CHAPTER 5

NONLINEAR PROPAGATION

The amplitude-steered array can achieve fast image formation because it equates received frequency with a target's vertical position, eliminating the need to transmit multiple pulses steered in the vertical direction. However, sound propagation in water and other media is nonlinear, so that harmonics of the transmitted signal are generated as the wave propagates. In this chapter, we consider the effect these harmonics have on the operation of the array. The concern is that frequencies generated through nonlinear propagation will be considered additional targets close to the array broadside direction.

In this chapter, we first describe the expected array response to nonlinearly generated harmonics, and then we describe our simulations, which confirm this response. Within this section, we show two methods for simulation of nonlinear propagation from a piston source, choosing one for generating data and the other for confirming the results. We also list the simplifying assumptions we used to perform our simulations. Next, we discuss the simulation of the array receiving a reflected signal that contains higher harmonics, and finally we give the resulting harmonic content of the received signal produced by the array.

5.1 Expected Array Response

The properties of the amplitude-steered array as a receiver are important for determining the response of the array to nonlinear propagation. The transmit signal contains only the fundamental frequency for each steering direction; however, due to the generation of harmonics, the received signal from a particular direction will also

contain higher frequencies. The array, in receive-mode, must reject these harmonics to avoid the appearance of false targets in the images.

We expect the received signal to approach the array at the same angle, or close to the same angle, that the emitted signal was transmitted. A flat target would produce the largest response if it is perpendicular to the insonifying beam. The reflection would then return to the array along the path of the incident beam, so for example, energy transmitted at a 30° angle relative to the array broadside will be received at a 30° angle relative to the array broadside. If the flat target were tilted to a small angle relative to the insonifying beam, the reflection might still be received by the array at a slightly different angle than the original angle of the transmitted signal. However, if the flat target were tilted to a large angle relative to the insonifying beam, then the reflection would be directed away from the array. Similarly, for a point target, which would scatter energy in all directions, only energy from a small range of angles close to the angle of the insonifying beam will be received by the array. This limit to the angle at which the received signal will approach the array is important because the amplitude-steered array accepts (or receives) different frequencies from different angles.

The amplitude-steered array acts as a spatial filter. We use a diagram of the array response in the temporal-frequency/spatial-frequency (ω–k) plane to show that the array will be able to reject higher harmonics that approach the array at the same direction as the fundamental frequency. First, Figure 5.1 shows the spatial and temporal frequency filtering characteristics for an array that uses time delays for beamforming. The figure shows ω versus k_y, where y is the coordinate in the vertical direction. We are interested in signals that approach the array at some vertical angle θ, where θ is measured from the array broadside direction. In the figure, red represents the main lobe of the array and ω is proportional to k_y ($\omega = \frac{k_y c}{\sin\theta}$, $\frac{c}{sin\theta}$ = constant), indicating that different frequencies are accepted at the same angle and attenuated at other angles of incidence.

In comparison, we show the spatial and temporal frequency filtering characteristics of the amplitude-steered array in Figure 5.2. The array used for this figure is 10 cm in

length, has an interelement spacing of 1.4 mm, and is steered to 30° at 100 kHz. From this figure, we see that the array accepts input of constant k_y value, and attenuates other values of k_y. Since k_y is proportional to frequency and the sine of the steering angle, constant k_y indicates that different frequencies are accepted at different angles of incidence. We make the angular dependence more clear in Figure 5.3 by showing the spatial and temporal frequency filtering properties of the amplitude-steered array as a function of steering angle, rather than of k_y. We can filter out signals of frequencies higher than our frequency range of interest (80–305 kHz), where grating lobes are most likely to occur. We conclude from this diagram that, in the absence of grating lobes, higher harmonics generated by nonlinear propagation will be rejected due to the spatial filtering of the array.

In the presence of grating lobes for frequencies within our range of interest, a narrow beamwidth allows the array to reject higher harmonics. With narrow beamwidths, the fundamental and higher frequencies accepted at a particular angle will have narrow bandwidths. Those higher frequencies generally will not be harmonics of the fundamental frequency steered to that direction. As an example, we have modified the array used for Figure 5.3. The interelement spacing has been increased by a factor of 4 and the total length of the array has been kept at 10 cm to generate Figure 5.4. In Figure 5.4 we show that although there are grating lobes present, the grating lobes do not correspond to harmonics of the fundamental frequency for any particular steering direction. For example, at 40°, the fundamental frequency is 72 kHz. The fourth and fifth harmonics are 288 kHz and 360 kHz, respectively. The grating lobe passes signals at 328 kHz.

Finally, when the array is small, and therefore the beamwidths are wide, higher harmonics may be accepted from the same direction as the fundamental frequency, and so the harmonics will affect the operation of the array. In Figure 5.5 we show an array with very wide beams. From the array used to generate Figure 5.3, the interelement spacing and total array length have been decreased by a factor of 8. At 30°, for example, the center frequency is 100 kHz, but the main lobe is wide enough that 200 kHz is also accepted, although it is somewhat attenuated. This example

shows a design problem that can be solved by increasing the array length. Therefore, we expect that nonlinear propagation will not cause the appearance of extraneous targets for well-designed arrays.

5.2 Field Simulation

We confirm our expected array response through simulation. No analytical solution currently exists to describe the nonlinear propagation of a sound beam. The difficulty in simulating nonlinear propagation arises from the combination of the effects of diffraction, interaction of harmonics, and attenuation. Many authors have considered numerical solutions to three-dimensional wave propagation in nonlinear media. Several algorithms exist for symmetrical sources with sinusoidal [39], [40] or pulsed excitation [41], [42], [43]. Some of the algorithms for sinusoidal excitation have been adapted for nonaxisymmetric sources with sinusoidal excitation [40], [44]. Those codes can also be used for pulsed excitation, but with huge computational cost due to the large number of harmonics needed to represent the distorted pulse.

We implemented two of the algorithms for simulation of nonlinear propagation. We use the algorithm by Christopher and Parker [40] to generate the data for our study. We use the algorithm by Aanonsen et al. [39] to compare our simulation results. We discuss Christopher and Parker's algorithm first, in order to introduce Burger's equation, and then we discuss Aanonsen's algorithm, which is based on an augmented Burger's equation.

5.2.1 Christopher and Parker's algorithm

Christopher and Parker use a two-substep algorithm, propagating the field over small steps, accounting for diffraction and attenuation in the spatial frequency domain in the first substep, and accounting for nonlinear effects in the temporal frequency domain in the second substep. The diffraction is solved using an exact solution to the Kirchoff-Helmholtz intergral, although there is some error associated with the finite transform size. Nonlinear effects are calculated using the frequency domain solution to Burger's equation. The results are accurate, even in the near field of the source.

Diffraction is included by convolving each harmonic field with an appropriate point spread function. This convolution is actually accomplished with multiplication in the spatial frequency domain. The point spread function is given by

$$h_n(\Delta z, r) = -\frac{1}{2\pi}\frac{\Delta z}{d^2}\left(jnk - \frac{1}{d}\right)e^{jnkd} \tag{5.1}$$

where Δz is the incremental distance in the direction of propagation, r is the lateral distance, $r = \sqrt{x^2 + y^2}$, $d = \sqrt{\Delta z^2 + r^2}$, k is the wavenumber, and n is the harmonic. If the source is axially symmetric, computational savings can be achieved by using the Hankel transform rather than the two-dimensional Fourier transform. The Hankel transform of the point spread function is

$$H_n(\Delta z, R) = \begin{cases} exp[j2\pi\Delta z\sqrt{(nf/c)^2 - R^2}\,] & |R| \leq nf/c \\ exp[-2\pi\Delta z\sqrt{R^2 - (nf/c)^2}\,] & |R| > nf/c \end{cases} \tag{5.2}$$

where R is the transform variable, the spatial frequency in the radial direction. Attenuation is included in this step in the spatial frequency domain by multiplying $H(\Delta z, R)$ by $exp[-\alpha(nf)^2(\Delta z/\cos(\sin^{-1}(Rc/f)))]$, where α is the attenuation coefficient and f is the fundamental frequency. The algorithm accounts for linear effects in the spatial frequency domain, which requires the repeated use of discrete transforms. To avoid wraparound error, the transform of the point spread function is appropriately windowed [45].

The nonlinear effects are accounted for in the temporal frequency domain. A plane wave assumption is made, which means that all phase fronts are assumed to travel in the same direction as the fundamental, allowing the use of Burger's equation for calculating the accretion and depletion of harmonics. This assumption is only loosely justified [40]. Burger's equation is given by

$$\frac{\partial u}{\partial x} = \frac{\beta\omega_0}{c_0^2}u\frac{u}{\partial\tau} + \Gamma\frac{\partial^2 u}{\partial\tau^2} \tag{5.3}$$

where u is the particle velocity, x is the direction of propagation, τ is the retarded, dimensionless time $\tau = \omega_0 t - kx$, and Γ is the Gol'dberg number, the ratio between absorption length and shock formation distance.

The solution to Burger's equation is given by Trivett and van Buren [46] and by Haran and Cook [47], [48]. We write the form given by Christopher in [40]:

$$u_n(z+\Delta z,i) = u'_n(z+\Delta z,i) + j\frac{\beta\pi f\Delta z}{c^2}\left(\sum_{k=1}^{n-1} k u'_k u'_{n-k} + \sum_{k=n+1}^{N} n u'_k u'^{*}_{k-n}\right) \qquad (5.4)$$

where u_n is the particle velocity for harmonic n, u'_n represents the output of the diffraction substep, Δz is the length that the field is advanced with each step of the algorithm, and f is the fundamental frequency.

This algorithm is implemented in Matlab, and the results are used to generate the reflected signal to be received by the amplitude-steered array.

5.2.2 Aanonsen's Algorithm

Several of the algorithms mentioned in the introduction to this section calculate the propagation of sound in a bounded beam using the Khokhlov-Zabolotskaya-Kuznetsov (KZK) equation. The KZK equation is a form of Burger's equation, augmented to include diffraction. The KZK equation was originally derived without absorption by Khokhlov and Zabolotskaya in 1969. Kuznetsov added absorption in 1971. Naze Tjøtta and Tjøtta later derived a nondimensional version of the equation using the method of multiple scales. The equation as written by Naze Tjøtta and Tjøtta is

$$\left(4\frac{\partial^2}{\partial\tau\partial\sigma} - \bar{\nabla}^2_{\perp} - 4\alpha r_0\frac{\partial^3}{\partial\tau^3}\right)\bar{p} = 2\frac{r_0}{l_d}\frac{\partial^2}{\partial\tau^2}\bar{p}^2 \qquad (5.5)$$

where τ is the normalized retarded time $\tau = (t - z/c)/T$ and the normalization time T for a periodic waveform is the period. The normalized axial distance σ equals z/r_0; the normalization distance r_0 is the Rayleigh distance, $r_0 = a^2/2cT$; α is the absorption coefficient; l_d is the plane wave shock formation distance; and $\bar{p}$ is the normalized acoustic pressure, $(p - p_0)/\rho_0 c u_0$ where u_0 is the source velocity. The $\bar{\nabla}^2_{\perp}\bar{p}$ term accounts for diffraction. Without the diffraction term, the KZK equation would reduce to Burger's equation.

The KZK equation is derived using the parabolic approximation, which means the acoustic variable, pressure or normal particle velocity, is assumed to vary slowly

in the direction of propagation. The solution is not correct very close to the source. The conversion between pressure and normal particle velocity is made using the plane wave impedance relation.

Aanonsen, Barkve, Naze Tjøtta, and Tjøtta formulated a solution to the KZK equation by assuming a Fourier series form for the pressure,

$$\bar{p} = \sum_{n=1}^{\infty} (g_n \sin n\tau + h_n \cos n\tau) \tag{5.6}$$

and substituting the solution into Equation (5.5). A coupled set of equations is obtained:

$$\begin{aligned} \frac{\partial g_n}{\partial \sigma} &= -n^2 \alpha r_0 g_n + \frac{1}{4n} \bar{\nabla}_{\perp}^2 h_n \\ &\quad + n \frac{r_0}{2 l_d} \left(\frac{1}{2} \sum_{p=1}^{n-1} (g_p g_{n-p} - h_p h_{n-p}) - \sum_{p=n+1}^{\infty} (g_{p-n} g_p + h_{p-n} h_p) \right) \end{aligned} \tag{5.7}$$

$$\begin{aligned} \frac{\partial h_n}{\partial \sigma} &= -n^2 \alpha r_0 h_n - \frac{1}{4n} \bar{\nabla}_{\perp}^2 g_n \\ &\quad + n \frac{r_0}{2 l_d} \left(\frac{1}{2} \sum_{p=1}^{n-1} (h_p g_{n-p} + g_p h_{n-p}) + \sum_{p=n+1}^{\infty} (h_{p-n} g_p + g_{p-n} h_p) \right) \end{aligned} \tag{5.8}$$

where n is the number of the harmonic and n ranges from 1 to ∞.

We program Equations (5.7) and (5.8) in Matlab using an implicit backward finite difference algorithm to integrate, for stability. The results are used for comparison with the Christopher and Parker algorithm.

5.2.3 Simplifying assumptions

In order to show the effect of harmonic generation on the operation of our array, we simulate nonlinear propagation using Christopher and Parker's algorithm. We make several simplifying assumptions to reduce the complexity of the problem while still illustrating the effects.

First, we simulate the low frequency amplitude-steered array, rather than the high-frequency array. The geometry of the low frequency transmitter is approximately circular, as opposed to the extremely long and narrow geometry of the high frequency transmitter. Despite the symmetry of the low frequency array, we use the

two-dimensional FFT rather than the Hankel transform to calculate diffraction and attenuation, as the Hankel transform is not readily available in Matlab. The circular geometry allows us to calculate the field on a smaller grid of points in the lateral directions because we do not have to account for the large spreading in one dimension. The lower frequency also allows a lateral grid with fewer points because we can sample spatially at a lower rate due to the longer wavelength. In the axial direction, we can calculate fewer points because the variation in the near field is slower than for the high frequency array. Also, we do not have to propagate the field as far to reach the range of interest.

Second, we simulate a case where the nonlinear effect is not very strong. The stronger the nonlinearity (the greater the source amplitude), the more harmonics will be generated, greatly increasing the cost of computation. However, we know *a priori* that frequencies outside the range 80 kHz to 305 kHz are not generated at the source, so we can use a lowpass filter to remove frequencies higher than 305 kHz. Therefore, we do not have to consider the problem where such high harmonic frequencies are generated. We simulate a signal that oscillates sinusoidally at 100 kHz at the source, as it will generate a large number of harmonics within our frequency range of interest. We generate harmonics up to 400 kHz.

In the simulation of the field, we do not include beamsteering. The source is assumed to be a plane circular piston, not a steered array. With this assumption, the grid on which field points are calculated does not have to be large to accomodate the steering angle of the beam. Nor does the grid have to shift to follow the beam. The implication is that a field point on the axis of the simulated beam pattern for the piston source corresponds to a field point on the axis of the 100-kHz beam of the amplitude-steered array, which is steered to 30°. This assumption is reasonable for axial distances greater than or equal to a^2/λ, the near-field/far-field transition, which occurs at 0.17 cm in our case.

5.2.4 Values used in our simulation

Christopher and Parker offer guidelines in the use of their algorithm, based on their experience [40]. If shocks are expected, then 30 to 50 harmonics may be required. If no shocks are expected, 5 to 10 harmonics should be enough to characterize the field. We choose a souce amplitude such that no shock is expected. We calculate the field to 1 m, and the plane-wave shock formation distance is 1.983 m. Satisfactory results are obtained by using five harmonics.

The radial sampling should be fine enough to show the structure of the field for the highest harmonic. For simulations requiring 50 harmonics, radial sampling of four times the Nyquist rate for the fundamental should be adequate. If no shocks are expected, one or two times the Nyquist rate for the fundamental should be adequate. Our fundamental frequency is 100 kHz, which corresponds to a wavelength of 1.5 cm. We use a lateral step size of 3 mm, which is smaller than half the Nyquist step size for the fundamental.

Finally, the step size in the axial direction is chosen to accurately represent the axial variations in the near field. The step size should also be chosen to account for the attenuation of the highest harmonic. That is, the step size should be made small enough that the magnitude of the highest harmonic at the end of the step does not decrease to less than 0.7 times the magnitude at the beginning of the step due to attenuation. In our case, the first condition is more restrictive. We use an axial step size of 1 mm.

5.2.5 Validation of simulation

Our simulation is validated two ways. First we compare our implementation of Christopher and Parker's algorithm with Ted Christopher's results for one case. Christopher and Parker have already satisfactorily compared their results to theory [40]. The case is a piston source with 5-cm radius. The source intensity is 20 W/cm^2, corresponding to an initial normal particle velocity of 51.6 cm/s. The medium is water, so the attenuation coefficient is 25×10^{-15} Np/m-Hz^2 and β is 3.5. The results of our simulation using Christopher and Parker's algorithm are compared with the

results from Christopher's implementation of the algorithm in Figure 5.6. For each harmonic, the maximum amplitude calculated by Ted Christopher is larger than the maximum amplitude calculated by our program. However, the differences are small. Our results are satisfactory for our purpose, which is to show how the amplitude-steered array will accept or reject harmonics.

Second, we compare our results to results calculated using Aanonsen's algorithm for the same case described above. The results for the first three harmonics are shown in Figure 5.7. The agreement between our calculated results and those using Aanonsen's algorithm are good, but not as good as the agreement between our implementation of Christopher's algorithm and his results, which is expected since in the previous comparison, the algorithms were the same. The results using Aanonsen's algorithm are not accurate close to the source, where the parabolic approximation fails. The results using Aanonsen's algorithm also show a higher peak particle velocity and then a faster decline in particle velocity for the fundamental frequency. Ted Christopher's results also showed a higher peak than we calculated; however, in both cases the agreement is reasonable.

5.3 Simulation of Array Operation and Results

Once the field generated by the array is simulated, the data are input into a simulation by Field II to generate the array response. Harmonic amplitudes and phases generated in the nonlinear field simulation using Christopher and Parker's algorithm are inverse Fourier transformed to generate the signal, which will be the signal reflected by a point target at that field point. The point reflector is taken from an axial position in the nonlinear simulation. The transmitted field is assumed to be reflected by a point target at a range of 80 cm along the 30° direction, corresponding to the steering direction for 100 kHz. The maximum array response is steered in the vertical direction but not in the horizontal direction. The reflected field is assumed to propagate back to the array linearly because only a small fraction of the incident energy will be reflected by a point target. Nonlinear effects in the reflected wave therefore will be insignificant.

The truncated signal that is radiated by the point source and the magnitude of its Fourier transform are shown in Figure 5.8. Four harmonics are used to form the signal. The signal has a length of 1 ms, but only one fourth of it is shown to give a clearer figure. Although the nonlinear effect is by design not strong, the peak positive and peak negative pressures are not symmetric.

The magnitude of the Fourier transform of the signal received by the amplitude-steered array is shown in Figure 5.9. The array is 10 cm in length, with an interelement spacing of 1.6 mm, and it is designed to steer to 30° at 100 kHz. The amplitude of the transform has been normalized, and for comparison the amplitudes of the signal harmonics before reception are shown as dots. The harmonics have been attenuated and are no longer visible.

For comparison in Figure 5.10, we show an array that has the same diameter as the original array, but with interelement spacing increased by a factor of 4, so that grating lobes are present. Again, dots in the figure represent the amplitudes of the harmonics before reception by the array. The harmonics are attenuated, but not as well as in the previous case.

Finally in Figure 5.11, we show the results of receiving the same signal with an array of smaller total size. The interelement spacing and total array length have been decreased by a factor of 8. The received beam pattern does not have any grating lobes, but the beamwidth is very wide. Here we see that the higher harmonics are not attenuated as much as in either of the previous two cases.

5.4 Conclusion

In this chapter, we have examined the effect of nonlinear propagation on the performance of the amplitude-steered array. The concern was that harmonics generated during propagation would appear as false targets in the image. We simulated the nonlinear propagation from a piston source, whose dimensions and operating frequency correspond to those of the low frequency amplitude-steered array. The waveform for a spatial point on the beam axis of the piston was formed and assumed to be a reflected waveform in a simulation of the operation of the amplitude-steered array. In

this second simulation, the position of the spatial point was shifted so that it corresponds to a point the same distance away, but along the beam axis of the 100-kHz beam, which was steered to 30° for the amplitude-steered array. The received signal was then formed. We have shown that in the received signal the higher harmonics are attenuated due to the spatial filtering properties of the array. For larger beamwidths, harmonics are accepted by the array, which means that they would appear as false targets in the image. However, nonlinear propagation does not have to cause false targets in the images for a properly designed array, that is, an array that is long enough.

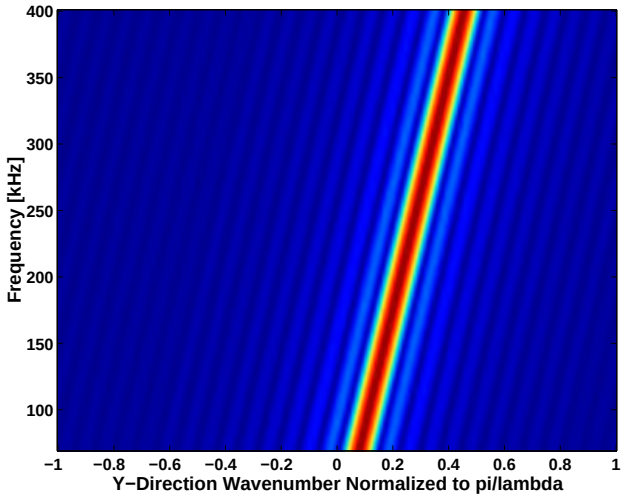

Figure 5.1: ω–k diagram for an array that uses time delays for beamforming.

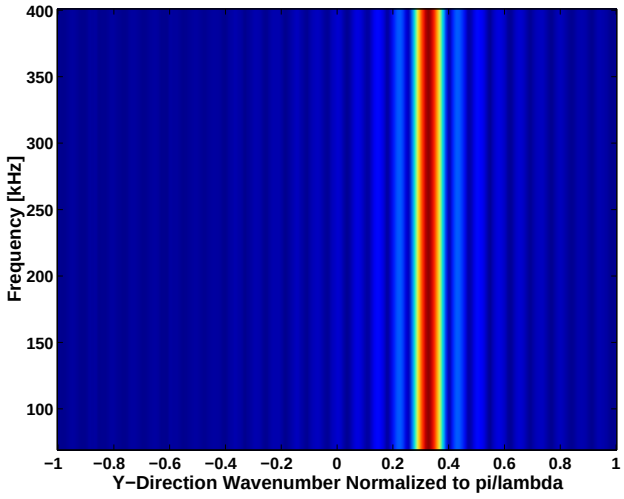

Figure 5.2: ω-k diagram for the amplitude-steered array, which uses amplitude weighting for beamforming.

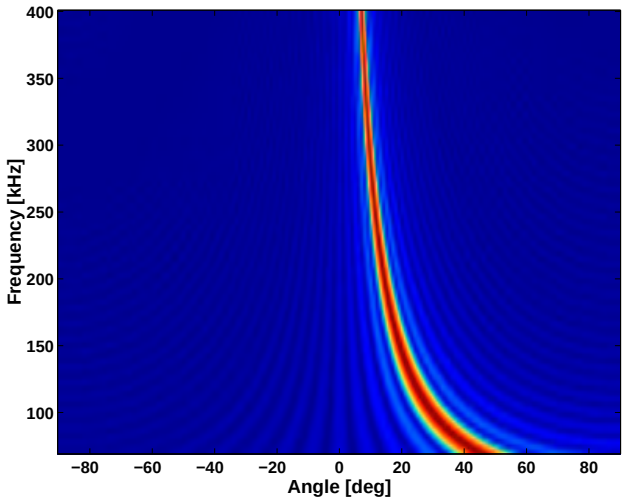

Figure 5.3: Spatial and temporal filtering properties of the amplitude-steered array.

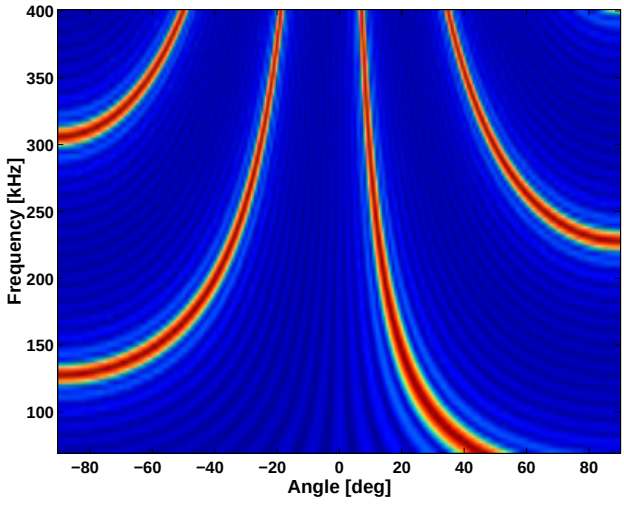

Figure 5.4: Spatial and temporal filtering properties of an amplitude-steered array, with the same length, but fewer elements.

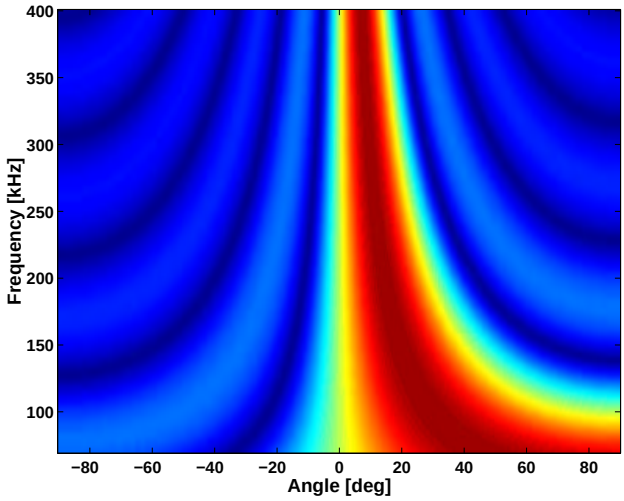

Figure 5.5: Spatial and temporal filtering properties of an amplitude-steered array with shorter length and fewer elements.

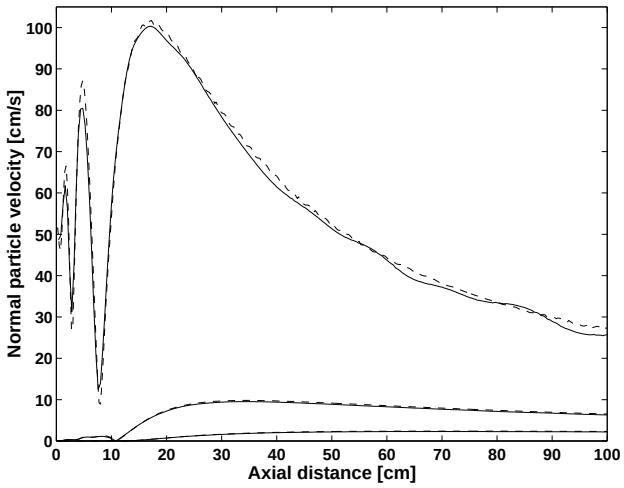

Figure 5.6: Comparison of simulated results (solid lines) with Christopher's results (dashed lines). The fundamental (considered the first harmonic) and two higher harmonics are shown.

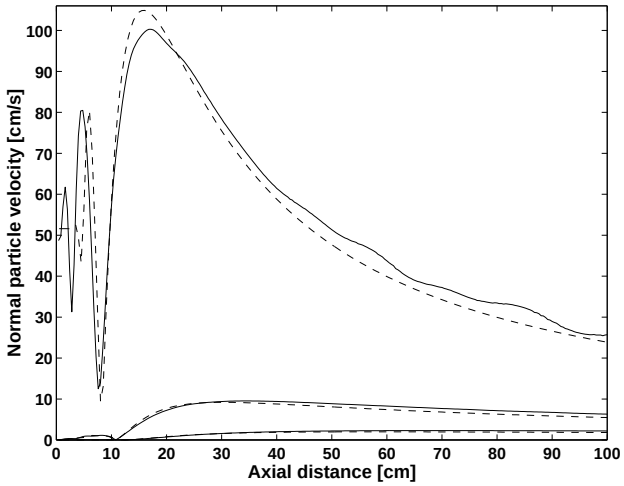

Figure 5.7: Comparison of simulated results (solid lines) with results from Aanonsen's algorithm (dashed lines). The fundamental (considered the first harmonic) and two higher harmonics are shown.

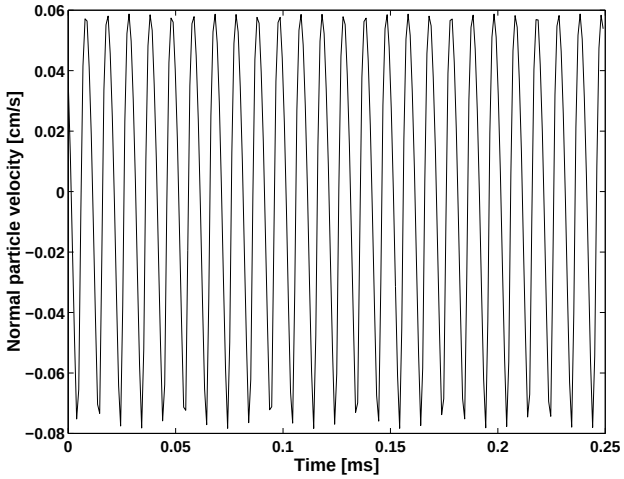

(a)

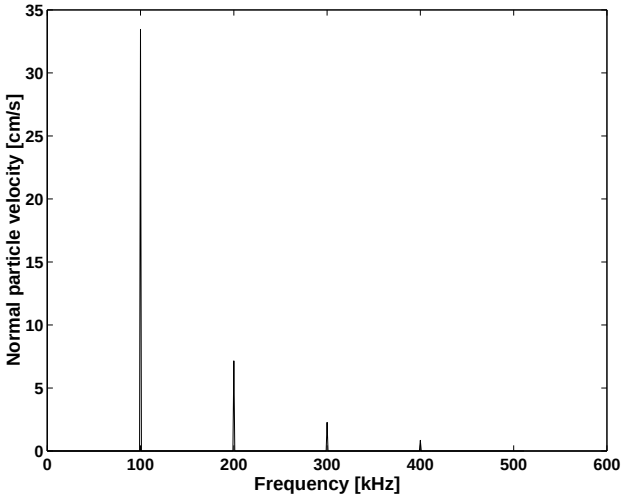

(b)

Figure 5.8: Signal generated by simulation of nonlinear field propagation: (a) time signal, (b) harmonic components.

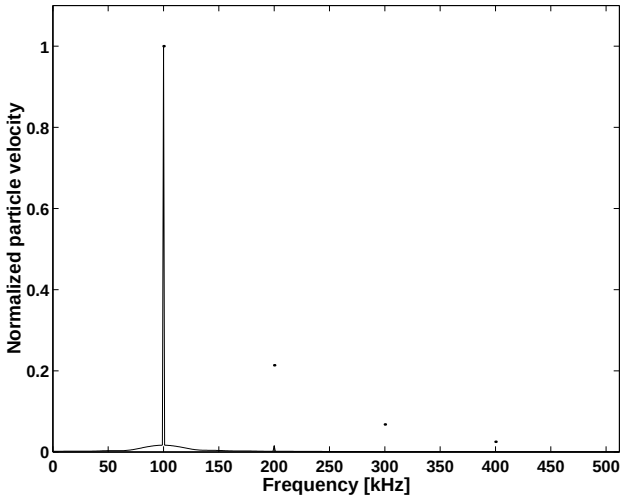

Figure 5.9: Fourier transform of the received signal using the low frequency amplitude-steered array.

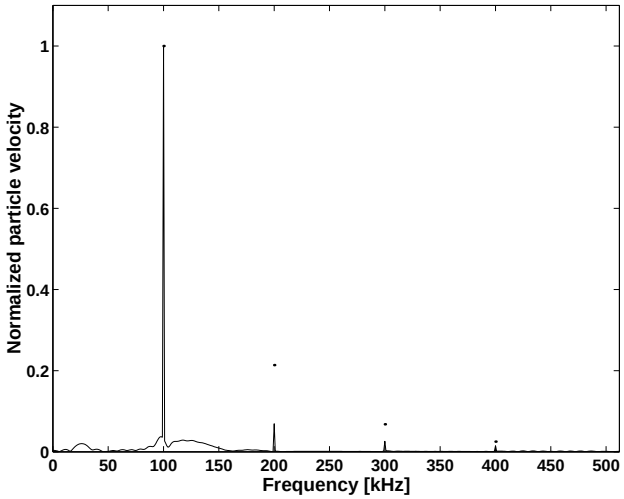

Figure 5.10: Fourier transform of the received signal using the low frequency amplitude-steered array. Interelement spacing has been increased without increasing the total length of the array, so that grating lobes are present.

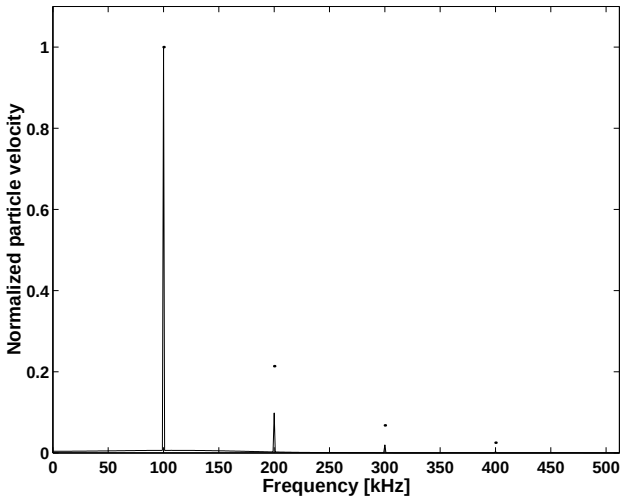

Figure 5.11: Fourier transform of the received signal using the low frequency amplitude-steered array. Interelement spacing has been decreased, changing the total length of the array by a factor of 8, so that beamwidths are greater.

CHAPTER 6

CONCLUSION

Acoustic volumetric imaging is difficult to accomplish in real time due to the speed of sound and the resulting time required to collect data. Conventional methods of data collection require the transmission and reception of at least one pulse per line in the image. For three-dimensional imaging, a reasonable set of data is 100 by 100 lines. The time to collect each line depends on the range of interest. For ranges more than a few centimeters, in order to be able to refresh the volume of data 30 times each second, innovative methods of data collection are required.

This thesis is a comprehensive study of the application of the amplitude-steered array to real-time volumetric acoustic imaging. Previously the linear amplitude-steered array had been used for narrow-band signal detection. The current project seeks to build an acoustic volumetric imaging system that is based on a two-demensional version of this array. The project has involved many groups working on various aspects of system design and implementation including crystal growth, electronics design, and building the array from the crystal (dicing and electrically connecting elements). The work presented in this thesis contributed to the analysis of the acoustic array operation and to the development of imaging algorithms.

6.1 Summary of Research Results

The specific research results produced by the author and presented in this thesis are listed here with some discussion.

- Demonstration of the tradeoff between axial and lateral resolution, which are related by the length of the array.

Initially it was unclear how two targets separated only in range could be distinguished with an array that spatially separates frequencies. If frequencies were perfectly separated, a point target would return a sinusoidal signal that was infinitely long. It was shown that targets separated only in range could be distinguished because frequencies were not separated perfectly, but rather overlapped. As a result, it was shown that there was a relationship between axial and lateral resolution. The width of the beam at each frequency decreases as the array length increases. Thus the overlap of frequencies decreases as the array length increases, leading to poorer axial resolution as lateral resolution improves.

- Analysis of time-frequency distributions for image formation.

  The spectrogram was initially proposed as the means to form an image from the signal received by the array. However, there is an inherent tradeoff between time (axial) and frequeny (lateral) resolution in the spectrogram that depends on the length of the window used for processing. Other distributions had different tradeoffs in time and frequency resolution and cross-term level. The smoothed pseudo-Wigner distribution (SPWD) produced the best combination of axial and lateral resolution, while reducing cross-term levels. Because the SPWD applies smoothing separately in the time and frequency directions, it has the most available parameters that can be adjusted to optimize axial and lateral resolution and to reduce cross-terms.

- Incremental development of the array layout to show that a sparse random array could be used.

  Each of the array designs originated at the Applied Research Laboratory at the Pennsylvania State University. The final array design was a random design that led to a narrow beamwidth with no inordinately high sidelobes. However, the fully sampled array may have too many elements to be implemented with current array fabrication technology, despite needing only four channels per stave rather than one channel per element. Therefore, it was shown that a

beam pattern with a narrow main lobe, but higher average sidelobe level could be achieved with a sparse random design. Average sidelobe level increases with the sparseness of the array.

- Development of projection algorithms to display data with reduced processing.

  Once the three-dimensional data set has been collected, it must be displayed on a two-dimensional screen. Therefore, several processing algorithms were presented for displaying the data. The three-projection method was adopted due to the amount of information displayed for the amount of processing. By forming only the three projection images, and not the entire volume of data, fewer time-frequency distributions had to be calculated. The projection through vertical angle used one Hilbert transform for each horizontal steering direction. The projection through range used one Fourier transform per horizontal steering direction. And the projection through horizontal angle used one calculation of the time-frequency distribution. In contrast, the surface rendering techniques presented showed only one view of the target. Forming slice images or trying to do more sophisticated volume rendering would require the calculation of one time-frequency distribution per horizontal steering direction.

- Demonstration that nonlinear propagation did not have to lead to false targets in the images.

  The processing of data from the amplitude-steered array equates frequency with target position, so it is possible for the generation of harmonics through nonlinear propagation to cause the appearance of false targets in the image. Although nonlinear propagation will cause the generation of harmonics, the spatial filtering properties of the array cause harmonics to be attenuated, as long as the beamwidths are narrow enough and as long as the number of grating lobes within the frequency range of interest is not extensive.

- Verification of simulations through comparison with experiment.

  The demonstration of all of the previously mentioned results was accomplished

through simulation of the array operation. In order to confirm the results, experimental data were collected at the Applied Research Laboratory at the Pennsylvania State University. Analysis of the data and comparison to simulation were accomplished by the author. Comparison between simulated and experimental results showed very good agreement, particularly for point-like targets.

6.2 Future Work

The amplitude-steered array has been characterized for imaging, and algorithms have been designed for processing data from a volumetric imaging system. The actual implementation of the system involves other issues which are currently being investigated. In addition, the use of the array for other applications should be investigated. Some of these issues are discussed below.

6.2.1 Environmental factors

The study of the array operation was carried out through simulation and through operation in a controlled water tank. The actual system will operate in a much more varied environment, which will induce more noise. The effects of current, temperature and pressure gradients, and water turbidity should be determined, although it is expected that temperature and pressure will not vary much over the desired operating range of 5 m. Also, the effect of noise on the image formation algorithms should be determined. The smoothed pseudo-Wigner distribution should not be greatly affected due to the the low-pass filter in each direction.

6.2.2 Full two-dimensional amplitude-steered array

The two-dimensional array discussed in this thesis is an amplitude-steered array in one dimension and a linear phased array in the other dimension. Three-dimensional data are collected with the transmission of a single pulse because the transmitter is narrow in one direction, leading to a broad beam in that direction. A full two-dimensional amplitude-steered array would steer a beam in two directions by changing

frequency. Most likely, to have such a system, amplitude-steering would have to be combined with pulse coding. Using frequency steering in two dimensions would cause a rotated line of data to be collected.

6.2.3 Medical imaging

Our array was discussed within the framework of an underwater imaging system. But the technology could also be used for medical imaging, for example, real-time imaging of the heart. Compared to underwater imaging, some of the requirements on our array could be relaxed for medical imaging, which would make the fabrication and operation easier. This is largely due to the much smaller range of interest. The array could be smaller, meaning many fewer elements. We could use a linear amplitude-steered array that is rocked or rotated to collect the data because there is time to transmit multiple pulses. Our transmit pulse would have to be much shorter to reduce the "dead zone" so that features close to the transducer could be imaged.