focusing. There were a lot more variations and the pressure was more diffuse. In the final images the differences were much smaller. The FWHM and FWTM were most of the time indistinguishable. However, there was slightly more artefacts due to the small timing errors and there was some modulation visible along the scan lines. Also, the CNR was lower. For the shallower depths, the CNR was only lower by about 5%, but at deeper depths the CNR was 20% lower with our method. This is most likely due to the more diffuse pressure distribution, which results in lower pressures deeper in the media. Overall though, the differences were relatively small and so it is clear that our new focusing method works very well.

The results also showed that by quantizing the phase to $N/4$ periods, as opposed to $N/2$ in the

previous paragraph, the results were almost indistinguishable from continuous delays. Only by very close examination, one could distinguish some modulation, but is much less severe than with the crudest quantization. So, it turns out that a quantization to $N/4$ periods is enough. However, this finer quantization might require an high voltage signal that cannot be included on the probe. A two-shot emulation was introduced to solve this. This will trade in frame rate to achieve the better results of a finer quantization. By using this emulation, only the CNR was found to be slightly lower for deeper depths. This is the result of the lower pressure in the separate shots. To increase the pressure we investigated a small overlap between the 2 shots, at the costs of some extra timing error, but this did deteriorate the results only and therefore an overlap is unusable to increase the per-shot pressure.

There are still some open ends, as discussed at the beginning of this chapter, but our results show that we can achieve similar results as continuous delays with our focusing method over the whole imaged range. This proves that we can achieve focusing with only a single channel in transmit and therefore that we are able to achieve a better penetration than an unfocused beam, which is an alternative to reduce the required transmit channels.

II

CYLINDRICAL STOLT MIGRATION

II.1

METHOD DEVELOPMENT

In this chapter we develop Stolt Migration in a cylindrical coordinate system that can be used with a matrix array doing a linear scan in one direction and a sector scan in the other. Stolt migration has already been developed for linear and phased arrays separately, as discussed in sections 2.5.1 and 2.5.2 and this method will in some way combine the two. A similar method as we will develop in this section has already been developed by Skjelvareid et al[39], but it uses a different approximation for the Hankel function than we will use. Furthermore, the whole derivation is skipped. Here, we will give the whole derivation and use approximation that is sufficient and better for our applications. After this we will discuss the comparison between this method and polar and linear Stolt. Finally, we describe the implementation of this algorithm.

II.1.1. DERIVATION

For the derivation of this method we will define the coordinate system in the following way

$$r = \sqrt{x^2 + z^2}, \quad \theta = \arctan x/z, \quad y = y. \quad (\text{II.1.1})$$

The derivation for Cylindrical Stolt Migration is very similar to the polar case. The main difference arises from the fact that there is now also a second derivative with respect to Y present in the Helmholtz equation that now reads

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \tilde{p}}{\partial r} \right) + \frac{1}{r^2} \frac{\partial^2 \tilde{p}}{\partial \theta^2} + \frac{\partial^2 \tilde{p}}{\partial y^2} + k^2 \tilde{p} = 0, \quad (\text{II.1.2})$$

where \tilde{p} is the Fourier Transform of the pressure. This equation can be solved by using separation of variables and therefore we separate the pressure as

$$\tilde{p} = R(r)\Theta(\theta)Y(y) \quad (\text{II.1.3})$$

and substitute it into the Helmholtz equation. After multiplying by $r^2/(R\Theta)$, we get

$$\frac{r^2}{R} \frac{d^2 R}{dr^2} + \frac{r}{R} \frac{dR}{dr} + \frac{1}{\Theta} \frac{d^2 \Theta}{d\theta^2} + \frac{r^2}{Y} \frac{d^2 Y}{dy^2} + k^2 r^2 = 0. \quad (\text{II.1.4})$$

The angular part can now be separated. Because the solution should be periodic, a negative separation constant has to be used

$$\frac{1}{\Theta} \frac{d^2 \Theta}{d\theta^2} = -m^2. \quad (\text{II.1.5})$$

The general solution to this equation is

$$\Theta(\theta) = A_m e^{im\theta} + B_m e^{-im\theta}. \quad (\text{II.1.6})$$

As θ is periodic every 2π , m is limited to integer values.

If we now plug equation II.1.5 back into equation II.1.4 and divide by r^2 we get

$$\frac{1}{R} \frac{d^2 R}{dr^2} + \frac{1}{rR} \frac{dR}{dr} - \frac{m^2}{r^2} + \frac{1}{Y} \frac{d^2 Y}{dy^2} + k^2 = 0. \quad (\text{II.1.7})$$

Now we can separate the y dependence. We will assume that y is periodic and use $-k_y^2$ as the separation constant. This means that the general solution is :

$$Y(y) = A_{k_y} e^{ik_y y} + B_{k_y} e^{-ik_y y}, \quad (\text{II.1.8})$$

where k_y is $2\pi n/d_y$, with n an integer and d_y the length over which y is assumed to be periodic.

Finally, we can solve the equation for R . This can be done by plugging the separation constant $-k_y^2$ into equation II.1.7 and multiplying the equation by R . This results in

$$\frac{d^2 R}{dr^2} + \frac{1}{r} \frac{dR}{dr} + \left(k^2 - k_y^2 - \frac{m^2}{r^2} \right) = 0, \quad (\text{II.1.9})$$

which is just a modified form of the Bessel differential equation with the following general solution:

$$R(r) = C_{mn} H_m^{(1)} \left(r \sqrt{k^2 - k_y^2} \right) + D_{mn} H_m^{(2)} \left(r \sqrt{k^2 - k_y^2} \right). \quad (\text{II.1.10})$$

In this equation $H_m^{(1)}$ and $H_m^{(2)}$ are the cylindrical Hankel functions of the first and second kind respectively. The first kind represent a wave propagating away from the origin, while the second kind represents a way propagating towards the origin.

With the above we can build our migration in cylindrical coordinates. As in the Polar case we first shift time axis to redatum the measured data to the focal points

$$t' = t - 2 \frac{r_f}{c}, \quad (\text{II.1.11})$$

where r_f is the radius at which the transducer is focused at each scan line. The data after redatuming represents the results of a pulse-echo measurement done with virtual elements at r_f . As we have described wave propagation in the frequency domain, the next step is to transform the data using a 3D Fourier Transform:

$$\tilde{p}(r_f, m, k_y, f) = \int_{-\infty}^{\infty} \int_0^{d_y} \int_0^{\pi} p(r_f, \theta, y, t') e^{i(m\theta + k_y y - 2\pi f t')} dy d\theta df. \quad (\text{II.1.12})$$

To form an image we use the ERM in which we first have to extrapolate the measured data using $\hat{c} = c/2$. For radii larger r_f we only have inward propagating waves. This means that only $H_m^{(1)}$ is present for these radii and we can use equation II.1.10 to get the pressure field at depth $r > r_f$:

$$\tilde{p}(r, m, k_y, f) = \tilde{p}(r_f, m, k_y, f) \frac{H_m^{(1)} \left(r \sqrt{\hat{k}^2 - k_y^2} \right)}{H_m^{(1)} \left(r_f \sqrt{\hat{k}^2 - k_y^2} \right)}. \quad (\text{II.1.13})$$

This has to be done for all variables appearing in the above equation.

Now an inverse transform has to be taken and, following the ERM, we have to evaluate it at $t' = 0$, which is the time at which the scatterers "exploded". This is described by

$$p(r, \theta, y, t' = 0) = \sum_m \sum_{k'_y} \int_{-\infty}^{\infty} \tilde{p}(r, m, k_y, f) e^{-i(m\theta + k_y y)} df. \quad (\text{II.1.14})$$

So equations II.1.12, II.1.13 and II.1.14 together form the whole migration algorithm for depths larger than r_f . For radii smaller than the transmit focus, $H_m^{(2)}$ has to be used instead of $H_m^{(1)}$ in equation II.1.13 and the time axis has to be reversed before migration.

II.1.2. COMPARISON TO POLAR AND LINEAR STOLT MIGRATION

What we have derived here in this section is in principle a combination of polar and linear Stolt. However, we see by comparing equation II.1.14 with 2.18 that the migration derived here is most similar to Polar Stolt. Furthermore, if we take only one slice in y , the terms corresponding to this direction drop out and we are left with the same equations as for Polar Stolt. It is therefore possible to use a similar implementation. This also means that both the near and far field can be processed in a single calculation and the same approximation of the Hankel function can be used.

There are however some complications due to the extra variable k_y . First of all, a square root has appeared in the Hankel functions which can turn complex if $k_y > k$. The waves described by this case are evanescent, which are not measured by the transducer and should therefore not be extrapolated. Secondly, the sign of the square root is important. A closer look at $\sqrt{k^2 - k_y^2}$ makes clear that this can be seen as the wavenumber in the radial direction k_r . Like in the polar case we should choose k_r positive if the frequency is positive and k_r negative for negative frequencies.

The computation time is also not positively affected by the extra variable, which could be expected as we now have to reconstruct a volume. The extrapolation and integration over all frequencies is most affected and becomes an even larger bottleneck. In the polar case it was implemented as a matrix multiplication for every value of m . Now we also have to do this for every value of k'_y , which increases computation time by at least a factor 100. Changing the implementation of this might positively affect the computation time. However, most of the time spent is on calculating the Hankel functions. Pre-calculation would help, but this is impossible due to memory limitations. Assuming that we have 1000 frequencies, 1000 radii, 100 values of k_y and also 100 values of m , this would easily require 160 GB to store the pre-calculated Hankel functions as complex double, which requires 16 B per value, in memory.

II.1.3. IMPLEMENTATION

The whole cylindrical migration algorithm described in this chapter has been implemented in Matlab (Matlab 2015b, The MathWorks, Inc., Natick, Massachusetts, USA). The implementation is very similar to Polar Stolt Migration and consists of the operations listed below:

1. Shift the time axis of all scan lines to $t' = t - \frac{2r_f}{c}$ to create an array of virtual point sources and receivers on a cylinder
2. Expand the number of scan lines. In θ to a full circle and enough in y .
3. Transform $p(r_f, \theta, y, t')$ to $\tilde{p}(r_f, m, k_y, f)$ by using a 3D-FFT and apply a band-pass filter for the frequencies
4. Extrapolate \tilde{p} to all radii of interest using equation II.1.13 and integrate over all frequencies. Repeat this for all m and k_y

5. Take a 2D-IFFT in θ and y
6. Discard scan lines outside of the scan sector

In step 2 the amount of scan-lines is extended with zero-padding. In the angular direction it is zero padded to a full circle, because it is periodic every 2π . Zero-padding in y is required because this direction is not periodic. A large enough padding will prevent any issues at the edges. At the end of the algorithm the padded scan lines are trimmed to their original size.

In the third step we apply a band-pass filter to reduce the amount of computations necessary. It is an ideal band-pass, meaning that any frequency outside are just discarded and the others are unchanged. We have chosen the pass-band equal to the -40 dB bandwidth of the excitation signal and therefore the ideal bandpass will not result in any problems.

Step 4 is implemented as a matrix multiplication. This has to be done for every m and k_y . While doing this, one has to make sure that the evanescent waves are excluded and the correct sign is taken before the square root as discussed in the section above.

II.2

SIMULATION SETUP

To test the performance of Cylindrical Stolt Migration we will only construct images from simulated data using FieldII. As stated the migration algorithm can be seen as a combination of polar and linear Stolt migration. Therefore, the transducer that will be used is a combination of a linear and phased array probe and consists of 80x128 elements. In the phased direction has a length of 80 elements that have a pitch of 150 μm . The linear direction, which lies along the y -axis, has a length of 128 elements. The pitch of these elements is 250 μm . The centre frequency of the transducer is 5 MHz, therefore the pitch in the phased and linear direction respectively will be about $\lambda/2$ and $5\lambda/6$. The excitation signal is a 3 cycle sinusoid with a Gaussian window and the impulse response is taken to be a delta-peak. The transducer is apodized in transmit and receive with a Hamming window.

The simulated scan consists of 256 scan lines in θ within a scan sector of 90° and 113 lines y within 32 mm, which is equal to aperture width in that direction. The transducer is focused at 13 mm. In y the focus will be accomplished by natural focusing, which means a subset of the elements is enabled in this direction and that there is no delay difference between the elements. In the other direction delays will be applied. This choice has been made to check the feasibility for a certain probe.

The phantom to be imaged consists of 3 lines of point scatterers. The first line consists of 10 points placed between a depth of 7.5 and 7.5+45 mm in the centre of the transducers. The second line runs in the x - z plane at a depth of 27.5 mm and the scatterers are spaced 5 mm. The final line runs in the y - z plane and is also placed at a depth of 27.5 mm and the scatterers are spaced 5 mm.

We will compare the image constructed by our developed algorithm to DRE. This comparison will be done by looking at the image and comparing the FWHM of the PSF as function of depth. Furthermore, to verify if the migration works correctly the results will also be compared to both polar and linear Stolt, which will process a single slice from the dataset. These two methods have already been experimentally verified and we can expect similar results.

Before display, first a depth correction is applied to compensate for the overall trend. This is done by fitting a low order polynomial through the line of point scatterers in the centre and then dividing each scan line through the fit. After this correction, the absolute value of the Hilbert transform is taken over the scan lines and finally the image is log compressed.

Table II.2.1: Simulation settings for Cylindrical Stolt Migration

	Parameter	Value
Transducer	Centre frequency	5 MHz
	Number of elements	80 x 128
	Pitch	150 μm \times 250 μm
	Element width	130 μm \times 230 μm
	Apodization	Hamming
	Excitation	3 cycle Gaussian windowed sinusoid
Simulation	Sampling Frequency	100 MHz
	Number of scan lines	256 x 113
	Scan sector in theta	$\pm 45^\circ$
	Scan length in y	32 mm
	Focal Depth	13 mm
	# elements in y per slice	16
Medium	Speed of sound	1540 m s ⁻¹
	Attenuation	Negligible

II.3

RESULTS

Cylindrical Stolt Migration reconstructs a volume. To be able to adequately show the results we will discuss two slices. In figure II.3.1 a slice is shown at $y = 0$ of the images generated with DRF, Cylindrical Stolt and Polar Stolt. As expected DRF shows the best resolution before and after the focus at 13 mm. At about a depth of 30 mm in the centre we can see various faint points appearing that are not part of the phantom. These artefacts most likely caused by points outside of this slice. DRF constructs the image line by line and thus does not take into account the connection between the scan lines.

Stolt migration does take this interaction into account and so these artefacts have disappeared, but on the other hand different ones have appeared. At the radius of every scatterer a ring appears at about -40 dB. This worsens at radii where multiple spots exist as is the case around a depth of 30 mm. The rings become more visible, but compared to the scatterers the intensity is still low. The reason why these rings appear is not clear. Apodizing the scan lines before migration, disabling the band-pass filter, using a better approximation and doing a depth correction before migration all do not help. Still, from figure II.3.2 we can see that the FWHM is up to 30% lower compared to DRF in this direction.

To verify the results of Cylindrical Stolt also an image generated with Polar Stolt has been included. It shows similar artefacts as DRF around a depth of 30 mm, but the rings do not appear. There are still things visible at the radii of the scatterers, but they are more spread out and therefore have a lower amplitude. This might indicate that the other slices that are included in Cylindrical Stolt cause these rings for some reason. The FWHM of the points are slightly better with Polar Stolt as compared to Cylindrical Stolt, indicating that those rings do negatively affect the final image.

The second slice is at $\theta = 0$, which is effectively a slice along the linear direction of the array. We can see the artefacts at about 30 mm deep that caused the faint point to appear in the slice at $y = 0$. The PSF width in this slice increases much steeper compared to the other slice, as can also be seen in figure II.3.4. This is the result of using the natural focus as only a subset of the elements can in this case be enabled.

Cylindrical Stolt again has a smaller PSF width, but this time the difference is larger. Except for the first reflector, Stolt always outperforms DRF. At the deepest reflector it is even 50% better. Furthermore, the width seems to approach an asymptote, which is also observed in literature.[4] The artefacts visible are again different compared to DRF. We can see the rings that appeared in figure II.3.1 at a depth 30 mm and just below it, but not the interference pattern as visible in the image constructed with DRF.

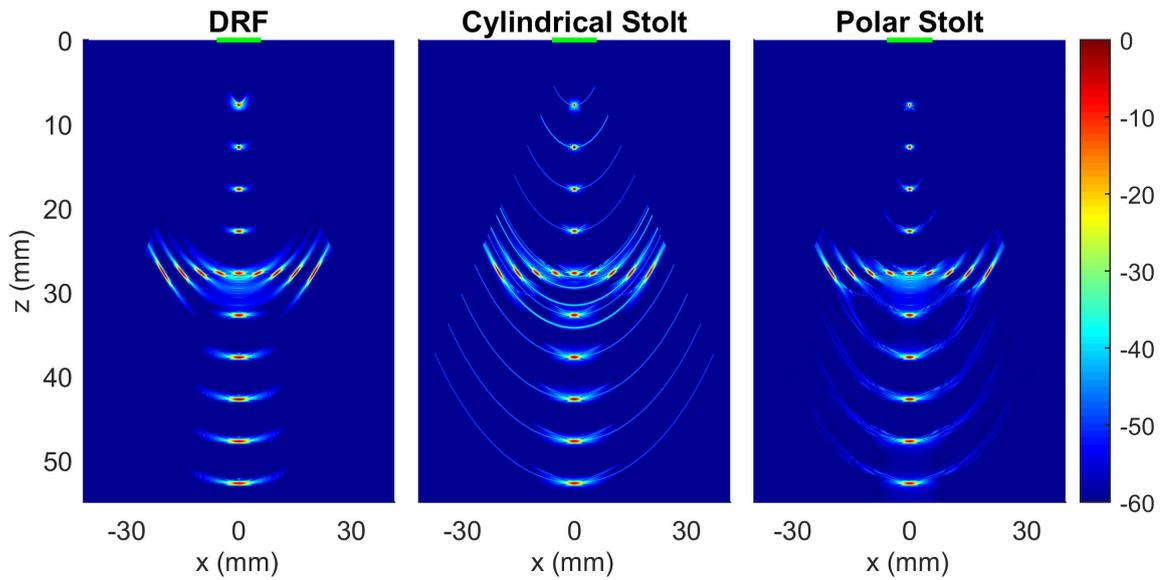


Figure II.3.1: Slice at $y = 0$ of the images constructed from the simulated data by DRE, Cylindrical Stolt and Polar Stolt.

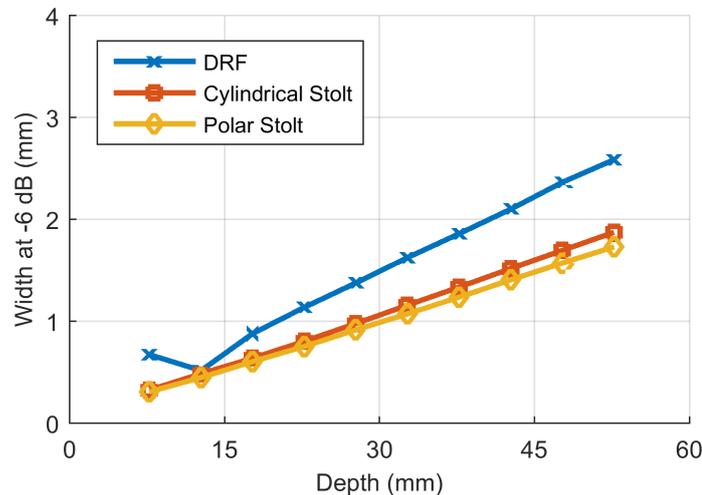


Figure II.3.2: The FWHM of the points visible in figure II.3.1

By looking at the image made with Linear Stolt we see that it is very similar to the image by Cylindrical Stolt. The only big difference are the extra artefacts that appear, but as they are mostly at -40 dB we can expect that this will not influence the result in a more realistic situation. The PSF width is this time not distinguishable from the width with Cylindrical Stolt.

So the results that linear and polar Stolt obtained in a single slice are similar to the results obtained with Cylindrical Stolt. This does partly verify that the results that we obtain here are correct. However, the advantage of Cylindrical Stolt would be that it can process all slices at once and obtain better results than if all slices are processed independently. We have not verified this. Furthermore, Cylindrical Stolt is much slower compared to Polar and especially to Linear Stolt. The computation time of Cylindrical Stolt easily is 100 to 200 times longer as compared to polar and linear Stolt. This is effectively the number of extra slices that are processed with Cylindrical Stolt and thus there is no computational advantage.

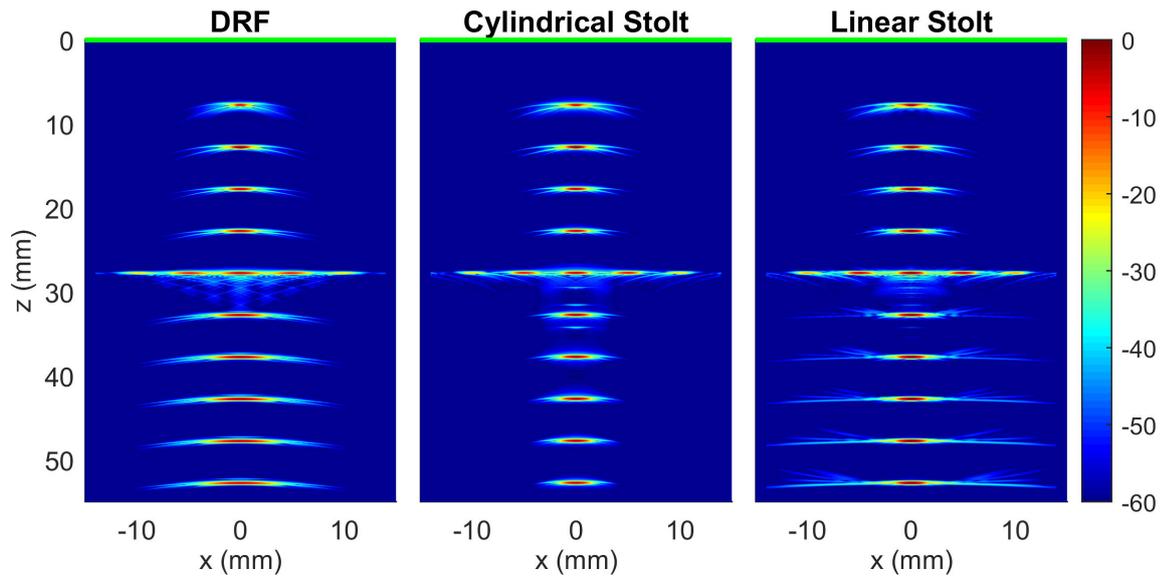


Figure II.3.3: Slice at $\theta = 0$ of the images constructed from the simulated data by DRE, Cylindrical Stolt and Polar Stolt.

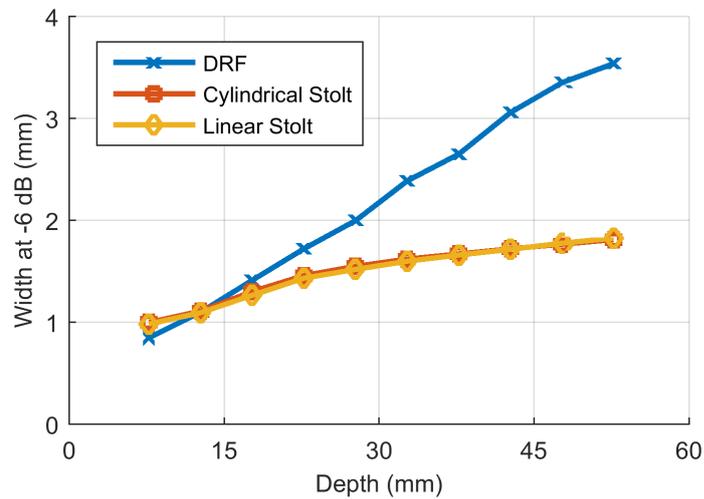


Figure II.3.4: The FWHM of the points visible in figure II.3.3

II.4

CONCLUSIONS

The goal of this part was to develop a frequency domain two-stage migration method in cylindrical coordinates. The results show that we have succeeded at developing such a method. The FWHM is up to 30% lower compared to DRF in the polar direction and up to 50% in the linear direction. Furthermore, the respective results obtained with Polar and Linear Stolt were very similar. This proves that our method works as expected. However, the computation time is very long and incompatible with real time imaging.

III

SPHERICAL STOLT MIGRATION

III.1

METHOD DEVELOPMENT

In this chapter we will develop a migration method in spherical coordinates. First we will derive it for a volumetric sector scan with a single focus in transmit and receive. Then we will discuss the algorithm and its implementation. At the end of this chapter we will modify the algorithm to work with spherical waves.

III.1.1. DERIVATION

During the derivation we will use the following convention for the coordinate system

$$r = \sqrt{x^2 + y^2 + z^2}, \quad \theta = \arctan x/z, \quad \varphi = \arccos y/r. \quad (\text{III.1.1})$$

This means that θ is the azimuthal angle and φ the inclination angle.

We start the derivation from the Helmholtz equation in spherical coordinates:

$$\frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \tilde{p}}{\partial r} \right) + \frac{1}{r^2 \sin \varphi} \frac{\partial}{\partial \varphi} \left(\sin \varphi \frac{\partial \tilde{p}}{\partial \varphi} \right) + \frac{1}{r^2 \sin^2 \varphi} \frac{\partial^2 \tilde{p}}{\partial \theta^2} + k^2 \tilde{p} = 0, \quad (\text{III.1.2})$$

where \tilde{p} is the Fourier Transform of the pressure. This equation can be solved by using separation of variables. First we separate the radial part from the angular part by writing the pressure as

$$\tilde{p}(r, \varphi, \theta) = R(r)\Omega(\varphi, \theta) \quad (\text{III.1.3})$$

and substitute it into the Helmholtz equation. After multiplying by $r^2/(R\Omega)$, we get

$$\frac{1}{R} \frac{d}{dr} \left(r^2 \frac{dR}{dr} \right) + \frac{1}{\Omega \sin \varphi} \frac{d}{d\varphi} \left(\sin \varphi \frac{d\Omega}{d\varphi} \right) + \frac{1}{\Omega \sin^2 \varphi} \frac{d^2 \Omega}{d\theta^2} + k^2 r^2 = 0. \quad (\text{III.1.4})$$

The above equation can be separated into a radial and angular part:

$$r^2 \frac{d^2 R}{dr^2} + 2r \frac{dR}{dr} + R(k^2 r^2 - l(l+1)) = 0, \quad (\text{III.1.5})$$

$$\frac{1}{\Omega \sin \varphi} \frac{d}{d\varphi} \left(\sin \varphi \frac{d\Omega}{d\varphi} \right) + \frac{1}{\Omega \sin^2 \varphi} \frac{d^2 \Omega}{d\theta^2} = -l(l+1), \quad (\text{III.1.6})$$

where $-l(l+1)$ has been used a separation constant.

To solve the radial differential equation we substitute $\gamma = kr$ to arrive at

$$\gamma^2 \frac{d^2 R}{d\gamma^2} + 2\gamma \frac{dR}{d\gamma} + R(\gamma^2 - l(l+1)) = 0 \quad (\text{III.1.7})$$

This equation is known as the Spherical Bessel Differential Equation of which the general solution is

$$R(\gamma) = a_l^{(1)} h_l^{(1)}(\gamma) + a_l^{(2)} h_l^{(2)}(\gamma), \quad (\text{III.1.8})$$

where $h_l^{(1)}$ and $h_l^{(2)}$ are the l -th order spherical Hankel functions of the first and second kind. The first kind represents a spherical wave propagating towards the origin, while the second kind represents a spherical wave travelling away from the origin. This can be derived from the asymptotic expansion of the Hankel functions in combination with the definition of the Fourier Transform.

We now turn to the angular differential equation in III.1.6. We again employ separation of variables and write the angular pressure as

$$\Omega(\theta, \varphi) = \Theta(\theta)\Phi(\varphi) \quad (\text{III.1.9})$$

and substitute it into the differential equation to get

$$\frac{\sin \varphi}{\Phi} \frac{d}{d\varphi} \left(\sin \varphi \frac{d\Phi}{d\varphi} \right) + \frac{1}{\Theta} \frac{d^2 \Theta}{d\theta^2} + l(l+1) = 0, \quad (\text{III.1.10})$$

which can be separated into an azimuthal and inclination part using the separation constant $-m^2$. After a bit of algebra, the separated parts read

$$\frac{1}{\Theta} \frac{d^2 \Theta}{d\theta^2} = -m^2, \quad (\text{III.1.11})$$

$$\frac{1}{\sin \varphi} \frac{d}{d\varphi} \left(\sin \varphi \frac{d\Phi}{d\varphi} \right) + \Phi l(l+1) - \frac{m^2}{\sin^2 \varphi} = 0. \quad (\text{III.1.12})$$

The general solution for the azimuthal angle θ is

$$\Theta = A_m e^{im\theta} + B_m e^{-im\theta}. \quad (\text{III.1.13})$$

This angle is periodic every 2π and therefore m can only be an integer value.

The differential equation for the inclination angle φ , shown in equation III.1.12, is more complex. To solve it we substitute $t = \cos \varphi$ into the equation to get:

$$(1-t^2) \frac{d^2 \Phi}{dt^2} - 2t \frac{d\Phi}{dt} + \Phi \left[l(l+1) - \frac{m^2}{1-t^2} \right] = 0. \quad (\text{III.1.14})$$

This equation is known as the general Legendre equation and the solution is the associated Legendre polynomial $P_l^m(t)$. It is required that l is a non-negative integer.[40] Furthermore, the polynomial is only non-zero if $|m| \leq l$.

We are now able to describe wave propagation in spherical coordinates and build the migration algorithm for imaging. First we have to redatum the data by shifting time axis of the line scans by

$$t' = t - 2 \frac{r_f}{c}. \quad (\text{III.1.15})$$

After redatuming the data represents a set of pulse-echo measurements done by an array of virtual sources and receivers on a sphere at radius r_f . The next step is to transform the wavefield to the

frequency domain using a Spherical Harmonic Transform for θ and φ , and a Fourier Transform for the time t' . This is mathematically described by

$$\tilde{p}(r_f, l, m, f) = \int_{-\infty}^{\infty} \int_0^{2\pi} \int_0^{\pi} p(r_f, \theta, \varphi, t') P_l^m(\cos \varphi) e^{i(m\theta - 2\pi f t')} \sin \varphi d\varphi d\theta dt'. \quad (\text{III.1.16})$$

The next step is to extrapolate the data. For $r > r_f$ there are only inward propagating waves and therefore only $h_l^{(1)}$ is of importance. Following the ERM this wave travels at velocity $\hat{c} = c/2$ and therefore the extrapolation for these depths can be described by

$$\tilde{p}(r, l, m, f) = \tilde{p}(r_f, l, m, f) \frac{h_l^{(1)}(\hat{c}r)}{h_l^{(1)}(\hat{c}r_f)}, \quad (\text{III.1.17})$$

which has been derived from equation III.1.8. The above equation has to be evaluated for all l , m and f . For $r < r_f$ the spherical Hankel function of the second kind has to be used instead of the first.

Next, an inverse transform has to be done. Furthermore, we know that the reflectors "exploded" at $t' = 0$, so we only have to evaluate that time. This results in the following inverse transform

$$p(r, \theta, \varphi, t' = 0) = \sum_{l=0}^{\infty} \sum_{m=-l}^{+l} \int_{-\infty}^{\infty} \tilde{p}(r, l, m, f) \bar{P}_l^m(\cos \varphi) e^{-im\theta} df \quad (\text{III.1.18})$$

So the whole migration algorithm to form an image is described by equations III.1.16, III.1.17 and III.1.18.

III.1.2. COMPARISON TO POLAR STOLT MIGRATION

Spherical Stolt Migration works for phased matrix arrays, which find their 2D counterpart in phased arrays. Therefore, it has many similarities to the Polar migration as described in section 2.5.2. For example the line density has to satisfy equation 2.24 in both φ and θ . Furthermore, the near and far field can be processed in a single calculation and also the ratio of two Hankel functions must be evaluated. In this case spherical Hankel functions, but they have a simple relation to the ordinary Hankel functions:

$$h_l^{(1)}(\gamma) = \sqrt{\frac{\pi}{2\gamma}} H_{l+1/2}^{(1)}(\gamma). \quad (\text{III.1.19})$$

This means that the spherical Hankel functions are just half integer cylindrical Hankel functions with a different modulus. The phase is unchanged. The approximations for large arguments, shown in equations 2.19 and 2.20, can therefore still be used.

As in the polar case, we still have Fourier Transforms over t' and θ , but in addition an extra transform over φ has appeared here. The transform over the angles together is called the Spherical Harmonic Transform(SHT), which was described in section 2.8.

III.1.3. IMPLEMENTATION

The whole spherical migration algorithm is written in Matlab (Matlab 2015b, The MathWorks, Inc., Natick, Massachusetts, USA) and consists of the operations listed below:

1. Shift the time-axis to $t' = t - \frac{2r_f}{c}$ to create virtual point sources and receivers on a semi-sphere
2. Transform $p(r_f, \varphi, \theta, t')$ to $\tilde{p}(r_f, \varphi, \theta, f)$ by using a FFT and apply a band-pass filter

3. Transform $\tilde{p}(r_f, \varphi, \theta, f)$ to $\tilde{p}(r_f, l, m, f)$ by using a SHT
4. Extrapolate \tilde{p} to all radii of interest using equation III.1.17 and integrate over all frequencies. Repeat this for all m and l
5. Take the inverse SHT for every radius

In step 2 the time domain signal is zero padded before the transformation to increase the density of frequencies and prevent wrapping errors. After the transformation, frequencies outside the -40 dB bandwidth of the transducer are discarded to speed up the algorithm. This does not have any negative effects as the discarded frequencies only have a very minor influence on the final image.

For the forward and inverse SHT in step 3 and 5 an implementation by Politis[41] has been used as basis and modified to our needs. The forward and inverse transforms are implemented as matrix multiplications. Due to this no explicit zero-padding in the angular directions is needed if Gauss-Legendre quadrature is used and the spherical harmonics do not have to be calculated for points outside of the scan sector. With the least squares transformation however, it is required to calculate the spherical harmonics at all points before the matrix inversion, but the values that do not relate to the scan sector can be discarded after the inversion.

For low order transformations this implementation works fine, but for large orders a very large amount of memory is needed for the spherical harmonics matrix. This limits the maximum order that can be processed. Furthermore, this is not the fastest kind of algorithm, but for purposes in this thesis it suffices.

Step 4 is implemented as a matrix multiplication. Because the spherical Hankel function only depends on l , all m for a given l can be multiplied at once and only a loop over all values of l is required.

III.1.4. COMPUTATION TIME

Currently, real time imaging cannot be achieved yet by this algorithm. In this section we propose some ways to decrease the computation time. A symmetry can be used and the extrapolation step can be changed. This will be investigated in the following two subsections.

III.1.4.1. SYMMETRY

One symmetry can be observed when looking at the forward transform of equation III.1.16. For $m \rightarrow -m$ and $f \rightarrow -f$ we get

$$\tilde{p}(r_f, l, -m, -f) = (-1)^m (\tilde{p}(r_f, l, m, f))^* , \quad (\text{III.1.20})$$

which can be derived by noting that $P_l^{-m} = (-1)^m P_l^m$ and that the measured data is real. We will use this to only keep the non-negative m during the calculation. This means that only the non-negative m have to be extrapolation in step 4 of the implementation described in the previous section. After the extrapolation and integration over all frequencies, the above symmetry leads to

$$\tilde{p}(r_f, l, -m, t' = 0) = (-1)^m (\tilde{p}(r_f, l, m, t' = 0))^* , \quad (\text{III.1.21})$$

The final step in the implementation is to do an inverse SHT. Due to the symmetry in equation III.1.21 it can be evaluated as follows:

$$\tilde{p}(r, \varphi, \theta, t' = 0) = \sum_{l=0}^{\infty} \tilde{p}(r, l, 0, t' = 0) \bar{P}_l^0(\cos \varphi) + \sum_{l=1}^{\infty} \sum_{m=1}^{+l} 2\Re \left(\tilde{p}(r, l, m, t' = 0) e^{-im\theta} \right) \bar{P}_l^m(\cos \varphi), \quad (\text{III.1.22})$$

where the \Re means taking the real part. This can be derived by noting that P_l^m is always real and thus

$$\begin{aligned} \tilde{p}(r, l, m, t' = 0)e^{-im\theta} P_l^m(\cos\varphi) + \tilde{p}(r, l, -m, t' = 0)e^{+im\theta} P_l^{-m}(\cos\varphi) = \\ \left[\tilde{p}(r, l, m, t' = 0)e^{-im\theta} + \left(\tilde{p}(r, l, m, t' = 0)e^{-im\theta} \right)^* \right] P_l^m(\cos\varphi) = \\ 2\Re \left(\tilde{p}(r, l, m, t' = 0)e^{-im\theta} \right) P_l^m(\cos\varphi) \end{aligned} \quad (\text{III.1.23})$$

By using this symmetry the computation time can most likely be cut down by a factor two. Furthermore, memory requirements for pre-calculation of the spherical harmonics or the ratio of the Hankel functions will also go down by a factor two.

III.1.4.2. FHT ALGORITHM

Step 4 listed in the implementation above is a time-consuming operation and is mathematically described for a single l and m by

$$b(r) = \int_{-\infty}^{\infty} \tilde{p}(r_f, f) \frac{h_l^{(1)}(kr)}{h_l^{(1)}(kr_f)} df. \quad (\text{III.1.24})$$

This then has to be repeated for all radii of interest. Effectively we do a transformation from $k(f)$ to r in this step. This is very similar to a Hankel Transform. The differences are that we have a ratio of two spherical Hankel functions instead of a Bessel function and that we also have negative frequencies. However, this does not hinder from using similar algorithms to the ones employed to do Hankel Transforms. In this subsection we will derive a QFHT like algorithm for the positive frequencies. The negative frequencies can be processed separately by a similar algorithm.

The first step is to employ a change of variables. Here we will use

$$f = f_0 e^{\alpha x}, r = r_0 e^{-\alpha y}. \quad (\text{III.1.25})$$

Substituting the above in equation III.1.24 and multiplying both sides by r gives us

$$r_0 e^{-\alpha y} b(r_0 e^{-\alpha y}) = \int_{-\infty}^{\infty} \frac{\tilde{p}(r_f, f_0 e^{\alpha x})}{h_l^{(1)}\left(\frac{2\pi}{\hat{c}} f_0 r_f e^{\alpha x}\right)} h_l^{(1)}\left(\frac{2\pi}{\hat{c}} f_0 r_0 \alpha e^{\alpha(x-y)}\right) f_0 r_0 \alpha e^{\alpha(x-y)} dx, \quad (\text{III.1.26})$$

which we simplify as

$$\bar{b}(y) = \int_{-\infty}^{\infty} \bar{f}(x) \bar{h}_l^{(1)}(x) dx, \quad (\text{III.1.27})$$

where $\bar{b}(y)$ contains the left hand side terms of equation III.1.26, $\bar{f}(x)$ the division on the right hand side and the other terms are absorbed in $\bar{h}_l^{(1)}(x - y)$. This equation has the form of a convolution with the kernel $\bar{h}_l^{(1)}$ of which the axis is flipped. A convolution like this can be efficiently evaluated using Fourier Transforms as

$$\bar{b}(y) = \mathcal{F}^{-1} \left[\mathcal{F}(\bar{f}(x)) \mathcal{F}^{-1}(\bar{h}_l^{(1)}(x - y)) \right], \quad (\text{III.1.28})$$

where we have used that if the axis are flipped we can use the inverse Fourier Transform instead of the forward.

The above equation, after dividing by r , does the same as step 4, but on an exponentially spaced grid. It can therefore not be used directly as the frequencies are on an equispaced grid. We have opted to solve this by replacing the FFT used to evaluate the transform from t' to f in equation III.1.16 by a non-uniform FFT (NUFFT). Furthermore, to arrive at an equispaced radial grid we have replaced the

outer Fourier Transform of the above equation with a non-uniform one. The NUFFT implementation that we will use for both transformations is the implementation by Fessler et al.[21]

The computational complexity of evaluating step 4 using this QFHT like algorithm scales with $O(N \log N)$ instead of $O(N^2)$, which was the case for straightforward evaluation, and could therefore be faster. However, the use of NUFFT might induce phase errors that will deteriorate the image quality. A bigger problem is the amount of frequencies that has to be used. Due to the exponential grid the spacing at high frequencies is much larger than between the lower ones. To make sure these are correctly taken into account the amount of frequencies has to be higher than required without this QFHT like algorithm. This increases the computation time for both the forward transform and step 4 in the implementation.

III.1.5. EXTENSION TO SPHERICAL WAVES

For phased arrays usually in the order of 100 scan lines are used. An extra angle means that the number of scan lines can easily increase to 100^2 lines. This is too many for real time imaging. Assuming that we can create 10^4 lines per second, a frame rate of 100 can be achieved with a phased array, but when reconstructing a volume the frame rate drops to 1.

It is possible to derive an algorithm which works with a single insonification by a spherical wave. This can be done by following the approach as described in section 2.5.3. The incident wavefield here is a spherical wave that is described by e^{-ikr} . The measured data can be extrapolated using equation III.1.8 and noting that the measured waves are always propagating towards the origin. This means that the imaging condition of equation 2.12 for this situation has to be evaluated as follows:

$$b(r, l, m, t' = 0) = \int_{-\infty}^{\infty} \tilde{p}(r_f, l, m, f) e^{ikr} \frac{h_l^{(1)}(kr)}{h_l^{(1)}(kr_f)} df. \quad (\text{III.1.29})$$

In principle a similar implementation scheme as described in section III.1.3 can still be used, replacing only step 4 by the above equation. However, the SHT can only be applied on data that lies on a sphere. This is usually not the case as we measure the data using a flat transducer. So the measured data lies on a flat surface instead of on the required sphere.

There are multiple ways to solve this issue. A simple and fast way is to focus the data in receive at multiple points on a sphere. Effectively we create an array of virtual receivers by doing this. If this method is used, the time-axis has to be shifted to $t' = t - \frac{r_f}{c}$ in step 1 instead of $t' = t - \frac{2r_f}{c}$.

Another way is to use the Rayleigh II integral which was presented in equation 2.8 to propagate the data to a set of virtual receivers on a sphere. This method is slower, but is most likely more robust. Step 1 in the implementation scheme is not required any more if this integral is used to transform the measured data.

III.2

SIMULATION AND EXPERIMENTAL SETUP

Both simulations, using FieldII, and measurements will be done to evaluate the performance of Spherical Stolt Migration (SSM) and Spherical Wave Stolt Migration (SWSM) with respect to DRF. SSM works with line scans, just like linear and polar Stolt as described in respectively sections 2.5.1 and 2.5.2. SWSM works with a single insonification by a spherical wave. We will use two versions of SWSM that create an array of virtual receivers in a different way. The first version uses receive focusing and will be named SWSM-RF, while the second version uses the Rayleigh II integral and will be named SWSM-RII. Both versions are shortly described at the end of section III.1.5.

Before displaying the final image of each method, first a depth correction is applied to compensate for the overall trend. This is done by first fitting a low order polynomial through a set of known scatterers and then dividing the scan lines through the fit. After this correction, the absolute value of the Hilbert transform is taken over the scan lines and finally the image is log compressed.

Beside a visual inspection of the images created using the 4 methods, we will use the FWHM and FWTM to quantify the spatial resolution. The FWHM will be compared to the far-field width obtained with the tool developed by De Wit, which is described in section 2.6.1.

Most of the evaluation will be done using simulations only in setup 1 described below. To verify the various simulations we will compare SSM with Polar Stolt Migration, which has been experimentally verified. Furthermore, we will compare measurements and simulations done on setup 2. A summary of the settings used in both setups can be found in tables III.2.1 and III.2.2. Both setups will be further described below.

III.2.1. SETUP 1

In this setup, which will be simulated only, a phantom consisting of point scatterers will be imaged. These point lie on three lines. The first line runs at the centre and contains 9 point scatterers between 20 and 110 mm. The second line lies in the x - z plane at a depth of 60 mm and contains 9 points spaced by 7.5 mm. The third line is similar to the second, except that it lies in the y - z plane.

The probe is based on a conventional phased array probe (ATL P4-1, Philips Ultrasound, Bothel, Washington, USA) used in phased array imaging. Instead of 96 elements, the simulated probe contains 48x48 elements with a pitch of 295 μm in both directions. The centre frequency is 2.5 MHz. The excitation pulse is assumed to be a 3 cycle Gaussian windowed sinusoid, while the impulse response is a delta-peak, i.e. infinite bandwidth in receive.

The sector scan, used as the first stage of SSM, consists of 128 scan lines in φ and θ both within 90 degrees. This number of scan lines satisfies the required line density given by equation 2.24. The scan angles in θ are equispaced, but in φ we have chosen the Gauss nodes as angles to allow for the faster Gauss-Legendre quadrature in calculating the forward Spherical Harmonics Transform. The probe is focused in transmit at 30 mm, which corresponds to an f-number of about 2.2. The choice for this relatively large f-number is related to the used approximation, as discussed in the appendix. For the first stage of SSM the probe was also focused at this depth in receive. During the sector scan a Hamming window was used both in transmit and receive.

For the insonification with a spherical wave, required for SWSM, the virtual focus behind the transducer was at 7 mm, which translates into a half width opening angle of about 45° . With the Rayleigh II integral the measured data was propagated to 128×128 points on a radius of about 20 mm as seen from the virtual focus. With receive focusing the measured data was also propagated to 128×128 points, but now on radius of 30 mm as seen from the virtual focus. During this scan a Tukey window was used in transmit with parameter 0.4, which makes it closer to a rectangular window than a Hann window, and a Hamming window in receive.

Table III.2.1: Settings used for the first setup.

	Parameter	Value
Transducer	Centre frequency	2.5 MHz
	Number of elements	48 x 48
	Pitch	295 μm \times 295 μm
	Element width	245 μm \times 245 μm
	Apodization in receive	Hamming
Line scan settings	Number of scan lines	128 x 128
	Scan angle	$\pm 45^\circ \times 45^\circ$
	Focal depth	30 mm
	Apodization in transmit	Hamming
Spherical wave scan settings	Virtual focus depth	-7 mm
	Opening angle	$90^\circ \times 90^\circ$
	Apodization in transmit	Tukey ($\alpha = 0.4$)
FieldII	Excitation	3 cycle Gaussian windowed sinusoid
	Sampling Frequency	100 MHz
Medium	Speed of sound	1540 m s^{-1}
	Attenuation	Negligible

III.2.2. SETUP 2

With this setup an image of the phantom shown in figure III.2.2 will be made. It consists of tiny needles, of which the longest is about 1 cm, with a small ball at the top with a diameter in the order of 1 mm, which is about 3.3 wavelengths. The centre row of needles is almost equispaced within 1 cm and should all be at a different depth. To measure the PSF over a larger depth range, the phantom was measured at different distances from the probe.

For these measurements the MICA probe is used. This matrix probe is under development for use in transesophageal echocardiography. It consists of a separate transmit and receive part as indicated in figure III.2.1. For our measurements only the transmit part is used, as it is impossible to get separate element data from receive part. The transmit part consists of 8×8 elements with a pitch of 150 μm in each direction. The centre frequency is 5 MHz. During the measurements the probe was connected

to a Verasonics Data Acquisition System (Verasonics Inc., Kirkland, Washington, USA) that samples the data at 20 MHz. The unfiltered pulse-echo signal of one of the elements and its Fourier Transform are shown in figure III.2.3. For display purposes it has been upsampled by a factor 5 to 100 MHz using a FFT method. This upsampled version is also used in the simulations as the excitation signal to allow for a good comparison to the measurements. Both simulations and measurements were bandpass filtered using a 4-th order Butterworth filter between 2.75 and 7.75 MHz. Furthermore, for Stolt migration an ideal bandpass filter has been used between 2 and 8 MHz.

In the first stage 31x31 scan lines were generated in φ and θ both within 90 degrees. In θ the angles are equispaced, while in φ Gauss nodes are used. The probe was focused at 1 mm, which corresponds with an f-number of 5/6. For SSM the same focus depth was used in transmit.

For SWSM the focus was 0.6 mm behind the transducer to allow for an opening angle of 90°. The measured data was then propagated using the Rayleigh II integral to 31x31 points at 15 mm from the virtual focus. This propagated data was the starting point for SWSM.

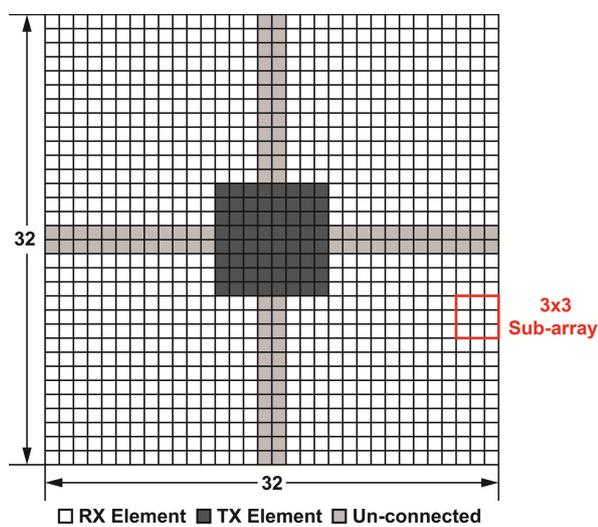


Figure III.2.1: Outline of the MICA probe. Only the TX elements are used for the measurements in this thesis.



Figure III.2.2: The phantom used in setup 2. It consists of tiny needles with a small ball at the top that has a diameter in the order of 1 mm.

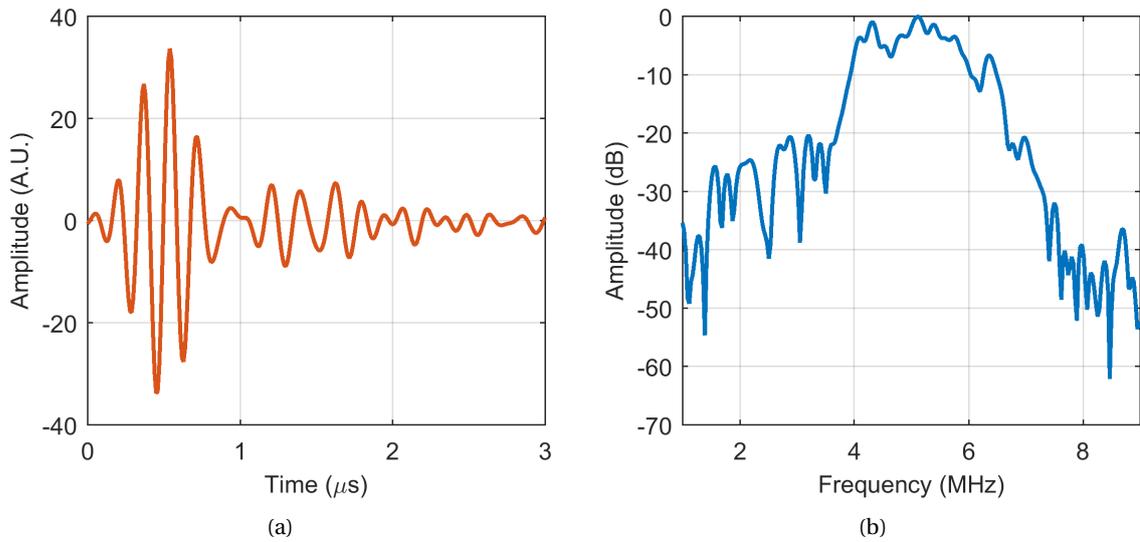


Figure III.2.3: In (a) the unfiltered pulse-echo signal from one of the transmit elements of the MICA probe and in (b) the Fourier Transform of this signal

Table III.2.2: Settings used for the second setup.

	Parameter	Value
Transducer	Centre frequency	5 MHz
	Number of elements	8x8
	Pitch	150 μm \times 150 μm
	Element width	130 μm \times 130 μm
	Apodization	None
Line scan settings	Number of scan lines	31x31
	Scan sector	$\pm 45^\circ \times 45^\circ$
	Focal Depth	1 mm
Spherical wave scan settings	Virtual focus depth	-1 mm
	Opening angle	$90^\circ \times 90^\circ$
	Receive focus radius	15 mm
Verasonics	Excitation	2 half cycles with equalization pulse
	Sampling frequency	20 MHz
FieldII	Excitation	Measured pulse-echo
	Sampling frequency	100 MHz
Medium	Speed of sound	1480 m s^{-1}
	Attenuation	Negligible

III.3

RESULTS

Here we will present various results obtained with Spherical Stolt Migration (SSM) and Spherical Wave Stolt Migration (SWSM). SSM works with line scans, while SWSM works with a single insonification by a spherical wave. We will use two versions of SWSM that create an array of virtual receivers in a different way. The first version uses receive focusing and will be named SWSM-RE, while the second version uses the Rayleigh II integral and will be named SWSM-RII. Both versions are shortly described at the end of section III.1.5.

First we will compare DRF applied to line scan data with the migration techniques. Next we will investigate if there are any visible differences between doing the SHT with Gaussian Quadrature on a grid with Gauss nodes or a least squares transformation on an equispaced grid. Then, the computation time is measured with and without the possible improvements described in section III.1.4. After this we will investigate SWSM a bit further and compare it with DRF applied to spherical wave data. In section III.3.5, we will verify the obtained results by comparing SSM with Polar Stolt Migration and compare simulations and measurements. Finally, the differences between PSASB and Polar Stolt will be investigated.

III.3.1. COMPARISON

In figure III.3.1 a slice of the volume at $\varphi = \pi/2$ is shown of the images generated with DRF applied on a line scan, SSM, SWSM-RII and SWSM-RE. The images generated with DRF and SSM look very similar. Both show an increasing PSF width and the same kind of artefacts. However, especially from the point scatterers at deeper depths it becomes clear that the points appear bigger in the DRF image. This can be more clearly seen in figure III.3.3(a) where the FWHM is shown. At transmit focus both methods produce similar widths, but the FWHM increases far more steeply as a function of depth for DRF than for SSM. At the deepest point scatterer in the image the improvement is 25% and this will most likely improve more even further from the focus. The result for DRF at deeper depths could be improved by using a transmit focus further from the transducer, but this will deteriorate the image closer to the transducer. This is thus a clear trade off that is not present in SSM. The same effects are visible in the FWTM, as shown in figure III.3.3(b). Besides this, SSM does also approach the theoretical far field width, indicating that it manages to focus the image over the whole depth range.

If we look further at the FWHM we see that it is slightly smaller in θ compared to the FWHM in φ for SSM, while there is no difference for DRF. This might be caused by the different grid in θ and φ . In θ an equispaced grid is used, while the points in φ are at the zeros of a Legendre polynomial. These

are almost equispaced, but not quite and might thus cause the FWHM to differ a bit. The FWTM is however unaffected, which should differ a bit too if the different kind of spacing was the cause. Besides this small difference in the FWHM, there are no clear differences visible in the image as can be seen when comparing figure III.3.1 and III.3.2. Therefore, it is not necessary to look at both angles the whole time.

The images created with both versions of SWSM appear a bit different from the other images. There are two differences caused by the different apodization. Two parabolas can be seen that cross each point and ghost points appear just below each point. A stronger apodization would help to reduce these artefacts, but this would also reduce the pressure transmitted into the medium and thus the signal. We note that the ghost points are also more visible because the focus is much further away from the natural focus than in the sector scan.

What is also visible if we compare the images generated by SWSM with the others, is that the points appear larger, which is confirmed by the lateral FWHM visible in figure III.3.3(a). Compared to SSM, the difference is larger than 25% in the whole image. One should note that only one transmit and receive event is used in SWSM and that for DRF and SSM we have used 128x128 transmit and receive events, so the difference could be expected. Interestingly, we see that for deeper depths the FWHM and FWTM of SWSM-RII almost becomes equal to the one obtained with DRF. This is not observed for SWSM-RF, as the FWHM is about 8% larger than obtained with SWSM-RII in the whole image.

There are more differences between both SWSM methods. SWSM-RF seems to result in a larger background noise level. This is especially the case if the lower frequencies, say beneath 0.8 MHz, are not filtered using a Butterworth filter. In this case the background noise level in the image obtained was at about -40 dB. It is not clear why this method suffers from this, while the other methods perform well even if these frequencies are not filtered. Fine tuning the filter might further reduce the background noise level. In this case both images will most likely be more similar, but we don't expect the resolution to improve. So, only by using the Rayleigh II integral, as is done SWSM-RII, a better resolution can be obtained.

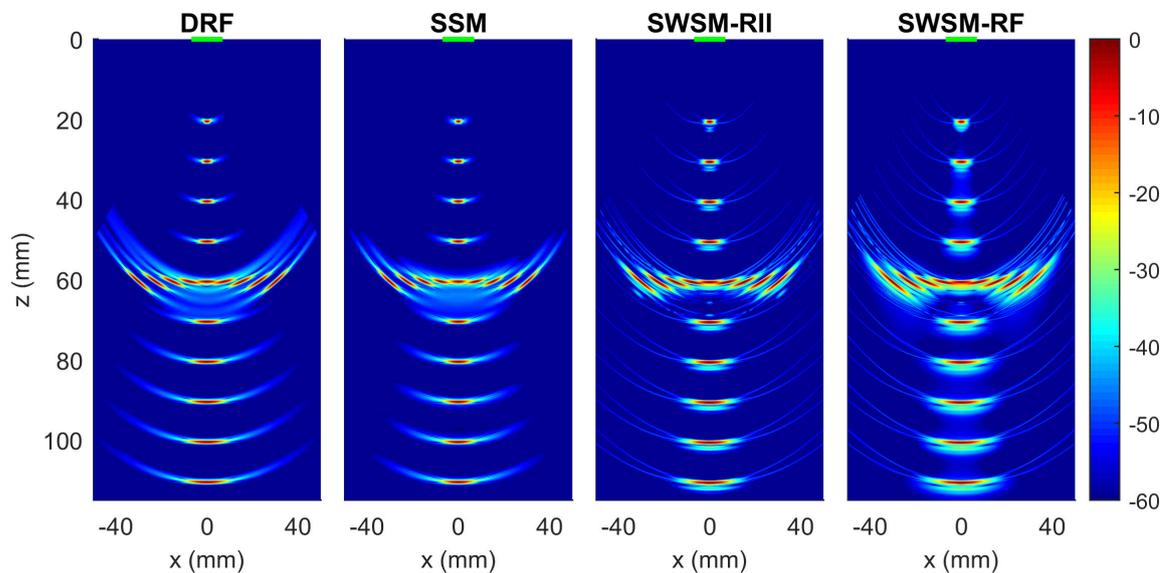


Figure III.3.1: Slice at $\varphi = \pi/2$ of the images generated with each of the four methods.

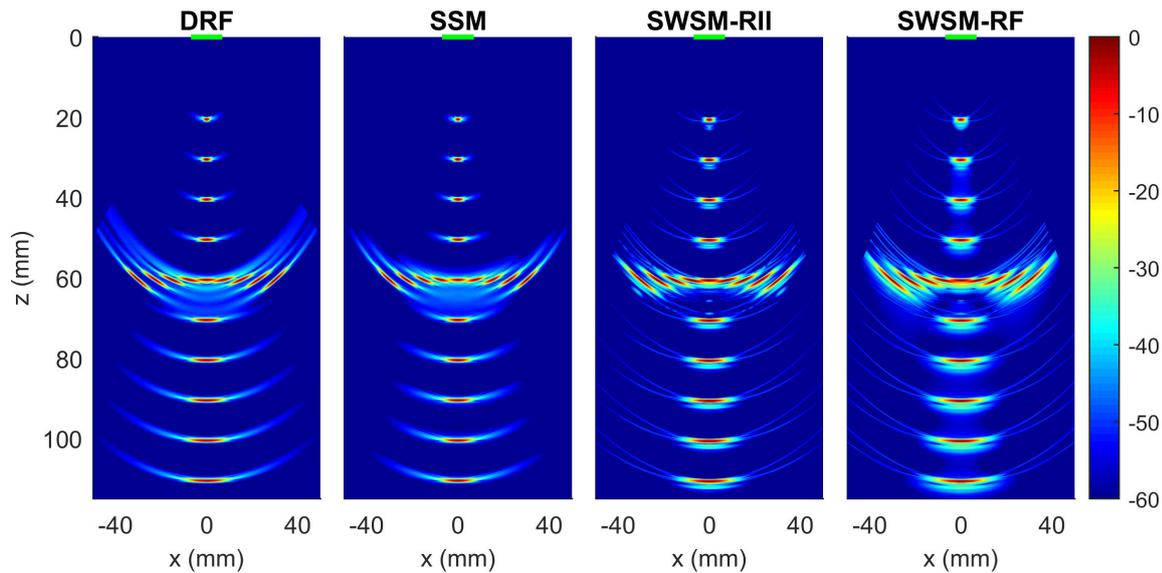


Figure III.3.2: Slice at $\theta = 0$ of the images generated with each of the four methods.

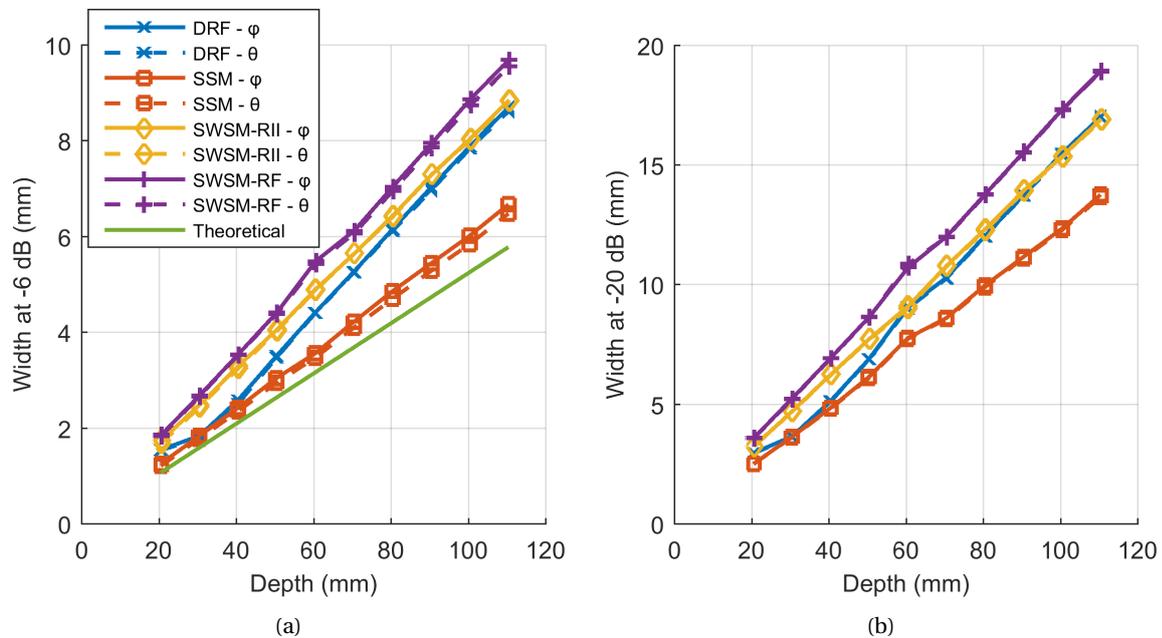


Figure III.3.3: In (a) the FWHM and in (b) the FWHM in θ and φ for the points visible in figures III.3.1 and III.3.2

III.3.2. LEAST SQUARES TRANSFORMATION

The results above obtained with the Stolt migration methods all used Gaussian quadrature to do the forward SHT. To do the quadrature the inclination angles have to be Gauss nodes. In some situations, data might be available that has been measured with a different set of angles. This data could still be processed with a least squares transformation instead of a Gaussian quadrature.

In figure III.3.4 we show the FWHM and the FWTM obtained with Gaussian quadrature on a grid with Gauss nodes or obtained with a least squares transformation on an equispaced grid. It is clear that we cannot distinguish the results. This is also the case for the images itself, which are not shown here. This proves two things. First, it does not seem to matter if a Gaussian quadrature is done or a least squares transformation. Furthermore, the results obtained on an equispaced grid do not differ from

the results on a grid with Gauss nodes as inclination angles.

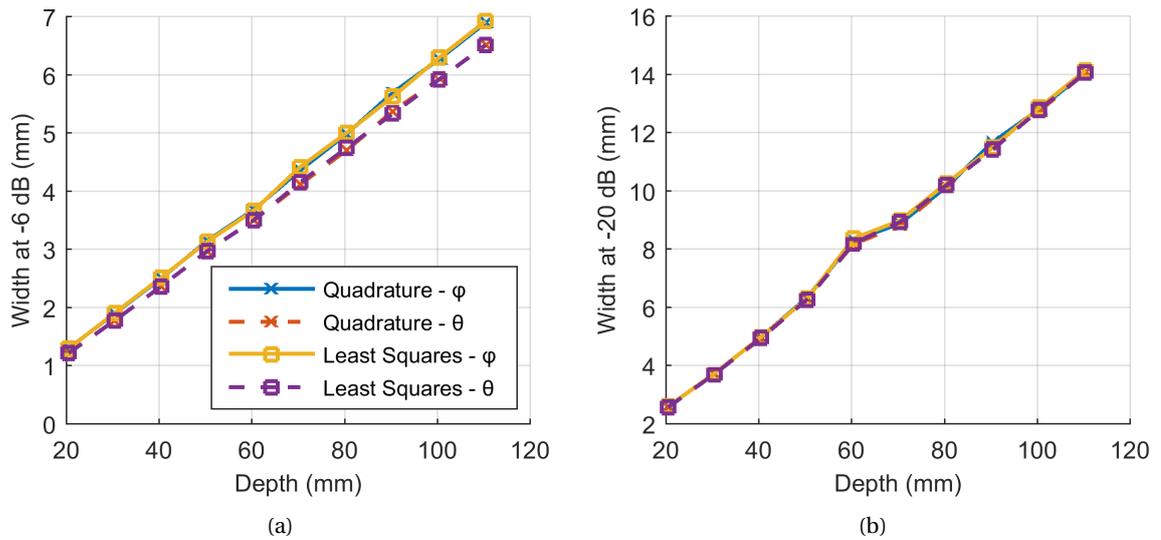


Figure III.3.4: In (a) the FWHM and in (b) the FWTM in θ and ϕ for the image created using Gaussian quadrature on a grid with Gauss nodes as inclination angles or using a least squares transformation on a fully equispaced grid.

III.3.3. COMPUTATION TIME

In this section we investigate the results if the symmetry or the FHT like algorithm, both introduced in section III.1.4, are used. For this we use SSM with only used 64x64 scan lines in a scan sector of 90°, instead of the 128x128 considered before due to memory limitations. We will compare the generated images and the computation time with and without the possible improvements. We will track the time of steps 2 to 5 as listed in section III.1.3. The time required for calculating the spherical harmonics matrix, used in the forward and inverse SHT, is excluded. However, the calculation of the Hankel functions, required for the extrapolation step, are included in the computation time.

A slice of the resulting images is shown in figure III.3.5. First of all, with or without the symmetry in f and m the image on the left is generated. There is thus no change in the image by using this symmetry, but there is a difference in speed as can be seen in table III.3.1. The computation time of both the forward and inverse transform almost halved if the symmetry is used. This improvement is not seen for the extrapolation step, as it is just 21% faster. The reason is that calculation of the Hankel functions is very time-consuming and the calculation time of these functions is not affected. The reason is that the Hankel functions do not depend on m . Still, the whole algorithm is 37% faster, without any effect on the image.

By using the FHT like algorithm for step 4 in the implementation the two rightmost images are generated. For the centre one, which we will consider first, a total of 8192 exponentially spaced frequencies were considered. This is significantly more than the about 1700 equispaced frequencies used in the normal algorithm. The resulting image contains a lot more artefacts, especially at the lower depths. Furthermore, the points appear wider in the radial direction. The differences are most likely caused by phase errors due to the use of a NUFFT. It is certainly not caused by the used approximation for the Hankel function. Using a higher order approximation or the exact function does not matter.

While the FHT like algorithm was introduced as a possible way to decrease the computation time, it seems to do the opposite. The total calculation for the centre most image has almost doubled. The

main issue is the many extra frequencies that have to be considered, which affect the calculation time required for the forward SHT the most. Furthermore, a NUFFT is slower than a FFT. This has the largest effect on the extrapolation step as using an FFT instead of a NUFFT in this step would reduce the computation time by a factor 3. However, this would have resulted in an exponentially spaced radial axis, which is unwanted.

The amount of frequencies can be reduced by a factor two to improve the speed. The time taken is shown in the table under FHT 2. The total calculation time is almost equal to the normal situation, but not faster. This does have its effects on the image as the artefacts are far more severe. Interestingly, most points do not appear elongated in the radial direction any more.

From this it is clear that the goal of a faster algorithm has not been reached with the FHT. Furthermore, the extra artefacts are unwanted. Fine tuning the parameters might reduce artefacts and a better implementation might improve the speed. From these results however, it seems that the use of a symmetry is the better way towards real time imaging.

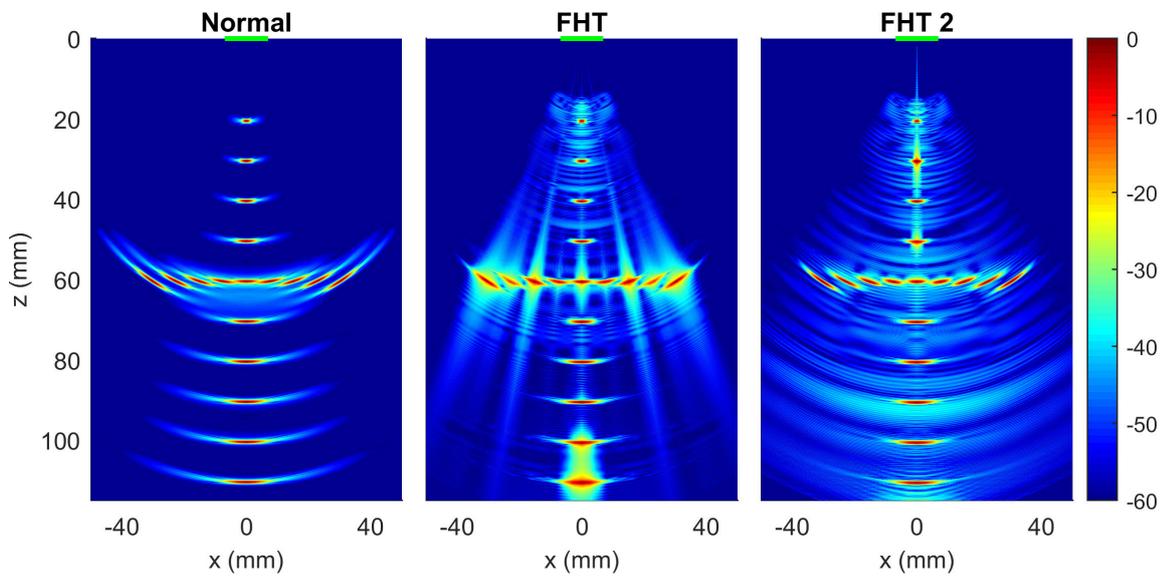


Figure III.3.5: Slice at $\varphi = \pi/2$ of the images. Doing the extrapolation in a straightforward manner results in the image on the left. With the FHT like algorithm the image in the centre and right are generated. The difference between these two is that fewer frequencies are considered in the right one.

Table III.3.1: The computation times relative to the computation time of the implementation without the possible improvements described in III.1.4. The difference between FHT and FHT 2 is that for the latter the frequency grid is less dense.

	Normal	Symmetry	FHT	FHT 2
(NU)FFT	1	1	6	3
Forward SHT	1	0,49	3	1,49
Extrapolation	1	0,79	1	0,5
Inverse SHT	1	0,57	1	1
Total	1	0,63	1,85	1,03

III.3.4. SPHERICAL WAVE STOLT MIGRATION

In this section we delve a bit deeper into Spherical Wave Stolt Migration(SWSM). First we compare the migration results with DRF applied to spherical wave data and after that the effect of the radius

at which the virtual receivers are located. There might be more things affecting the final result with Stolt Migration, but except for the virtual focus depth, which did not have much impact, no other things were investigated.

III.3.4.1. COMPARISON WITH DRF

In the first section of this part we only showed DRF applied to data from a sector scan. In this section, we will apply DRF to spherical wave data and compare it with the result of SWSM.

A slice of the reconstructed volume using 3 different methods is shown in figure III.3.6. We see that DRF produces a very similar result to SWSM-RF. The points only appear a bit different in the radial direction, but the lateral width appears equal. SWSM-RII does appear more distinct. The points appear more confined both in the radial and lateral direction.

This last point is confirmed by the FWHM and FWTM as visible in III.3.7. The PSF width in the lateral direction is smallest for SWSM-RII. The other two methods produce similar results. So it seems that SWSM-RF can at best produce an image similar to DRF. Better resolution can only be obtained in this case if the Rayleigh II integral is used, as is done in SWSM-RII.

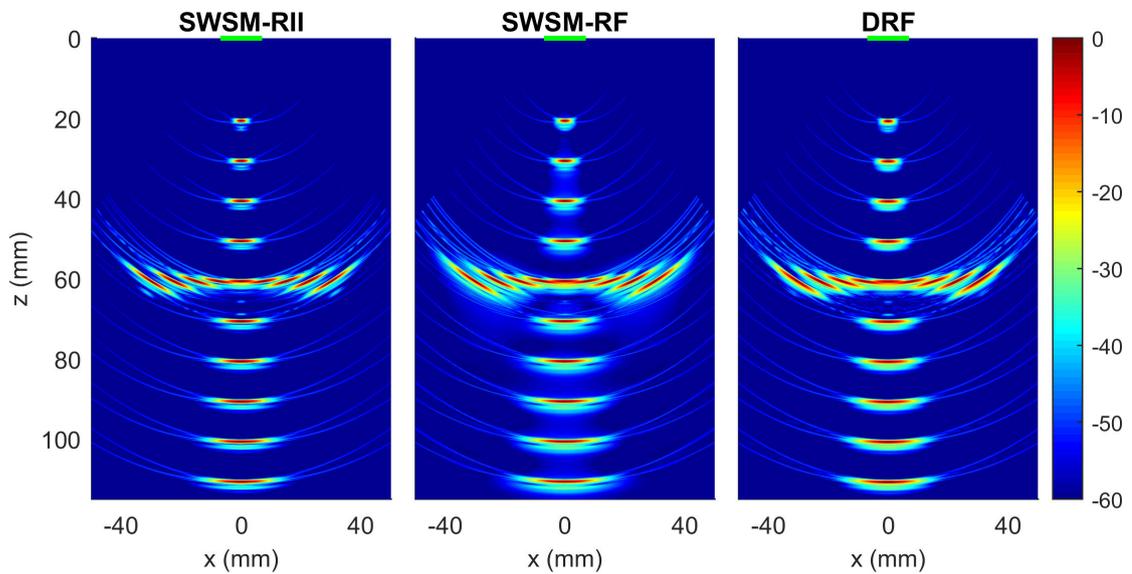


Figure III.3.6: Slice at $\varphi = \pi/2$ of the images generated from spherical wave data.

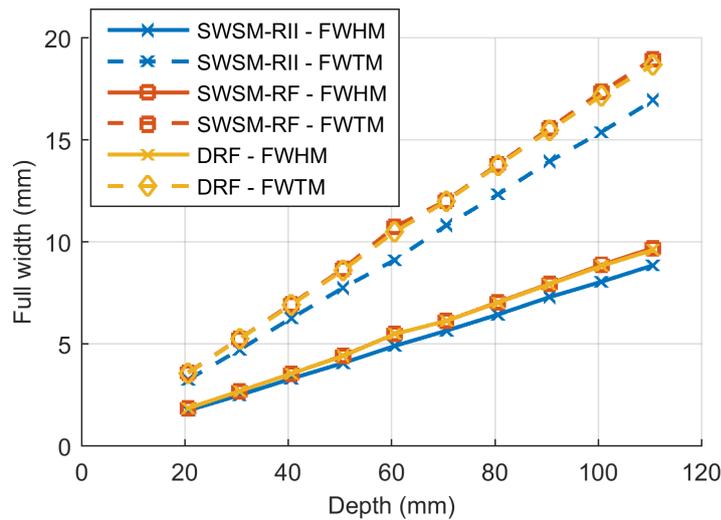


Figure III.3.7: The FWHM and FWTM of the points visible in figure III.3.6

III.3.4.2. VIRTUAL RECEIVER RADIUS

The radius as seen from the virtual focus at which the virtual receiver array is placed in SWSM does affect the final image. This can be seen in figure III.3.8, where we plotted the FWHM and FWTM for different radii as obtained with SWSM-RII. The results are similar for SWSM-RF. If the radius is too small, the resolution suffers and, although not shown, extra artefacts appear. The transition is somewhere around 10 mm, close to the point at which all virtual receivers lie above the transducer. The centre of the transducer is found at a radius of 7 mm as seen from the virtual focus, but the edges of the transducer are at a larger radius.

This behaviour is caused by multiple effects. First of all, the approximation that is used for the Hankel function results in very high phase errors if the radius is too small, as can be seen in the appendix. Using a higher order approximation or the exact values reduces the artefacts if the radius is 10 mm, but the image with a virtual receiver radius of 5 mm is almost unchanged. Most likely the small effective aperture of the virtual array causes this. All virtual receivers are so close together if the radius is small that they effectively measure the same, making it very difficult to retrieve useful information from the migration. The same kind of behaviour is also observed in line scans for low f -numbers.[3][6] On the other hand, if a too large f -number is used, the resolution has also been shown to suffer in line scans, but that is not observed here.

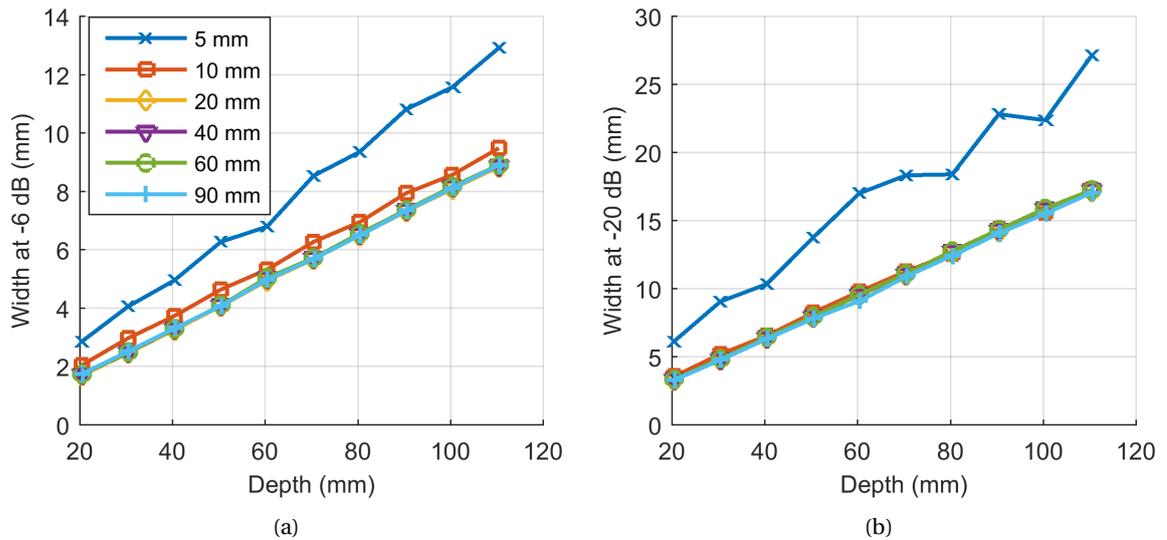


Figure III.3.8: In (a) the FWHM and in (b) the FWTM obtained after using SWSM RII for different virtual receiver radii.

III.3.5. VERIFICATION

To verify the results presented above we have done two things. First the simulation results of SSM are compared with Polar Stolt, which already has been experimentally verified. After that we look at measurements done with the needle phantom and compare them with simulations.

III.3.5.1. POLAR AND SPHERICAL COMPARISON

For as slice at $\varphi = \pi/2$, similar results are expected for SSM and Polar Stolt. The images shown in figure III.3.9(a) do seem to indicate that this is true. The shape of each point scatterer in both images is very similar. The "wings" that appear at the sides of each point are different though. In SSM these wings appear at the same radii as the points. Polar Stolt also has this kind of wings, but another pair of wings seems to stay at constant depth. Both wings together in Polar Stolt also appear to be dimmer compared to the wings in SSM. The largest effect of this is visible at a depth of about 60 mm where multiple points appear. At this depth the background level for SSM is much higher there as compared to Polar Stolt.

The FWHM obtained with both methods, shown in figure III.3.9(b), seems to be consistently a bit lower with Polar Stolt. The difference in FWHM is almost not visible, but the FWTM is clearly lower, especially at deeper depths. This might be caused by the different "wing" shape that each point has.

These small differences might be explained by the small differences in the setup. For example the probe used for SSM was apodized in both x and y during transmit and in receive. For Polar Stolt there was no apodization in y and there is no active focusing in y . Furthermore, the other slices in φ , which are not considered in Polar Stolt, can also influence the slice that is shown. Overall, the differences are small and we can conclude that both methods produce very similar results in a single slice. Therefore, the results that De Wit[6] obtained with Polar Stolt are also relevant for SSM.

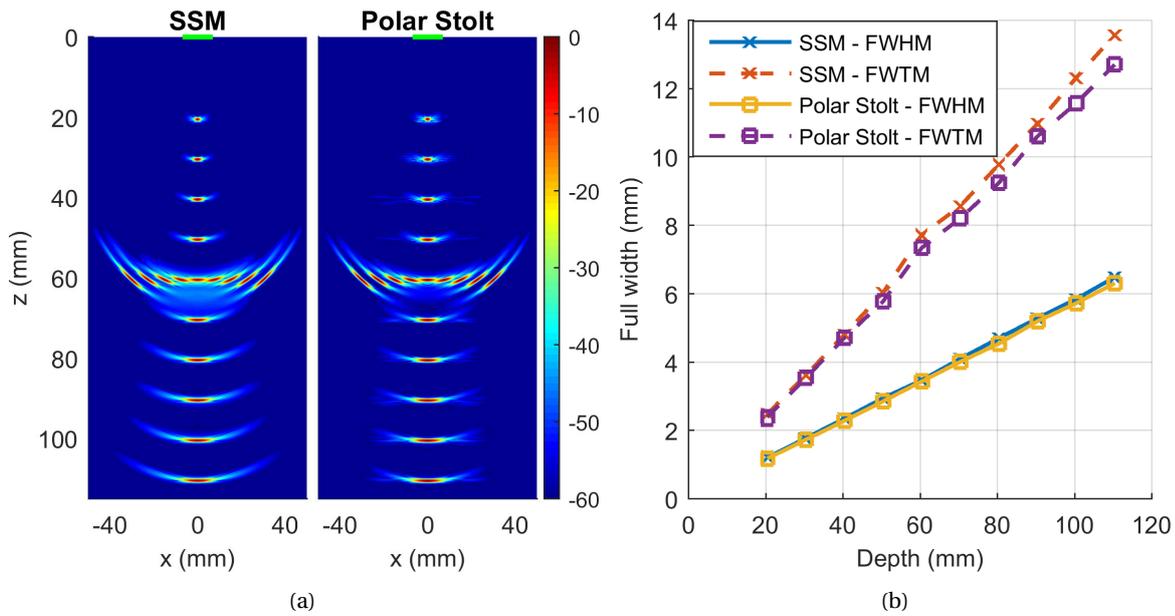


Figure III.3.9: In (a) images created with SSM and Polar Stolt. In (b) the FWHM and FWTM of the PSF as a function of depth for the points visible in (a).

III.3.5.2. MEASUREMENTS

LINE SCAN

In figure III.3.10 the FWHM of the point scatterers in the measurements and the corresponding simulation are shown. Two measurements were done to obtain these values, however all points were simulated at once. The difference between measurements and simulation is very small for DRE. SSM produces a smaller width than DRE, but the width obtained with SSM in the simulation is clearly lower than in the measurements. The difference might be caused by the objects used in the measurements. As mentioned they are rather large compared to the wavelength and therefore they cannot really be considered as point objects as is done in the simulation. Interestingly, the FWHM of SSM in the measurement is almost right on top of the theoretical width, but in the simulation it is consistently below this. We note however, that the theoretical width is obtained with a far-field approximation and that we cannot consider the imaged depths as being in the far field.

To further investigate the behaviour we look at figures III.3.11 and III.3.12 where the deepest reflector of the phantom is shown in two different measurements. First of all we see that the simulated DRF images are almost equal in both figures. This is expected because from Fraunhofer diffraction we can derive that in the far-field the PSF as a function of angle is independent of depth[34], which can be observed here. Furthermore, the PSF of a square aperture obtained with continuous waves is very similar to the one that appears here. The difference is explained from the fact that instead of a single frequency, pulsed waves contain multiple frequencies and thus the PSF is effectively a weighted sum of the different frequencies. SSM has a different PSF because the effective aperture, containing the virtual point sources and receivers, is a spherical surface.

The slices constructed by DRF of the measurement data look very similar to the slices of the simulation. The noise levels for the measured results seem to be a bit higher and reverberation of points above the shown slice are also visible. What we can also see in the measurements is that the central spot appears more round than in the simulation. This is because we do not image a point scatterer, and thus the PSF, in the measurements, but a finite size scatterer. Therefore, a convolution of the PSF with the measured object is visible.

The differences for SSM are much larger. It seems to be more affected by the fact that we do not have a point scatterer. The spot also appears to be larger than obtained in DRF, even though the FWHM with SSM is lower. However, from the simulation images we can determine that SSM has a larger FWTM. This could explain why the spot in the measurements appears larger than with DRF.

So, the difference between simulation and measurements can be explained. Therefore, even though the spot does not appear smaller in SSM, the results prove that the method also works on measured data.

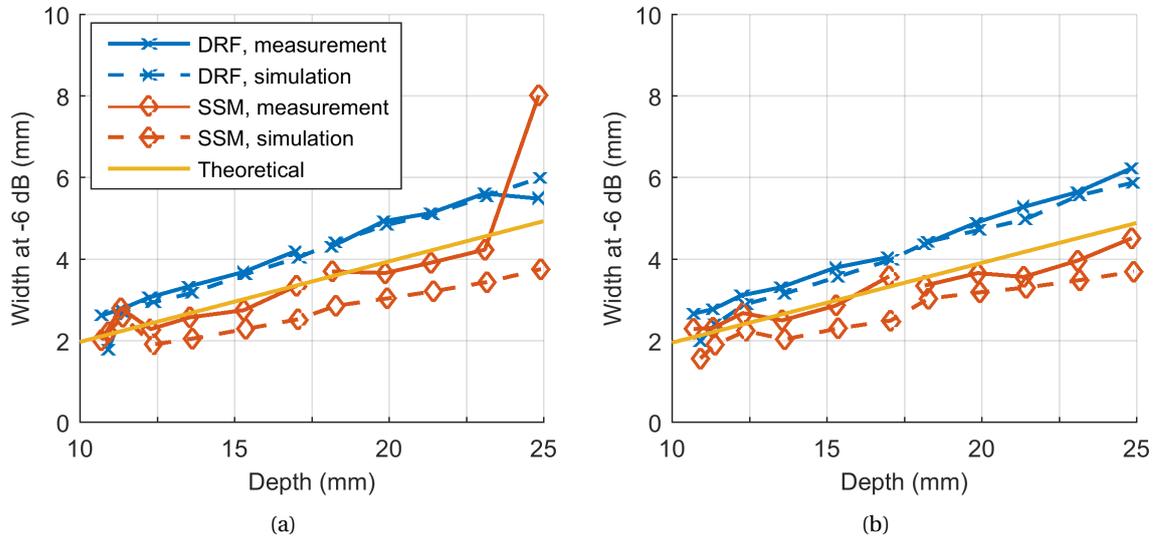


Figure III.3.10: In (a) the FWHM and in (b) the FWTM in θ and φ obtained with the different methods. Two measurements were done with the same phantom, but at a different distance from the transducer, hence the discontinuous lines.

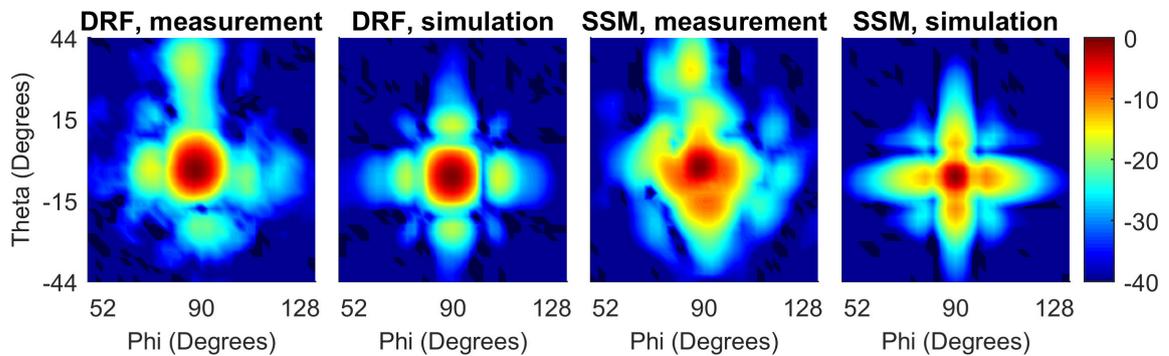


Figure III.3.11: A slice at constant radius of the deepest reflector at about 17 mm. This is the same point as in III.3.12, but imaged in another measurement and at a different depth.

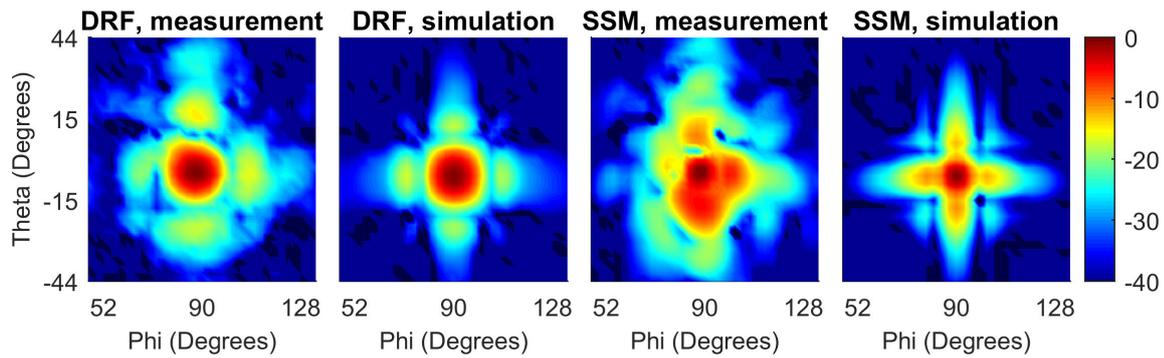


Figure III.3.12: A slice at constant radius of the deepest reflector at about 25 mm. This is the same point as in III.3.11, but imaged in another measurement and at a different depth.

SPHERICAL WAVE

In figure III.3.13 the FWHM in θ is shown for measurements and simulation of DRF applied on spherical wave data, SWSM-RF and SWSM-RII. In the simulations, the FWHM is similar for each method. We expected that DRF and SWSM-RF would show no difference, but SWSM-RII did produce a smaller FWHM as could be seen in figure III.3.7.

The measurements however show something slightly different. First of all, we see that all methods show a larger FWHM than in the simulation. This could be explained by the relatively large objects used in the measurements. The second thing is that the FWHM obtained with SWSM-RII is consistently the smallest, although the difference is very small.

In figure III.3.14 and III.3.15 we show a single slice of respectively the simulation and measurements. It is clear that the measurements seem to have more artefacts, which might be caused by reverberation of points above this slice. What can also be determined from the figures is that there are no differences visible between the DRF images and the images generated with SWSM-RF in both the simulation and measurements. This proves that it is not possible to obtain better resolution with receive focusing as the propagation step. SWSM-RII produces different results. The PSF visible in the simulation seems to be more confined, but this is not clearly visible in the measurements. However, we can see that the background level around the central spot appears to be lower, which might be partially explained by the more confined PSF as visible in the simulation.

So, the results prove that SWSM-RF performs equal to DRE, but the results do not show a clear advantage of using SWSM-RII.

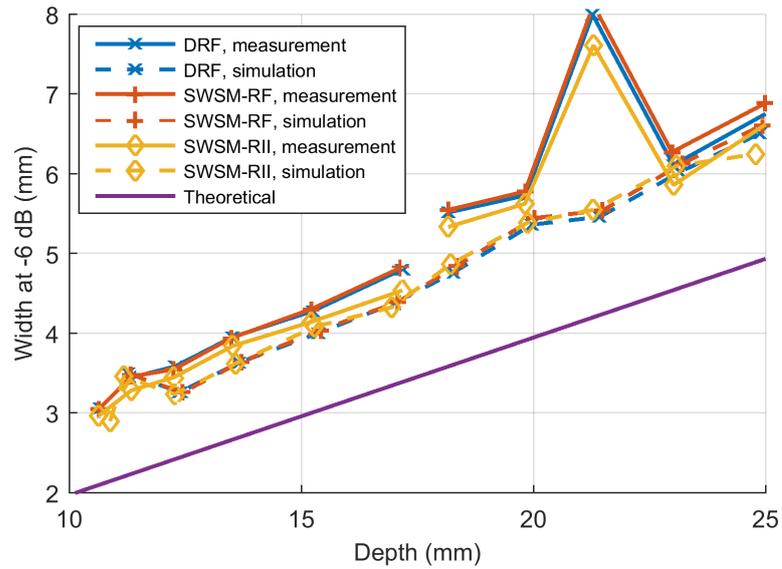


Figure III.3.13: The FWHM in θ obtained with the different methods. Two measurements were done with the same phantom, but at a different distance from the transducer, hence the discontinuous lines..

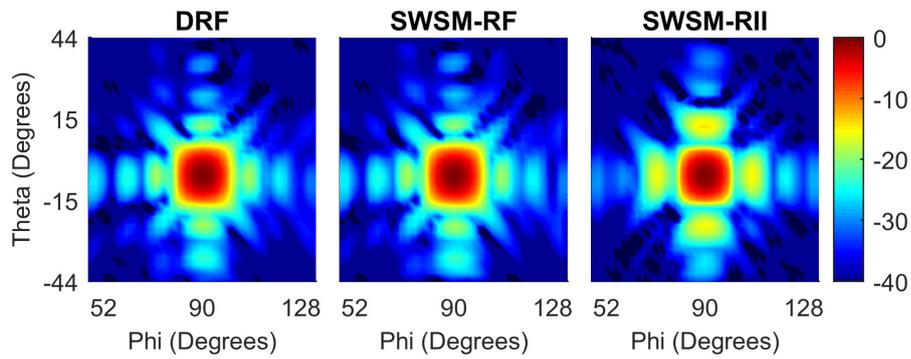


Figure III.3.14: A slice of simulations at constant radius of the deepest reflector at about 25 mm for each of the three methods.

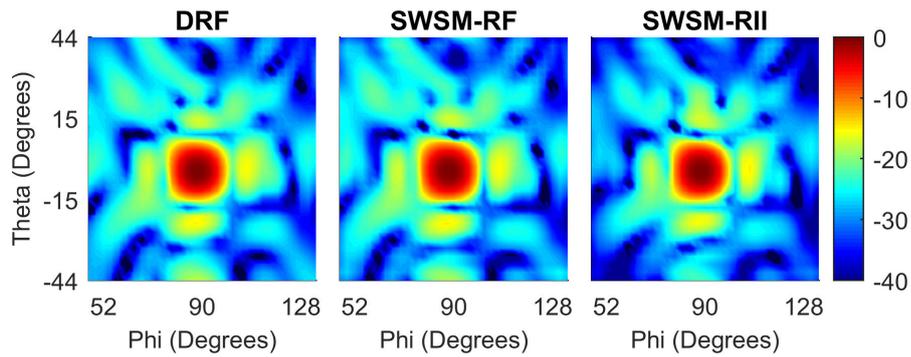


Figure III.3.15: A slice of measurement at constant radius of the deepest reflector at about 25 mm for each of the three methods.

III.3.6. PSASB VS STOLT

In this chapter we have compared DRF with SSM. This is not a completely fair comparison as DRF constructs each line from a single transmit. SSM on the other hand combines data from multiple lines to construct the final image. It would be more fair to compare SSM with SASB, but currently no 3D implementation of the latter exists. We would still like to point out some important differences between the two. Therefore, we will compare PSASB with Polar Stolt Migration in this section. We will look at the simulation results obtained with a phased array with similar dimensions as the MICA probe. No apodization is used in transmit and receive and furthermore, no apodization is used in PSASB while combining multiple low resolution lines to one high resolution line.

III.3.6.1. RESOLUTION

In figure III.3.16 the PSF width for DRE, PSASB and polar Stolt is shown for two f-numbers. Stolt Migration obtains the smallest lateral width for an f-number of 0.8. A smaller or larger f-number results in a wider PSF, as is for example illustrated for a f-number of 1.5. This behaviour is a bit different from what De Wit[6] observed when using the P4-1 probe, as in his case the performance clearly deteriorated for f-numbers below 1. This indicates that the behaviour cannot be fully caught with the f-number.

PSASB also performs best for the lowest f-number shown, but the PSF for both f-numbers is clearly broader than what can be obtained by Stolt. Furthermore, we see that similar results can be obtained with DRF. It is not completely unexpected that PSASB does not perform well with this aperture. The difference in time delays between each element to combine multiple low resolution lines are smaller than half a period. This means that effectively all delays can be considered as equal. Also, because of these small delays, similar results can also be obtained by using a single focus in receive for each scan line, which makes both DRF and SASB unnecessary in this case. To obtain better resolution it is thus necessary to use a frequency domain method like Stolt for this kind of array.

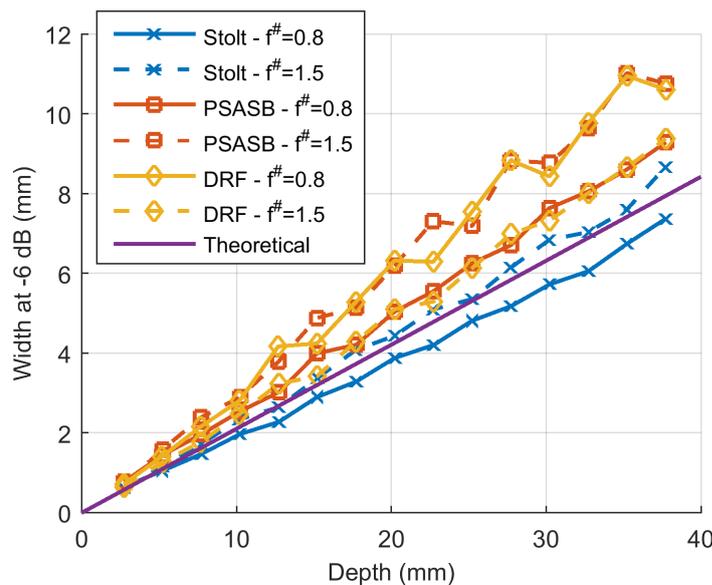


Figure III.3.16: FWHM of the PSF for different f-numbers

III.3.6.2. ARTEFACTS

While Stolt thus manages to provide a better resolution than time based methods, it does suffer from artefacts. This can be seen by comparing the image generated by SASB in figure III.3.17 and the one by Stolt visible in the leftmost panel of figure III.3.18. With SASB the points only become wider with depth, but there are no other changes visible. With Stolt the lateral resolution is better, but the artefacts become more severe deeper into the medium.

It turns out that the artefacts are in fact points that are imaged a second time. This happens because of the periodic nature of the FFT. Points that are already imaged will wrap-around and appear a second time. This time the points are unfocused as each frequency is affected differently and hence they appear as arcs. Because they are unfocused, they have a low amplitude. However, the signal from the points far from the transducer also has a low amplitude. Therefore, both the real points and artefacts can be seen. We note that we always apply a depth correction before display, and therefore points that appear further from the transducer have equal amplitude to points closer to the transducer.

There are multiple ways to get rid of the artefacts. Extending the time-signal by extensive zero-padding reduces the artefacts. Using a better approximation for the Hankel function also works, but in both cases the artefacts are still visible. Another way is to apply a depth correction before migration, instead of after the migration and just before display. The result is shown in the rightmost panel of figure III.3.18. The artefacts have almost completely disappeared. The ones that appear have an amplitude of -40 dB as compared to the surrounding points. Therefore, in a medium with speckle the artefacts will be hidden.

Depth correction is very common in ultrasound machines, so it should be possible. Still, SASB does not suffer from this and other frequency domain related issues and is thus more robust. However, in the end the best results can be obtained by Stolt migration.

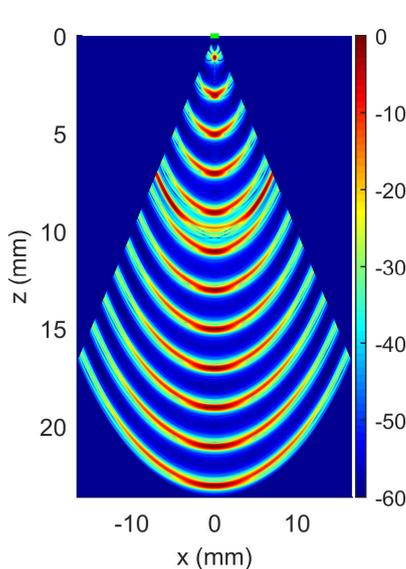


Figure III.3.17: Image created with PSASB

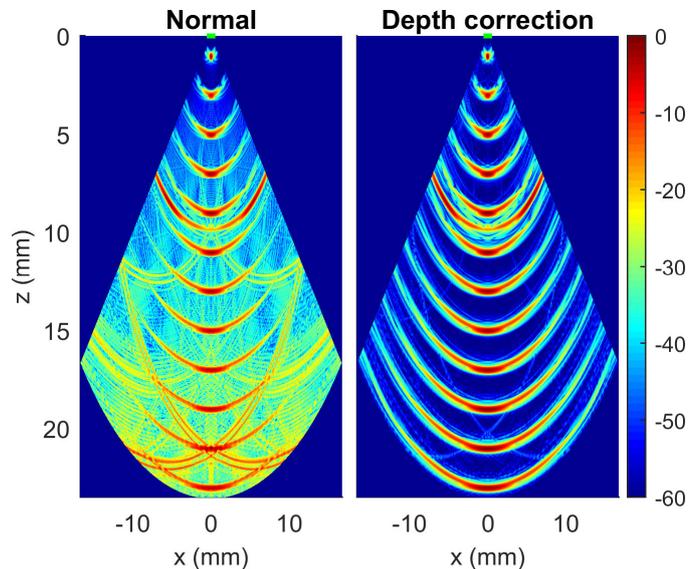


Figure III.3.18: Images created with Polar Stolt. For the leftmost image no depth correction is applied prior to the migration. This has been done however for the rightmost one.

III.4

DISCUSSION AND CONCLUSIONS

In this chapter we will first discuss a few open ends that were not addressed in this part and after that we will give our conclusion.

III.4.1. DISCUSSION

In this thesis we have developed SSM, without considering how the on-chip electronics would have to be to merge all element data into a single channel. This is the first thing that we will bring up in this section. Also, we think there is still a possibility to get even better results from the same data, which will be discussed the second subsection. In the next subsection, we will discuss our implementation of Spherical Wave Stolt. In the last subsection we will look at ways to improve the computation time.

We note that some discussion points are already brought up by De Witc[6] and are also applicable here. We will not repeat his discussion about the point source assumption and the resolution that exceeds the theoretical measure. If necessary, one could read them in his thesis. The computation time will however be discussed.

III.4.1.1. COMPLEXITY

It sounds great that only a single output from a scan line is required for SSM, instead of individual element data as is the case for DRF. However, the on-chip electronics will become very complex to achieve this as there are numerous different time delays that have to be applied and they differ per steering direction. This means that it might be impossible to implement Spherical Stolt as how it has been presented here. This is of course beside the fact that it would take too long to do a full sector scan.

A simple way of simplifying the receive electronics is to not apply time delays any more and just use the natural focus of a set of elements. This is a very coarse way of focusing and it will most deteriorate the resolution as compared to time-delay focusing, but it might be better than DRF. Still, the scan time will be very long.

Spherical Wave Stolt is very fast. It requires only a single insonification to form a full image. The results showed however that the resolution suffers in comparison with the sector scan. This might be solved by angular compounding, which might improve the resolution at the cost of frame rate. The major problem with this method is that it requires all channel data. This is not directly possible, but

it might be possible to read out a subset of the elements per insonification and thus reduce the frame rate.

Spherical and Spherical Wave Stolt as presented in this part can be seen as extremes. There is also a sort of middle way, which is presented by Bera et al.[42] This 3D beamforming method consists of three stages. In the first stage groups of receive elements are pre-steered to an angle and a diverging transmit beam is steered at a direction close to that. The pre-steering effectively reduces the channel count by the number of elements in the receive groups. In the second stage of this method scan lines are formed by delay-and-sum beamforming on the measured data. Finally, selections or combinations are made of scan lines from different transmit and receive events to form the final volume. In the paper two methods are described, which have the above two steps as basis, and can reconstruct a full volume in 25 or 169 transmit and receive events. With the 169 events the results were almost equal to a full sector scan, which in the case of Bera required 65x65 transmit and receive events and is thus much slower. We have shown that it is possible to get a better resolution in a sector scan with SSM as compared to DRF and therefore Stolt migration might also improve the methods presented in this paragraph, but this has to be investigated.

III.4.1.2. FULL MATRIX CAPTURE

The ERM is a simple way to obtain an extrapolation scheme with which we can correctly describe the phase and thus image all scatters at the correct point. Obtaining this scheme via a different procedure is more complex. A simple guess could have been to assume that the reflected waves can be extrapolated via the Hankel function of the first kind, which represent waves propagating towards the origin, and that the incident wave is described by the Hankel functions of the second kind. For an array of elements on a sphere, just like we created for SSM, the imaging condition would look like

$$b(r, l, m, t' = 0) = \int_{-\infty}^{\infty} \tilde{p}(r_f, l, m, f) \frac{h_l^{(1)}(kr)}{h_l^{(1)}(kr_f)} \left(\frac{h_l^{(2)}(kr)}{h_l^{(2)}(kr_f)} \right)^* df. \quad (\text{III.4.1})$$

This is very different from the ERM that contains just a single ratio of Hankel functions and the ERM wavenumber \hat{k} . Clearly, these two ratios together do not describe a pulse-echo experiment. Instead, also the cross paths are included, which means that the transmitted signal is not just measured by the same element, but also all others. This is called a full matrix capture and with this propagator we thus have a full matrix imaging algorithm.

Multiple full matrix imaging algorithms already exist, like the one by Hunter et al[43] for flat arrays. The geometry is different from our case, but otherwise the situation is very similar. We do a single transmission for each line or in a different perspective excite the elements on the virtual array independently. However, we do not measure the reflected signal with the other virtual elements. This is possible, assuming we have access to all channel data. Instead of just focusing in receive at the transmitting element, we can also receive focus on the other virtual elements. After doing all transmissions, we just sum up all the different contributions. With this scheme, we do have the situation which can be imaged by using the imaging condition described above.

This scheme has some similarities the method of Bera et al, presented in the previous subsection. Proving that one of the methods works will make it more likely that the other works as well. Using this method, we can then expect a better resolution than can be achieved with just the pulse-echo data.

III.4.1.3. SPHERICAL WAVE STOLT

We saw that the resolution obtained with Spherical Wave Stolt is 20% worse than with Spherical Stolt. Part of it is due to the fact that we just have a single insonification, instead of 128x128 scan

lines. It might be that spatial compounding imaging might help, but this of course reduces frame rate. We think however, that it is possible to increase the resolution without more insonifications. For all results with Spherical Wave Stolt we have used the cross-correlation imaging condition of equation 2.12. We know however that a better resolution can be obtained with a deconvolution imaging condition. So by using a deconvolution we are certain that the resolution can be improved.

There are some downsides however. For the deconvolution to work well, the transducer bandwidth has to be known. A mismatch between the expected source pulse and the real pulse will result in the image being convolved with the function that converts the true pulse to the assumed pulse.[10]. This indicates that if the assumed pulse is close enough to the real pulse, the resolution might still be better than with a cross-correlation imaging condition.

It might also be possible to combine a deconvolution with the ERM used in Spherical Stolt, but we have not found any examples of this in literature. If it could be combined, this would mean that beside the transducer bandwidth, it also needs to be known how the source differs as a function of the angular frequencies.

III.4.1.4. COMPUTATION TIME

Currently, the whole migration takes 80 seconds to complete with 8 CPU cores(Intel Xeon(R) CPU E5430, Intel, Santa Clara, California, USA) and the symmetry described in section III.1.4. This is incompatible with real time imaging, which is common in medical ultrasound, and therefore the algorithm has to become a lot faster.

Part of the improvement in computation time can be achieved by tweaking the settings. For example, as De Wit already mentioned, the radial grid can be made coarser. For the images in this part we have always used a spacing larger than $2\lambda_{\min}$, but this could be reduced $0.5\lambda_{\min}$ without any problems. For grids that are less dense aliasing may occur. Also, if the region of interest is known, the migration only has to be done in this depth range. Another thing is that the angle between subsequent scan lines might be chosen smaller. This might introduce aliasing if it does not satisfy equation 2.24, but for the higher frequencies this may be allowed. The final setting to be tweaked is the amount of frequencies propagated. We propagated all frequencies that were in the -40 dB bandwidth of the transducer, but it might be that less frequencies can be propagated without any detrimental effects.

The algorithm itself could also be improved. We already used a symmetry to cut down calculation time by almost half, but there might be other symmetries that can be exploited, but that are less obvious. Changing the algorithm used for certain steps might also give improvements. For example, the transforms take almost half of the computation time. The used algorithm is certainly not one of the fastest and it uses a massive amount of memory for larger transformations. A comparison of multiple algorithms can be found in the paper by Reinecke and Seljebotn[44, Table 1]. A faster SHT would mean that the bottleneck becomes the propagation part. The FHT like algorithm did not work as expected. Maybe tweaking the parameters could help or a different kind of scheme could be found.

Of course, for the final implementation dedicated hardware can be used. This will certainly speed up the whole algorithm, but the possible improvements listed above will certainly also help.

III.4.2. CONCLUSION

In this part we have developed and investigated the results of both Spherical Stolt Migration (SSM) and Spherical Wave Stolt Migration (SWSM). By using SSM we could achieve a better resolution than DRF. In the simulated images made with the first setup, the observed difference in FWHM was up to 25%, while there were no additional artefacts. The FWHM did approach the far field PSF FWHM, which indicates that the image is focused over the whole depth range.

To verify our simulated results, we first compared the result obtained with SSM and Polar Stolt Migration. The FWHM and FWTM were consistently slightly lower in the image obtained with Polar Stolt, but the difference in FWHM was less than 1%. The artefacts in both images were also slightly different. Overall, the differences were very small and therefore we concluded that the results obtained by De Wit for Polar Stolt are also relevant for SSM. The second way to verify the results was to compare measurements and simulations done with the second setup. As the measured objects could not be considered as point scatters, there were some clear differences between measurements and simulations. However, both measurements and simulations showed a lower FWHM for SSM as compared to DRF. Therefore, this seems to prove that SSM works.

SWSM gave worse results as compared to DRF applied on line scan data and SSM. The FWHM with SWSM was about 25% larger than SSM in the whole image. This is not strange, as SWSM constructs a whole image from a single transmit and receive event, while DRF and Stolt did use 128x128 events. DRF could also be applied to the spherical wave data. In this case we could show that better results than DRF could only be obtained if SWSM is combined with the Rayleigh II integral (SWSM-RII). The FWHM was about 8% smaller at all imaged points. With receive focusing (SWSM-RF) as intermediate step the results were equal to DRF applied on spherical wave data.

The measurements achieved with the second setup verified that SWSM-RF can only produce results equal to DRF applied to spherical wave data. SWSM-RII managed to obtain a smaller FWHM in the measurements, but the difference was very small. In the simulations there was no difference in FWHM. So, we could not get a clear answer.

For real time imaging, computation time is very important. By using a symmetry in the data we could reduce the computation time by 37% as compared when no symmetry was used. We also tried a different algorithm based on the Fast Hankel Transform, but this resulted in many artefacts and longer computation times. Still, real time imaging is not yet possible.

Comparing Stolt with DRF was not deemed completely fair, as the modus operandi are quite different. Therefore, we went back to 2D and compared Polar Stolt and PSASB in similar settings as the second setup. We found out that PSASB could not improve results if a very small probe is used as compared to using no focus in receive, which produces similar results as DRF for these configurations. Stolt managed to achieve a lower resolution than this, but also again better as compared to the theoretical measure. However, Stolt is less robust as compared to PSASB as there are also various frequency domain related problems, possibly resulting in extra artefacts. Luckily, in many cases it is possible to work around these problems, but PSASB works almost directly out of the box.

We have shown that SSM and SWSM is able to get better results than the current golden standard used for imaging in medical ultrasound: Dynamic Receive Focusing. Theoretically our method requires a single output channel and many transmit and receive events, in case of SSM, or all channel data and a single shot, in case of SWSM. Practical implementation is more complex, as discussed at the start of this chapter. There are other schemes that work around these problems, but they do decrease resolution. If we could show that Stolt works together with these implementation schemes, we will increase the resolution, while reducing channel count and complexity. This is one step forward towards the use of matrix arrays in a medical setting.

ACKNOWLEDGEMENTS

First of all, I would like to thank Martin Verweij for being my supervisor during this project. The two research subjects, Fresnel Zone Beamforming and Stolt Migration, he provided me with were both very interesting subjects that have easily kept me engaged during the almost 9 months of research. Furthermore, the discussions that we had gave me more insight into both subjects and the feedback on my report really helped make it more clear for people who are not me.

I would also like to thank Rik Vos for the various discussions about the problems I ran into, the valuable feedback on my report and for sharing the scripts of the Linear Stolt algorithm.

I am grateful to Deep Bera for helping me with the experiments for Stolt Migration, the various discussions about beamforming and for providing his PSASB code that I used to prove that Stolt is clearly better.

My thanks also go to Jos de Wit who did some further preparation work on Fresnel Zone Beamforming. Also, his research on Polar Stolt Migration kick-started my research on both Cylindrical and Spherical Stolt Migration.

Finally, I am also grateful to all the people that participated in playing sjoelen. Without them, I would never have got such a high score. At the time of writing, my score of 123 points is still the highest score achieved to date, but I have to admit that I was not the first to achieve that score.

BIBLIOGRAPHY

- [1] CBS, *Zorguitgaven stijgen langzamer*, (2016), last visited on 14/1/2017. Available on <https://www.cbs.nl/nl-nl/nieuws/2016/20/zorguitgaven-stijgen-langzamer>.
- [2] B. W. Orenstein, *Ultrasound First — AIUM Initiative Points Out the Diagnostic Benefits and Cost Savings of Using Ultrasound as the Primary Imaging Modality in Many Cases*, (2013), last visited on 14/1/2017. Available on <http://www.radiologytoday.net/archive/rt0213p14.shtml>.
- [3] J. Kortbek, J. A. Jensen, and K. L. Gammelmark, *Sequential beamforming for synthetic aperture imaging*, *Ultrasonics* **53**, 1 (2013).
- [4] H. J. Vos, P. L. van Neer, M. M. Mota, M. D. Verweij, A. F. van der Steen, and A. W. Volker, *F-k domain imaging for synthetic aperture sequential beamforming*, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **63**, 60 (2016).
- [5] D. Bera, J. G. Bosch, N. de Jong, and H. J. Vos, *Synthetic aperture sequential beamforming for phased array imaging*, in *Ultrasonics Symposium (IUS), 2015 IEEE International* (IEEE, 2015) pp. 1–4.
- [6] J. de Wit, *Sequential Beamforming for Phased Array Imaging using Stolt Migration*, Master's thesis, TU Delft, the Netherlands (2016).
- [7] H. J. Vos, P. L. van Neer, M. D. Verweij, N. de Jong, and A. W. Volker, *3D post-processing of pre-beamformed RF data in the frequency-wavenumber domain*, in *Ultrasonics Symposium (IUS), 2015 IEEE International* (IEEE, 2015) pp. 1–4.
- [8] D. Gisolf and E. Verschuur, *The principles of quantitative acoustical imaging* (EAGE Publ., 2010).
- [9] T. Szabo, *Diagnostic Ultrasound Imaging: Inside Out*, Academic Press series in biomedical engineering (Elsevier Academic Press, 2004).
- [10] T. L. Poole, A. Curtis, J. O. Robertsson, and D.-J. van Manen, *Deconvolution imaging conditions and cross-talk suppression*, *Geophysics* **75**, W1 (2010).
- [11] R. Stolt, *Migration by fourier transform*, *Geophysics* **43**, 23 (1978).
- [12] A. Volker, *Plane wave imaging using phased array*, (2014).
- [13] M. Abramowitz and I. A. Stegun, *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*, Vol. 55 (Courier Corporation, 1964).
- [14] D. Garcia, L. Le Tarnec, S. Muth, E. Montagnon, J. Porée, and G. Cloutier, *Stolt's fk migration for plane wave ultrasound imaging*, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **60**, 1853 (2013).
- [15] M. Van Wijk and J. Thijssen, *Performance testing of medical ultrasound equipment: fundamental vs. harmonic mode*, *Ultrasonics* **40**, 585 (2002).

- [16] B. D. Steinberg, *Principles of aperture and array system design: Including random and adaptive arrays*, New York, Wiley-Interscience, 1976. 374 p. (1976).
- [17] M. F. Rasmussen and J. A. Jensen, *Comparison of 3-d synthetic aperture phased-array ultrasound imaging and parallel beamforming*, IEEE transactions on ultrasonics, ferroelectrics, and frequency control **61**, 1638 (2014).
- [18] M. Cree and P. Bones, *Algorithms to numerically evaluate the hankel transform*, Computers & Mathematics with Applications **26**, 1 (1993).
- [19] A. Siegman, *Quasi fast hankel transform*, Optics Letters **1**, 13 (1977).
- [20] Q. H. Liu and Z. Q. Zhang, *Nonuniform fast hankel transform (nufht) algorithm*, Applied optics **38**, 6705 (1999).
- [21] J. A. Fessler and B. P. Sutton, *Nonuniform fast fourier transforms using min-max interpolation*, IEEE Transactions on Signal Processing **51**, 560 (2003).
- [22] N. Schaeffer, *Efficient spherical harmonic transforms aimed at pseudospectral numerical simulations*, Geochemistry, Geophysics, Geosystems **14**, 751 (2013).
- [23] F. G. Tricomi, *Sugli zeri dei polinomi sferici ed ultrasferici*, Annali di Matematica Pura ed Applicata **31**, 93 (1950).
- [24] E. W. Weisstein, *Halley's method. From MathWorld—A Wolfram Web Resource*, Last visited on 15/9/2016. Available on <http://mathworld.wolfram.com/HalleysMethod.html>.
- [25] K. K. Sharma, *Optics: principles and applications* (Academic Press, 2006).
- [26] J. C. Wiltse, *The fresnel zone-plate lens*, in *1985 Technical Symposium East* (International Society for Optics and Photonics, 1985) pp. 41–47.
- [27] Lawrence Berkeley National Laboratory Center For X-Ray Optics, *Fresnel zone plate theory*, (2014), last visited on 15/1/2017. Available on <http://zoneplate.lbl.gov/theory>.
- [28] D. Peterson and G. S. Kino, *Real-time digital image reconstruction: A description of imaging hardware and an analysis of quantization errors*, IEEE transactions on sonics and ultrasonics **31**, 337 (1984).
- [29] M. M. Nguyen, J. Mung, and J. T. Yen, *Fresnel-based beamforming for low-cost portable ultrasound*, IEEE transactions on ultrasonics, ferroelectrics, and frequency control **58**, 112 (2011).
- [30] P. A. Magnin, O. Von Ramm, and F. Thurstone, *Delay quantization error in phased array images*, IEEE Transactions on Sonics Ultrasonics **28**, 305 (1981).
- [31] R. R. Entekin, J. R. Jago, and S. C. Kofoed, *Real-time spatial compound imaging: technical performance in vascular applications*, in *Acoustical Imaging* (Springer, 2002) pp. 331–342.
- [32] J. A. Jensen, *Field: A program for simulating ultrasound systems*, in *10th Nordic-Baltic Conference on Biomedical Imaging, Vol. 4, Supplement 1, Part 1: 351–353* (Citeseer, 1996).
- [33] J. A. Jensen and N. B. Svendsen, *Calculation of pressure fields from arbitrarily shaped, apodized, and excited ultrasound transducers*, IEEE transactions on ultrasonics, ferroelectrics, and frequency control **39**, 262 (1992).
- [34] J. W. Goodman, *Introduction to Fourier optics* (Roberts and Company Publishers, 2005).

- [35] F. Fool, *Transmit focusing in a plane - Fresnel zone focusing of continuous wave ultrasound*, Bachelor's Thesis, TU Delft, the Netherlands (2015).
- [36] I. Beekers, *Modeling the amount of backscattered pressure from the collection of red blood cells in a vessel*, Master's thesis, TU Delft, the Netherlands (2015).
- [37] L. Bohs, B. Geiman, M. Anderson, S. Gebhart, and G. Trahey, *Speckle tracking for multi-dimensional flow estimation*, *Ultrasonics* **38**, 369 (2000).
- [38] L. van der Graaf, *Frequency- and angle dependency of the scattering of ultrasound as a contrast mechanism for imaging blood in the neonatal brain*, Master's thesis, TU Delft, the Netherlands (2015).
- [39] M. H. Skjelvareid, Y. Birkelund, and Y. Larsen, *Synthetic aperture focusing of outwardly directed cylindrical ultrasound scans*, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **59**, 2460 (2012).
- [40] R. Haberman, *Applied partial differential equations: with Fourier series and boundary value problems*, Vol. 4 (Pearson Prentice Hall Upper Saddle River, 2004).
- [41] A. Politis, *polarch/spherical-harmonic-transform*, (2015), last visited on 31/1/2017. Available on <https://github.com/polarch/Spherical-Harmonic-Transform>.
- [42] D. Bera, H. Vos, S. Raghunathan, C. Chen, Z. Chen, M. Verweij, M. Pertijs, N. de Jong, and J. Bosch, *Three-dimensional beamforming combining micro-beamformed rf datasets*, in *Ultrasonics Symposium (IUS), 2016 IEEE International* (IEEE, 2016) pp. 1–4.
- [43] A. J. Hunter, B. W. Drinkwater, and P. D. Wilcox, *The wavenumber algorithm for full-matrix imaging using an ultrasonic array*, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **55** (2008).
- [44] M. Reinecke and D. S. Seljebotn, *Libsharp–spherical harmonic transforms revisited*, *Astronomy & Astrophysics* **554**, A112 (2013).
- [45] M. A. Haun, D. L. Jones, and W. D. O'Brien, *Efficient three-dimensional imaging from a small cylindrical aperture*, *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **49**, 861 (2002).
- [46] J. Gardner and R. Collin, *An accurate closed-form approximate representation for the hankel function of the second kind*, *IEEE Transactions on Antennas and Propagation* **48**, 1699 (2000).



HANKEL FUNCTION APPROXIMATION

Exactly calculating the Hankel function using an integral representation or a series can be time consuming. Furthermore, pre-calculating the Hankel functions used may not be feasible due to memory limitations, as was shown in section II.1.2 This is why the Hankel function is frequently approximated. In this thesis we have used the same approximation for the Hankel function as De Wit used. This is the asymptotic expansion for large arguments found in Abramowitz and Stegun.[13] While we can of course expect this approximation to work best for large arguments, it is unknown how this approximation handles large orders. In this appendix we will look at this and also investigate some other approximations and compare them with the solution as obtained by Matlab's built in functions, which we assume to produce exact results. For the imaging algorithms we always require a ratio of two Hankel functions and therefore we will focus on this here.

A.1. APPROXIMATIONS

In this section we will look at three different approximations. First the large argument approximation that de Wit also used. Next we look at the approximation by Haun, which is derived from the Rayleigh-Sommerfeld diffraction formula. Finally, we look at the approximation by Gardner and Collin.

A.1.1. LARGE ARGUMENT APPROXIMATION

This approximation for large arguments can be found in Abramowitz and Stegun[13, eqs.9.2.28-29]. The first few terms in the expansion of the square modulus and the phase are:

$$|H_m^{(1)}(\rho)|^2 \approx \frac{2}{\pi\rho} + \frac{4m^2 - 1}{4\pi\rho^3} + \dots, \quad (\text{A.1})$$

$$\arg(H_m^{(1)}(\rho)) \approx \rho - \left(\frac{m}{2} + \frac{1}{4}\right)\pi + \frac{4m^2 - 1}{8\pi\rho} + \dots \quad (\text{A.2})$$

De Wit made all his images with just the first three terms of the phase. This third term was very important. Not including this resulted in an unfocused image and an extra term did have negligible effects. For the modulus, only the first term was taken as it is of less importance than the phase. This results in the following approximation of the ratio of two Hankel functions:

$$\frac{H_m^{(1)}(kr)}{H_m^{(1)}(kr_f)} = \sqrt{\frac{r_f}{r}} e^{ik(r-r_f)} \left(1 - \frac{4m^2 - 1}{8k^2 r r_f}\right). \quad (\text{A.3})$$

A.1.2. HAUN APPROXIMATION

The approximation by Haun et al[45] does not approximate the Hankel function itself. Instead, it approximates the ratio of two Hankel functions, which is used as the transfer function in their paper. This approximation reads

$$\frac{H_m^{(1)}(kr)}{H_m^{(1)}(kr_f)} = \sqrt{\frac{r}{r_f}} e^{i(r-r_f)\sqrt{k^2 - \frac{m^2}{r_f r}}}. \quad (\text{A.4})$$

Compared to the large argument approximation there are a few clear differences. First of all, all variables appear under the square root. This is very inconvenient, as calculating a square root is computationally expensive. The large argument approximation does not have all arguments under the square root, so it is possible to pre-calculate and reuse parts while looping over the other variables. Furthermore, the modulus is exactly opposite. Instead of yielding an approximation to the Hankel function, this modulus compensates for the spreading of energy as a function of depth. Finally, this approximation will break down for wide beams as it approximates a cosine with the first two terms in its Taylor series.

A.1.3. GARDNER APPROXIMATION

The last approximation that we will introduce is the closed-form approximate representation for the Hankel function of the second kind by Gardner and Collin[46] that reads

$$H_m^{(2)}(\gamma) = \sqrt{\frac{2}{\pi\sqrt{\gamma^2 - m^2}}} e^{-i\left(\sqrt{\gamma^2 - m^2} - \frac{\pi}{2}\left(\frac{1}{2} + m\right) + m \arctan\left[m/\sqrt{\gamma^2 - m^2}\right]\right)}. \quad (\text{A.5})$$

Using the fact the Hankel function of the first kind is just the complex conjugate of the second kind for real arguments, we can write the ratio as

$$\frac{H_m^{(1)}(kr)}{H_m^{(1)}(kr_f)} = \left(\frac{(kr_f)^2 - m^2}{(kr)^2 - m^2}\right)^{1/4} e^{i\left(\sqrt{(kr)^2 - m^2} - \sqrt{(kr_f)^2 - m^2} + m\left(\arctan\left[m/\sqrt{(kr)^2 - m^2}\right] - \arctan\left[m/\sqrt{(kr_f)^2 - m^2}\right]\right)\right)} \quad (\text{A.6})$$

This is clearly the most complex approximation of the three. Again all variables are found under the complex root, but also an arctangent is included. These two things combined make it not very computationally efficient. However, from the paper it becomes clear that this approximation is quite accurate, but things will go wrong if the order approaches the argument. Effectively, this is thus also a large argument approximation.

A.2. COMPARISON

We have compared the above three approximations with Matlab's built in function in a setting similar to the one used for Cylindrical Stolt Migration in part II. This thus also includes the use of the ERM velocity ($\hat{c} = c/2$). For each order we evaluate the ratio of two Hankel functions for 10^5 radii between 3 and 5 mm. Other settings are noted in table A.1. To check how the large argument approximation improves with more terms, we have also included a version with the first 6 terms in the phase and the first 4 in the modulus.

Table A.1: Settings used for comparing the different approximations

Parameter	Value
r_f	13 mm
r	3 to 50 mm
m	0 to 256
Evaluation points per order	10^5
Frequency	5 MHz
Speed of sound	1540 m s^{-1}

A.2.1. MODULUS

The absolute difference between the modulus of Matlab's built in function and the 4 approximations is shown in figure A.1. For the large argument approximations, the minimum error always stays relatively low. It only slowly increases for higher orders. The mean and maximum error slowly increase until about order 125 and then jump up very fast. It is clear that with more terms the approximation is better, but both still show this behaviour. We will investigate in the where exactly the large errors come from in the next subsection.

The approximation by Haun does not approximate the modulus very well, but this is expected. As we could see in equation A.4, the modulus in front was already flipped. It thus clearly does not approximate the ratio itself, but compensates for the spreading of energy as a function of depth.

The approximation by Gardner seems to work best of the four approximations shown. The minimum error is lowest of all and the mean and maximum error slowly increase towards order 123. At this point things the behaviour changes. If we look back at equation A.5 we see that it includes a square root with both the argument and order in it. At order 123, the order becomes larger than the smallest argument, and the approximation of the modulus fails. Interestingly, the other error of other approximations also start to rise very fast around this point.

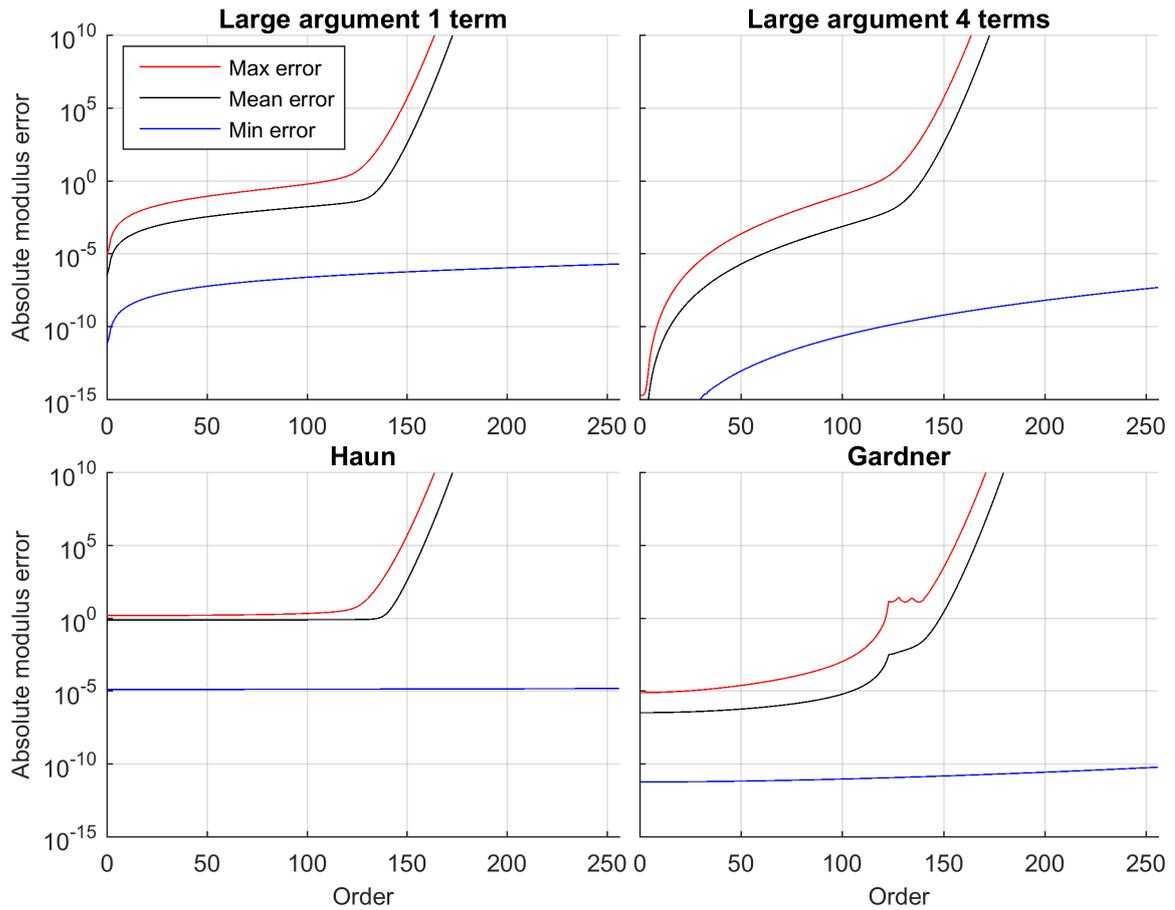


Figure A.1: Absolute modulus error between the approximations and the value obtained with Matlab's built in function for a range of orders.

A.2.2. PHASE

The absolute phase error with respect to the exact value obtained with the above approximations can be seen in figure A.2. We see that the large argument approximation, visible in the upper row, works very well for low orders, but the error increases fast. At order 100 the maximum error with 3 terms has gone to the maximum value of π and the mean error is there already about 0.1 rad. This is also the point at which the approximation failed to approximate the modulus, as could be seen in the previous subsection. It seems therefore that the approximation only works well as long as the argument is much larger than the order. This is not the case any more for the higher orders, as the argument for values close to the origin becomes relatively small.

Including more terms in the approximation does improve the accuracy. The order at which the maximum error is reached does not shift very much, but the mean error is reasonable for most of the range. This will of course come at the cost of computation time, as we will see later.

The approximation by Haun is very bad. It reaches a mean error of almost 1 rad at order 50. After that, it starts oscillating a bit, but the mean error is always above 1. The minimum error also varies a lot, because there are a few well approximated values, but this is most likely just by chance as we have observed that the phase evolves completely different with depth as compared to the exact values.

The closed form approximation by Gardner has a good overall performance. For the first few orders, the large argument approximation is better, but the approximation by Gardner is very constant. Even after order 123, the point at which the square root becomes imaginary for the smallest arguments,

the maximum error stays beneath 0.6 rad and the mean error beneath 2×10^{-3} rad. Clearly, the approximation by Gardner is the best approximation for the phase of all discussed approximations.

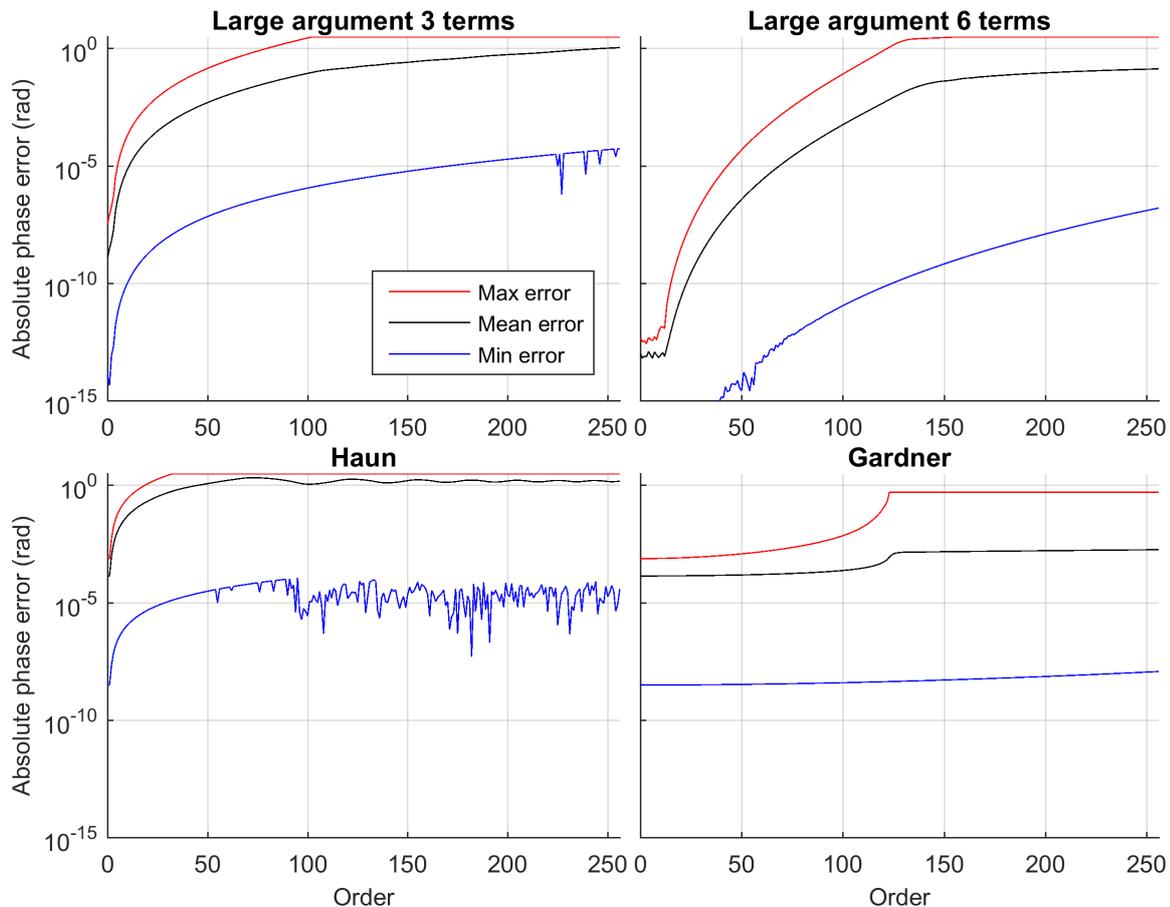


Figure A.2: Absolute phase error between the approximations and the value obtained with Matlab's built in function for a range of orders.

To see where the largest errors are obtained with the approximations, we look at the phase error as a function of depth for order 100. This is visible in figure A.3 for all approximations except for the approximation by Haun, which gave results far away from the exact values. We can see that all approximations are perfect at a depth of r_f , which is as expected as two equal Hankel functions are divided at that depth. The error increases rapidly to both sides of this depth. For depths shallower than r_f we see that the phase error increases rapidly for all approximations. This is not very strange as the arguments for these depths are relatively small, invalidating the approximations. For depths larger than this all approximations go to an asymptotic value. However, we should expect that large argument approximations start working better for larger arguments.

The reason is simple. The Hankel function in the denominator is the source of error. For a depth of r_f the large argument approximations give a relatively larger error, which affects the whole depth range. If we replace this denominator by the exact value we get the results visible in figure A.3(b). Especially the 3 and 6 term approximations benefit. This figure also clearly shows that all approximations work better for large arguments, as the error becomes less the larger the argument.

The effect of using the exact denominator on figure A.2 would be that both the mean and minimum error will be lower. However, they will still increase slowly for increasing order. The maximum error will not change, as for small arguments the approximations are still bad. The same effects would be observed on the modulus in the previous section.

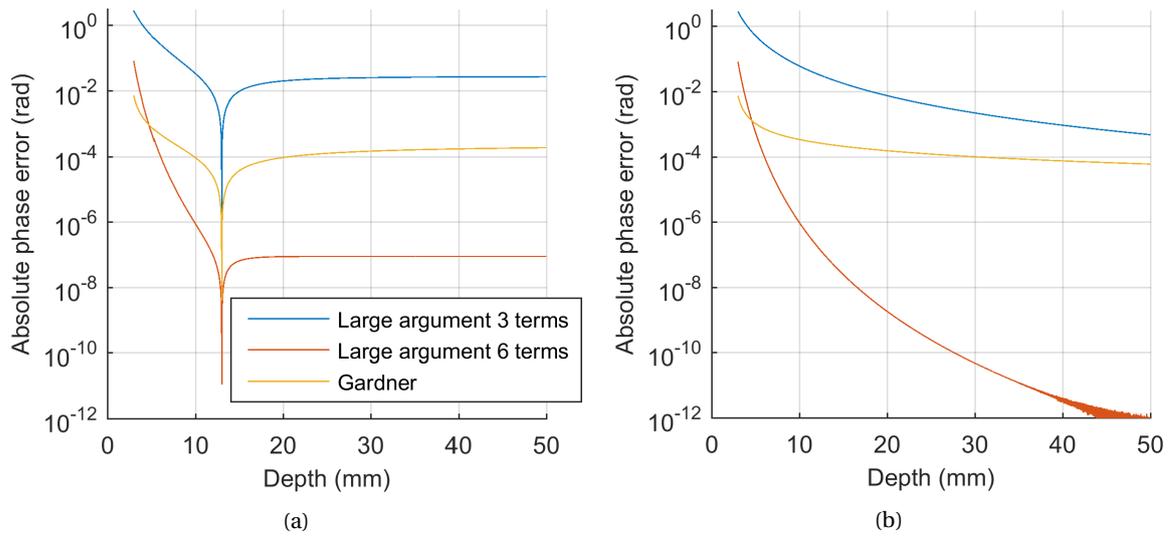


Figure A.3: In (a) the absolute phase error between the approximations and the value obtained with Matlab's built in function as function of depth for order 100. In (b) the same thing is done, except that the Hankel function has been replaced by the exact value.

A.2.3. COMPUTATION TIME

In table A.2 the relative computation time as compared to Matlab's built-in function for calculating the ratio of two Hankel functions for all values of m and r . It is clear that all approximations are way faster. The fastest approximation is the large argument approximation with 1 term for the modulus and 3 terms for the phase, which is shown in equation A.3. Increasing the number of terms in both the modulus and phase does increase accuracy as we have seen in the sections above, but also clearly affects the computation time. We would like to note that only the 1-3 version has been optimized, so there is some improvement possible.

The approximations by Haun and Gardner are both slower than the fastest large argument approximation. The difference between Haun and the large argument approximation is quite small, but Gardner takes double the amount of time. There might be some improvement possible, but we do not expect both of them to be faster than the large argument approximation. Especially in case of the migration, as all arguments appear under the square root and in case of the approximation by Gardner also in the arctangent. Calculating the square root and arctangent is time-consuming and has to be done for every variable. This is not required in the large argument approximation, which is therefore the fastest.

Table A.2: Computation time relative to the computation time required for calculating the exact value.

	Relative computation time
Large argument 1-3 terms	0,045
Large argument 2-4 terms	0,144
Large argument 3-5 terms	0,192
Large argument 4-6 terms	0,285
Haun	0,051
Gardner	0,090
Matlab	1

A.3. EFFECT ON RESULTS

In the previous section we saw that the error for higher orders becomes very large with the approximation that we used during this thesis. For the array of virtual sources and receivers at $r_f = 13$ mm orders larger than 100 had very large errors. In part III, we used the MICA Probe and the virtual array was placed at 1 mm. This will result into even larger errors and even less usable orders. We determined that the maximum possible error for the phase was already reached just above order 20. However, this will not affect the final image significantly. The reason can be found in figure A.4, where we have summed the absolute value of frequencies for each individual order of measurement and simulation data obtained with the MICA probe. Almost all content can be found for $l < 30$, especially in the case of the simulation where there is no noise. This means that although the higher orders are not propagated correctly, the influence on the final image is limited.

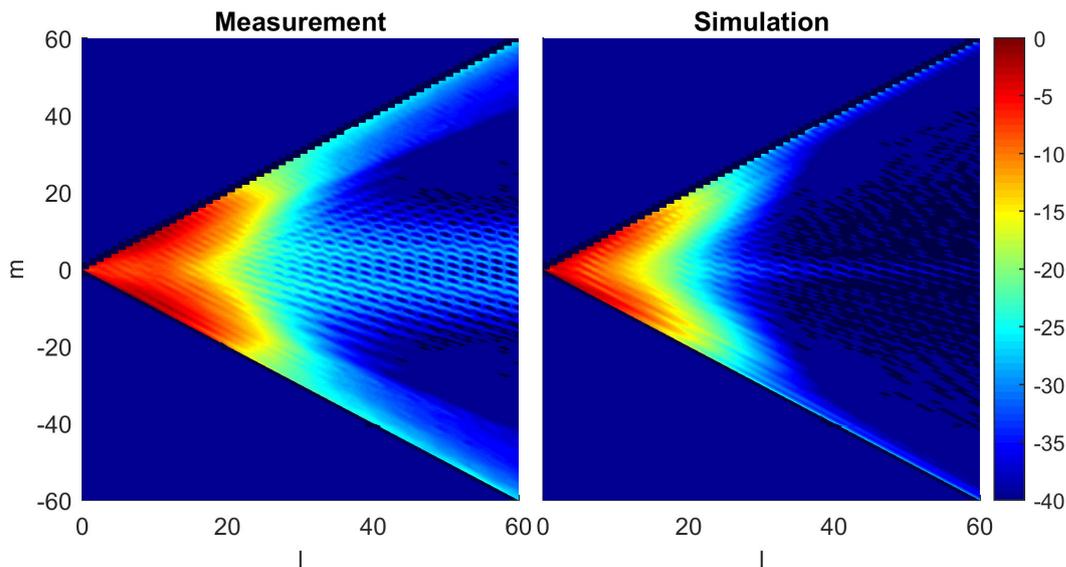


Figure A.4: The total frequency amplitude for each of the orders plotted on a log scale. This has been obtained by summing the absolute values of the frequencies for each individual order. The black lines are on the outside of where all non-zero values lie.

For all the results shown in part II and III we have checked if a better approximation makes any difference. We found out that the difference was negligible for the results we obtained. Partly this is because there is almost no content in the higher orders that are not correctly propagated as discussed above, but also due to some choices. For example, we used an f-number of 2.2 for the first setup in part III. This was because an f-number of 1.1, which is much more common, gave more artefacts without improving the resolution. We found out only later that this was caused by the phase errors that play a larger role if a lower f-number is used, and thus a smaller r_f .

A.4. CONCLUSION

We have seen that the approximation that have used during this thesis does only work well if the argument is much larger than the order. This is usually the deeper into the medium, but close to the transducer this might not be the case. Placing the virtual array deep enough decreases the overall error, but this does have negligible effect on the image close to the transducer. For these depths a better approximation is required, possibly the approximation by Gardner, or the exact values should be pre-calculated.