NEW APPROACHES TO ABERRATION CORRECTION IN MEDICAL ULTRASOUND IMAGING

BY

MARK ALDEN HAUN

B.S.E., Walla Walla College, College Place, 1996
M.S., University of Illinois at Urbana-Champaign, 1999

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2003

Urbana, Illinois

ABSTRACT

Medical ultrasound imaging is in widespread use today due to its low cost, portability, lack of side effects, and unique ability to probe the mechanical properties of tissue. As transducer apertures and operating frequencies grow, however, ultrasound's resolving power continues to be limited by variations in the speed of sound within tissue. Concepts borrowed from other imaging disciplines can provide new insights into the aberration problem in medical ultrasound; in particular, some of the same issues have been studied for many years to improve acoustic imaging of the nonhomogeneous Earth.

The seismic imaging community has understood for some time that layered media may be approximated by constant-sound-speed media for beamforming purposes, leading to so-called time-migration algorithms. This raises the possibility of medical ultrasound applications—for example, brain imaging through the adult human skull. While our simulation results have been encouraging, experiments with animal skulls have been inconclusive due to the high attenuation of ultrasound in skull bone. Further research may validate the time-migration concept for brain imaging.

Complete data sets are composed of the raw reflection data from every combination of one transmit element and one receive element in an array. The tremendous redundancy of a complete data set can be exploited for aberration correction by analyzing the time shifts on common-midpoint gathers. Until now, however, the wide-angle, random-scattering nature of medical ultrasound targets has limited the accuracy and robustness of this approach, particularly when estimating azimuth-dependent aberration profiles. Prefiltering the data with two-dimensional fan filters largely solves this problem and produces an aberration-correction algorithm (OFF) that outperforms the most popular existing algorithms in almost all cases.

The concept of focusing-operator updating, recently popular in seismic imaging, provides insight into iterative aberration-correction algorithms using a transmit focus. We

develop a new updating procedure based on dynamic programming. With careful selection of initial focus points, the resulting algorithm outperforms existing algorithms in some experiments.

Our aberration-correction results imply that imaging with single-valued focusing operators may be able to correct for most of the aberration encountered in soft tissues; that increasing aperture should not be viewed merely as a source of aberration, but as an opportunity to more fully correct it; and that the noise penalty for using complete data sets may not be as serious a problem as commonly assumed.

A different kind of image-formation challenge is posed by small-diameter cylindrical imaging platforms. We derive a fast, three-dimensional, frequency-domain imaging algorithm for this geometry by making suitable approximations to the point spread function for wave propagation in cylindrical coordinates and obtaining its Fourier transform by analogy with the equivalent problem in Cartesian coordinates. For the most effective use of limited aperture, the focus of a transducer is treated as a virtual source, and the synthetic-aperture algorithm then forms images on deeper cylindrical shells. Computer simulations and experimental results show that this imaging technique attains the resolution limit dictated by the operating wavelength and transducer characteristics.

ACKNOWLEDGMENTS

During my time at UIUC, I have been fortunate to have Prof. Doug Jones as my thesis adviser for both the M.S. and Ph.D. degrees. His incurable optimism and many helpful suggestions have helped keep my research on track, and I have benefited tremendously from our honest and frank conversations about research, careers, and life in general.

I wish to thank everyone in the Bioacoustics Research Lab for enabling and supporting my experimental work and counting me as part of their research group. In particular, thanks go to Prof. Bill O'Brien for his sage advice on all things related to ultrasound; Jim Blue for his persistence in obtaining all manner of animal specimens for me to scan; Rita Miller for her assistance in laboratory procedures; Dennis Matthews for writing the software application to collect data sets with our 64-element array; and Mark Johnson for measuring the acoustic properties of the silicone rubber aberrating phantoms.

One of the most enjoyable aspects of this research has been the opportunity it brings for cross-discipline collaboration. In May 2001 I was privileged to travel to Houston, TX, and discuss ultrasound imaging with geophysicists at BP and the University of Houston. I am grateful to Joe Dellinger for arranging this visit and to all the researchers who kindly took hours away from their busy schedules to meet with me: At BP, Ray Abma, Sverre Brandsberg-Dahl, Joe Dellinger, John Etgen, Paul Garossino, Ken Matson, and Simon Shaw; at UH, Joongmoo Byun, Kurt Marfurt, Arthur Weglein, and Hua-wei Zhou.

Hal Kunkel at Philips/ATL helped us procure the 64-element array transducer used for the aberration-correction experiments. He faithfully answered my seemingly endless questions while the array was being integrated into our data-acquisition system.

Financial support for this research was provided by the University of Illinois Research Board and the National Institutes of Health (grant no. CA79179).

Finally, I thank my parents, who have patiently supported and encouraged me throughout my education—especially during these uncertain and often difficult years of doctoral research.

TABLE OF CONTENTS

CHAPTER **PAGE**

LIST OF TABLES

LIST OF FIGURES

Figure **Page**

CHAPTER 1

INTRODUCTION

"Today, I get a weird feeling whenever a physician uses a sonogram to image my gizzard. Analog data, two-dimensional image, no pulse compression; it's what we used in Wyoming in 1952. I'm just glad the doctor isn't using dynamite. It isn't trivially easy to translate modern seismic technology into medical practice. Deep seismic reflections in the oil patch take three and four seconds to return. In my body, everything is over in one millisecond. It will require fast computers and some tough engineers."

—Kenneth S. Deffeyes, in *Hubbert's Peak: The Impending World Oil Shortage*, p. 87 [1].

1.1 Medical Ultrasound Imaging

Virtually every pregnant woman in the western world today will have her fetus clinically evaluated with the aid of ultrasound—high frequency acoustic waves in the 1–100 MHz range. Ultrasound images of the heart are used to examine almost anyone suffering from chest pain. In many parts of the body, suspected tumors are routinely scanned with ultrasound.

This widespread use is evidence of the unique advantages offered by ultrasound, which complements other important diagnostic imaging technologies like MRI, PET, and X-rays. Compared to most of these, an ultrasound scanner is small, lightweight, and inexpensive, requiring only a handheld transducer head and the processing and display electronics. Acoustic radiation is nonionizing and produces few or no side effects. It also has the advantage of being able to probe the mechanical properties of human tissue directly, while other techniques must infer this indirectly.

The simplest method of obtaining an ultrasound image is the pulse-echo *B-mode scan* (Figure 1.1), in which a focused source of pulsed ultrasound is translated parallel to the skin of the

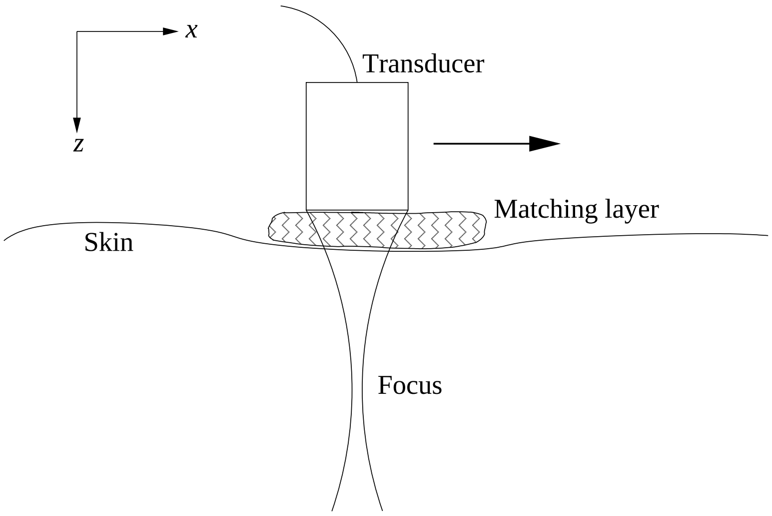

Figure 1.1 A B-mode ultrasound scan.

patient or rotated to achieve a sector scan. A gel-like substance ensures good coupling between the transducer and the skin. In soft tissues, a weak scattering assumption is used to relate the backscattered echoes in time to the tissue reflectivity in depth. The received radio-frequency (RF) signals are envelope-detected and displayed side by side in the (x, t) plane, forming the desired image. To preserve good lateral (x-direction) resolution over a range of depths, the ultrasound transducer is designed with an f-number (that is, the *focal ratio* of focal length to diameter) high enough to yield the desired depth of focus. Single-element transducers suitable for this task usually contain a flat disk of piezoelectric material coupled to a mechanical lens which creates a fixed point focus and also helps match the transducer to the acoustic impedance of the medium.

In modern clinical use, single ultrasound transducers have been replaced with phased arrays. These allow the ultrasound beam to be steered over a range of angles and focused at any depth without the need to physically move the transducer. 1-D arrays, by far the most common, incorporate a number of elements which are narrow in the x direction and wide in the y direction (see Figure 1.2). This permits flexible focusing in azimuth, while a mechanical lens covering the array provides a weak fixed focus in elevation. Fully sampled 2-D arrays are rare because of the wiring and signal processing challenges for even a small aperture. Multirow arrays with full sampling along the x direction and a handful of elements along the y direction are a compromise between these extremes, allowing for limited beamforming in the elevation plane. These are nicknamed with fractional dimensions between one and two, e.g., a 1.5-D array [2].

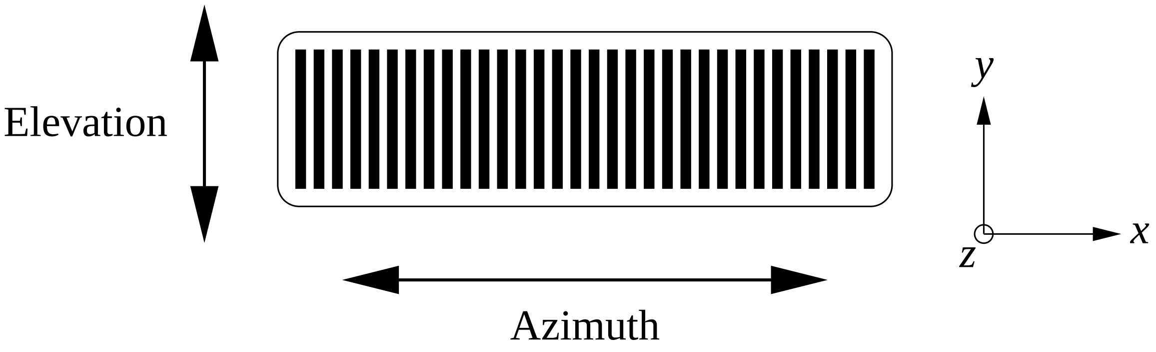

Figure 1.2 Bottom view (facing up from the skin) of a 1-D ultrasound array showing typical element shape and plane definitions.

Ultrasound arrays are also increasingly being designed with novel geometries for specialized applications inside the body. Circular arrays are now being placed on catheters for close-up imaging of blood vessels [3]. A proposed application of very-high-frequency ultrasound would place tiny transducers on a needle for high-resolution *in vivo* imaging of suspected tumors [4]. Chapter 7 presents an efficient, Fourier-domain, three-dimensional imaging technique for pulse-echo data collected on cylindrical apertures.

Sound waves propagating in the body are subject to attenuation, scattering by impedance discontinuities, and refraction by propagation speed variations. Attenuation in soft tissues is approximately 0.5 dB cm^{-1} MHz^{-1}. For this reason, clinical applications must seek a compromise between resolution and penetration depth, usually staying below 10 MHz. Scattering in an idealized model is weak and caused primarily by tissue microstructure at subwavelength scales. Multiple scattering effects are neglected due to the weak scattering assumption. The bulk acoustic properties are assumed constant at large scales (with respect to a wavelength), with the transmitted wave remaining mostly intact and traveling at a constant speed. In reality, there is significant variation in these properties; in particular, the speed of sound varies from 1470 m/s in fatty tissues to over 1600 m/s in muscle and up to 3700 m/s in bone [5, 6]. The resulting effects are collectively termed "aberration" and are the focus of Chapters 2 through 6.

Medical ultrasound images are formed assuming a constant speed of sound, usually 1540 m/s. Straightforward "delay-and-sum" beamforming is used almost universally. Dedicated hardware applies geometrically determined delays across the array so that transmitted signals focus to a point and received signals are added coherently for a point source at the focus. To aid the sonographer in positioning the array and to make the image less susceptible to tissue-motion artifacts, a high frame rate is desirable. Using transmit and receive foci to interrogate every

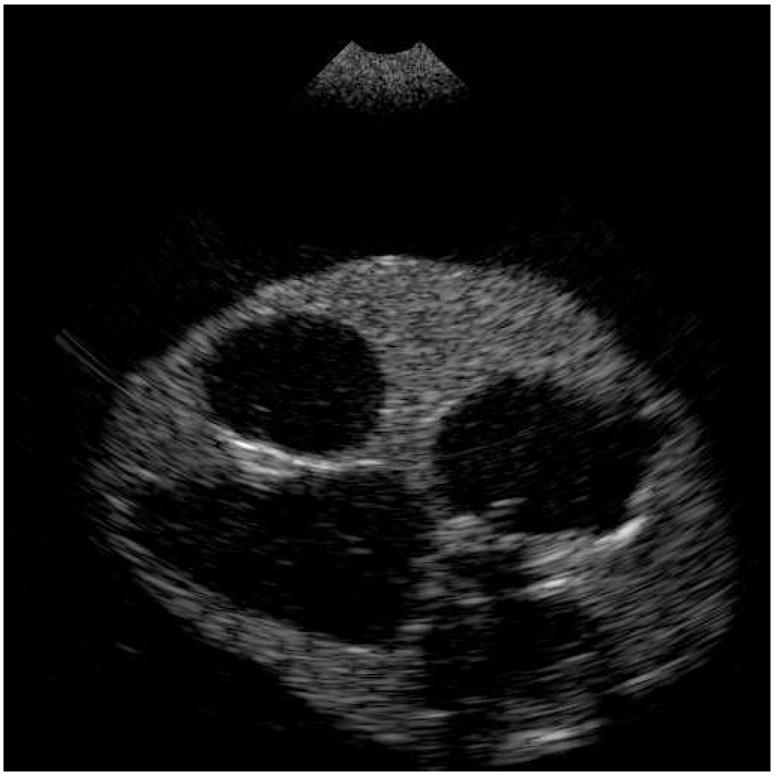

Figure 1.3 An ultrasound image of a "phantom" target designed to mimic the appearance of a heart. (Image courtesy U. Michigan Biomedical Ultrasonics Laboratory.)

point in an image, however, would require too much time for sound propagation, even if the processing circuitry was arbitrarily fast. The usual compromise is similar to the B-mode scan described above: Pulses are transmitted at various azimuth angles using a large depth of focus, so that one transmission covers the entire range of depths in the image. For reception, instead of using the same fixed focus, the focus can be scanned along the range line as the echoes arrive. This is called *dynamic focusing* on receive. Common enhancements to this scheme use more than one transmit pulse per azimuth, segmenting the image in both azimuth angle and range.

One striking feature of medical ultrasound images, like the simulated heart in Figure 1.3, is the grainy background texture. This is not noise, but *speckle*, a phenomenon observed in any coherent imaging system when the target contains many random, subwavelength scatterers per resolution cell. The brightness of a single pixel in the image is determined by the vector sum of the contributing scatterers' complex reflectivities. Modeling this as a random walk in the complex plane [7], the central limit theorem predicts circular Gaussian statistics as the number of scatterers per resolution cell becomes large. The pixel intensities are thus Rayleigh distributed, making the background appear grainy. Because speckle-producing regions are so common in tissue, speckle plays an important role in aberration correction; this will be discussed in more detail in Chapter 2.

The resolving power of an ultrasound scanner is not evaluated solely by its ability to resolve closely spaced point targets, as is the case in many other imaging systems. For diagnostic purposes, it is more important to be able to pick out small, anechoic (scatter-free) regions surrounded by speckle. Thus, sidelobe level is at least as important as mainlobe width.

1.2 Exploration Seismic Imaging

Because some of the results in this dissertation derive from techniques used in seismic imaging, it is convenient to summarize briefly the current practice, concepts, and terminology from that field. Many similarities with medical ultrasound will be noted, as well as some differences.

1.2.1 Acquisition

Geophysicists and petroleum prospectors have been using reflection seismology to image underground features in search of oil and gas since the 1920s [8]. The technology required for a rudimentary survey on land is simple: A dynamite charge in a borehole serves as the energy source, and a string of geophones (devices which record vibration) record the echoes. Most of the echo energy lies between 10 and 100 Hz and takes seconds to arrive back at the geophones, leading to modest sampling and recording requirements.

The chemical explosives used by early surveys have mostly been abandoned in favor of vibrating trucks (on land) or airguns (at sea). Unlike ultrasound, the sources and receivers are not interchangeable, and the receivers are less expensive than the sources. The firing of a source is recorded over an array of receivers simultaneously, then the source is moved to a different position (along with the receivers, in a towed-array marine survey) and the experiment is repeated. For a single survey line, this leads to a collection of signals $d_{s,g}(t)$, where s is the *shot* (source) position and g is the *geophone* (receiver) position along the linear aperture. Of vital importance later in this dissertation is the fact that the data are recorded separately for each source and receiver position. This is called *prestack* data, for reasons which will become clear shortly, and should be contrasted with the conventional ultrasound practice, where echo data are implicitly summed by the transmit focusing operation due to a number of sources being fired at about the same time.

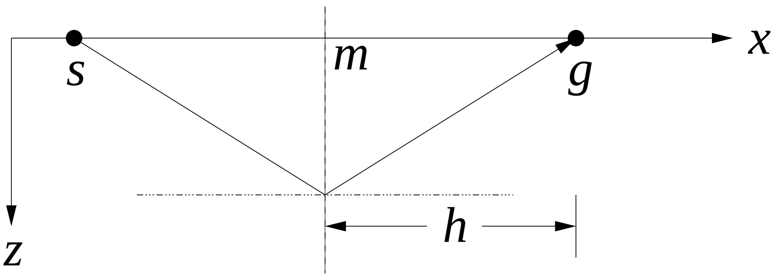

Figure 1.4 A simple reflection seismology experiment showing the shot point, geophone point, midpoint, and offset.

For most common processing techniques, it is best to think of the raw data in *midpoint–offset* (m, h) space, defined by the coordinate transformation [9]

$$m = \frac{g+s}{2} \quad \text{(midpoint)} \tag{1.1}$$

$$h = \frac{g-s}{2} \quad \text{(offset)}. \tag{1.2}$$

This geometry is depicted in Figure 1.4. Various slices through the raw data cube are termed *gathers* or *sections*. For example, the traces in a *common-midpoint gather*, $\{d_{m+h,m-h}(t);\ m \text{ constant}\}$, will all contain reflections from a common subsurface point if the medium consists of flat, uniform layers. A *common-shot gather*, $\{d_{s,g}(t);\ s \text{ constant}\}$, is the data recorded along the aperture for one source firing. The *zero-offset section*, $\{d_{m,m}(t)\}$, is just the monostatic acquisition case common to SAR (synthetic aperture radar), where the source and receiver are always colocated.

1.2.2 Simple sound-speed estimation

The speed of sound in rock is similar to that in bone—in the low thousands of meters per second for longitudinal waves. If the case of bone is excluded, the typical range of sound speeds encountered in the Earth is much greater than in biological tissue. Because of this, a constant-c assumption has never worked well in seismic imaging, and geophysicists have had to develop techniques for estimating the unknown speed of sound.

In the Earth, there is usually a gradual increase in wave speed with depth. Dramatic variations can be superimposed on this general trend; for example, salt deposits may well up through less-dense sedimentary layers, forming "salt domes" with much higher wave speed than their surroundings. Gas pockets, on the other hand, can create areas where sound propagates even more slowly than in water [10]. Usually, though, a layered pattern is seen, with interfaces that dip or become curved near interesting geological structure. This layered structure is the basis for many sound-speed-estimation techniques.

Consider a seismic experiment in which a flat reflector is located at depth z' in a medium with a constant speed of sound c. An echo will be received at time t depending on the shot and geophone positions as

$$t = \frac{2}{c}\sqrt{\left(\frac{g-s}{2}\right)^2 + z'^2} \tag{1.3}$$

$$= \frac{2}{c}\sqrt{h^2 + z'^2}, \tag{1.4}$$

which is a hyperbola in (h, t) space, that is, on a common-midpoint (CMP) gather. The same holds true in a point-scatterer model for scatterers directly below the midpoint, as may be seen in the left panel of Figure 1.5. The delay with increasing offset h is called *normal moveout*—NMO for short. In the hypothetical experiment, the location and shape of the hyperbola allow one to determine both c and z'. Note that this would not be possible with zero-offset data—the echo would arrive at the same time on every trace, leaving a depth-speed ambiguity. The prestack data contains redundancy which permits sound-speed estimation.

Normal moveout may be "corrected," and the CMP hyperbolas flattened, by an axis stretch [9],

$$t \to t', \text{ where } t' = \sqrt{t^2 + \frac{4h^2}{c^2}} \tag{1.5}$$

(Figure 1.5, right-hand side). Simple sound-speed analysis consists of performing this NMO correction for many trial values of c on selected CMP gathers in the prestack data set, then summing (stacking) over offset in each case. At some "stacking sound-speed" (which may be different at different t, if the true speed of sound in the medium is depth-varying) the hyperbolas will be maximally flat and the sum over offset will be maximized. Applying the NMO correction at the stacking sound-speed to every CMP gather in the data and then summing over offset has the additional benefits of (1) reducing noise, and (2) reducing the amount of data that must be processed in the imaging algorithm. The prestack data becomes poststack data and is treated like a zero-offset section.

Speed-of-sound estimation by NMO correction and stacking has serious limitations; for example, it gives the wrong answer when the reflecting layers are not flat. Modern practice is increasingly to take prestack data all the way into the imaging process. Some of the termi-

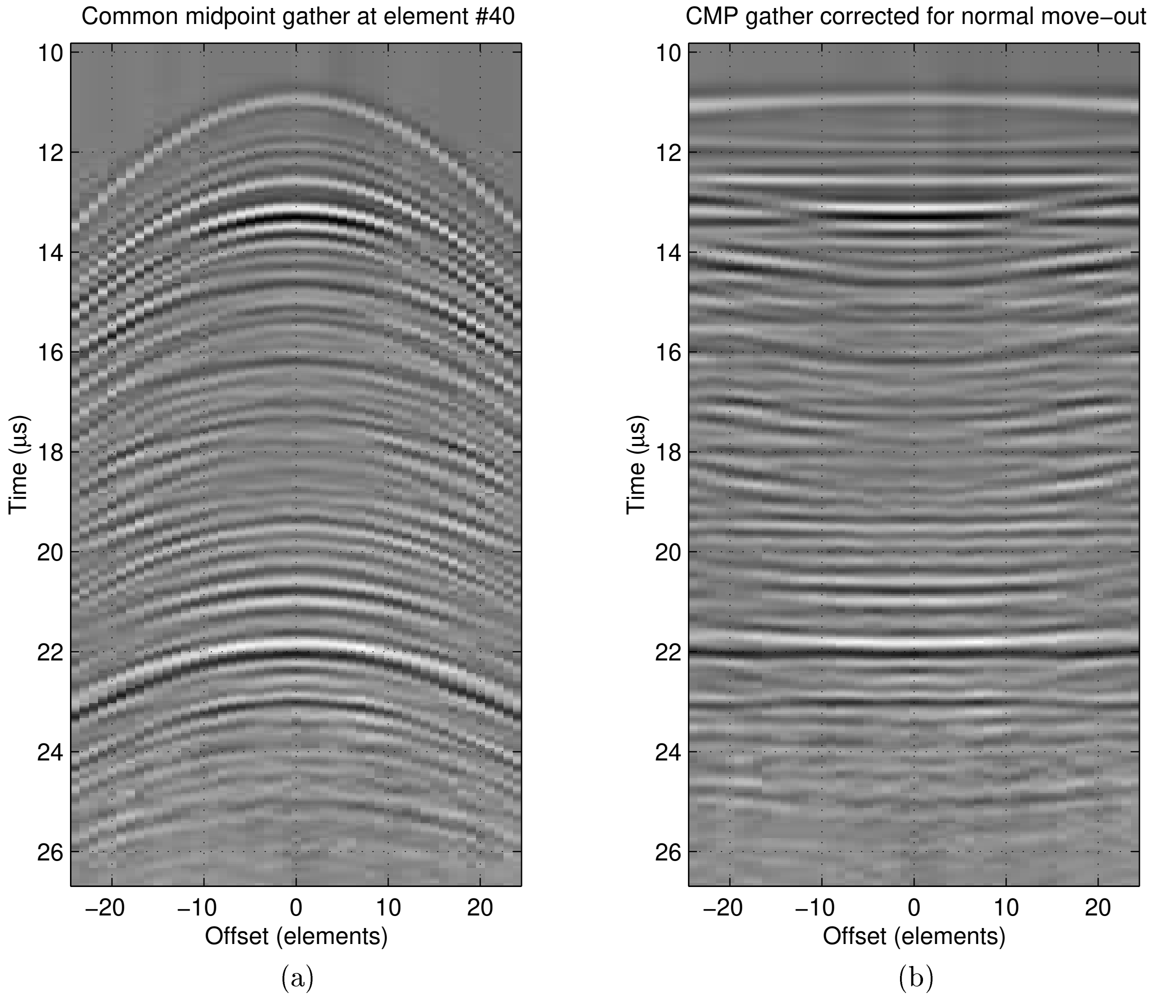

Figure 1.5 The normal moveout (NMO) correction applied to simulated ultrasound data from a speckle-producing region. (a) Common-midpoint gather. (b) After NMO correction, echoes from scatterers directly below the midpoint are perfectly flattened. Other echoes are overcorrected.

nology ("prestack" and "poststack") remains, however, and sound-speed estimation via NMO correction and stacking also illustrates why the redundancy in prestack data is important.

1.2.3 Imaging

Algorithms for image formation (termed *migration*) in geophysics are more varied than in ultrasound [11]. *Kirchoff migration* is a point-by-point correlation of the expected point target response with the raw echo data. For poststack or zero-offset data, this involves summing the raw data over hyperbolic curves; for prestack data, the summation is over the surface

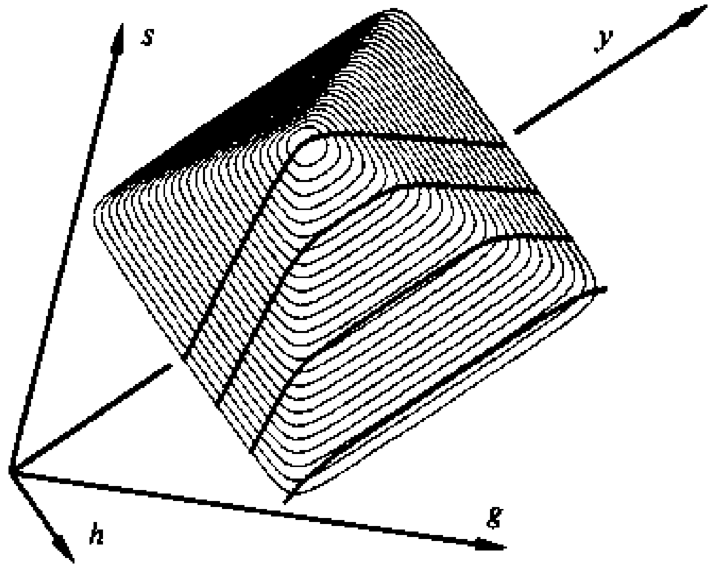

Figure 1.6 "Cheops' Pyramid"—the signature of a point target in prestack data. The time axis points down into the page. (From Claerbout, *Imaging the Earth's Interior*, p. 165 [9].)

in Figure 1.6, dubbed "Cheops' Pyramid" [9]. Kirchoff migration is essentially the same as traditional delay-and-sum beamforming in ultrasound except that it also takes into account the proper obliquity factor and phase predicted by wave theory. *Phase-shift migration* [12] is a more efficient, Fourier-domain approach which back-propagates the recorded wavefield in increments. The still more efficient *Stolt*, or *f-k migration* [13] focuses an entire image at once by interpolating the Fourier-domain image data from the Fourier-domain raw data; it was brought to the SAR community in [14] and called the *wavenumber* or $\omega - k$ algorithm. All of these, and many others, are still used in seismic processing because there is an inherent trade-off between speed and flexibility. Stolt migration is strictly valid only for media having a constant speed of sound c; phase-shift migration can handle a depth-variable $c(z)$ model; Kirchoff migration and other time-and-space-domain approaches work in a general $c(x, z)$ model (or, because most modern seismic surveys are three-dimensional, a $c(x, y, z)$ model). All of these migration algorithms have prestack and poststack variants.

From the foregoing, one can see that dynamic receive focusing in ultrasound is similar to a prestack, constant-c Kirchoff migration. The main differences are as follows: (1) Part of the stacking is wrapped up in the data acquisition—it happens when the array transmits a focused pulse using a particular set of delays. Thus, potentially useful information is discarded right

Table 1.1 Comparison of approximate physical parameters in exploration seismic and medical ultrasound imaging.

	Medical ultrasound	Seismic imaging
Frequency	$10^6 \sim 10^7$ Hz	$10^1 \sim 10^2$ Hz
Sound speed	1.5×10^3 m/s	1–5 $\times 10^3$ m/s
Maximum depth	10^{-1} m	10^4 m

at the outset. For example, the entire experiment would need to be repeated if an image at a different speed of sound were desired. (2) As discussed in Section 1.1, transmitting a focused pulse onto every resolution cell in the imaging area would be prohibitively time-consuming for real-time imaging, so short-cuts are used on the transmit side.

Commercial ultrasound scanners transmit only a limited number of times per frame, focusing on a subset of points spread out over the imaging area. The main reason why they do not acquire a prestack data set is noise: By forming a transmit focus during the acquisition process, one obtains a signal-to-noise (SNR) power gain of N, the number of array elements. Prestack data sets are sometimes used by the ultrasound research community, however. They are useful not only because of the redundancy in the data, but also because, neglecting the effect of noise, any imaging or aberration correction algorithm may be emulated using the prestack data set without the need to acquire more data. If algorithms requiring prestack data significantly outperform those which do not, there may be an increased motivation to collect prestack data in commercial scanners and solve the SNR problem by some combination of (1) transmitting longer pulses, with pulse compression on reception, and (2) averaging signals over multiple transmissions, possibly relaxing the high frame-rate requirement. In the ultrasound literature, prestack data sets are called *complete* data sets; this terminology will be used for the remainder of this dissertation.

1.3 Summary

In this introduction, the similarities and differences between medical ultrasound and seismic imaging have been noted. Table 1.1 presents one more comparison, in terms of the approximate frequencies, typical sound propagation speeds, and penetration depths in each discipline. Notice that although the frequencies differ by a factor of about 10^5, the media sound speeds and the ratio of depth to wavelength are very similar for both forms of acoustic imaging. This suggests that each application can learn from the other's accumulated knowledge and tools.

The rest of this dissertation is organized as follows: Chapter 2 provides additional background and a description of existing solutions to the aberration problem in medical ultrasound imaging. The next three chapters introduce three new approaches to the problem: Chapter 3 discusses the concept of rms sound speed and applies this idea to imaging situations where the aberrating structures can be approximated as planar layers. Chapter 4 describes a method for estimating angle-dependent aberration profiles directly from complete data sets using 2-D fan filters and a least-squares solution. Chapter 5 considers the application of dynamic programming to iterative focus-updating techniques. Chapter 6 details the experimental apparatus and procedures used for collecting complete data sets and presents performance comparisons for the algorithms described in the previous two chapters. Finally, Chapter 7 describes a novel Fourier-domain algorithm for efficient, three-dimensional imaging from cylindrical apertures. This work was previously published in the *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* [4].

CHAPTER 2

ABERRATION CORRECTION

Much attention has been given to improving the spatial resolution of ultrasound by using higher frequencies and/or increasing the size of the transducer aperture. The attenuation properties of tissue impose an upper limit on frequency, yet under this constraint, diffraction theory still permits a significant improvement in resolution over current systems.

The main obstacle to improved spatial resolution in medical ultrasound, and one problem which prevents its use in bone-shielded areas like the brain, is the variation of sound propagation speed in the body. As was discussed in the Introduction, the true speed of sound deviates from the assumed value by up to 10% in soft tissues and much more in bone. Although medical ultrasound imaging depends on the echoes generated by scattering from impedance discontinuities, including changes in the speed of sound, current systems assume a bulk homogeneity—that is, straight raypaths and a constant speed of sound at scales above a wavelength. When this assumption fails, imaging resolution suffers as the point spread function broadens. Higher sidelobes are especially damaging because a primary benchmark of an ultrasound system is its ability to resolve small, negative-contrast (scatter-free) regions surrounded by speckle. Thus, aberration seriously harms the diagnostic usefulness of ultrasound [15]. The effects of moderate aberration are illustrated in Figure 2.1.

Many researchers have studied the effects of aberration experimentally in various parts of the human body. A few areas of particular interest are the abdominal wall [16, 17], the female breast [18, 19], and the brain [20]. Of the three, the abdominal wall is considered the least challenging because the aberrating fat and muscle is close to the surface and oriented parallel to it. In the breast, sonography would be preferred over X-rays for cancer screening, but a large propagation-speed contrast between the irregularly distributed fatty and glandular tissues has kept ultrasound from achieving the necessary resolution. The brain has so far been almost completely off-limits to ultrasound due to the attenuation, variable thickness, and high speed of sound in the skull bone.

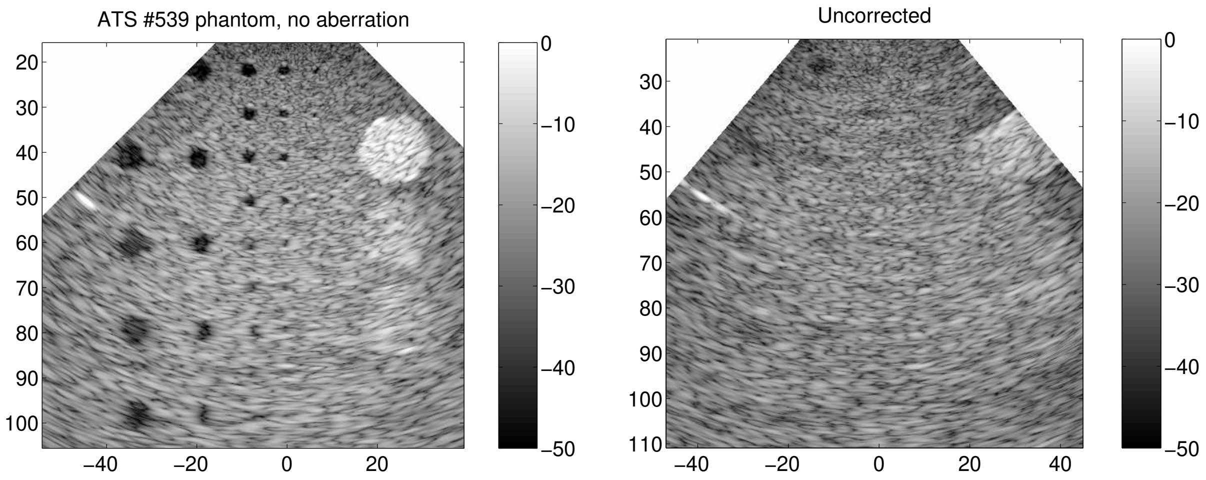

(a) (b)

Figure 2.1 Effects of aberration. (a) A "phantom" target designed to simulate the acoustical properties of biological tissue, imaged with a 64-element 1-D array at 2.6 MHz in a water bath. (b) The same target, imaged through an aberrating layer of rippled silicone rubber. (Axes are labeled in millimeters.)

Current approaches to aberration correction have been only mildly successful at improving resolution. In many parts of the body, tissue inhomogeneities still degrade the resolution of ultrasound images to that which can be obtained with relatively low frequencies and small apertures. An improved technique for aberration correction would have life-saving diagnostic implications. For example, it might become possible to image potentially cancerous lesions when they are as small as 1 mm in diameter. At this size, the lesion is too small for a blood supply to develop, and successful diagnosis can prevent a terrible disease. However, even a small improvement in resolution would be significant in clinical use.

2.1 Past Work

2.1.1 Screen methods

Virtually all of the published correction algorithms that would be practical *in vivo* (that is, with live subjects) are based on a screen model of tissue aberration. In it, the effects of aberration are modeled entirely by variable time delays on the received and transmitted signals due to an infinitesimally thin screen at the transducer surface. If these delays can be estimated correctly, perfect focusing is easily accomplished by adding compensating delays to the geometrically determined, hyperbolic focusing operators. Note that the terms *phase screen*, *phase aberration*, and *phase aberration correction* are common in the literature, despite the wideband nature of the pulses used in ultrasound.

The screen model has a long history of success in many fields. It has been applied to seismic imaging [21, 22], where a "weathered layer" of variable thickness and slow sound-propagation speed lies between the Earth's surface and deeper, higher-speed rock layers. The effects of the weathered layer are sometimes well-approximated by pure, angle-independent time delays; these are known as *surface-consistent statics*. In synthetic aperture radar, small deviations from a straight flight path may be corrected by applying the right set of phase shifts across the aperture. Similar ideas are used in radio astronomy to correct for phase aberration caused by the inhomogeneous atmosphere, which, over the typical small field of view, is angle-independent and in the near field.

In ultrasound, unlike the examples above, the screen model clearly fails to describe reality. Although the principal aberrators may be located close to the array in some cases (as in the abdominal wall and the skull), they are not "thin" with respect to the array-to-target distance. Nevertheless, one may consider a set of delay corrections to be valid for a small region of the imaging area called an *isoplanatic patch*. By calculating a different set of delay corrections for each isoplanatic patch, the entire image may be corrected. This still assumes that multiple scattering can be neglected, that is, that the focusing operators are single-valued. In this dissertation, only single-valued focusing operators will be considered; all of the proposed algorithms may be viewed as screen methods.

2.1.1.1 Iterative algorithms using a transmit focus

One class of aberration-correction algorithms based on a screen model is exemplified by the nearest-neighbor cross-correlation (NNCC) approach of Flax and O'Donnell [23, 24]. A focused transmit pulse is aimed at some region of interest, and the received signals on neighboring array elements are cross-correlated to estimate the relative time shifts. These time shifts are integrated across the array and taken as an improved focusing operator. Since the original transmit focus is degraded by uncompensated aberration, the procedure is iterated with the expectation of convergence to a correct focusing operator.

Algorithms of this type have the advantage of estimating aberration profiles from small regions of the target; thus, small isoplanatic patch size should not be a problem. If a single scatterer dominates the focal region, the received signals are highly correlated and the algorithm performs well. Unfortunately, dominant point scatterers are scarce in ultrasound images.

Speckle-generating tissues, which contain many random scatterers within the focal region of an ultrasound transducer, are far more common. In this case, the spatial correlation of the received signals can be predicted with a form of the van Cittert-Zernike theorem [25, 26], which treats the random scattering region at the focus as a source of incoherent radiation. For a linear, nonapodized aperture, the expected correlation between elements as a function of the distance separating them is a triangle function. Even for neighboring elements, in the ideal aberration-free case, decorrelation limits the accuracy of time-shift estimates. A more serious problem is that aberration drastically reduces the correlation between adjacent element signals. If it becomes severe enough, the algorithm may not converge. Some ideas to improve the performance of transmit-focus-based aberration correction algorithms are presented in Chapter 5.

2.1.1.2 Algorithms using common-midpoint signals

Another approach to delay estimation exploits the strong correlation between common-midpoint signals when single array elements are used for transmitting and receiving. As noted in Section 1.2, complete data sets contain redundant information about the target to be imaged; under a screen model of aberration, this information can be applied to find the unknown delays [21, 22]. Signals in a common-midpoint gather are redundant in the sense of sampling the same portion of the target's Fourier transform [27]. After moveout correction, the signals are highly correlated over a range of source-receiver offsets, regardless of the target composition. This is the "signal redundancy principle" of Li [28] and is clearly visible in the simulated common-midpoint gather of Figure 1.5.

In their published algorithms, both Rachlin [27] and Li [29] perform cross-correlations of common-midpoint signals. Rachlin uses data at many offsets to form a robust, over-determined matrix system, whereas Li only uses offsets of zero and one element. On the other hand, Li performs moveout correction on the common-midpoint gathers; this was neglected in Rachlin's far-field approximation. To date, there are no published comparisons between these algorithms and the NNCC-based algorithms.

Aberration correction based on common-midpoint signals is potentially more robust than methods which depend on echoes from a focal point. Because common-midpoint signals are highly correlated, even in the presence of aberration, there is no need for iteration and no "boot-strapping" problem as with methods that require a good transmit focus. Moreover, if many different offsets per midpoint are used for time-shift estimates, the problem becomes highly overdetermined, reducing error propagation across the array and allowing a robust, least-squares solution. The penalty for this robustness is reduced targeting ability: With no transmit focus, finding different aberration profiles for different steering angles is difficult. A new technique to overcome this problem is the subject of Chapter 4.

The translating apertures (TA) algorithm [30] may be seen as a compromise between the extremes of full-aperture transmission in NNCC algorithms and single-element transmission in the common-midpoint methods. In TA, a subaperture of the array transmits from two slightly shifted positions, focusing each time on a spot at the midpoint of the shifted apertures. In general, there may be a trade-off between the potentially more accurate time-shift estimates of the common-midpoint algorithms (due to over-determinancy) and the better SNR and targeting ability of the NNCC algorithms (due to a transmit focus).

2.1.1.3 Algorithms based on an image quality metric

A third fundamental approach to screen-based aberration correction is to adjust the element delays based on an image quality metric [31–33]. The integrated value of some power of the image intensity has been used in radar and astronomy as a quality metric [34]. In [31], a widening of the imaging point spread function is shown to decrease the average speckle brightness in the image. By sequentially adjusting the delay at each array element for maximum speckle brightness, image quality was improved in many cases. As with the previously mentioned algorithms, however, no performance comparison with other screen methods is available.

An extension to the basic screen model was proposed in [35], in which the received wavefield is back-propagated (assuming a homogeneous medium) some distance before the time delays are estimated. This sometimes improves the correlation between neighboring element signals. In principle, this extension could be combined with any of the screen methods.

Very few of the screen-based algorithms for aberration correction have gone on to be implemented in commercial scanners. For those that have, the resolution improvement in clinical situations has not been as dramatic as expected. There is reason to believe that the disappointing performance of screen algorithms in ultrasound is due to the difficulty of estimating the time shifts with sufficient accuracy, rather than any fundamental failure of the screen model. If the effects of multiple scattering may be neglected, that is, if the best focusing operators are single-valued, then the aberration problem in medical ultrasound is solvable by such an algorithm, at least in theory. It is interesting to note that an optimal choice of single-valued focusing operators has been shown to be effective in seismic imaging, where the aberrations are much more severe than in soft tissue [36].

2.1.2 Time reversal

Suppose a point source is buried deep inside a complicated, nonhomogeneous medium. A pulse from this source propagates through the complex structure and is recorded on an array at the surface. If the received array signals are time-reversed and retransmitted, the invariance of the wave equation under time reversal implies that the wavefront will refocus on the source point. This is time-reversal focusing, which implements the filter matched to the propagation channel [37].

Since attenuation is the same regardless of the propagation direction, time reversal requires that attenuation be negligible. Also, since the recording aperture intercepts only a portion of the energy from the point source, the resolution is, in most cases, naturally limited by the aperture size. However, if the medium exhibits strong multiple scattering, a superresolution phenomenon may be observed, allowing time-reversal focusing to beat the diffraction limit imposed by the recording aperture [38].

With iteration, time-reversal focusing can "lock on" to a passive point reflector, so it may be applied to medical ultrasound in the few cases where dominant point targets are available [39]. Short of implanting a target in the body, however, this cannot solve the general aberration problem.

2.1.3 Direct inversion

Aberration correction could be stated formally as the task of solving the scalar wave equation in inhomogeneous media [40],

$$\left[\nabla^2 + k^2(\vec{r})\right] p(\vec{r}) = 0, \tag{2.1}$$

written here for the time-harmonic pressure field at position $\vec{r}$. In homogeneous media, k is independent of position and the solution can be obtained as a superposition of Green's functions, but for inhomogeneous media, no closed-form solution is known [41].

2.1.3.1 Diffraction tomography

Under certain weak scattering conditions, approximate solutions to (2.1) may be obtained via the Born or Rytov approximations. These lead to a tomographic formulation in which the well-known projection-slice theorem is generalized to the case of diffracting waves. In this *Fourier diffraction theorem*, the target's Fourier data are obtained on curved paths which change with the projection angle [40, 41]. There have been some encouraging results using diffraction tomography in ultrasound (see [42], for example), and it may have a place in applications where the projections can be collected over a large range of angles, e.g., in breast scans. Both the Born and Rytov approximations, however, place severe limits on the allowable strength and physical extent of the inhomogeneities being imaged. Beyond these limits, the accuracy of the reconstruction degrades very quickly [40]. Diffraction tomography has received relatively little attention in the past decade.

2.1.3.2 Nonlinear inverse scattering

An infinite-series solution to the full inverse-scattering problem has been considered recently [43]. In principle, it may be possible to perform a complete inversion, from reflection data to the target's physical properties, without any prior knowledge of the propagation medium. The computational demands for this approach are extreme, and there is as of yet no proof of convergence. Simulations suggest that the inverse scattering series will converge when the sound-speed variation is less than about 14%, so in the long run it may be a solution to ultrasound imaging in soft tissues [44].

2.2 Focusing Operators and Notational Conventions

Table 2.1 lists the notation that will be used throughout the rest of this dissertation. The spatial coordinate system is fixed with respect to an array transducer, with $(x, z) = (0, 0)$ at the array center. The x, or *lateral*, axis is parallel to the array. The k^{th} array element (out of N) is located at $(x_k, 0)$; elements are assumed to be infinitesimal points unless otherwise indicated. The z, or *depth*, axis is the array normal, positive "downward." Where polar coordinates are used, θ, or *azimuth*, is measured from the z-axis, i.e., $\theta = \arctan(x/z)$.

In accordance with the assumptions of the screen model for aberration, the concept of focusing operators will be used extensively. The focusing-operator terminology in Table 2.1 is illustrated here with an example. Consider the one-way travel time for waves moving between the point (x', z') and a position x on the array aperture. In a medium with constant sound speed c,

$$t(x; x', z', c) = \frac{1}{c}\sqrt{(x - x')^2 + z'^2}, \tag{2.2}$$

which is one branch of a hyperbola in the (x, t) space of the recorded data:

$$c^2t^2 - (x - x')^2 = z'^2 \quad \text{(constant)}. \tag{2.3}$$

Table 2.1 Notation used in this dissertation.

c	Speed of sound
c_{rms}	Sound speed of the optimal constant-c focusing operator
$g(x,z)$	A target scatterer distribution or its image
x_k	Center of the k^{th} element
m	Midpoint
h	Offset
δ_t	Sampling period in time
δ_x	Sampling period in space; array pitch
$p(t)$	Transmitted pulse
$d_{s,g}(t)$	Signal from a complete data set: transmit element s and receive element g
$d_g(t;s)$	Common-shot gather: received signals after a single-element transmission
$d'_g(t)$	Once-focused data (advanced and summed over s); the echoes received after a focused transmit pulse
$d''_g(t)$	Shifted, once-focused data; $d'(g,t)$ that has been advanced again in preparation for summing over g
$t_k(x',z',c)$, $t(x;x',z',c)$	Focusing operator for a constant-sound-speed medium; delay at element k or as a continuous function across the aperture
τ_k, $\tau(x)$	Aberrating delay at the k^{th} element, or as a continuous function across the aperture
$\tilde{t}_k(x',z')$, $\tilde{t}(x;x',z')$	True focusing operator for the point (x',z'); delay at element k or as a continuous function across the aperture

Firing the array elements $\{x_k\}$ at corresponding times $\{-t_k\}$ will generate a wavefront which focuses on the point (x',z') at time zero. Similarly, an "exploding point" (x',z') at time zero will leave a hyperbolic signature in the data at the times $\{t_k\}$. Therefore, $t_k(x',z',c)$ is termed the *focusing operator* for the point (x',z'), given a medium of constant sound propagation speed c. Since we are interested mainly in large-scale aberrations with respect to the acoustic wavelength, this geometric acoustics approximation suffices for the discussion. In addition, amplitude variations will not be considered. While amplitude fluctuations across the aperture do change the imaging point spread function, the principal challenge in aberration correction is to add the elemental signals coherently.

Dynamic focusing on transmit and receive (a.k.a. prestack migration, or imaging with complete data) may now be defined as a sequence of two focusing steps [45–47]. Transmit focusing is expressed as the summation over transmit elements of time-shifted signals in the complete data set d. Neglecting any apodization,

$$d'_g(t) = \sum_{s=1}^{N} d_{s,g}(t + \tilde{t}_s(x', z')), \tag{2.4}$$

where $\tilde{t}_s(x', z')$ is the true focusing operator (possibly including the effects of time-shift aberration) for a target at (x', z'). The second focusing step is the summation over receive elements,

$$d''(t) = \sum_{g=1}^{N} d'_g(t + \tilde{t}_g(x', z')). \tag{2.5}$$

Notice that the same focusing operator is used on receive as was used on transmit. This follows from reciprocity: The delays necessary to focus a wavefront onto a point are the same as those observed following an "explosion" at that point. (The amplitudes are not generally the same, however.) If the focusing operator is correct, the image point corresponding to (x', z') is obtained by evaluating the doubly focused data at zero time:

$$g(x', z') = d''(0). \tag{2.6}$$

Since ideal focusing operators in a medium with a constant speed of sound are always hyperbolic, the *aberration profile* $\tau(x)$ (τ_k at element k) will be defined as the difference between a true focusing operator and a hyperbolic operator which approximates it:

$$\tau(x) = \tilde{t}(x; x', z') - t(x; x', z', c'). \tag{2.7}$$

Some aberration correction techniques are able to estimate τ values directly, but others provide only the estimated focusing operator $\tilde{t}$. To estimate the aberration profile in this case, the best-fit hyperbola must first be determined. Casting this as a least-squares problem, we have

$$\begin{aligned}
\min_{A,B,C} E &= \frac{1}{2}\sum_{k=1}^{N} \|\tilde{t}_k^2 - Ax_k^2 - Bx_k - C\|^2 && (2.8)\\
\frac{\partial E}{\partial A} &= -\sum_{k=1}^{N} (\tilde{t}_k^2 - Ax_k^2 - Bx_k - C)x_k^2 = 0 && (2.9)\\
\frac{\partial E}{\partial B} &= -\sum_{k=1}^{N} (\tilde{t}_k^2 - Ax_k^2 - Bx_k - C)x_k = 0 && (2.10)
\end{aligned}$$

$$\frac{\partial E}{\partial C} = -\sum_{k=1}^{N}(\tilde{t}_k^2 - Ax_k^2 - Bx_k - C) = 0 \tag{2.11}$$

leading to the matrix equation

$$\begin{bmatrix} \sum x_k^4 & \sum x_k^3 & \sum x_k^2 \\ \sum x_k^3 & \sum x_k^2 & \sum x_k \\ \sum x_k^2 & \sum x_k & N \end{bmatrix} \begin{bmatrix} A \\ B \\ C \end{bmatrix} = \begin{bmatrix} \sum x_k^2 \tilde{t}_k^2 \\ \sum x_k \tilde{t}_k^2 \\ \sum \tilde{t}_k^2 \end{bmatrix} \tag{2.12}$$

from which the parameters of the best-fit hyperbola are found as

$$c_{rms} = \sqrt{\frac{1}{A}} \tag{2.13}$$

$$x' = \frac{-B}{2A} \tag{2.14}$$

$$z' = \sqrt{\frac{C}{A} - \frac{B^2}{4A^2}}. \tag{2.15}$$

This procedure is also used to estimate the location of a focal point when only the [non-hyperbolic] focusing operator is known. The 3-by-3 matrix is poorly conditioned if all of the $\{x_k\}$ are far from 1, so it is important to scale the data appropriately before solving.

Sometimes we wish to find an operator's focal point location (x', z') under the assumption of some fixed sound speed c_0. This is useful, e.g., when an ensemble of focusing operators are to be used to form an image. Since the focal point location depends on the assumed speed of sound, and the operators generally have different best-fit hyperbola c values, all of the operators' positions should be translated to a common reference sound speed used for the (x, z) coordinates in the image. To do this, force $c = c_0$ in the solution by letting $A = 1/c_0^2$; then

$$\begin{bmatrix} \sum x_k^2 & \sum x_k \\ \sum x_k & N \end{bmatrix} \begin{bmatrix} B \\ C \end{bmatrix} = \begin{bmatrix} \sum x_k \tilde{t}_k^2 - A \sum x_k^3 \\ \sum \tilde{t}_k^2 - A \sum x_k^2 \end{bmatrix} \tag{2.16}$$

where x' and z' are obtained as before.

CHAPTER 3

FOCUSING WITH AN RMS SPEED OF SOUND

In Section 1.2.2, a simple sound-speed estimation technique was presented, based on the sound-speed-dependent flattening of common-midpoint gathers. A better way to use the same principle and avoid the restriction to horizontal reflectors is to use an imaging algorithm which operates directly on complete data sets (*prestack migration* in seismic terminology). The data are imaged repeatedly, each time with the assumption of a different constant speed of sound. The set of outputs at zero offset is like the set of outputs from stacked CMP gathers that were moveout-corrected for those same propagation speeds [9]. The coherent 2-D summation implicit in the imaging process discriminates in favor of features that are consistent with the imaging sound-speed assumption. This type of sound-speed estimation is known as *migration velocity analysis* in the geophysics community [48]. As discussed in the next section, the technique may be applicable to medical ultrasound imaging when the aberrating tissues have a layered structure—for example, when imaging the adult human brain through the skull.

3.1 $c(z)$ Models, RMS Sound-Speed, and Time Migration

In practical seismic imaging scenarios, the speed of sound is not constant, but "migration velocity analysis" is still useful. The sound-speed estimates obtained in this case do not correspond to the true speeds of sound in the medium, but they are related [49, 50]. This relationship may be understood by considering the case of a horizontally layered medium in which the speed of sound is a function of depth z only. Figure 3.1 shows one such model with a reflecting target beneath a layer of "bone." An expression for travel time versus offset h (compare Equation (1.4)) may be written as

$$t = \frac{2}{c_1}\sqrt{l_1^2 + d^2} + \frac{2}{c_2}\sqrt{l_2^2 + (h-d)^2}, \tag{3.1}$$

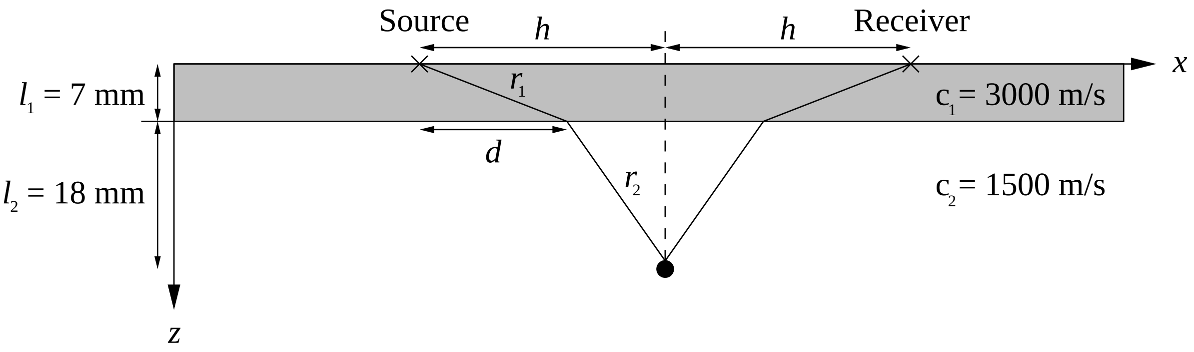

Figure 3.1 A simple two-layer skull-like model used to validate the rms sound-speed approximation.

where d parameterizes the set of possible ray-paths due to refraction. The correct ray-path may be determined by finding the minimum of the travel time with respect to d, following Fermat's principle of least time which he used to derive Snell's law,

$$\frac{dt}{dd} = \frac{2d}{c_1\sqrt{l_1^2 + d^2}} - \frac{2(h-d)}{c_2\sqrt{l_2^2 + (h-d)^2}} = 0. \tag{3.2}$$

Obtaining an expression for the travel time solely in terms of h, l_1, l_2, c_1, and c_2 is surprisingly difficult. It involves solving the quartic

$$(c_1^2 - c_2^2)d^4 - 2h(c_1^2 - c_2^2)d^3 + (c_1^2 l_1^2 - c_2^2 l_2^2 + h^2(c_1^2 - c_2^2))d^2 - 2hc_1^2 l_1^2 d + h^2 c_1^2 l_1^2 = 0, \tag{3.3}$$

which does not have a compact solution.

It was shown in [49] that for the general case of any number of flat layers of differing sound speeds and thicknesses (thus, in the limit, for any $c(z)$ medium),

$$t^2(h) = k_1 + k_2 h^2 + k_3 h^4 + k_4 h^6 + \cdots, \tag{3.4}$$

where $k_1 = t_0^2$ (the squared vertical two-way travel time) and the other $\{k_i\}$ are functions of $c(z)$. The second-order approximation is of special interest, because then the travel time is a hyperbola just like the constant-sound-speed case in Equation (1.4). It turns out k_2 has a nice interpretation in terms of the *rms speed of sound* c_{rms}:

$$k_2 = \frac{1}{c_{rms}^2} \quad \text{where} \quad c_{rms}^2 = \frac{1}{t_0}\int_0^{t_0} c^2(t)\,dt. \tag{3.5}$$

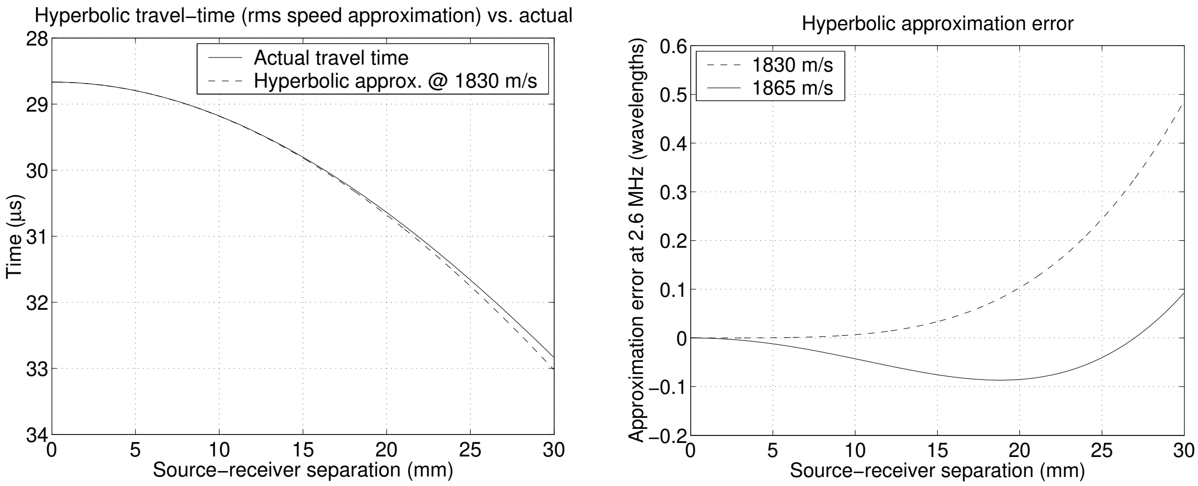

(a) (b)

Figure 3.2 The travel times for a scatterer beneath the "skull" are modeled almost perfectly by assuming propagation in a homogeneous medium with a higher propagation speed. (a) The actual travel time and its hyperbolic approximation. (b) The approximation error, in fractions of a wavelength, at c_{rms} = 1830 m/s and c_{best} = 1865 m/s.

The integral of squared wave speed is taken over the round trip of the vertical ray; for an N-layer model where t_n is the vertical two-way time and c_n is the speed of sound in the nth layer,

$$c_{rms}^2 = \frac{\sum_{n=1}^{N} c_n^2 t_n}{\sum_{n=1}^{N} t_n}. \tag{3.6}$$

The hyperbolic approximation is surprisingly accurate for source-receiver separations equal to or less than the target depth, a fact that was observed as long ago as 1955 [51]. Figure 3.2 shows the approximation (rms sound-speed of 1830 m/s) for the skull-like model in Figure 3.1. For this example and a 20-mm array aperture, the worst-case phase error at 2.6 MHz is only about 40°. Note that a best-fit hyperbola at $c_{best} \approx c_{rms}$ does even better by spreading the error more evenly over offset; for simplicity, the notation c_{rms} will still be used to refer to this best-fit hyperbola.

This result implies that, subject to the limitations of the hyperbolic (second-order) travel-time approximation, *targets in a horizontally layered $c(z)$ medium may be accurately focused by assuming a constant propagation speed of c_{rms} in the imaging algorithm.* Of course, the correct "constant speed" will vary from point to point in the image. This is the essence of *time migration* [11], so called because, since the true target depths are unknown, the "depth" axis of the final image is left in units of time.

This understanding has the potential to allow ultrasound imaging of the brain—a most difficult challenge for aberration correction. It is true that even in its flat regions, the skull bone is not a simple, homogeneous medium with a high speed of sound. There are anisotropic effects to consider, as well as the diploe, a layer of sponge-like bone with high attenuation. Promising experimental results were obtained, however, by Smith et al. [52]. They modeled the skull as a planar aberrating layer and derived static focusing corrections from prior knowledge of the layer's thickness and speed of sound, but they did not note the suitability of the hyperbolic approximation. This approximation is crucial, for it allows an imaging algorithm to determine c_{rms} blindly using some focusing metric, without the need for *a priori* knowledge of the intervening layers' sound speeds and thicknesses.

The range of aberration scenarios in tissue for which time migration will yield good results is unknown. Although generally associated with a $c(z)$ assumption, in seismic imaging it is commonly used to image strata which are not flat [11]. At some level of structural complexity and lateral variation in the speed of sound, time migration starts to fail, but this is subjective. In any case, it should always perform at least as well as the current practice of assuming one constant sound-speed (1540 m/s) for the entire image.

Note that this technique does not account for multiple scattering because it presumes a single-valued travel time from any reflecting target to the array. Thus, it can be viewed as a screen method, much like the previous approaches for aberration correction reviewed in Chapter 2. A key difference, however, is a reduction in the number of unknown variables from N (a time shift on each array element) to one (the assumed speed of sound for best focusing).

3.2 Measures of Focusing Quality

A practical time-migration algorithm starts with a series of images for a range of constant sound propagation speeds, the assumed speeds spanning the expected range of rms sound-speed for all the targets in the image. The final output is a composite image carved from the set of constant-sound-speed images, where the best speed of sound at each point is chosen based on some quality measure.

A number of published criteria for choosing the best speed of sound derive from the following conceptual model of the prestack migration process (imaging from complete data sets) [9]:

1. Move the receivers downward to some depth z.

2. Move the sources downward to the same depth as the receivers.

3. Evaluate the wavefield at zero source-receiver offset and zero time.

The first step is just a backpropagation of the received wavefield. By reciprocity, the second step is computationally very similar. The third step is known as the *imaging condition*: As the sources and receivers are pushed toward the depth of a reflector, the travel time for that reflector decreases until, at the reflector depth, the response is at zero time for a colocated source and receiver.

Normally, the nonzero-offset and nonzero-time data are discarded at the end of the imaging process. (In delay-and-sum beamforming, a.k.a. Kirchoff migration, nonzero-offset outputs are usually not calculated at all.) If these data are kept for analysis, various schemes for sound-speed selection become possible [48, 53, 54]. For example, one can test the correspondence of the zero-time and zero-offset conditions; a mismatch indicates faulty focusing [48].

In the present research of applying these ideas to medical ultrasound imaging, one technique which showed some promise in early simulations was to pick the imaging sound-speed that minimizes the energy at all offsets greater than some threshold h^*:

$$\hat{c}_{\mathrm{rms}} = \arg\min_{c} \int_{h^*}^{h_{\max}} [\text{Migrated prestack data}]^2 \, dh. \tag{3.7}$$

Although this worked well in bright regions of the image, it was prone to bad estimates in fainter or anechoic regions. Ultimately, a different, simpler scheme proved superior—maximizing the energy at zero offset at selected bright points in the image.

Picking the rms sound-speed based on a maximization of the energy at zero offset parallels the basic NMO-correction-and-stack technique. Applying this criterion at every point in the image would be a bad idea because ultrasound images have strong brightness contrasts at the edges of scatter-free regions like cysts; in the dark regions this criterion will select the wrong sound-speed due to energy from defocused bright targets nearby. Picking a set of equally spaced, bright targets throughout the image, determining c_{rms} at these points, and interpolating between them works better because the c_{rms} estimates are more reliable and the best focusing sound-speed is expected to vary smoothly over the image.

3.3 Time-Migration Algorithm

1. Form images, dynamically focused on transmit and receive, at a range of constant speeds of sound. Gain control should be applied so that the average brightness is similar in different parts of the image.

2. From the ensemble of constant-speed images, pick the brightest point (magnitude g^*; it may lie in any one of the images). Set c_{rms} for this point to the speed of sound assumed when forming its parent image.

3. Find the next-brightest point. If it is a sufficient distance (at least several resolution cells) from the last point, set its c_{rms} to the speed of sound assumed when forming its parent image, otherwise skip.

4. Continue picking points in decreasing order of brightness, subject to a minimum spacing constraint. Each time, use the already-established estimates to interpolate c_{rms} over the region of interest, and discard the new point if its c_{rms} estimate differs from the interpolated value by more than Δ (typically a few tens of meters per second).

5. Stop when all of the points brighter than αg^* have been scanned. α should be small enough to insure even coverage of the image by the selected points; a typical value is 0.1.

6. Use the final, interpolated map of c_{rms} to carve out the finished, composite image.

Dynamic focusing on both transmit and receive (that is, prestack migration) highlights reflections consistent with a given speed of sound. Rather than using a complete data set, it could be performed "online" as in conventional ultrasound beamforming; however, acquiring many complete scans at various c would be time-consuming.

3.3.1 Prestack Stolt migration

A computationally efficient algorithm for generating the constant-c images for the initial stage of time migration is a prestack variant of Stolt migration [13], known in SAR as the wavenumber or $\omega - k$ algorithm [14]. A simplified derivation follows: Let $d_{s,g}(t)$ be the raw prestack data in transmitter and receiver coordinates and time. Neglecting spherical spreading losses and angle-dependent scattering, the data may be expressed as

$$d_{s,g}(t) = \int\int g(x', z')p\left(t - \sqrt{(x'-s)^2 + z'^2} - \sqrt{(x'-g)^2 + z'^2}\right) dx'dz'. \tag{3.8}$$

Fourier transforming the time variable, we have

$$D_{s,g}(f) = P(f) \int\int g(x', z') \exp\left[-j2\pi f\left(\sqrt{(x'-s)^2 + z'^2} + \sqrt{(x'-g)^2 + z'^2}\right)\right] dx'dz'. \tag{3.9}$$

The spatial Fourier transforms $s \to f_s$ and $g \to f_g$ may be approximated using the principle of stationary phase [55] or derived in closed form, followed by a large-argument approximation to the Hankel function [56]. Only the phase terms are preserved here:

$$\begin{aligned} D(f_s, f_g, f) \quad \approx \quad & P(f) \int\int g(x', z') \exp\left[-j2\pi z'\left(\sqrt{\frac{f^2}{c^2} - f_s^2} + \sqrt{\frac{f^2}{c^2} - f_g^2}\right)\right] \\ & \times \exp(-j2\pi(f_s + f_g)x')dx'dz' \end{aligned} \tag{3.10}$$

$$D(f_s, f_g, f) \approx P(f)G(f_x, f_z), \tag{3.11}$$

where

$$f_z = \sqrt{\frac{f^2}{c^2} - f_s^2} + \sqrt{\frac{f^2}{c^2} - f_g^2} \tag{3.12}$$

and

$$f_x = f_s + f_g. \tag{3.13}$$

The image data in the Fourier domain are obtained on a Cartesian grid by interpolation from the Fourier-transformed raw data. The process is the same as in the monostatic (colocated source and receiver) wavenumber algorithm, except here the $f \to f_z$ interpolation is performed at every (f_s, f_g), i.e., N^2 times for data from an N-element array.

This algorithm is much more efficient than standard, delay-and-sum beamforming, but applying it to ultrasound data sets involves some complications. The main problem is that the final inverse Fourier transform will result in space-domain aliasing if there were scatterers located beyond the ends of the imaging array. This is almost always the case in clinical applications—the target area directly below the array is tiny compared to the angular sector that is typically imaged. The aliasing can be avoided by zero-padding the input data in space, on both the s and g axes, but this quickly becomes prohibitive in terms of memory usage. For this reason, the Fourier-domain algorithm is used only for the finite-difference simulation data presented in this chapter; a standard delay-and-sum algorithm is used elsewhere (Section 6.2.4).

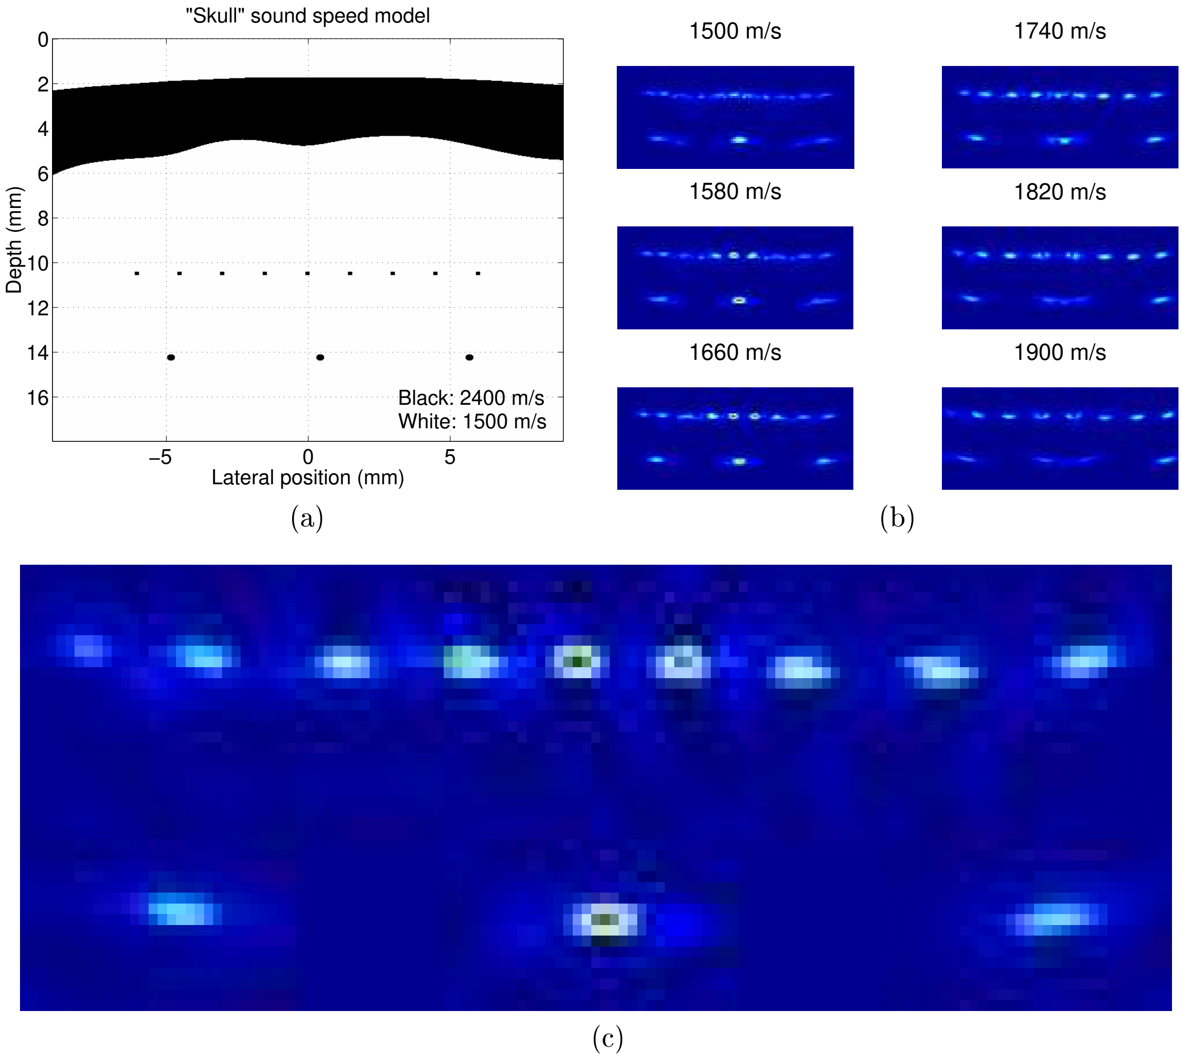

Figure 3.3 Time-migration results for a simple skull model simulated with a finite-difference code. (a) The skull model. (b) Images formed with various constant-sound-speed assumptions. (c) Corrected image using the time-migration algorithm. (Targets in the upper row are 2λ apart.)

3.4 Simulations

Figures 3.3 and 3.4 present simulation results designed to validate the rms sound-speed concept. The second-order finite-difference code `sufdmod2_pml` from version 35.3 of the *Seismic Unix* package [57] is used to simulate acoustic propagation through constant-density, variable-sound-speed media. The simulation domain is sampled at 30 μm, corresponding to $\lambda/25$ in soft tissues at the center frequency of 2 MHz. A map of the sound speed at every point in the domain serves as a "virtual target" and is constructed by a MATLAB program. An impulsive

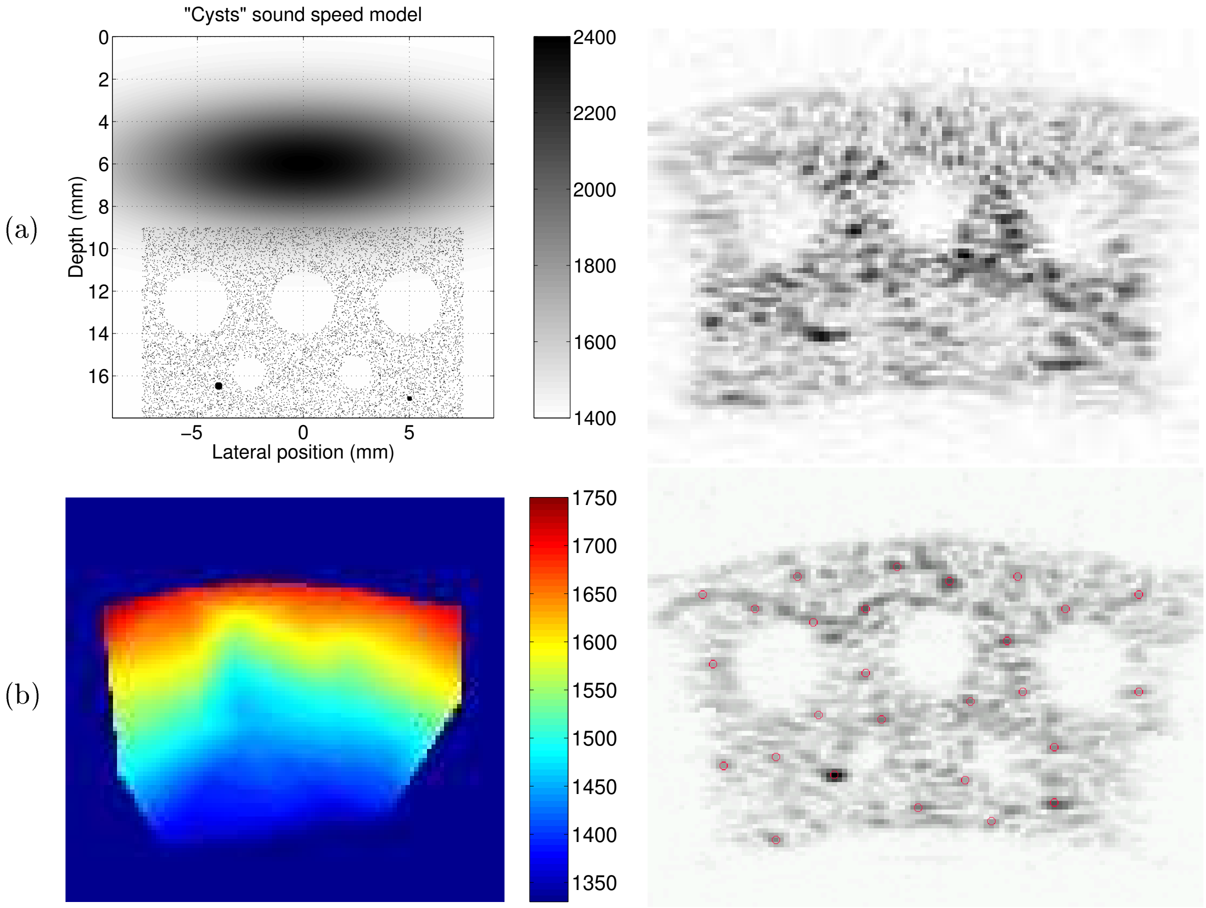

Figure 3.4 Time-migration results for a cysts-and-speckle model simulated with a finite-difference code. (a) Cysts and speckle model; conventional 1500 m/s image. (b) Map of c_{rms} from the time migration algorithm; corrected image with red circles indicating the bright spots picked for c_{rms} estimation. (The two small cysts are 2λ in diameter.)

spherical wave is launched from the top of the domain, which has perfectly absorbing boundary conditions on all sides. During each simulation run, the up-coming reflections are recorded at 64 positions along a 15-mm aperture. By running the simulation 64 times with 64 source positions across the aperture, a complete data set is obtained, approximating the data acquired with a real 64-element phased array. After these data are read into MATLAB, images are formed using the Fourier-domain algorithm in Section 3.3.1, each assuming a different, constant speed of sound. These are then combined into a single composite image using the map of c_{rms} obtained in the time-migration algorithm.

The model for a simulated brain imaging problem is shown at the left of Figure 3.3. In this simple approximation, the skull bone is modeled as a homogeneous medium with 2400 m/s speed of sound. The background speed is 1500 m/s. The nine points in the upper row are two wavelengths apart. In the first image, focused assuming a 1500 m/s speed of sound, the targets are not cleanly resolved—we cannot even tell how many there are. At higher speeds approaching 2000 m/s, different targets come into focus, indicating the success of the hyperbolic travel-time approximation. In the corrected composite image, all of the targets are reasonably well-focused, even though the sound-speed variation has a lateral component.

Figure 3.4 shows simulation results for a speckle-producing target with embedded scatter-free regions (cysts). The background speed of sound is 1400 m/s with a Gaussian-shaped, higher-speed anomaly above the target. In the conventional image, formed assuming a 1500 m/s speed of sound, the large upper cysts are poorly defined and the lower small cysts are almost invisible. In the corrected composite image, all of the cysts are much improved.

3.5 Experimental Results

In order to test these ideas on real data, the complete data set `ats_slab` was collected using a 64-element array transducer and a tissue-mimicking "phantom" target in a water tank. (See Chapter 6 for details of the experimental system and procedures.) An 18-mm-thick piece of silicone rubber (GE RTV615, 1100 m/s speed of sound) was placed between the array and the phantom target. The speed of sound in the phantom is 1450 m/s; we thus expect the rms focusing speeds to be near 1100 m/s at the top of the phantom, growing toward 1450 m/s at large depths.

Constant-sound-speed images of the upper 7 cm of the phantom (Figure 3.5) illustrate the depth-dependence of c_{rms}. The effect is particularly noticeable on the smaller cyst targets. The shallower cysts focus best near 1200 m/s, while the deepest cysts appear best at 1350 m/s.

The time-migration algorithm from Section 3.3 was applied to these images, but the results were inconsistent. Some targets provided "correct" sound-speed estimates, but numerous bad estimates harmed the quality of the c_{rms} map. This was still the case after numerous attempts with different Δ and α parameters. Another time-migration algorithm could probably be devised that would work better on this data set.

Because the primary motivation for this work was the brain-imaging problem, a number of complete data sets were collected using rat, pig, and sheep skulls. In most of the scans, wire targets were placed through the skull cavity; in several, the intact brain was scanned. Both medial and sagittal scans were performed at a variety of skull locations and stand-off distances.

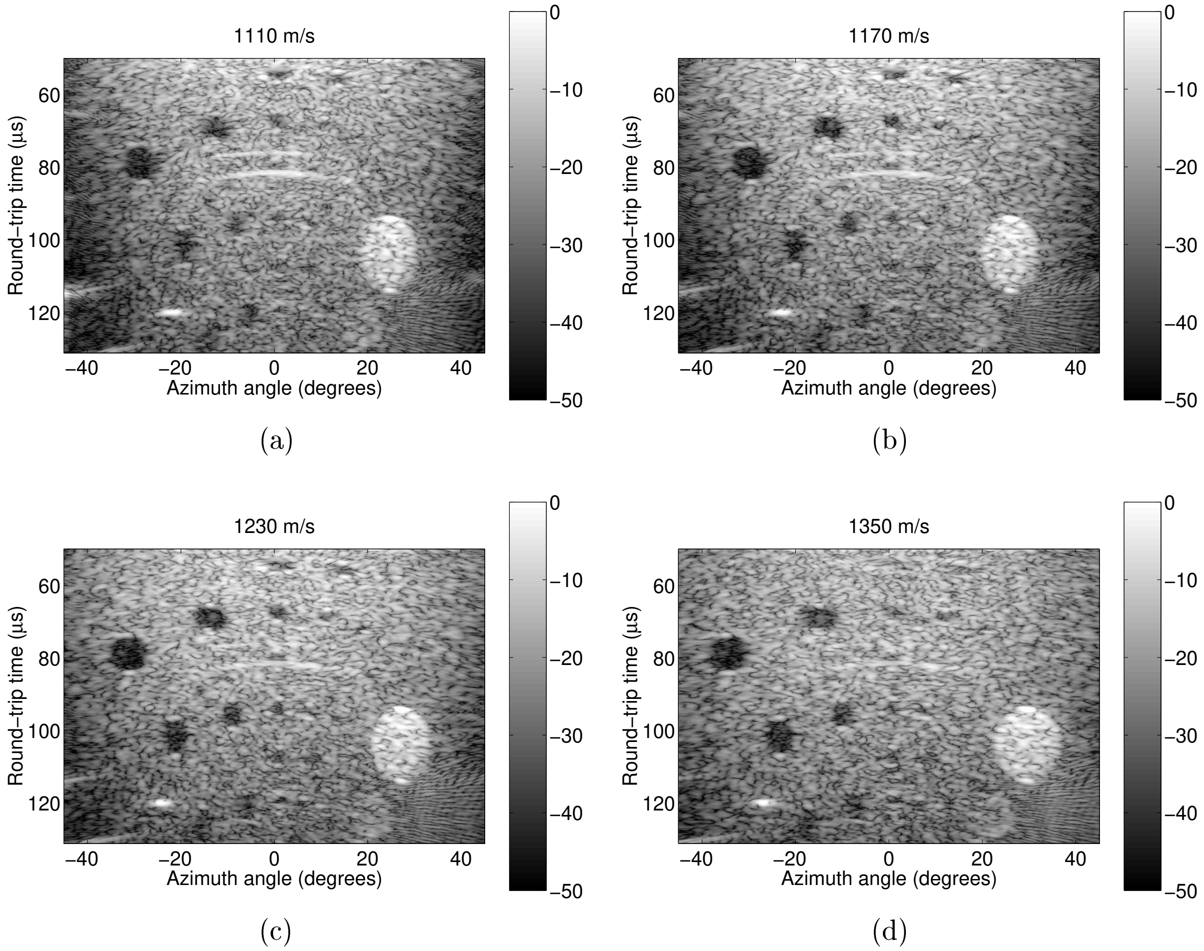

Figure 3.5 Constant-sound-speed images of data set `ats_slab` at (a) 1110 m/s, (b) 1170 m/s, (c) 1230 m/s, and (d) 1350 m/s reveal the c_{rms} trend due to an 18-mm-thick layer of silicone rubber (1100-m/s speed of sound) between the transducer and the target. (The cysts at 0° are 7λ in diameter, spaced 1 cm apart.)

Unfortunately, the results were inconclusive. In most cases, the intended target could not be found in the data. The reason may be that the 2.6-MHz transducer frequency is too high; the images in [52] were obtained using a 1.3-MHz transducer. The signal-to-noise ratio also suffered because only one element was fired at a time. In the one or two data sets that appeared to show a wire target, the wire appeared to focus best at or near the speed of sound in water. This may have been due either to the skull being too thin, or to a lensing effect of nonuniform skull thickness, or a combination of these.

3.6 Future Work

The time-migration algorithm based on an automatic selection of bright spots did not perform well on noisy, experimental data. More work is needed to find alternatives that will produce a more reliable c_{rms} map.

The applicability of time migration to medical ultrasound imaging is still uncertain. The best prospect remains brain imaging through the adult skull, but severe attenuation in the skull is the first obstacle that must be overcome. More experiments are needed at lower frequencies to establish how well the skull may be modeled as a flat layer of high sound-propagation speed.

CHAPTER 4

OVERDETERMINED LEAST-SQUARES ABERRATION ESTIMATES

As discussed in Chapter 2, the high degree of correlation between common-midpoint signals in a complete data set makes them attractive for estimating aberration profiles (Figure 4.1). The starting point for these calculations is a collection of estimated time-shifts between signals within the common-midpoint gathers. Usually, to estimate the time shift between two signals, the displacement from zero of the peak in the signals' cross-correlation function is used.

Algorithms based on common-midpoint signal analysis are less susceptible to misconvergence than those requiring an initial transmit focus because common-midpoint signals remain highly correlated even in the presence of aberration. Looking at it another way, the process of transmit focusing (summing $d_{s,g}(t)$ over s) discards potentially useful information; by analyzing the common-midpoint gathers before this step, the extra information can be exploited, leading to a more robust solution.

An automated method for estimating screen-model parameters from complete (prestack) data was first developed in the context of the *statics* problem in seismic imaging [21]. This work modeled the aberrating delays (the statics) as a rapidly varying, zero-mean random sequence. By placing the time-shift estimates into a matrix with shot and geophone position (s and g) axes, crude estimates of the aberrating delays were obtained by averaging individual rows or columns. Later, the problem was re-examined and cast into a least-squares framework [22].

Rachlin [27] was the first to recognize the potential of common-midpoint signals in medical ultrasound imaging. He explained the redundancy of common-midpoint signals using a far-field approximation in which these signals sample the same portion of the target's Fourier transform. His algorithm estimates the pair-wise time shifts between many signals in each common-midpoint gather. These time-shift estimates define a system of linear equations in the unknown aberrating delays $\{\tau_k\}$. In a complete data set collected using an N-element array, there are N midpoints at the element centers and $N-1$ midpoints half-way between elements.

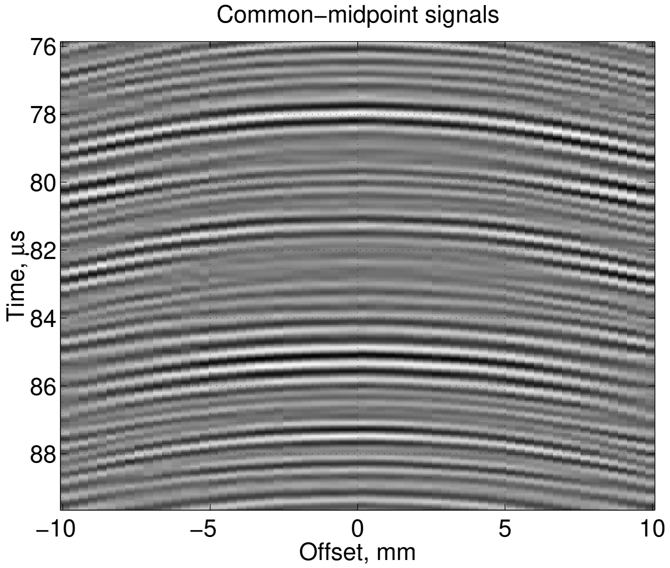

Figure 4.1 Common-midpoint signals are highly correlated.

For N even, at midpoint m there are

$$n_x(m) = \begin{cases} \lfloor m \rfloor & 1 \leq m \leq \frac{N+1}{2} \\ N + 1 - \lfloor m + 1/2 \rfloor & \frac{N}{2} + 1 \leq m \leq N \end{cases} \tag{4.1}$$

signals in the common-midpoint gather, from which

$$n_e(m) = \frac{n_x(m)(n_x(m) - 1)}{2} \tag{4.2}$$

independent, pair-wise time-shift estimates may be obtained. The total number of equations that may be written in the unknown τ_k's is then

$$N_e = \sum_{m=1}^{N} n_e(m) + \sum_{m=1}^{N-1} n_e(m + 1/2). \tag{4.3}$$

This simplifies to

$$\begin{aligned} N_e &= 4 \sum_{m=1}^{N/2} \frac{m(m-1)}{2} - \frac{N}{4}\left(\frac{N}{2} - 1\right) && (4.4) \\ &= \frac{N^3}{12} - \frac{N^2}{8} - \frac{N}{12}, && (4.5) \end{aligned}$$

or 21 328 equations for a 64-element array! Clearly, the system is highly overdetermined, even in the typical case of throwing out all but the highest-quality shift estimates.

Rachlin based his algorithm on a far-field assumption in which common-midpoint signals are theoretically identical. As we have seen in Section 1.2.2, this is not the case for reflectors in the near field—their echoes trace out hyperbolas. Normal moveout (NMO) correction,

$$t \to t', \text{ where } t' = \sqrt{t^2 + \frac{4h^2}{c^2}}, \tag{4.6}$$

should be applied before attempting to cross-correlate the signals. Even this is a simplification, however, because it assumes that reflections occur directly below the midpoint, as they would for a flat reflector. This is the major difficulty in applying these ideas correctly to medical ultrasound imaging: It is usually easy to find flat reflectors in seismic data, but they are rare in biological tissue. In the body, the bulk of the reflection energy comes from countless subwavelength scatterers. This is why medical ultrasound scanners can image a sector 90° wide and several inches deep from an aperture less than an inch long—echoes return to the transducer from every part of the medium reached by the transmitted wavefront.

Li recognized the need to extend NMO correction to scatterers at angles other than $\theta = 0°$, the array normal [28]. His generalized moveout correction for scatterers at angle θ, called the "dynamic near-field delay correction" in [28], defines the change of variable $t \to t'$, where

$$t' = \sqrt{\left(\frac{t}{2}\right)^2 + \frac{h^2}{c^2} - t\frac{h}{c}\sin\theta} + \sqrt{\left(\frac{t}{2}\right)^2 + \frac{h^2}{c^2} + t\frac{h}{c}\sin\theta}. \tag{4.7}$$

Unfortunately, this operation by itself does nothing to suppress echoes coming from other directions. These echoes will not be flattened in the "corrected" common-midpoint gather, and this will lead to bias in the time-shift estimates. The moveout correction grows rapidly with increasing offset, and so does the bias from uncorrected, off-axis scatterers. The problem is particularly severe for small-element arrays having wide radiation patterns in azimuth.

Due to the wide-angle nature of ultrasound imaging, it is frequently the case that one aberration profile cannot correct the entire image. This is the concept of an isoplanatic patch, discussed in Section 2.1.1. Multiple, independent profiles must be derived for different azimuth angles (and, less frequently, for different ranges). Thus, the lack of directionality in the processing discussed so far is a compound problem: First, how can the echoes from different azimuths be separated in order to derive independent aberration profiles? Second, even if the aberration profile is not θ-dependent, how can bias be eliminated in the time-shift estimates, given that the generalized moveout correction (Equation (4.7)) only corrects the common-midpoint gathers for one azimuth at a time, but echoes from many azimuths are present?

Li's near-field signal redundancy (NFSR) algorithm [28, 29] only considers the time shifts between signals of offset 0 and 1 element (after applying the moveout correction). This largely succeeds in avoiding time-shift estimation bias, and the resulting set of $N-2$ equations is sufficient to find the aberration profile to within an unknown linear tilt; however, the algorithm has to ignore the wealth of redundant information contained in the longer offsets. Furthermore, it has difficulty finding different aberration profiles at different azimuth angles when used with a small-element array.

A modified version of NFSR has been introduced [58, 59] using steered subarrays to provide some directionality. This is effective, but has the drawback of reducing the resolution of the estimated aberration profile because the elements of each subarray are assumed to share the same time shift. It also still precludes the use of a large offset spread to achieve a robust, overdetermined system.

In this chapter, the directionality problem is addressed in a new way using two-dimensional fan filters. After introducing fan filters as an aid to azimuth angle selection in the aberration-free case, the next section analyzes the effect of aberration on their performance. Fan filtering is then integrated into a new aberration-correction algorithm that uses common-midpoint signals at many offsets to derive robust, least-squares aberration profiles. Results in Chapter 6 show that this overdetermined, fan-filtering (OFF) algorithm is effective at estimating multiple, azimuth-dependent aberration profiles.

4.1 Angle Preselection Using 2-D Fan Filters

Consider a point source located at (x', z'), representing the reflection from a point target at time zero (Figure 4.2). Let $p(t)$ be the transmitted pulse and $\tau(x)$ be the time-shift aberration profile across the aperture. In a medium with a constant speed of sound, and neglecting amplitude factors, the aberrated signal received across the array is

$$\tilde{d}(x,t) = p\left[t - \frac{1}{c}\left((x-x')^2 + z'^2\right)^{1/2} - \tau(x)\right]. \tag{4.8}$$

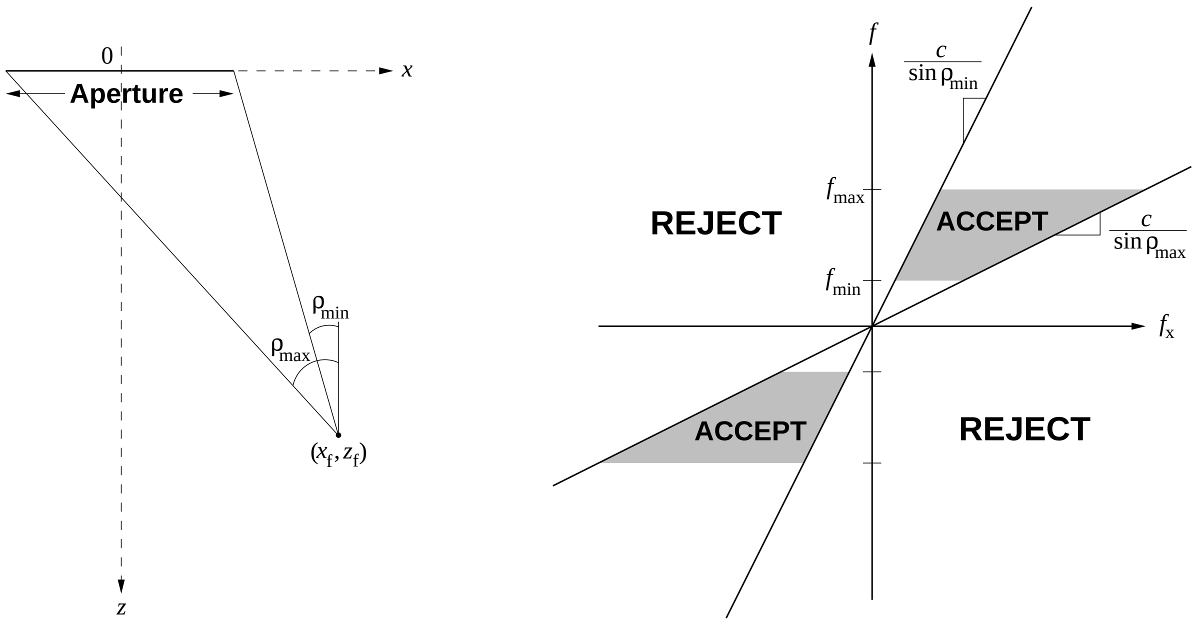

Figure 4.2 The signals received along a linear array from a point source or point reflector occupy a fan-shaped region in the two-dimensional Fourier domain.

Fourier transforming in time, we have

$$\tilde{D}(x,F) = P(F)\exp\left[\frac{-j2\pi F}{c}\left((x-x')^2+z'^2\right)^{1/2} - j2\pi F\tau(x)\right]. \tag{4.9}$$

Because only the phase of D depends on x, the instantaneous spatial frequency F_x may be found from the derivative of this phase with respect to x. For the case of no aberration,

$$\begin{aligned} F_x &= \frac{1}{2\pi}\frac{\partial \phi_D}{\partial x} && (4.10)\\ &= -\frac{F}{c}\frac{x-x'}{\left((x-x')^2+z'^2\right)^{1/2}} && (4.11)\\ &= \frac{F}{c}\sin\rho && (4.12) \end{aligned}$$

where ρ is defined in Figure 4.2. Thus, the unaberrated spectrum from a point target is contained within the fan-shaped regions determined by the temporal bandwidth and the spatial extent of the aperture. If δ_t and δ_x are the sampling period and array pitch, respectively, then using f and f_x as the frequency variables in the discrete-time and discrete-space Fourier domain,

$$f_x = f\frac{\delta_x \sin\rho}{\delta_t c}. \tag{4.13}$$

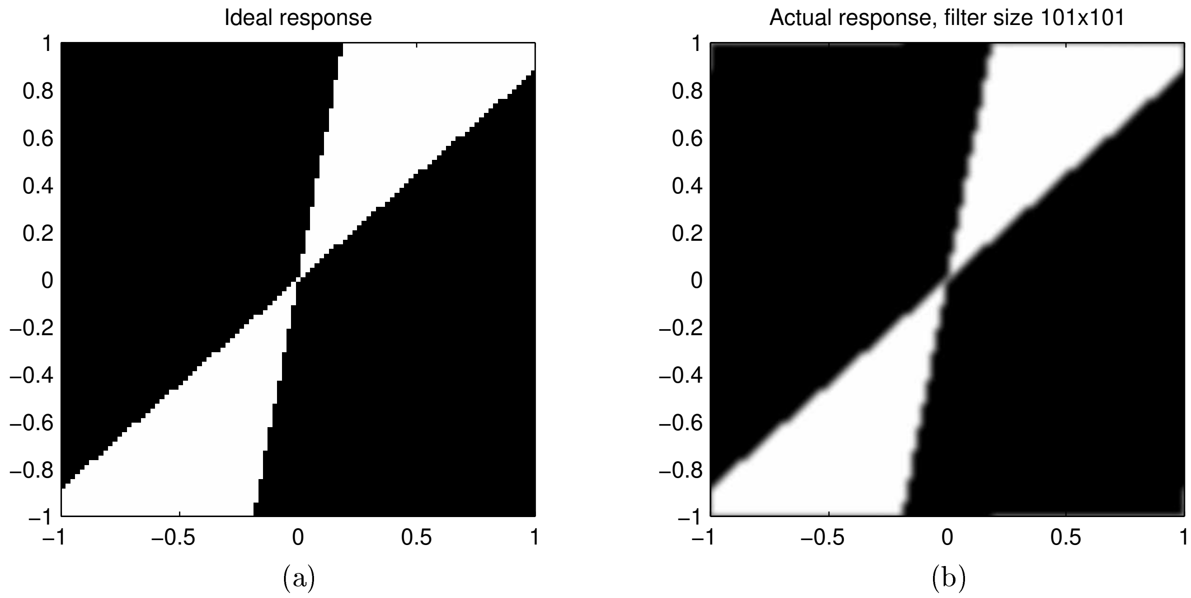

Figure 4.3 (a) Ideal and (b) actual FIR fan-filter responses in the digital frequency domain for removing all plane wave components except those from $\theta = 2°$ to $\theta = 12°$. White is one and black is zero on this linear scale.

4.1.1 Implementation

In geophysical imaging, where the subsurface is commonly modeled as a series of linear interfaces between rock layers, digital fan filters defined by Equation (4.13) are called *dip filters* [9]. This is because echoes from a specular reflecting interface come from an angle equal to the dipping angle of that interface. Ultrasound targets, by contrast, are better viewed as random distributions of point scatterers. Thus, fan filters can be used to discard the echoes from all but one image region, such as that being targeted for aberration correction.

While there are various techniques available for designing 2-D FIR fan filters, a design method based on circularly symmetric windows [60] was found to give good results for a wide range of filter angles (Figure 4.3).

Figure 4.4 illustrates the application of fan filters to data obtained experimentally from a "phantom" target (a tissue-mimicking physical model) with no aberration and moveout-corrected for two different correction angles: 0° and 40° from the array normal. The bottom panels are obtained as follows: First, fan filters are applied to the raw data by separately filtering each set of received signals across the array when transmitting from one element. N different fan filters are used on these N wavefields, each one designed to pass a fixed angular range centered on the angle determined by the transmitting element and the center of the scattering

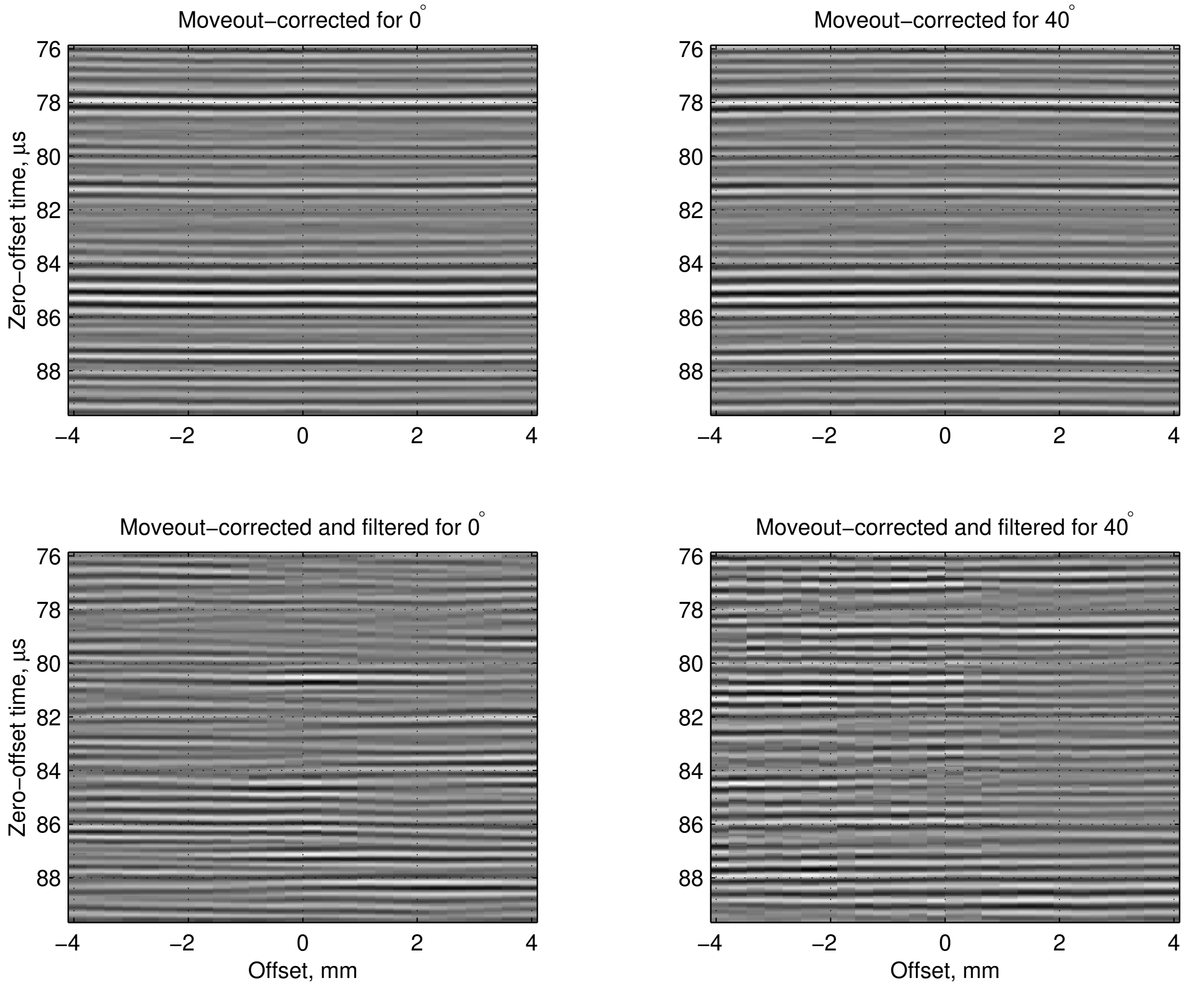

Figure 4.4 Common-midpoint signals corrected for scatterers at 0° and 40° from the array normal, and constructed from filtered and unfiltered wavefields. Only when using the appropriate 2-D FIR fan filters are different echoes observed at each angle.

region. Then, a set of common-midpoint signals are selected from the filtered complete data set. Finally, the generalized moveout correction is applied.

Notice that in the unfiltered, upper panels, the same echoes are dominant; the effect of the moveout correction is merely to flatten some echoes and under- or overcorrect others. After passing the complete data set through fan filters targeting scatterers at 0° and 40° with an angular bandwidth of 10°, the common-midpoint signals appear at first to be of lower quality; notice, however, that *different echoes are dominant* at each angle. If angle-dependent aberration were present, different time shifts would be observed in each panel.

4.1.2 Effect of aberration

The claims made for the benefits of fan filtering have thus far neglected the effects of aberration. Now suppose that $\tau(x)$ is not zero. Letting $\tilde{D}$ and D denote the received signals' 2-D spectra in the aberrated and unaberrated cases, respectively, application of the frequency-convolution property yields

$$\tilde{D}(F_x, F) = D(F_x, F) * U(F_x, F), \tag{4.14}$$

where

$$U(F_x, F) = \mathcal{F}_x \left[e^{-j2\pi F \tau(x)} \right] \tag{4.15}$$

and the convolution is in F_x. The effect of aberration, then, is to convolve the original spectrum in spatial frequency with that of a phase-modulated signal. The spectrum of phase-modulated signals cannot, in general, be found analytically; some approximate results are known, however [61].

Suppose the aberration profile, $\tau(x)$, can be modeled as a Gaussian random process obtained by passing white Gaussian noise through the low-pass filter defined by

$$h(x) = e^{-\pi a^2 x^2}. \tag{4.16}$$

The autocorrelation function of $\tau(x)$ is then

$$r_\tau(x) = \frac{1}{a\sqrt{2}} e^{-\pi a^2 x^2/2} \tag{4.17}$$

with power spectral density (PSD)

$$T(F_x) = \frac{1}{a^2} e^{-2\pi F_x^2/a^2}. \tag{4.18}$$

Defining the aberrator correlation length x_0 as the width of $r_\tau(x)$ at half maximum, we have

$$a = \frac{1}{x_0}\sqrt{\frac{8\ln 2}{\pi}}. \tag{4.19}$$

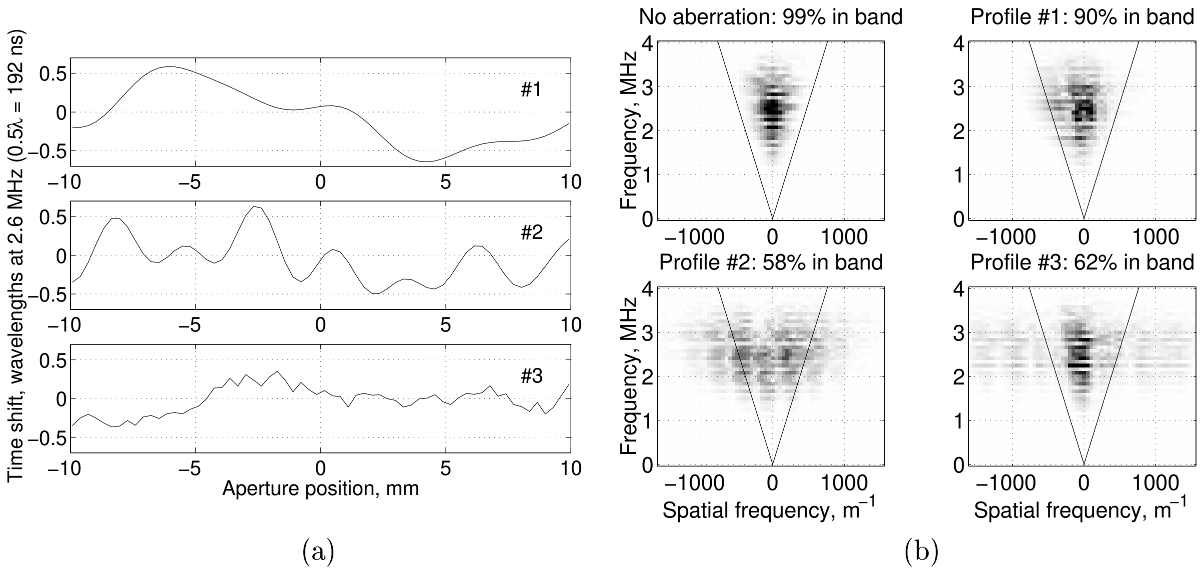

Figure 4.5 The effect of aberration on a target's spectrum in the 2-D Fourier domain is to convolve the unaberrated spectrum in spatial frequency with a spreading function—the spectrum of a sinusoid phase-modulated by the aberration profile. Here, the effect of three aberration profiles (a) is simulated in the 2-D Fourier domain (b). The spreading becomes worse as the aberration profile fluctuates more rapidly and as its amplitude increases. The angled lines show the theoretical bounds derived in Section 4.1 in the absence of aberration.

Following the derivation in [61] for phase modulation by a Gaussian random process, define the mean-square bandwidth of the aberration process $\tau(x)$:

$$W_T = \sqrt{\frac{\int_{-\infty}^{\infty} F_x^2 T(F_x) dF_x}{\int_{-\infty}^{\infty} T(F_x) dF_x}}. \tag{4.20}$$

The mean-square bandwidth of the phase-modulated signal U in spatial frequency is then

$$W_U = \frac{F}{x_0}\sqrt{\frac{8\ln 2}{L}\int_{-L/2}^{L/2} \tau^2(x)dx}, \tag{4.21}$$

where L is the aperture length. The spatial-frequency spreading effect of aberration is thus inversely proportional to the aberrator correlation length and proportional to the rms aberration.

There is qualitative agreement between this result and the examples considered in Figure 4.5. Simulated speckle data were collected from a 10°-wide swath of scatterers, broadside to and 55 mm from a 64-element array. When no aberration is present, the spectrum of the received signals falls cleanly within the bounds dictated by the recording geometry. A smooth, long-correlation-length aberration profile (profile 1) broadens the spectrum only slightly. Profile 2, with similar amplitude but a shorter correlation length, smears more energy outside the boundary, but more than half of the signal is still inside. Under profile 3, with its rapid, low-amplitude fluctuations, most of the signal remains within the original band while the remainder is spread widely across the rest of the spatial frequencies.

4.2 Overdetermined, Fan-Filtering (OFF) Algorithm

For data from an N-element array, let $d_{s,g}(t)$ denote the signal received on element g when firing element s. $d_g(t; s)$ is the set of N signals received on all elements after firing element s. x_s is the position of array element s. h_{min} and h_{max} are the minimum and maximum offsets to consider at each midpoint (subject to the array edges for midpoints near the ends). Figure 4.6 presents the steps needed to derive the aberration profile for a region of interest centered at $(x_{\text{ROI}}, z_{\text{ROI}})$.

Different 2-D fan filters are applied to each signal set $d_g(t; s)$. The angular acceptance range is centered for the zero-offset signal $d_{s,s}(t)$ from the region of interest, and $\Delta\rho$ is chosen at least large enough to pass signals at h_{max}. Following this, at each midpoint, the common-midpoint signals are corrected for geometric path-length differences according to Equation (4.7). Within the corrected common-midpoint signals, all pairwise cross-correlations are performed. Time-shift estimates from those pairs having cross-correlation coefficients above some threshold x_{thresh} are incorporated into the overdetermined linear system $\mathbf{A}\tau = \mathbf{b}$. After every common-midpoint set has been processed, the $\mathbf{A}$ matrix is regularized using the singular-value decomposition, and the least-squares solution for the aberration profile τ is obtained.

The maximum cross-correlation lag, Δ_{max}, must be at least twice the expected peak-to-peak amplitude of τ. (To see why, imagine a cross-correlation between signals at small and large offsets, where the midpoint is located at a sharp dip in the aberration profile and the maximum offset reaches to the two adjacent peaks.) To prevent bad time-shift estimates due to "cycle skipping," however, Δ_{max} should be kept as small as possible. Fortunately, a small number of bad estimates in this highly overdetermined system will not perturb the least-squares solution very much.

For cases where Δ_{max} must be set high enough that cycle-skipping in the time-shift estimates becomes problematic, the following correction procedure has been used: First, calculate the error vector $\mathbf{e}$ and sort it by magnitude. Starting with the largest element of $\mathbf{e}$, add or subtract $1/f_c$ from the corresponding element of $\mathbf{b}$ for negative or positive error, respectively. Continue

For $s = 1 \ldots N$ {

  Filter $d_g(t; s)$ with fan filter passing angles $\rho \pm \Delta\rho$, where $\rho = \arctan\left(\frac{x_{\mathrm{ROI}} - x_s}{z_{\mathrm{ROI}}}\right)$.

}

For $m = 1 \ldots N$ {

  For $h = h_{min} \ldots h_{max}$ {

    Interpolate $d_{m+h,m-h}(t) \;\rightarrow\; d_{m+h,m-h}(t')$ (Equation (4.7)).

  }

  For $h_1 = h_{min} \ldots h_{max}$ {

    For $h_2 = (h_1 + 1) \ldots h_{max}$ {

      Cross-correlate $d_{m+h_1,m-h_1}(t')$ and $d_{m+h_2,m-h_2}(t')$ at lags $-\Delta_{max} \ldots \Delta_{max}$, yielding shift estimate $\Delta_{h_1 h_2}$.

      If peak correlation $> x_{thresh}$ {

$$\mathbf{A} = \begin{array}{c} \begin{array}{ccccccccccccc} & \scriptstyle(m-h_2) & & & \scriptstyle(m-h_1) & & & \scriptstyle(m+h_1) & & & \scriptstyle(m+h_2) & & \end{array} \\ \left[\begin{array}{ccccccccccccc} & & & & & & \mathbf{A} & & & & & & \\ \cdots\ 0 & -1 & 0 & \cdots & 0 & 1 & 0 & \cdots\ 0 & 1 & 0\ \cdots\ 0 & -1 & 0 & \cdots \end{array}\right] \end{array}$$

$$\mathbf{b} = \begin{bmatrix} \mathbf{b} \\ \Delta_{h_1 h_2} \end{bmatrix}$$

      }

    }

  }

}

Regularize $\mathbf{A}$ and solve $\mathbf{A}\tau = \mathbf{b}$ using the SVD.

Figure 4.6 The OFF algorithm: Overdetermined, least-squares solution for the aberration profile at $(x_{\mathrm{ROI}}, z_{\mathrm{ROI}})$.

for elements of $\mathbf{e}$ larger than β times the mean error, then stop, calculate an updated error vector, and start the process over. Repeat until the mean error stops decreasing.

For the results presented in Chapter 6, $\Delta\rho$ was 10° and x_{thresh} was 0.5. The offset range h_{min} to h_{max} was 1 to 12 elements. (Zero-offset data was excluded for simplicity because the data-acquisition procedure routes the N zero-offset, pulse-echo signals through different analog hardware than the other $N^2 - N$ signals.) For the $\mathbf{b}$-adjustment procedure in the preceding paragraph, β was 3.

Figure 4.7 shows a typical segment of the $\mathbf{A}$ matrix. Using the above parameters and data from a 64-element array, the number of rows is commonly 2000 or more.

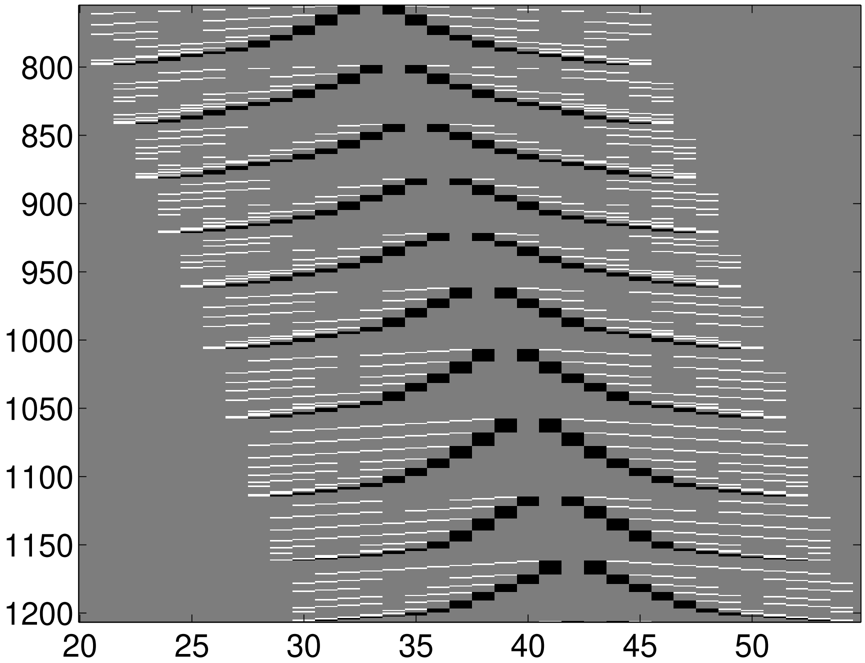

Figure 4.7 A portion of the **A** matrix. Each row corresponds to a time-shift estimate which defines an equation in four unknown aberrating delays. The pattern is due to the ordering of the equations by midpoint; equations for 10 midpoints form this segment.

4.3 Discussion and Future Research

The proposed algorithm is very similar to the approach described by Taner et al. for the seismic statics problem [22], the key addition being the angle-selectivity afforded by fan filtering. Like the linear system described in that work, the **A** matrix in OFF turns out to be rank-deficient. In this formulation, the rank is always $N-2$, implying that the solution for τ is indeterminate by a linear component. Although surprising at first, this fact has an intuitive explanation.

First, note that because the system is built up from pair-wise relative shift estimates between signals, the solution for $\{\tau_k, k = 1 \ldots N\}$ is clearly insensitive to an overall constant shift. Thus, $\{\tau_k + C,\ k = 1 \ldots N\}$ is also a solution for any constant C because this does not change the relative time-shifts between signals. Now consider any single row in **A**. Because the pair of signals has the same midpoint, the contribution of any linear component in τ is canceled out by the symmetry of the equation—the displacement of transmitters is equal and opposite to the displacement of receivers. (If we were not restricted to common-midpoint signals, this would not be true, and **A** would have rank $N-1$.)

The rank-deficiency of $\mathbf{A}$ is not a problem in practice. All of the robustness benefits from a highly overdetermined system still apply. Using the singular-value decomposition yields the minimum-norm solution, which will be the true aberration profile with any linear component subtracted out. To first order, a linear tilt of the focusing delays is equivalent to a steering-angle change, and the addition of a constant is equivalent to a change in depth. Because these terms are always small (at worst, of the same order as the peak-to-peak aberration), the quality of focusing will not be affected.

Experimental results in Chapter 6 indicate that the greatest error in the derived aberration profiles occurs near the ends of the aperture. This is to be expected, since there are fewer offsets and hence fewer equations available to constrain the profile there. It may be beneficial, then, to employ a larger aperture for aberration correction, then throw away a handful of elements at each end and image using the central part of the aperture where the aberration profile is more accurate.

As noted in the previous section, the least-squares solution is robust to a small number of bad time-shift estimates. Despite this, experiments suggest that this algorithm's performance is ultimately limited by "cycle-skipping"—one-period errors in the time-shift estimates. As aberration becomes more severe, the maximum cross-correlation lag (Δ_{max}) must be increased, and one-period errors become more likely. If the solution scheme could be modified to exploit the fact that most of the time-shift estimates are still correct modulo one period, the overall performance might be significantly improved. The post-processing, $\mathbf{b}$-adjustment step is an improvement, but not an optimal solution. A different strategy would be to view this as a nonlinear optimization problem. In the context of the seismic statics problem, a form of simulated annealing has been shown to reach good solutions when cycle-skipping would normally be troublesome [62].

The effect of iteration on OFF remains to be investigated. In cases of severe aberration, a partially correct estimated profile from the first iteration could be applied to the raw data prior to the second iteration. Correcting some or most of the aberration prior to fan filtering would reduce the spreading effect of aberration in spatial frequency, possibly leading to a better aberration estimate.

The ability of OFF to find an approximate solution even in the presence of severe aberration suggests its use as a bootstrapping method for other algorithms, particularly those requiring a transmit focus. The degree to which OFF might complement iterative focusing approaches should be explored further.

CHAPTER 5

FOCUSING-OPERATOR UPDATING VIA DYNAMIC PROGRAMMING

Recent work in the geophysics community has recognized the utility of viewing the prestack migration process as a sequence of two focusing steps: focusing on transmission followed by focusing on reception[1] [45–47]. The intermediate result of this process is a collection of N received signals resulting from a focused transmit pulse. (Note that with a complete data set, the two focusing steps can be interchanged; the intermediate result then becomes the set of beam-sum signals resulting from N single-element transmissions.) These signals have been widely used for aberration correction in ultrasound, e.g., in the nearest-neighbor cross-correlation (NNCC) algorithm [23]. In [45], the intermediate signals are called a common-focus-point (CFP) gather; the two focusing steps are also formalized in terms of focusing operators. It will be seen that this formalism yields a better understanding of aberration correction algorithms utilizing a transmit focus and also provides ideas for improvements.

5.1 Focusing Operators and the Principle of Equal Travel Time

As explained in Section 2.2, transmit focusing is expressed as the summation over transmit elements of time-shifted signals in the complete data set d. Neglecting amplitude factors,

$$d'_g(t) = \sum_{s=1}^{N} d_{s,g}(t + \tilde{t}_s(x', z')), \tag{5.1}$$

where $d'_g(t)$ is the CFP gather and $\tilde{t}_s(x', z')$ is the true focusing operator (one-way travel times) for a target at (x', z'). ($\tilde{t}_s(x', z')$ equals the hyperbola $t_s(x', z', c)$ in the absence of aberration.)

[1] This may seem strange, given that ultrasound researchers have long tended to view the imaging process in precisely this way. Geophysicists, however, have used other equivalent but conceptually different models—for example, "moving" sources and receivers via backpropagation to the same depth, then evaluating the result at coincident source and receiver location and zero time.

The second focusing step, this time advancing and summing the signals in the CFP gather, yields

$$d''(t) = \sum_{g=1}^{N} d'_g(t + \tilde{t}_g(x', z')). \tag{5.2}$$

If the focusing operator was correct, the image point (x', z') is given by $d''(0)$. Notice again that the same focusing operator is used for both focusing operations. By reciprocity, the operator that focuses on (x', z') at $t = 0$ is the same as the operator that coherently sums the arrivals from (x', z') following an "explosion" at $t = 0$. (The amplitudes are not generally the same, however.) If we define $d''_g(t)$ to be the shifted signals from the CFP gather prior to summation,

$$d''_g(t) = d'_g(t + \tilde{t}_g(x', z')), \tag{5.3}$$

then a properly focused event (a point reflector, for example) will appear as a straight line at $t = 0$. This is termed a differential time-shift (DTS) panel in [47]. Note that defining these focusing steps as time shifts rather than convolutions presumes single-valued focusing operators and sharp transmit pulses.

Consider the example of a point reflector focused using two different operators—one correct and one assuming a speed of sound that is 150 m/s too high (Figure 5.1). Using the correct operator, the echo in $d'_g(t)$ (the common-focus-point gather) from the point reflector exactly matches the focusing operator; upon shifting for the second receive-focusing step, this echo is perfectly flattened at $t = 0$ and ready for coherent summation. Using the incorrect operator, the echo and the operator do not match and the response is displaced from $t = 0$ in $d''_g(t)$. Notice that the correct focusing operator lies somewhere between the incorrect operator and the received echo in the CFP gather. In [46] it is shown that this will always be the case, leading to an iterative procedure for operator updating.

5.1.1 Focusing operator updating

The true focusing operator for a given target is always bounded in the CFP gather by the target's echo and the [generally incorrect] operator being used. A reasonable updating strategy [46] is to locate the target's echo and take the difference, Δt_k, between it and the current operator at each array element. This is best visualized in the differential time-shift (DTS) panel $d''_g(t)$, where the once-focused data has already been shifted by the focusing operator and any difference between the operator and the target echo shows up as a deviation from $t = 0$. Add

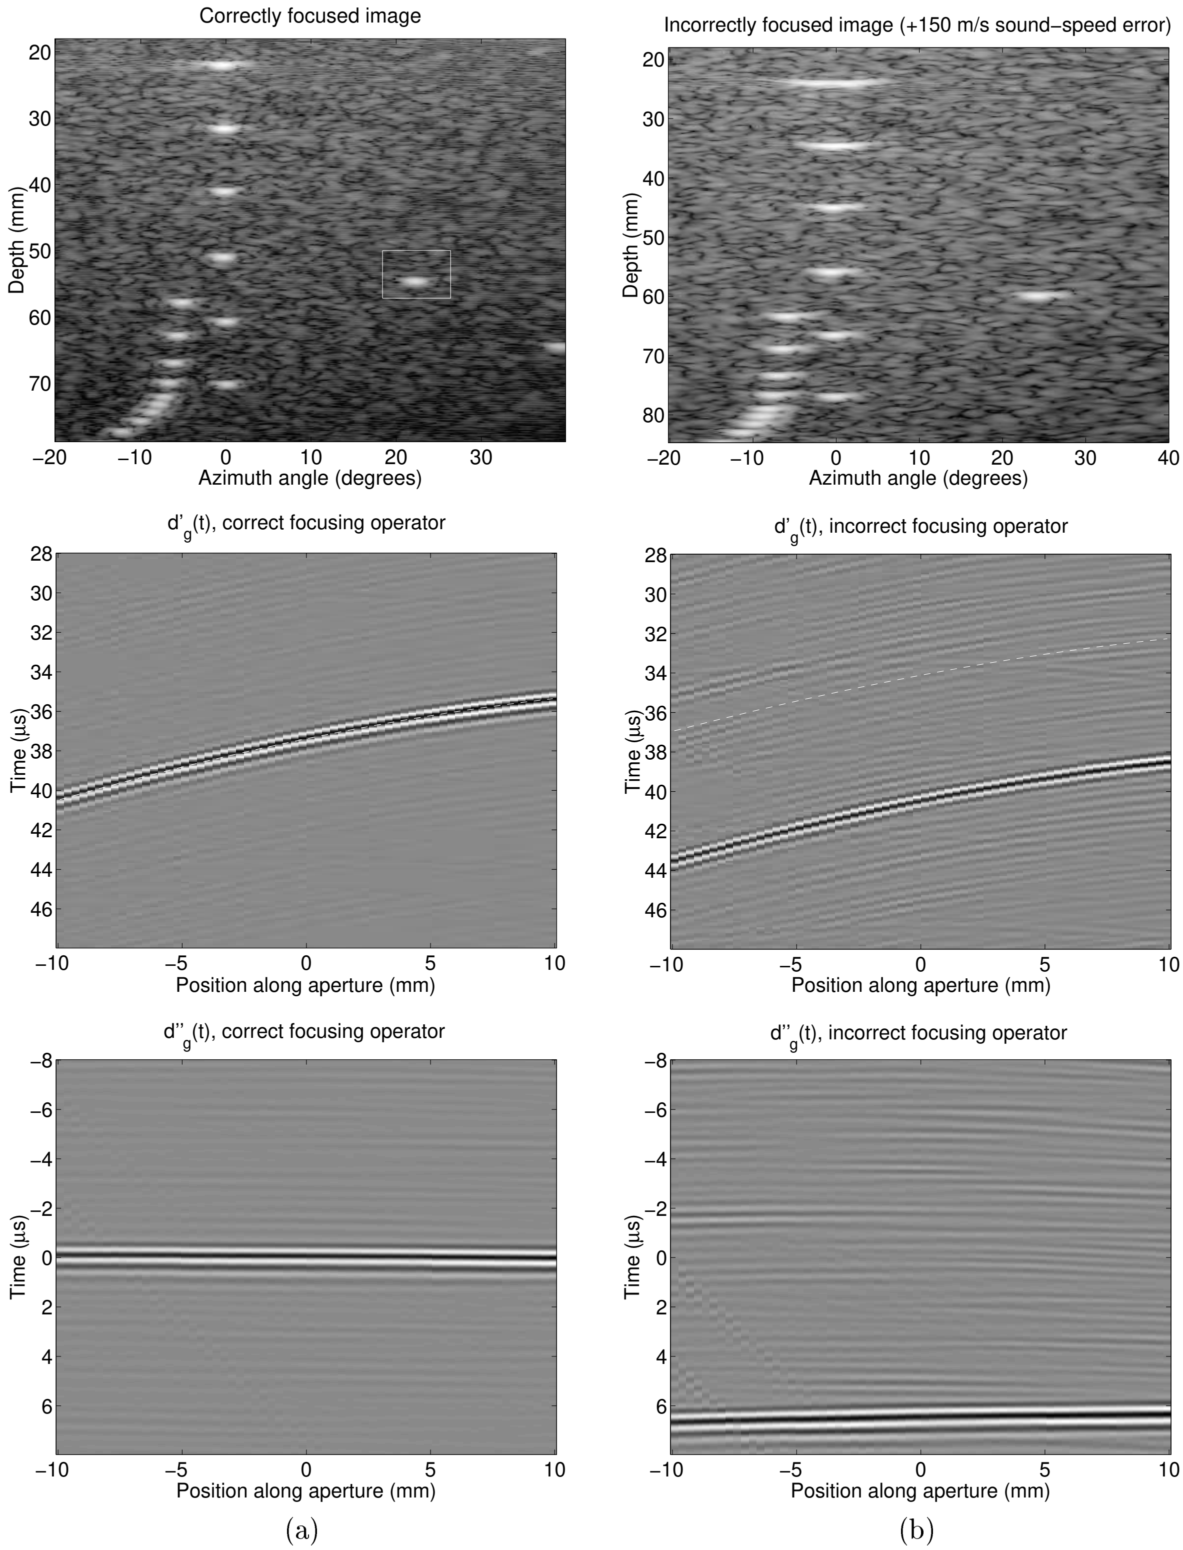

Figure 5.1 Common-focus-point (CFP) gathers and differential time-shift (DTS) panels for a correct focusing operator and for a 150 m/s sound-speed error. (a) Correct operator for focusing on the boxed target. (b) Operator calculated for a speed of sound that is 150 m/s too high.

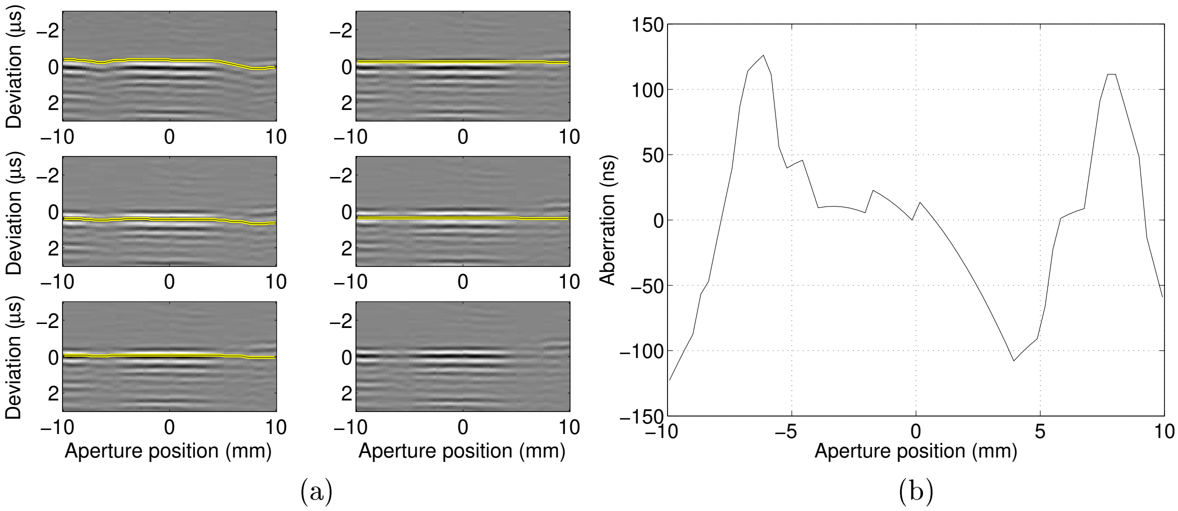

Figure 5.2 (a) DTS panels and picked target echoes (in green) for five updates to the focusing operator. (b) The resulting aberration profile.

half of this difference to the current operator:

$$t_k^{i+1} = t_k^i + \frac{\Delta t_k}{2}, \quad 1 \leq k \leq N. \tag{5.4}$$

An example of the updating process is shown in Figure 5.2. In this case, the aberrator was a section of a rabbit's chest wall oriented so that the scan plane cut across the ribs. The target was a point-like reflector at the top of an anechoic region in a "phantom"—a physical model designed to simulate the ultrasonic properties of real tissue. The reflection is similar to what would be observed at the top of a blood vessel. The green line shows the times picked for the echo at each iteration; the picking is automatic, using a technique described in the next section. After five iterations, the target echo appears flat. The aberration profile derived from the final operator shows two peaks at the rib locations.

5.1.2 Comparison to NNCC

The updating technique just presented is very similar to the nearest-neighbor cross-correlation (NNCC) algorithm of Flax and O'Donnell [23]. Translating their algorithm into the framework used here, it may be summarized as follows:

1. Transmit a focused pulse using the ideal, hyperbolic focusing operator $t_k(x', z', c)$.

2. Record the received signals from all array elements; this is the common-focus-point gather, $d'_g(t)$.

3. Cross-correlate $d'_k(t)$ and $d'_{k+1}(t)$ for $k = 1 \ldots N-1$ to estimate the time shifts $\Delta t_{k,k+1}$ between adjacent elements. The time window used should span many resolution cells in depth.

4. Form the first updated focusing operator by integrating the time-shift estimates across the array: $\tilde{t}^1_1 = t_1$, $\tilde{t}^1_2 = \tilde{t}^1_1 + \Delta t_{1,2}$, $\tilde{t}^1_3 = \tilde{t}^1_2 + \Delta t_{2,3}$, etc.

5. Adjust the new operator as needed to prevent focal point drift with iteration; for example, removing the best-fit linear component from the update term prevents a change in the steering angle.

6. Iterate by using the new operator $\tilde{t}^1$ on transmission and performing the cross-correlations and the update step again to obtain $\tilde{t}^2$. Continue until the operator converges.

Clearly, both NNCC and the algorithm proposed here attempt to update a focusing operator using the information in $d'_g(t)$. Both are iterative. There are three significant differences, however:

1. In its original form, the proposed updating procedure assumes that a specific reflection event (the target echo) can be identified and tracked in $d''_g(t)$. The focusing operator is updated until this echo is flattened at $t = 0$. By contrast, NNCC cross-correlates fairly long portions of the received signals, corresponding to a 1 or 2 cm range at typical ultrasound bandwidths, without attempting to "lock on" to any particular feature.

2. NNCC updates the focusing operator in a relative sense, without an absolute reference point. The event-tracking algorithm uses the deviation of an event from its expected arrival time to update the operator in an absolute sense—it exploits the principle of equal travel-time.

3. NNCC estimates the element-to-element time shifts using simple cross-correlations, with no memory. A tracking method to pick the target echo in the proposed algorithm has not been specified yet; in seismic imaging it would typically involve some human input. An automatic technique with path memory, based on dynamic programming, is discussed in the next section.

Regarding (1), it could be asked whether "events" (to use the geophysical term) even exist in typical ultrasound data. In other words, are there coherent reflections which a tracking algorithm could exploit in the updating process? While seismic images are filled with interfaces between rock layers, ultrasound depends on volumes of random scatterers for much of the reflected energy. (This is why it can image such a wide swath in azimuth from a comparatively small aperture.) Previous work in the medical ultrasound literature has concentrated on speckle because it is prevalent in ultrasound images and is a limiting factor in time-shift estimation accuracy [30]. The spatial correlation of the N received echoes following a transmission focused onto a speckle-generating tissue patch can be predicted with a form of the van Cittert-Zernike theorem [25, 26]. This analysis treats the random scattering region at the focus as a source of incoherent radiation. When there is no aberration present, the correlation between signals in $d_g''(t)$ is a triangle function of the distance between the elements under consideration. The rate of decorrelation with element separation increases as the aperture shrinks, or as aberration is introduced. To obtain better time-shift estimates, longer time windows should be used. Under severe aberration, the decorrelation is too great for algorithms like NNCC to converge.

While fully developed speckle may serve as a useful worst-case test for these algorithms, clinical ultrasound images exhibit varying degrees of coherent and incoherent scatter. If enough locations returning coherent echoes can be found in the image, the operator-updating strategy presented here may be successful. This ultimately depends on the content of "typical" ultrasound images and the ability of the correction algorithm to pick the best locations for operator updating.

Because complete clinical data sets are not publicly available and the present research has been largely restricted to synthetic aberrators and tissue-mimicking phantom targets, further research is necessary to determine if the proposed operator-updating algorithm will be useful in practice.

5.2 An Automatic Operator-Updating Algorithm

Figure 5.3 outlines the implementation of the overall algorithm for picking initial points, iteratively updating the focusing operator at those points, and evaluating the converged operators. In this section, each part is described in detail. Because there are so many adjustable parameters in the algorithm, the values chosen for the experiments in Chapter 6 are listed in Table 5.1 for reference.

```
Form an initial image (assume no aberration).

Choose bright spots subject to a minimum spacing d_min1 (see text). For each bright
point, compute d''_g(t) and run the Viterbi algorithm once. Store the point's
coordinates (r_b, θ_b) and Q_max in 𝔹.

For each φ ∈ {−30°, −20°, ..., 30°} {
    For each b ∈ 𝔹 where (φ − 5°) ≤ θ_b ≤ (φ + 5°) {
        Initialize t̃ with hyperbolic operator for (r_b, θ_b), or other initial operator in b.
        Δt_k = 0,  k = 1...N
        Do {
            t̃_k = t̃_k + Δt_k/2,  k = 1...N
            Compute new d''_g(t) using t̃ and run Viterbi algorithm.
            Set Δt = highest-quality path through d''_g(t), Q = corresponding quality.

            If (Q < Q_min, or Q < C_1 Q_max (stored in 𝔹), or iteration count > N_i)
              Reject this operator: Remove entry from 𝔹 and terminate iteration.
            Else {
              p = S_1 Σ_{i=2}^{N} |Δt_i − Δt_{i−1}| + S_2 |Δt_N − Δt_1| + S_3 Σ_{i=1}^{N} |Δt_i|
              Find t, the best hyperbolic fit to t̃ in the least-squares sense.
              If (p < 1/f_c and max|t̃_k − t_k| < S_4 and max|t̃_k − t_k − t̃_{k−1} + t_{k−1}| < S_5)
                Accept this operator and terminate iteration.
            }
        } While (iteration not terminated)

        If (operator was accepted) {
          If (t̃ located ≥ d_min2 from all existing operator locations)
            Add to list of operators.

          For k = 1...3 {
              t̂_{k+} = t̃ + kΔr ∂t/∂r,   t̂_{k−} = t̃ − kΔr ∂t/∂r     (Equation (5.14))
              Compute d''_g(t) using t̂_{k+}, t̂_{k−} and run Viterbi algorithm once on each.
          }
          For the t̂_{k+} and t̂_{k−} with the highest-quality paths {
              If (t̂ located ≥ d_min2 from all existing points in 𝔹)
                Store location and initial operator t̂ in 𝔹.
          }
        }
    }
}
```

Figure 5.3 Dynamic-programming, operator-updating aberration-correction algorithm.

Table 5.1 Parameter values used in the dynamic-programming algorithm for the results in Chapter 6.

d_{min1}	8 mm
d_{min2}	4 mm
V.A. N_{states}	101
V.A. d_{max}	4
$w(t)$	Gaussian, 5.6 μs width at half-max
α	0.2
Q_{min}	0.16
S_1	0.8
S_2	1
S_3	0.006
S_4	3×10^{-7} s
S_5	$(0.35/f_c)$ s
C_1	0.9
N_i	12
Δr	$N_{states}\delta_t c$ (6 mm, for typ. $\delta_t = 40$ ns)

5.2.1 Dynamic programming

The task of picking the target echo in $d_g''(t)$ can be recast as that of finding an optimal path through a trellis. This trellis has N stages, and the time samples of $d_g''(t)$ constitute the discrete set of states at each stage. We wish to select the path through $d_g''(t)$ that traces a constant phase along the reflection event. Finding a theoretical basis for optimality will be difficult or impossible; in many situations the necessary information may not even exist within $d_g''(t)$. We can, however, define the costs associated with each step along a path so that the "optimal" path will approximate the one chosen by a skilled human observer.

This is an example of a deterministic dynamic programming problem [63]. These problems require the evolution of the system (the construction of a path) to be in discrete stages, with a cost that is additive over the path. Then, the dynamic programming principle of optimality holds: Consider the optimal path through the trellis from start to finish. The optimal sub-path from any point along this path to the finish must follow the same route as the overall optimal path. (If this were not true, a new path from start to finish could easily be constructed that would be "better than optimal.")

The principle of optimality makes it possible to find the optimal path efficiently. A well-known algorithm for doing so is the Viterbi algorithm [64, 65]. Trial paths are built through the trellis, one stage at a time, by connecting each trellis state at stage k with the path terminating at stage $k-1$ that minimizes the total cost-thus-far. At the last stage, the best path is kept. The Viterbi algorithm was originally used for the maximum-likelihood decoding of convolutional codes, but has been used in many other settings since then. For example, in seismic processing, it has been used to automate the matching of compressional-wave and shear-wave seismograms [66].

5.2.2 Quality function

The Viterbi algorithm is implemented so as to maximize an additive quality value, $Q = \sum q_{k-1,k}$, which is defined for each transition between stages and is composed of two parts,

$$q_{k-1,k} = q_{k-1,k}^{coherence} + q_{k-1,k}^{wiggle}. \tag{5.5}$$

In order to make the optimal path relatively insensitive to amplitude fluctuations, Q incorporates a measure of echo coherence [67]. This is defined as the ratio of coherently summed energy to incoherently summed energy over a time-windowed path through $d_g''(t)$ up to stage k:

$$\Gamma_k = \frac{\int_t w(t) \left| \sum_{g=1}^{k} d_g''(t + \Delta t_g) \right|^2 dt}{\int_t w(t) \sum_{g=1}^{k} |d_g''(t + \Delta t_g)|^2 dt}. \tag{5.6}$$

(The same measure is used extensively in seismic processing, where it is called *semblance* [68].) Here Δt_g is the path deviation from $t = 0$ at stage (and element) g. $w(t)$ is a Gaussian time window. The quality measure must be additive over every path, so the change in coherence is used:

$$q_{k-1,k}^{coherence} = \Gamma_k - \Gamma_{k-1}. \tag{5.7}$$

To encourage smooth paths, the quality is discounted by a term which grows with the length of the transition between stages:

$$q_{k-1,k}^{wiggle} = -\alpha(\Delta t_k - \Delta t_{k-1})^2 f_c^2, \tag{5.8}$$

where f_c is the array-transducer center frequency and α is a tunable parameter.

The number of trellis states, $N_{nstates}$, is a trade-off between computation time and the likelihood of finding a coherent echo; 101 states were used in the experiments. The maximum jump at each stage, d_{max}, depends on the expected maximum slope of the aberration profile; a value of 4 was used in the experiments.

If the window $w(t)$ is long, the picking algorithm will behave more like NNCC in the sense that it will not "lock on" as readily to individual coherent events. In contrast to NNCC, the operator is still updated in an absolute sense because the algorithm picks one particular path through $d_g''(t)$, but this choice of optimal path will tend to bounce around from one iteration to the next. If $w(t)$ is short (on the same order as the length of the transmitted pulse), the algorithm will settle down quickly once it finds a coherent echo. The choice depends on the expected target properties. Since the main target used in our experiments is primarily speckle-generating and returns very few coherent echoes, the half-magnitude width of $w(t)$ was set to 15 cycles at $f_c = 2.6$ MHz. This is equivalent to about 4 mm in range, still less than the 20 mm recommended for NNCC in [24].

The picking algorithm is not phase-sensitive; that is, it may choose any constant phase along the picked event that maximizes the quality measure. The picked phase will sometimes differ from one iteration to the next, until the quality maximum is centered. Aligning a random phase at $t = 0$ is not a significant problem because adding a small constant (less than the length of the transmit pulse) to a focusing operator is equivalent to a tiny change in depth.

5.2.3 Selection of initial points

In order for the operator-updating algorithm to converge, $d_g''(t)$, the DTS panel, must contain an echo or echoes with some minimal degree of coherence. As aberration increases, the coherence of echoes from speckle regions rapidly degrades to the point where convergence is unlikely—even a human operator could not hope to pick the correct path through Figure 5.4, for example.

Under severe aberration, the success of this algorithm is critically dependent on the selection of initial focal points that can provide echoes of higher coherence than those from speckle. If convergence can be achieved at just a handful of these locations, the resulting aberration profiles can be used for the initial transmit foci at other places in the image. In this way, many aberration profiles can eventually be found.

The approach taken here is to search an initial, conventionally focused image for isolated bright spots that will serve as the initial focal points in the operator-updating process. The following algorithm has been used in the experimental results:

1. Convolve the image with a circular 2-D filter having 1 at the center and a small negative value over the rest of the area, such that the sum of the elements is zero. The filter is at least several resolution cells in diameter.

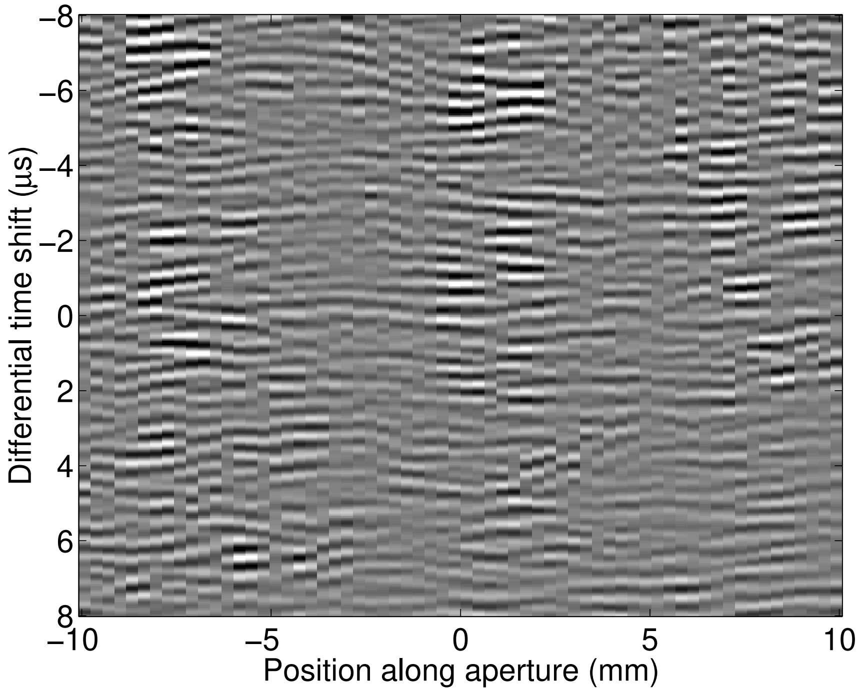

Figure 5.4 A DTS panel with poor coherence for which the dynamic-programming algorithm is unlikely to converge.

2. Sort the points in the processed image by decreasing magnitude. Store the brightest point as the first in the list of initial points for operator updating.

3. Consider the next-brightest point. If it is at least d_{min1} distance away from the brightest point, add it to the list also.

4. Consider the rest of the points down to some threshold in magnitude. For each one, add it to the list of initial points for operator updating only if it is d_{min1} distance away from all of the points already on the list.

After the initial points are chosen, the algorithm sorts them into groups by azimuth angle and the operator-updating process begins on the first point.

5.2.4 Evaluating convergence and terminating the iteration

Because misconvergence occurs easily (see Section 5.3), tests are performed at each iteration of the operator-updating process to reduce the chance of an incorrect focusing operator inadvertently being accepted by the algorithm. Any of the following conditions will terminate the iteration and cause the algorithm to move on to the next target point:

1. The quality, Q, of the last chosen path through the DTS panel is less than Q_{min}, an adjustable parameter.

2. The quality, Q, of the last chosen path through the DTS panel is less than C_1 times the quality of the path chosen in the first iteration at this point. (C_1 is an adjustable parameter.)

3. The maximum number of iterations, N_i, has been exceeded, and convergence has not been reached.

Convergence is detected using a test statistic that evaluates how well an event is aligned at $t = 0$ in the DTS panel. It is calculated each time the Viterbi algorithm is run to pick an event in $d''_g(t)$. Convergence is declared when

$$S_1 \sum_{i=2}^{N} |\Delta t_i - \Delta t_{i-1}| + S_2 |\Delta t_N - \Delta t_1| + S_3 \sum_{i=1}^{N} |\Delta t_i| < \frac{1}{f_c} \tag{5.9}$$

where S_1, S_2, and S_3 are all adjustable parameters. The first term measures the bumpiness of the picked event, the second term measures tilt, and the third term evaluates how well centered at $t = 0$ the event is.

If convergence is detected in the DTS panel, further checks are made. The aberration profile is computed as the difference between the converged operator $\tilde{t}$ and the best-fit hyperbolic operator t,

$$\tau_k = \tilde{t}_k - t_k, \quad k = 1 \dots N. \tag{5.10}$$

Then, requiring

$$\max_k |\tau_k| \quad < \quad S_4 \quad \text{and} \tag{5.11}$$

$$\max_k |\tau_k - \tau_{k-1}| \quad < \quad S_5 \tag{5.12}$$

ensures that the aberration profile lies within reasonable bounds and contains no large jumps (determined by adjustable parameters S_4 and S_5). Finally, the converged operator is added to the final list as long as its estimated location is at least d_{min2} away from every other operator in the list.

5.2.5 Stepping in range using first-order differential updates

Once an accurate focusing operator is determined for one focal position, it is desirable to use the associated aberration profile as a first estimate for nearby foci, since aberration profiles are expected to change slowly across the image, especially in range. Once a correct operator is found, the dynamic-programming algorithm will often converge successfully at many locations where it would have failed starting from an initial hyperbolic operator.

Estimated operators for deeper and shallower locations in the image are formed using first-order perturbations of the current operator. Taking the derivative of the hyperbolic travel-time with respect to range, we have

$$t(x; r', \theta', c) = \frac{1}{c}\sqrt{(x - r'\sin\theta')^2 + r'^2\cos^2\theta'} = \frac{1}{c}\sqrt{x^2 - 2r'x\sin\theta' + r'^2} \tag{5.13}$$

$$\frac{\partial t}{\partial r'} = \frac{-x\sin\theta' + r'}{c\sqrt{x^2 - 2r'x\sin\theta' + r'^2}}. \tag{5.14}$$

Trial operators are calculated for three range steps of size Δr in each direction from the current operator. For each one, a DTS panel is calculated and the Viterbi algorithm is run to pick the optimal path. The operators yielding the highest-quality picks in each direction are saved to the list of points still to process, as long as their estimated locations are at least d_{min2} away from all of the others in the initial-points list.

5.3 Discussion and Future Research

The dynamic-programming aberration-correction algorithm is prone to misconvergence when a sufficiently coherent echo is not present in the DTS panel. A typical case is shown in Figure 5.5. The algorithm appears to be flattening the speckle echoes over disconnected pieces of the aperture. From comparisons with the OFF algorithm in Chapter 4, the aberration profile in the lower right-hand corner is known to be incorrect.

In some cases, most of the aperture appears well-corrected, but a large discontinuity forms where the picking algorithm has jumped cycles. Once this happens, it becomes self-reinforcing and the algorithm will not reverse its mistake. If a method for detecting and fixing these sharp jumps could be implemented, the performance would improve somewhat.

More research is needed to optimize the selection of initial focal points for the algorithm. This would be broadly applicable to other aberration-correction algorithms using a transmit focus, like NNCC. The bright-spots technique in Section 5.2.3 was easy to implement, but it

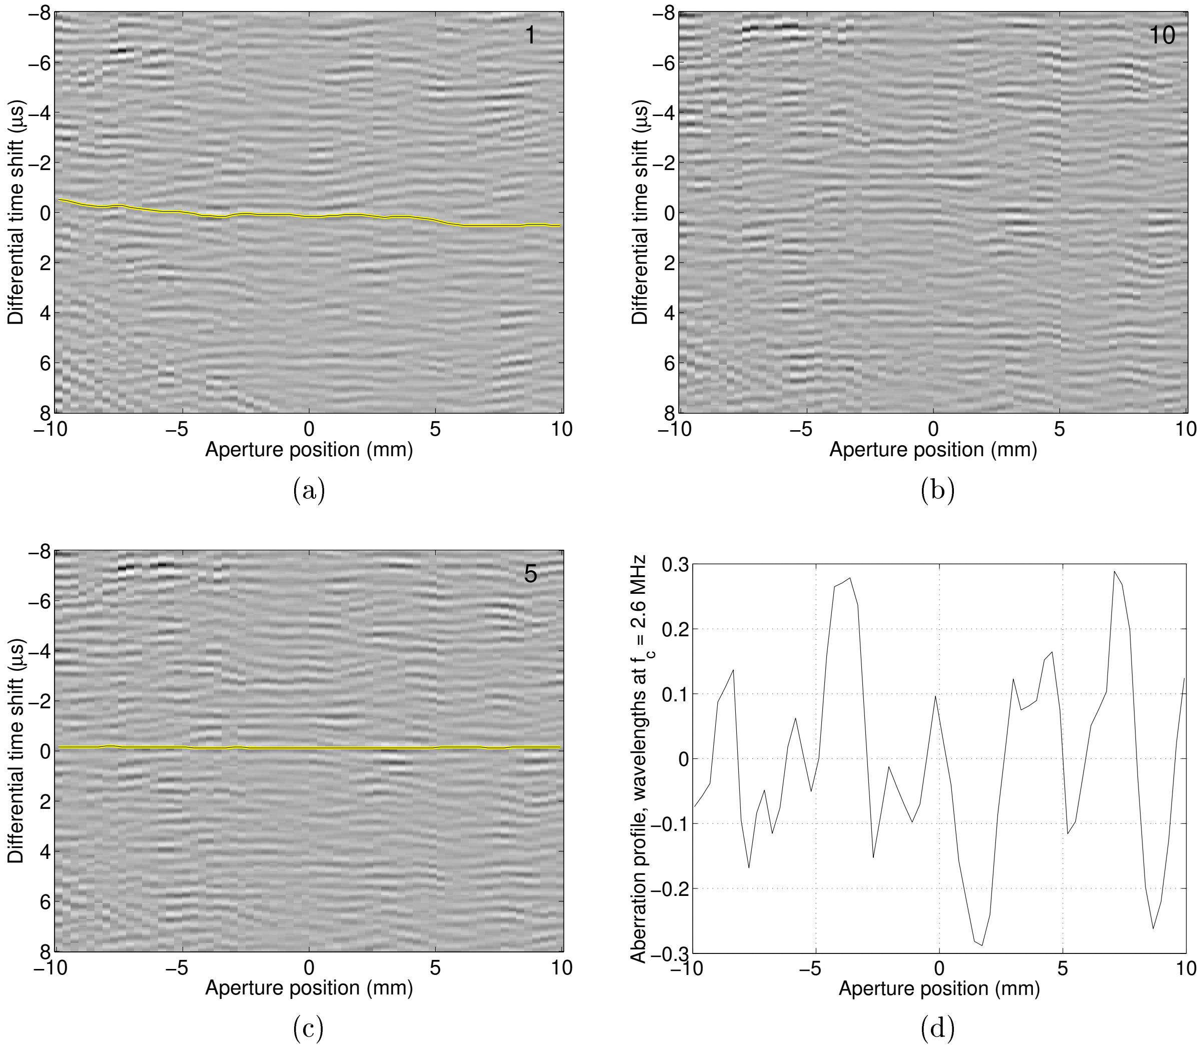

Figure 5.5 An example of misconvergence with the dynamic-programming algorithm. (a) The initial DTS panel. (b) and (c) After 5 and 10 iterations the dynamic-programming algorithm appears to have converged. (d) The derived aberration profile is incorrect. Note that the echoes above and below $t = 0$ in the DTS panels still appear jumbled.

could probably be improved. When isolated bright spots are present, picking them is a good strategy, but in a more uniform field the brightest points may not necessarily coincide with the most coherent echoes. In addition, when the aberration is severe, the initial, conventionally focused image becomes less useful for finding these points. Examining the received wavefields in $d_g''(t)$ over the entire image, using the Viterbi algorithm to estimate the quality of the echoes, and then going back to promising locations for iteration is one possibility. It would greatly increase the computational requirements, however.

It is unclear under what conditions the DTS panel $d_g''(t)$ contains sufficient information to correct for a given level of aberration. Apart from studies of speckle and its influence on correlation-based time-shift estimates [25, 26, 30], little is known. It is possible that NNCC-like algorithms (including the dynamic-programming algorithm presented here) cannot be significantly improved, and further progress will require the use of complete data prior to any focusing.

CHAPTER 6

EXPERIMENTAL PROCEDURES AND RESULTS

Water-tank experiments have been conducted to evaluate the aberration-correction performance of the overdetermined, fan-filtering algorithm (OFF, Chapter 4), the dynamic-programming algorithm (Chapter 5), and three representative published algorithms. This chapter describes the equipment and procedures for collecting complete data sets, the particular implementations of the three published algorithms, and the imaging methodology, followed by an extended discussion of the results.

6.1 Experimental Apparatus

6.1.1 Array imaging system

Most of the complete data sets used in this research were acquired in the Bioacoustics Research Laboratory using a Philips/ATL P4-2 64-element array transducer (2.6-MHz center frequency, 315-μm pitch) and a custom data-acquisition system (Figures 6.1 and 6.2). RF multiplexers (Matrix Systems, Calabasas, CA) permit the independent selection of single elements on receive or transmit. Transmit pulses (approximately 100 volts peak-to-peak, loaded) are provided by a Panametrics 5800 pulser/receiver. A GPIB-connected digital oscilloscope carries out the tasks of digitization and averaging. Scans were conducted in a tank of degassed water at room temperature.

The receiver portion of the Panametrics pulser/receiver is unused except for the 64 pulse-echo signals comprising the monostatic subset of the complete data set. Signals from the other 4032 bistatic transmit-receive element combinations are routed to the data-acquisition system through four low-noise RF preamplifiers (Appendix A). These were designed, built, and tested by hand for optimum noise performance at the frequencies and input impedances presented by the array transducer. The circuit uses a National Semiconductor CLC5509 chip which was designed specifically for medical ultrasound scanner front-ends. The preamps achieve a peak

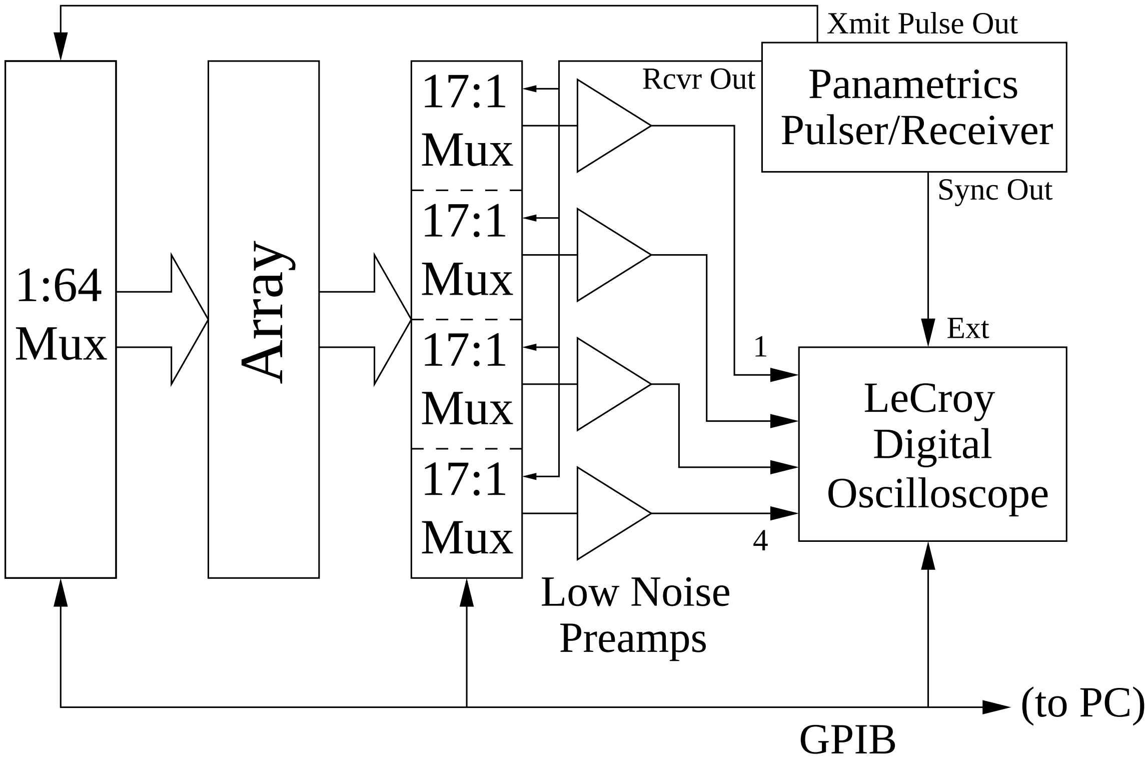

Figure 6.1 Apparatus for collecting complete data sets from a 64-element array transducer.

voltage gain of 33 dB at 3 MHz, with less than 1 nV/$\sqrt{\text{Hz}}$ of input-referred noise. In terms of noise performance, the custom preamps are far superior to the front end of the pulser/receiver, but the latter must still be used for the special case of transmitting and receiving on the same element because the custom preamps are not capable of transmit/receive switching.

Individual array elements are selected for receive or transmit by a GPIB-controlled RF multiplexer system. The switch configuration was designed locally and custom-built for this application. It allows parallel acquisition of up to four elements on receive and provides spare inputs for routing the same-element (monostatic) signals to the receiver side of the pulser/receiver. The switches can handle standard high-voltage transmit pulses and maintain better than 90-dB isolation between channels over the transducer bandwidth. To connect the array to the multiplexers, an adapter box was constructed by carefully soldering 64 miniature coaxial cables with BNC connectors on one end to the mating high-density array connector.

A software application commands the multiplexers through all 4096 transmit/receive element combinations and acquires four waveforms at a time from the digital oscilloscope. The oscilloscope averages each signal over many transmit pulses (commonly 100–300) to improve SNR. The data are saved as a $64 \times 64 \times N_t$ complete data cube for later processing in MAT-

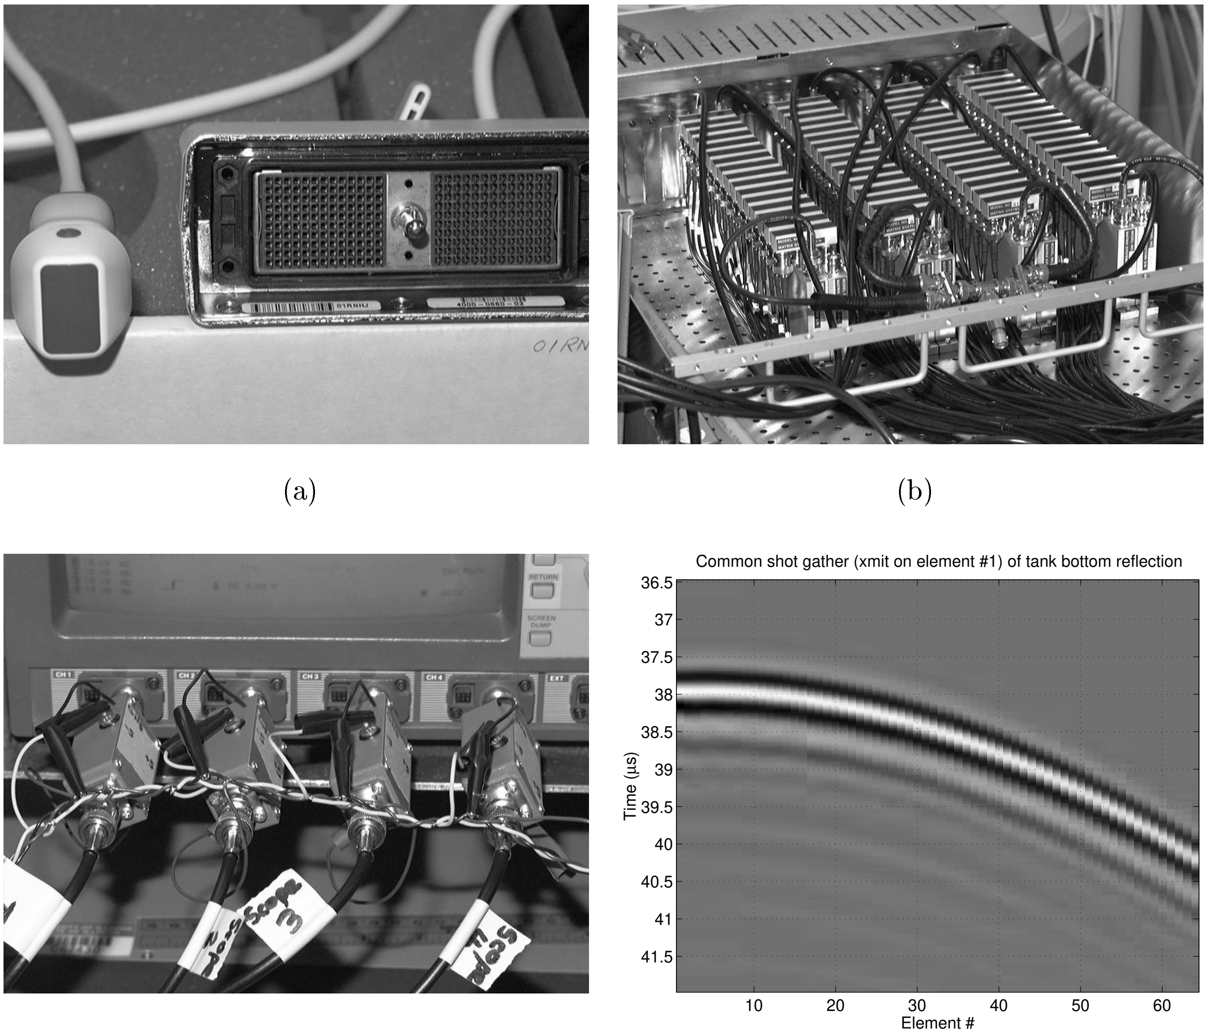

(a) (b) (c) (d)

Figure 6.2 Pictures of the experimental apparatus. (a) The 64-element phased-array transducer. (b) The computer-controlled RF switching system used for selecting single array elements on transmit and receive. (c) The low-noise preamps in use. (d) A common-shot gather showing the expected hyperbolic moveout of the reflection from the bottom of the water tank.

LAB. N_t, the number of time samples, is usually about 5000. Depending on the length of the waveforms and the amount of averaging, a complete data set takes between 1 and 4 hours to acquire.

6.1.2 Targets for imaging

A tissue-mimicking phantom (ATS model #539, Figure 6.3) served as the target for most of the data sets. The phantom is a physical model approximating the scattering and attenuation

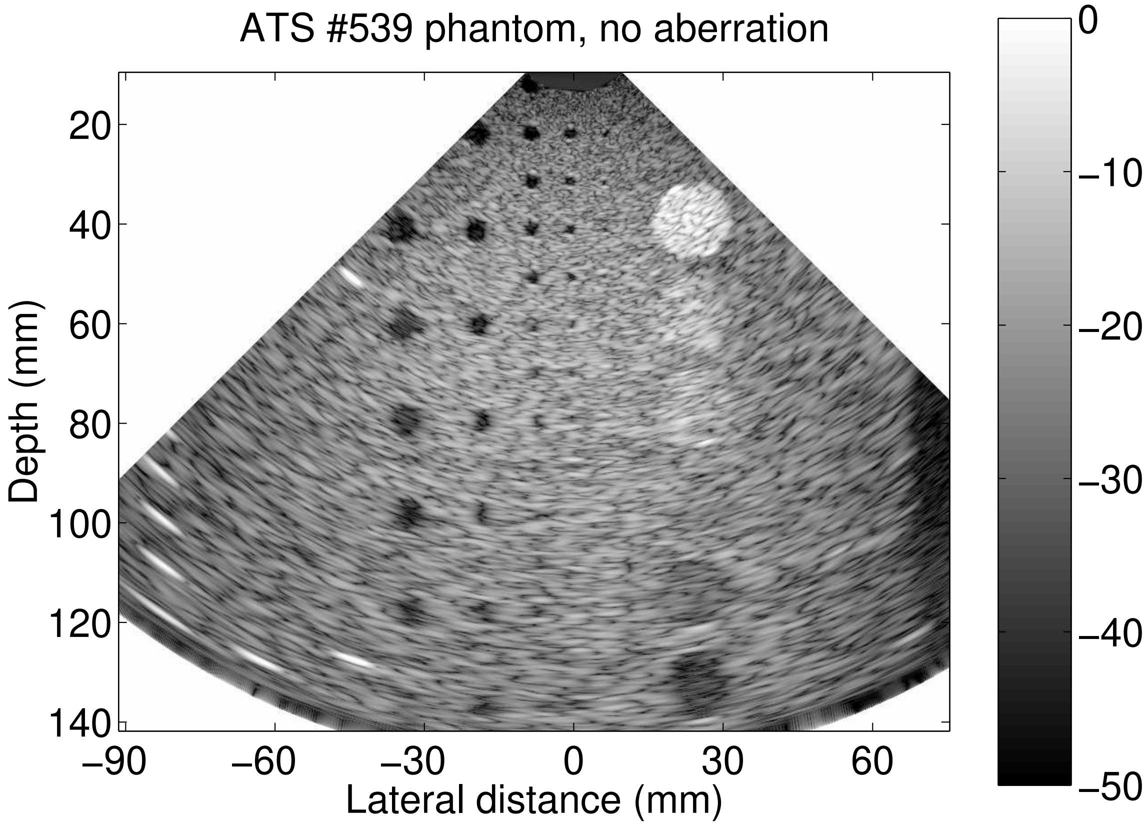

Figure 6.3 The ATS 539 phantom target.

properties of biological tissue. It uses wires and cylindrical structures to simulate point targets and cysts of various sizes placed within speckle-producing background scatter. Anechoic (scatter-free) cysts with diameters of 8, 6, 4, 3, and 2 mm are arranged in five columns. In Figure 6.3, the 2-mm cysts are just barely visible near the top of the image. Another column of cysts are 15 mm in diameter, with scattering gradations from 15 dB above to 15 dB below the background level. The background has a 1450-m/s speed of sound.

The $\pm 40^\circ$ field of view for our locally collected data sets covers the part of the phantom containing cysts with a speckle background. There is only one point target, far to one side, in most of the images. This choice of target region is intended to increase the difficulty of the aberration correction challenge. Strong, coherent scatterers (such as wire targets) make aberration correction much easier for some algorithms by serving as clear "beacon signals" [23]. In our target region, however, only the echoes from the top and bottom of each cyst are more coherent than fully developed speckle.

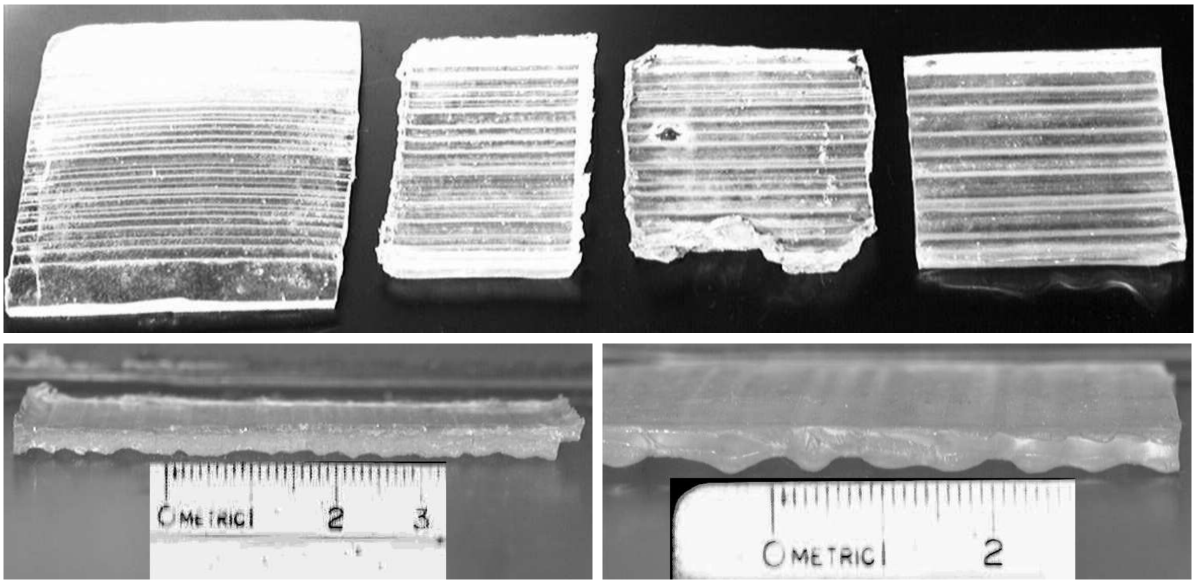

Figure 6.4 The RTV silicone aberrators; profile views of the "thin" and "thick" aberrators. (The "slab" aberrator, not pictured here, is a thick piece of silicone with no ripples.)

Table 6.1 Synthetic aberrator properties.

Name	Maximum thickness (mm)	Peak-to-peak ripple (mm)	Correlation length (mm)
`slab`	18	0	—
`thin`	3	1	2.4
`thick`	4	2	2.9

6.1.3 Silicone aberrators

Artificial aberrating structures (Figure 6.4, Table 6.1) were constructed from GE RTV615, a two-part, room-temperature-vulcanizing silicone rubber. This material has an 1100-m/s speed of sound and an acoustic impedance similar to that of soft tissue, providing strong refraction effects without much reverberation. The aberrators have a rippled surface on the underside, created by pouring the liquid silicone into a rippled mold. The mold material was Ivory soap, carved using pieces of sheet metal that had template curves cut into their edges with a rotary tool.

During scans in the water tank, the aberrators' ripples were carefully aligned with the array elements. By thus restricting the aberration to one dimension, a fair test of aberration correction was made possible with the one-dimensional array. (It is well-established that aberration in real tissue can have characteristic correlation lengths in the elevation direction which are much

Table 6.2 Complete data sets used for the aberration-correction experiments.

Name	Source	N	f_c (MHz)	f_{sample} (MHz)	Averaging
`ats`	BRL	64	2.6	50	200
`ats_syn`	BRL	64	2.6	50	200
`ats_slab`	BRL	64	2.6	40	300
`geabr_2`	BUL	64	3.3	17.76	none
`ats_2abl`	BRL	64	2.6	50	160
`ats_4abl`	BRL	64	2.6	50	200

Name	Target	Aberrator	Stand-off dist. (mm)	Peak-to-peak aberratio τ_{p-p} (ns, approx.)
`ats`	ATS phantom	none	—	—
`ats_syn`	ATS phantom	synthetic delays	0	440
`ats_slab`	ATS phantom	"slab"	4	—
`geabr_2`	AIUM phantom	grooved RTV	$\sim 0?$	280
`ats_2abl`	ATS phantom	"thin"	6	180
`ats_4abl`	ATS phantom	"thick"	9	~ 300

shorter than the typical 1-D array element size [17]. To properly correct for these aberrations, 1.5-D or 2-D arrays will be needed [69, 70], along with 2-D versions of the aberration-correction algorithms. While computationally daunting, these extensions are reasonably straightforward.)

6.2 Procedures

6.2.1 Data sets

The complete data sets used for the results in Section 6.3 are listed in Table 6.2. They are listed in approximate order of aberration-correction difficulty. `ats` is the control, obtained with no aberrator interposed between the array transducer and the ATS phantom (Figure 6.3). `ats_syn` is the same data set with its signals time-shifted in software; the synthetic aberration profile (Figure 6.5) is low-pass-filtered Gaussian noise. This data set is used to evaluate the aberration-correction algorithms' performance under conditions approximating the ideal screen model, in which the aberration profile is the same over the entire image. In other words, the isoplanatic patch is the entire image; this is the ideal "surface-consistent statics" problem in seismic imaging [21, 22].

All of the data sets were collected locally (BRL) except for `geabr_2`, which comes from the Biomedical Ultrasound Laboratory (BUL) at the University of Michigan [71]. `geabr_2` is

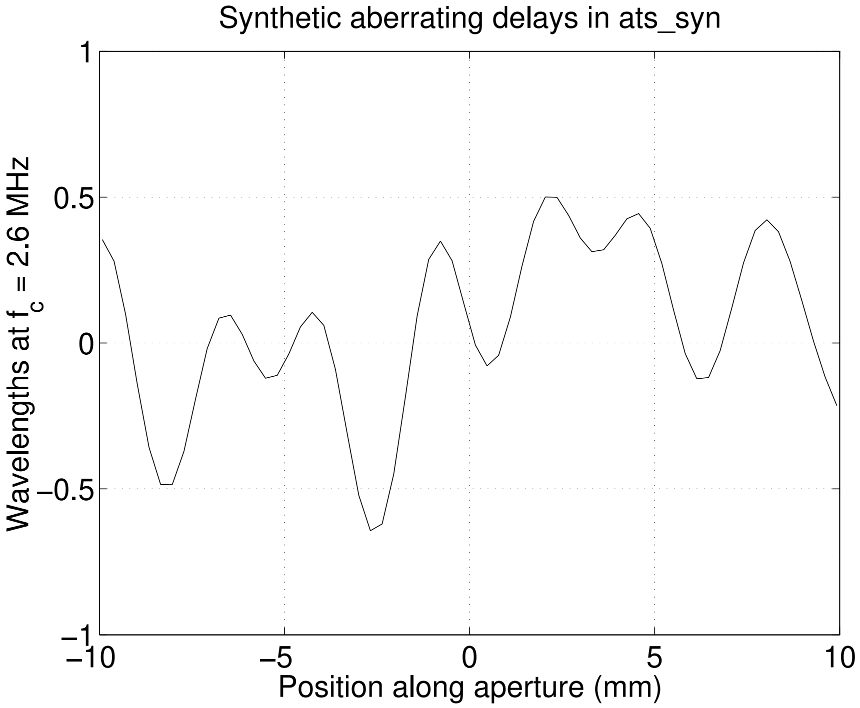

Figure 6.5 Synthetic aberration profile used for data set `ats_syn`.

the "2X" aberrator example from the original work on the NNCC algorithm [24]. It is used here to verify the implementation of NNCC and to provide a comparison with the locally collected silicone aberrator data sets. Because the RTV aberrators in [24] were molded onto the transducer surface, the aberration in `geabr_2` conforms much more closely to a screen model than the aberration in `ats_2ab1` and `ats_4ab1`. Another important attribute of this data set is the lack of averaging; this allows the evaluation of algorithms at low SNR.

The "thin" and "thick" aberrators used for data sets `ats_2ab1` and `ats_4ab1`, respectively, are pictured in Figure 6.4. The nonzero stand-off distance from the silicone aberrators in these scans, combined with the significant thickness of the aberrators themselves, means that the aberration in these data sets is more angle-dependent (smaller isoplanatic patch sizes). Although the "thick" aberrator appears smoother and not much thicker than the "thin" aberrator, it has proven the most difficult aberrator to correct for, by far.

Data set `ats_slab` was collected using the "slab" aberrator, an 18-mm-thick layer of silicone with flat surfaces. This was intended as a test of rms sound-speed concepts and time migration. Results using this data set can be found in Chapter 3.

6.2.2 Initial data processing

After the complete data set is imported into MATLAB, the 64 same-element pulse-echo signals are corrected for passage through the Panametrics pulser/receiver's amplifier. Advancing these signals by 80 ns and multiplying them by -0.65 causes them to blend in visually on common-shot and common-receiver gathers, and this was deemed sufficient for the subsequent processing. To improve the SNR, each of the 4096 signals is then bandpass filtered from 250 kHz to 5 MHz with a length-251 FIR filter. Finally, a 2-D FIR fan filter (Section 4.1) with a typical angular passband of $0 \pm 40°$ is applied to each of the 64 received wavefields. This further improves the SNR and strongly attenuates undesirable events in the data. These events are either slow waves (such as surface waves) or strong reflections from far-off-axis objects—typically the hardware used to hold the target in position.

Note that the data are not demodulated to complex baseband prior to processing. The center frequency is low enough, and the fractional bandwidth large enough, that very little would be accomplished by doing this.

6.2.3 Derivation of aberration profiles

The processed data are passed to a collection of programs implementing different algorithms for aberration correction. These take a complete data set as input, and provide a collection of aberration corrections as output. Each correction contains the length-N aberration profile τ and three parameters from the least-squares, best-fit hyperbola: the sound speed and (x, z) coordinates. Among other parameters, each algorithm is provided a set of target regions of interest (ROIs) for which to estimate aberration profiles, except for the dynamic programming algorithm, which finds its own "targets of opportunity" in the image. For most of the experiments, between 30 and 60 ROIs were specified, with azimuth angle θ covering the range $\pm 30°$ or $\pm 40°$ in steps of $10°$, and depth r between 40 and 100 mm in steps of 10 or 20 mm.

Two of the algorithms, the overdetermined fan-filtering (OFF) algorithm and the dynamic-programming, operator-updating algorithm, have been described in detail in Chapters 4 and 5, respectively. Three previously published algorithms have also been implemented, representative of existing techniques using a full-aperture transmit focus, an image quality metric, or common-midpoint signals. The implementation details of each will now be described.

$\tilde{t}^1_k = t_k(r'_{\mathrm{ROI}}, \theta'_{\mathrm{ROI}}, c), \quad k = 1 \ldots N$

For $i = 1 \ \ldots \ M-1$ {
    Transmit a pulse using $\tilde{t}^i$ and record the received signals from all array elements:
    $d'_g(t) = \sum_{s=1}^{N} d_{s,g}(t + \tilde{t}^i_s)$

    $\tilde{t}^{i+1}_1 = \tilde{t}^i_1$
    For $k = 1 \ \ldots \ N-1$ {
        Cross correlate $d'_k(t)$ and $d'_{k+1}(t)$ over a time window equivalent to 2 cm in range, centered on the times of the focusing operator.
        Estimate the relative time shift Δ from the cross-correlation peak, searching the range of lags corresponding to $(\tilde{t}^i_{k+1} - \tilde{t}^i_k) \pm 1/(2f_c)$.

        $\tilde{t}^{i+1}_{k+1} = \tilde{t}^{i+1}_k + \Delta$
    }

    Find t^{i+1}, the best hyperbolic fit to $\tilde{t}^{i+1}$ in the least-squares sense, yielding parameters (r_e, θ_e, c_e).
    Find the best hyperbolic fit, having sound speed c_e, to the initial hyperbolic operator $t(x; r'_{\mathrm{ROI}}, \theta'_{\mathrm{ROI}}, c)$, yielding parameters (r_f, θ_f).

    $\tilde{t}^{i+1} = \tilde{t}^{i+1} - (\theta_e - \theta_f)\frac{\partial t}{\partial \theta} - (r_e - r_f)\frac{\partial t}{\partial r}$ (Equations (6.2) and (6.3))
}

Find t^M, the best hyperbolic fit to $\tilde{t}^M$.
$\tau = \tilde{t}^M - t^M$

Figure 6.6 Implementation of the nearest-neighbor cross-correlation (NNCC) algorithm.

6.2.3.1 Nearest-neighbor cross-correlation (NNCC)

The NNCC algorithm ([23, 24], Figure 6.6) transmits a focused pulse onto a target point and records $d'_g(t)$, the N echoes received across the array. With a complete data set, this is emulated by advancing and summing $d_{s,g}(t)$ over the transmit elements. The initial focusing operator is the ideal hyperbola for the targeted point.

An updated focusing operator begins with the current operator's time value at array element 1. Adjacent signals in $d'_g(t)$ are cross-correlated, yielding time-shift estimates which are integrated along the array to construct the updated operator at elements 2, 3, $\ldots N$. A time window equivalent to 2 cm in range and centered on the target was used for the correlations in these results, following [24]. The maximum time-shift permitted between elements is the expected shift, as seen in the current focusing operator, plus or minus one-half the period of the transducer center frequency; this corresponds to a critically sampled aberration profile. Depending on the sampling rate used, interpolation may be desirable prior to the cross-correlation; for these results an interpolation factor of eight was used.

At this point, because the new operator was derived from relative time-shifts only, its focus may have drifted relative to the original target. To recenter the operator, its current focus location is estimated from the parameters of a best-fit hyperbola; then, first-order perturbations are applied using derivatives of the hyperbolic travel-time with respect to range and angle.

$$t(x;r',\theta',c) = \frac{1}{c}\sqrt{(x - r'\sin\theta')^2 + r'^2\cos^2\theta'} = \frac{1}{c}\sqrt{x^2 - 2r'x\sin\theta' + r'^2} \tag{6.1}$$

$$\frac{dt}{dr'} = \frac{-x\sin\theta' + r'}{c\sqrt{x^2 - 2r'x\sin\theta' + r'^2}} \tag{6.2}$$

$$\frac{dt}{d\theta'} = \frac{-r'x\cos\theta'}{c\sqrt{x^2 - 2r'x\sin\theta' + r'^2}} \tag{6.3}$$

The recentering step was not explicitly described in [24]. Fortunately, the authors have made their complete data sets available online [71]. A comparison of the images in the original paper with those obtained using this implementation shows good agreement.

The NNCC algorithm is iterated by transmitting another focused pulse, this time using the updated focusing operator, then performing the update steps again. After a small number of iterations (four, for these results), the aberration profile is calculated as the difference between the final operator and its least-squares hyperbolic fit.

6.2.3.2 Speckle brightness

The aberration profile in the speckle brightness algorithm ([31–33], Figure 6.7) is constructed one element at a time. Each τ_k is adjusted in small increments to achieve a local maximum in the average magnitude of a small image region. In this implementation, the region size is 2 cm in range (0.4-mm resolution) by 10° in azimuth (0.5° resolution). The profile adjustment step size, Δ, is $1/(40f_c)$. The delay-and-sum beamforming algorithm (Section 6.2.4) performs dynamic focusing on transmit and receive; this is probably not necessary, but it allows code reuse.

At each element k, three initial images are used to determine whether to adjust τ_k and, if so, in which direction. If increasing or decreasing τ_k increases the average magnitude, τ_k continues to be adjusted in the same direction until the local maximum is found; then, the next element is considered. Only one pass through the elements is performed, in order from 1 to N.

```
τ_k = 0,   k = 1...N
Choose Δ, the profile adjustment step size (1/(40 f_c) for these results).

For k = 1...N {
    Form images of region (r'_ROI ± 1 cm, θ'_ROI ± 5°), dynamically focused on transmit
    and receive, using three aberration profiles: τ_1...τ_N, τ_1...(τ_k + Δ)...τ_N, and
    τ_1...(τ_k − Δ)...τ_N.
    Calculate the average magnitude of each image: b_0, b_+, b_−
    If max(b_0, b_+, b_−) = b_+ {
        τ_k = τ_k + Δ
        While b_+ > b {
                b = b_+
                τ_k = τ_k + Δ
                Form new image using profile τ; b_+ = average magnitude
        }
        τ_k = τ_k − Δ
    }
    Else if max(b_0, b_+, b_−) = b_− {
        τ_k = τ_k − Δ
        While b_− > b {
                b = b_−
                τ_k = τ_k − Δ
                Form new image using profile τ; b_− = average magnitude
        }
        τ_k = τ_k + Δ
    }
}
```

Figure 6.7 Implementation of the speckle-brightness algorithm.

6.2.3.3 Near-field signal redundancy (NFSR) with subarrays

Li's subarray NFSR algorithm for small-element arrays ([58, 59], Figure 6.8) derives an aberration profile based on the similarity of common-midpoint signals, as does the least-squares algorithm of Chapter 4. It only uses offsets of zero and one, however, and also uses a different strategy to achieve directionality. The array is divided into subarrays of size p (16 subarrays of four elements each, in this implementation), which function as single "elements" in the cross-correlations and the system of linear equations.

Partition array into N/p length-p subarrays; let $d_{S,s,G,g}(t)$ be the signal received on subarray G, element g after transmitting on subarray S, element s ($s, g = 1 \ldots p$).

For $m = 2 \ \ldots \ N/p - 1$ {
  Perform moveout correction $d(t) \to d(t')$, relative to center of subarray m, for subarrays $m-1$, m, and $m+1$ at angle θ'_{ROI} (Equation (6.4)).

  $y_0(t) = \sum_{k=1}^{p} \sum_{l=1}^{p} d_{m,k,m,l}(t')$
  $y_1(t) = \sum_{k=1}^{p} \sum_{l=1}^{p} d_{m+1,k,m-1,l}(t')$
  Cross-correlate $y_0(t)$ and $y_1(t)$ over a time window equivalent to 2 cm in range, centered at r'_{ROI}.
  Estimate the relative time shift Δ_m from the cross-correlation peak, searching lags $-0.8/f_c \ \ldots \ 0.8/f_c$.
}

$$\mathbf{b} = [\ \Delta_2 \ \ \Delta_3 \ \ \cdots \ \ \Delta_{N/p-1}\ ]^{\mathrm{T}}; \ \mathbf{A} = \begin{bmatrix} 2 & -1 & 0 & \cdots & 0 \\ -1 & \ddots & \ddots & & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & & \ddots & \ddots & -1 \\ 0 & \cdots & 0 & -1 & 2 \end{bmatrix}$$

$\mathbf{A}^{-1}\mathbf{b}$ is the subarray-by-subarray estimated aberration profile for the central $N/p - 2$ subarrays. Let $\tau = [\ 0 \ \ (\mathbf{A}^{-1}\mathbf{b})^{\mathrm{T}} \ \ 0\ ]$, interpolated by a factor of p.

Figure 6.8 Implementation of the near-field signal redundancy (NFSR) subarray algorithm.

To calculate the zero-offset signal $y_0(t)$ at midpoint m, all combinations of transmit and receive elements in subarray m are summed, but only after the moveout correction

$$\begin{aligned} t' &= \sqrt{\frac{t^2}{4} + \frac{\left((S-m)p + s - \frac{p+1}{2}\right)^2 \delta_x^2}{c^2} - \frac{t}{c}\left((S-m)p + s - \frac{p+1}{2}\right)\delta_x \sin\theta'_{\text{ROI}}} \\ &+ \sqrt{\frac{t^2}{4} + \frac{\left((G-m)p + g - \frac{p+1}{2}\right)^2 \delta_x^2}{c^2} - \frac{t}{c}\left((G-m)p + g - \frac{p+1}{2}\right)\delta_x \sin\theta'_{\text{ROI}}} \end{aligned} \qquad (6.4)$$

has been applied to the individual signals in the complete data set. (S and G are the subarray numbers and s and g are the element numbers within a subarray, for transmit and receive elements respectively.) Likewise, at midpoint m, $y_1(t)$ is the sum of all the moveout-corrected signals with a transmit element in subarray $m+1$ and a receive element in subarray $m-1$.

Signals $y_0(t)$ and $y_1(t)$ are cross-correlated over a 2-cm range-equivalent time window for each m to estimate $N/p-2$ time-shift estimates (14, in this implementation) which comprise the vector $\mathbf{b}$. The maximum correlation lag is $0.8/f_c$ in this implementation; this is a trade-off between the knowledge that the aberration profile can change quite a bit over a subarray, and the possibility of "cycle skipping" if the search range is too wide. The matrix $\mathbf{A}$ is always the same, as given in the table.

Solving $\mathbf{Ax} = \mathbf{b}$ gives the aberration profile for the central $N/p-2$ subarrays under the assumption that the aberration values are zero for the two subarrays at the ends of the array. To generate a profile compatible with the other algorithms used here, interpolation by p is used.

6.2.4 Imaging

Images are formed from the complete data sets and the derived aberration profiles using delay-and-sum beamforming at each point on transmit and receive, i.e., dynamic focus on transmit and receive. Hamming apodization is used on the transmit aperture only. Linear interpolation is included in the imaging code; at the typical sampling rate of 25 MHz, however, it does not make a large difference in image quality when compared to a simpler, nearest-neighbor approach. Each range line of the beamformed image is envelope-detected by applying the Hilbert transform and taking the magnitude of the analytic signal. Finally, an approximate gain correction is applied by multiplying each pixel by $(r/r_{min})^2$, the squared ratio of the current range to the minimum range in the image.

The imaging process is decoupled from aberration correction. If an image at a constant speed of sound is desired, the ideal hyperbolic operators are used. Otherwise, the various aberration correction codes supply a collection of estimated profiles, and the locations for which they were derived, to the imaging code. For each point in the image, the algorithm selects two aberration profiles which are valid at the point's depth and located on either side of it in azimuth angle. If the profiles differ primarily by a lateral shift, as evidenced by a high cross-correlation value at some small lag, an interpolated profile is derived: The average of the lined-up profiles is used, shifted by a linearly interpolated distance based on the azimuth angles of the two aberration profile locations and the image point between them. If, instead, the two aberration profiles do not appear correlated, or the image point lies outside the region populated by profiles, the profile nearest the point's location is used. This occasionally leads to obvious boundaries in the finished image, especially far off-axis, when the quality of the aberration profiles is poor or the profiles give insufficient coverage.

Where multiple images of the same data set are compared, the method of presentation has been carefully chosen in an effort to ensure fairness. All of the images are drawn over a 50-dB dynamic range. Within the set of images being compared, the 0-dB point is chosen subjectively to place the background speckle in the upper quarter of the brightness range. This "burns out" concentrated point reflectors but provides the best contrast between the scatter-free regions and background speckle. Most importantly, the white and black points in each image represent the same magnitudes; they are not dependent on a peak pixel value which changes from one image to the next.

6.3 Results

6.3.1 Data set `ats`

This data set is an experimental control. It provides a benchmark image (Figure 6.3) with which to compare the results of aberration correction on the other data sets. It also assists in verifying that the aberration-correction algorithms are unbiased in the absence of aberration, i.e., their aberration profile estimates are close to zero.

In Figure 6.9, about 45 aberration profiles are plotted for each of the five algorithms under consideration. The profiles with the greatest deviations from zero usually correspond to deep and/or far off-axis image regions. Aside from these, most of the algorithms keep the error to 1/20-wave or below. The true error may be even lower if there is minor aberration present in the ATS phantom. The most obvious possibility is the difference in sound speed between the water and the phantom; however, this would mostly lead to an overall effective sound-speed change (Chapter 3) which would be subtracted out of the profile by the hyperbolic fit.

Two other interesting features may be observed here. In a few of the plots, the aperture appears divided into four pieces by small jumps at ± 5 mm and 0 mm. Closer inspection shows that the jumps occur exactly between elements 16–17, 32–33, and 48–49. This is undoubtedly due to small delay variations between the four channels used to acquire the complete data sets. Notice also that the error in the profiles usually reaches a maximum at the endpoints. This effect is one reason to argue for the use of larger apertures in aberration correction, even if a smaller aperture is used for image formation.

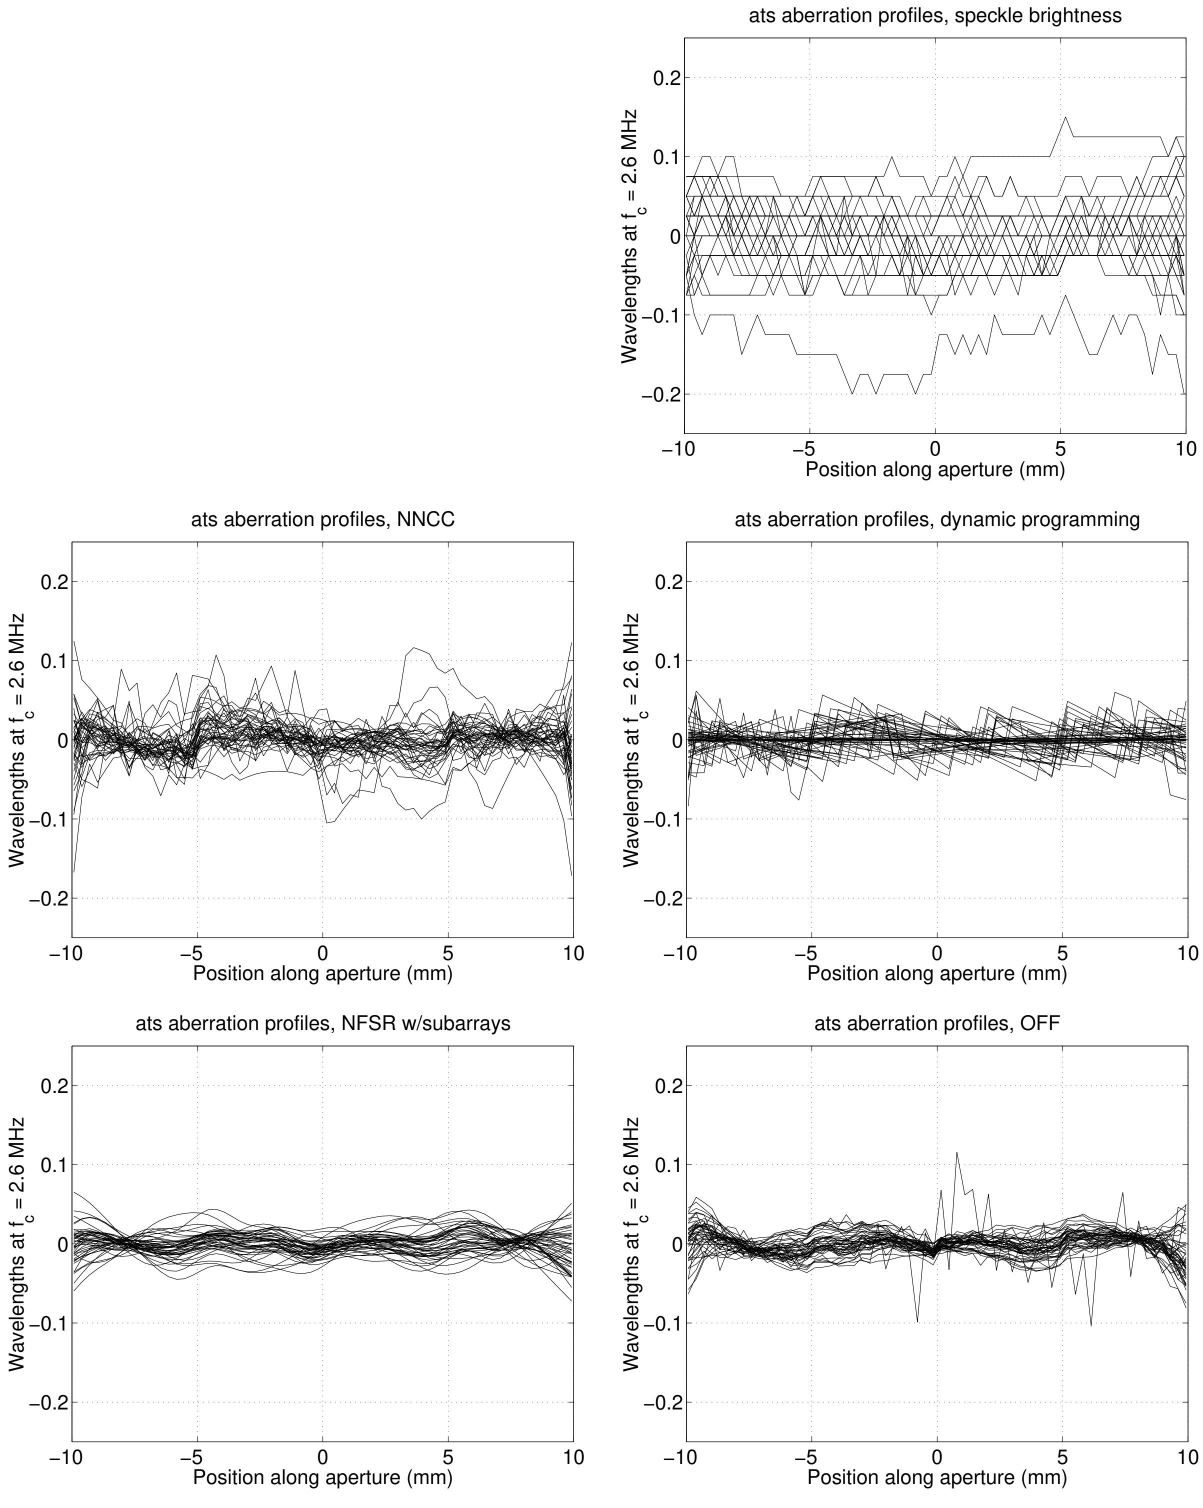

Figure 6.9 Residual aberration profiles estimated by the five algorithms for data set `ats`, which is assumed to have no aberration. The profiles were estimated at image locations from $-40°$ to $40°$ in azimuth and 40 to 120 mm in range.

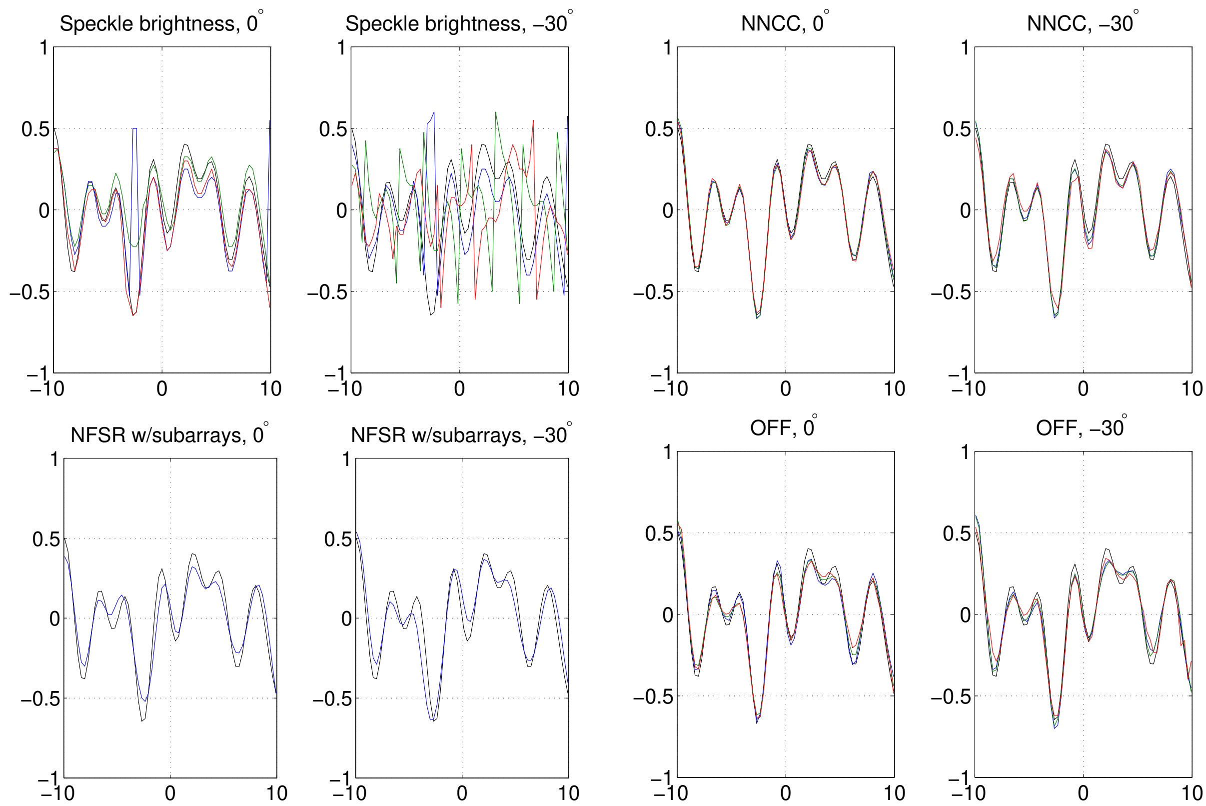

Figure 6.10 Estimated profiles compared to the known synthetic aberration profile for data set `ats_syn`. Profiles estimated at 40, 80, and 120 mm (blue, green, and red curves, respectively) match the actual profile (black curve) closely for the NNCC and OFF algorithms. The NFSR subarray algorithm fails below 40 mm and the speckle brightness algorithm suffers from multiple cycle-skipping errors. (For all plots, the horizontal axis is in mm and the vertical axis is in cycles at the center frequency of 2.6 MHz.)

6.3.2 Data set `ats_syn`

Adding the pure time-shift aberration in Figure 6.5 to data set `ats` provides a good test of aberration-correction performance in the absence of the amplitude variations, diffraction artifacts, and azimuth steering-dependence associated with real-world aberrators. Figure 6.10 compares the known aberration profile with the estimates of the speckle brightness, NNCC, NFSR subarray, and OFF algorithms at $\theta = 0°$ and $-30°$ and $r = 40$ mm, 80 mm, and 120 mm. (The dynamic programming algorithm is not included because it selects its own targets rather than targeting a preset list of ROIs.) Prior to these comparisons, the least-squares best-fit line was subtracted out of the known aberration profile. This is necessary because small linear tilts in a focusing operator correspond to steering angle changes, so a linear term in the derived profile is subsumed into the best-fit hyperbola.

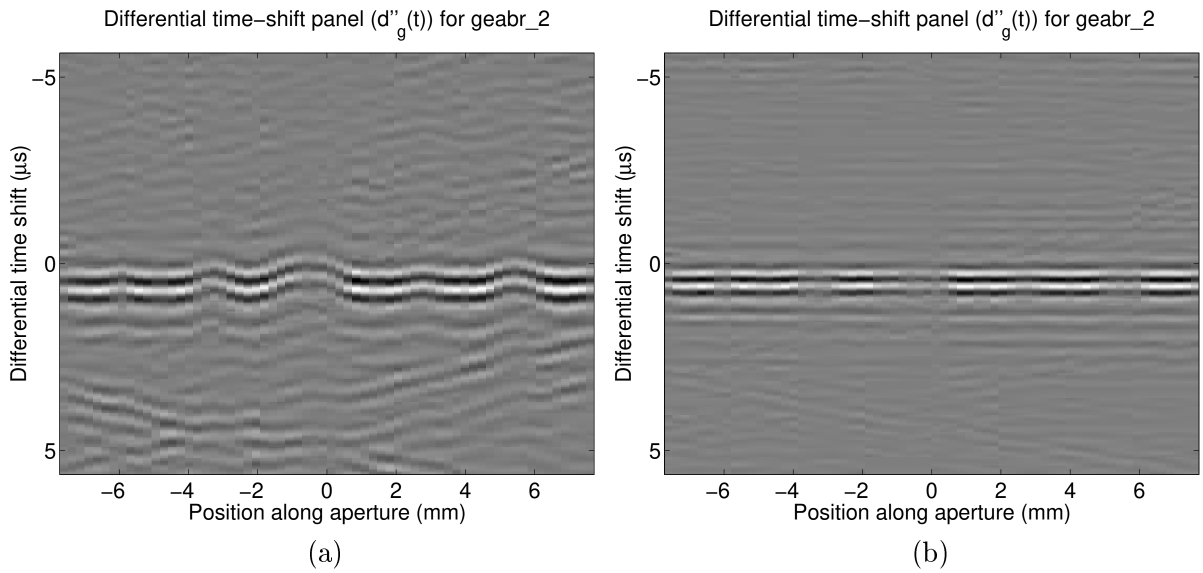

Figure 6.11 The differential time-shift (DTS) panels, $d''_g(t)$, for a point target in data set `geabr_2` reveal minor amplitude fluctuations along the aperture, but no other aberration-induced artifacts. (a) Initial focus and (b) after operator updating.

The NNCC algorithm appears best in this test, with OFF also a close match at both angles and all ranges. NFSR with subarrays fails at 80 and 120 mm, perhaps because the aberration varies rapidly enough across the aperture to disrupt the steering of individual subarrays. Most interesting is the behavior of the speckle brightness algorithm. The magnitude of the aberration (more than one cycle peak-to-peak) is sufficient to cause cycle-skipping artifacts which become worse as the focus is moved off-axis.

In the following sections, it will be seen that other data sets pose a much more difficult aberration correction problem, despite having a much smaller τ_{p-p}.

6.3.3 Data set `geabr_2`

As evidenced by the differential time-shift panels in Figure 6.11, the aberration in this data set causes minor amplitude fluctuations along the aperture. Otherwise, it appears as easy to correct as the synthetic time-shifts in data set `ats_syn`. A look at the OFF-derived aberration profiles for a wide range of steering angles (Figure 6.12) confirms that the aberration is in the extreme near-field; there is very little change in the profile with θ. (The claim that this is not due to the OFF algorithm itself is supported by the observation that the OFF-derived profiles are able to correct a wide range of azimuths in the image, and also by the strong θ-dependence

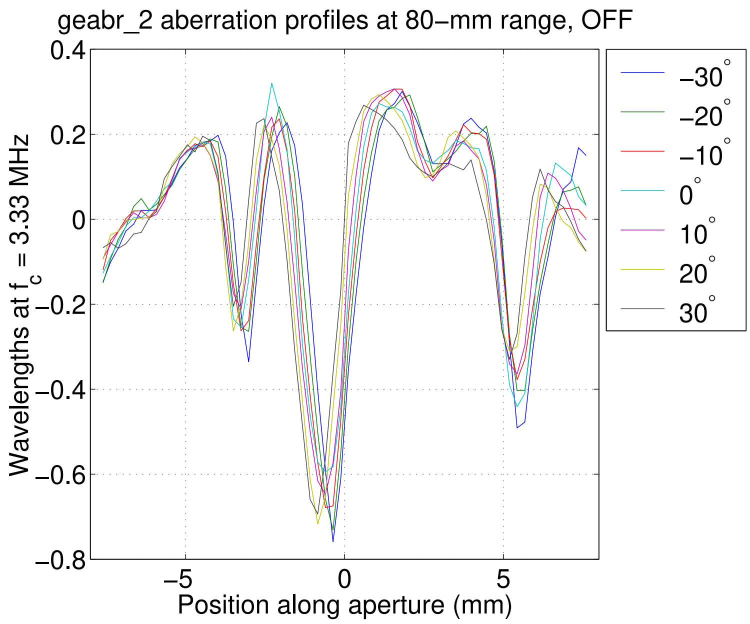

Figure 6.12 The OFF-estimated aberration profiles for data set `geabr_2` are almost angle-independent, indicating very-near-field aberration.

observed for data set `ats_2ab1`.) This was expected, given the description in [24] of a thin silicone aberrator attached directly to the surface of the array transducer.

An uncorrected image and five corrected images using aberration estimates from the speckle brightness, NNCC, dynamic programming, NFSR subarray, and OFF algorithms are shown in Figure 6.13. Forty-five ROIs were specified for $\theta = -40°, -30°, \ldots, 40°$ and $r = 40, 60, \ldots, 120$ mm. After weeding out bad profiles (those for which the associated best-fit operator was not hyperbolic, and those having a peak value above a reasonable cutoff threshold), there were 27 profiles used for the speckle brightness and OFF images, 18 for the NFSR image, and 14 for the NNCC image. The dynamic programming algorithm found 17 profiles.

The best images are provided by the speckle brightness, dynamic programming, and OFF algorithms. All of these algorithms apparently had trouble finding correct operators for the far-right edge of the image, perhaps because of the bright point-targets nearby. The dynamic programming algorithm found at least one operator with an erroneous linear tilt (probably the result of cycle-skipping), and this caused the lowest bright cyst and a nearby point target to be displaced. The NNCC image matches the published image [24] except for the anechoic cyst at 70-mm depth. With *a priori* information that the aberration profiles are mostly independent of θ, the $\theta = 0°$ profiles could have been used at all angles, instead of attempting to correct for different angles. In this case the image would improve.

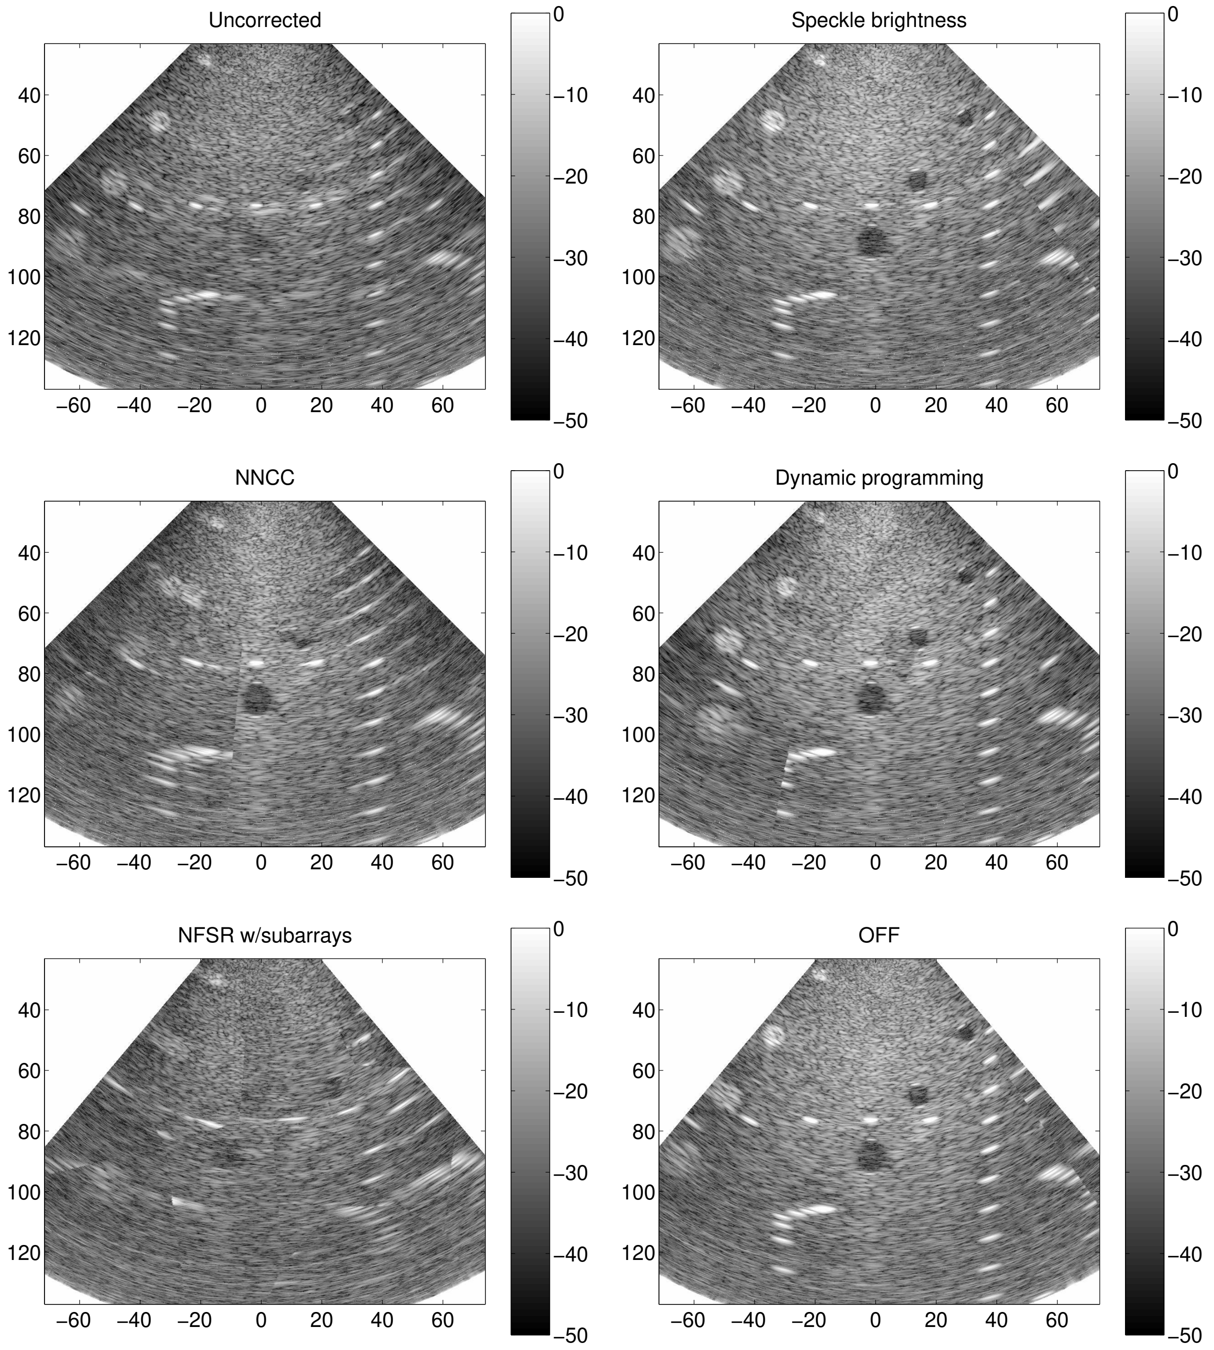

Figure 6.13 Images of data set geabr_2—uncorrected, and corrected with aberration profiles supplied by five different algorithms. (Axis labels are in millimeters.)

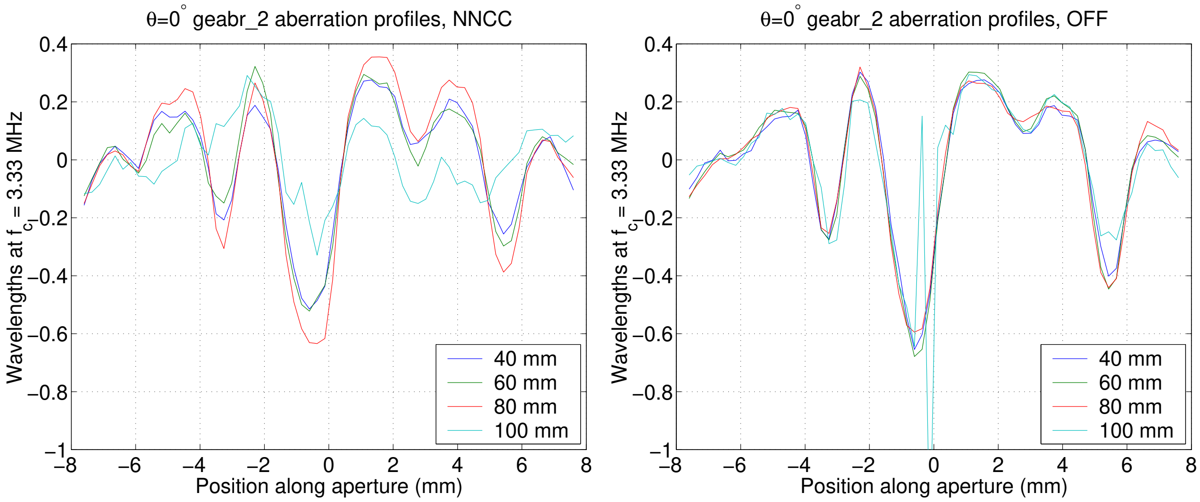

(a) (b)

Figure 6.14 Comparing the estimated aberration profiles at several depths using NNCC (a) and OFF (b) shows that OFF performs better than NNCC at large depths, even broadside to the array. Note that data set `geabr_2` was collected using no averaging, so the available SNR should be similar to a clinical situation.

Because this data set was acquired without any averaging [71], it offers a good opportunity to compare the low-SNR performance of the various correction algorithms in a setting similar to that which might be encountered in clinical practice. It is often believed that aberration correction using complete data sets is infeasible due to the SNR penalty incurred by transmitting and receiving on single elements. As seen in Figure 6.14, this is not necessarily the case. The performance of both algorithms does indeed suffer at large depths, but OFF does much better than NNCC. An authoritative comparison would require NNCC to be implemented with a true transmit focus, rather than the synthetic aperture used here. Still, this experiment suggests that the redundancy in complete data sets may be exploited to compensate for low SNR.

6.3.4 Data set `ats_2ab1`

The received signals across the array (shown in Figure 6.15 after correction) for this data set exhibit greater amplitude fluctuations than those in `geabr_2`. Some diffraction artifacts may also be visible as the slanting lines below the bright areas. (The line at lower-left is from a strong, off-axis scatterer.)

The aberrator in this case was not located at the array surface, but about 6 mm away from it. Consequently, the aberration profiles are strongly θ-dependent (Figure 6.16, compare with Figure 6.12). Notice that the profiles are almost completely out of phase in the left and right

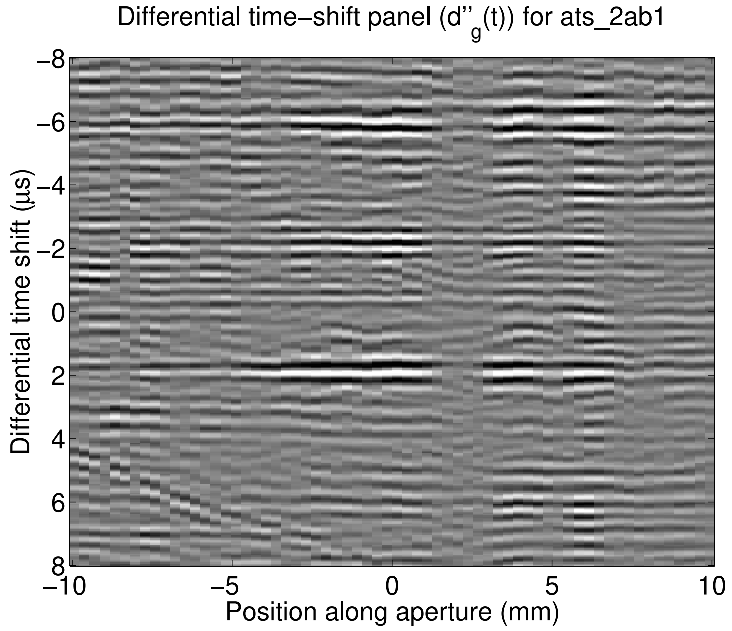

Figure 6.15 A differential time-shift (DTS) panel, $d''_g(t)$, obtained using a correct focusing operator for data set `ats_2ab1` shows large amplitude fluctuations along the aperture. Some diffraction artifacts may also be present.

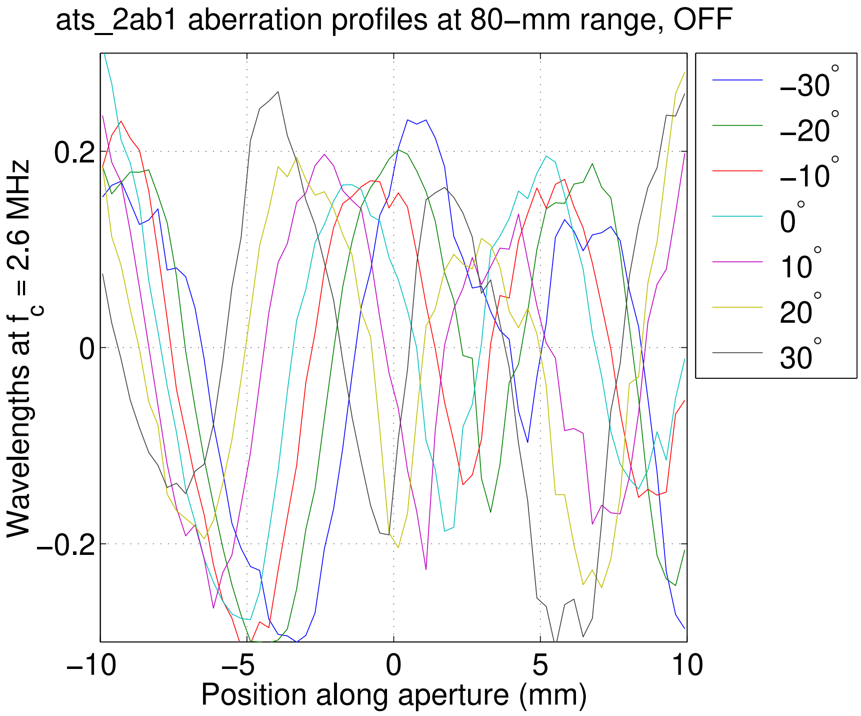

Figure 6.16 In contrast to `geabr_2`, the aberration profiles for `ats_2ab1` vary considerably with azimuth steering; accurate estimates at many azimuth angles are thus needed for good imaging performance.

halves of the image. Profiles over the full range of angles are clearly necessary to achieve good correction throughout the image.

Figure 6.17 shows, for the five aberration correction algorithms, estimated profiles at $\theta = 0°$ for five depths ranging from 40 to 120 mm. "Bad" profiles (lack of convergence, or grossly incorrect) are not plotted. The aberrator is close enough to the array that there should be very little change in the true aberration profile with depth. As the SNR decreases with increasing depth, the consistency of the estimated profiles starts to suffer. Of the algorithms tried, OFF maintains the best consistency over depth.

The uncorrected image and five corrected images from the different algorithms are shown in Figure 6.18. Twenty-eight ROIs were specified for $\theta = -30°, \; -20°, \; \ldots \; 30°$ and $r = 40, \; 60, \; \ldots \; 100$ mm. After deleting the bad profiles, there were 28 profiles used for the speckle brightness and OFF images, 24 for the NNCC image, and 13 for the NFSR image. The dynamic programming algorithm found 64 profiles.

For this data set, OFF is the clear winner. (Examine the visibility of the deeper cysts, the definition of the shallower cysts, and the correction of the large extra-scattering cysts on the right-hand side.) The dynamic programming and speckle brightness algorithms perform well also. NFSR and NNCC are both hampered by an inability to find the correct aberration profiles at large angles from the array normal.

6.3.5 Data set `ats_4ab1`

This data set was obtained by scanning the ATS phantom through the "thick" silicone aberrator at 9-mm stand-off distance from the transducer. Notice the strong amplitude fluctuations and obvious diffraction artifacts in the DTS panel (Figure 6.19) for the return from a point target. Out of the data sets considered here, this is by far the most difficult to form images with. Because none of the aberration correction algorithms had excellent performance on this data set, the true aberration profiles are not known, but an examination of the various estimates suggests that the peak-to-peak aberration, τ_{p-p}, is about 300 ns (Table 6.2), not dramatically larger than the 180 ns for data set `ats_2ab1`. The difference may be due to the increased distance (9 mm) from the array.

The uncorrected image and five corrected images from the different algorithms are shown in Figure 6.20. Twenty-eight ROIs were specified for $\theta = -30°, \; -20°, \; \ldots \; 30°$ and $r = 40, \; 60, \; \ldots \; 100$ mm. After deleting the bad profiles, there were 28 profiles used for the speckle

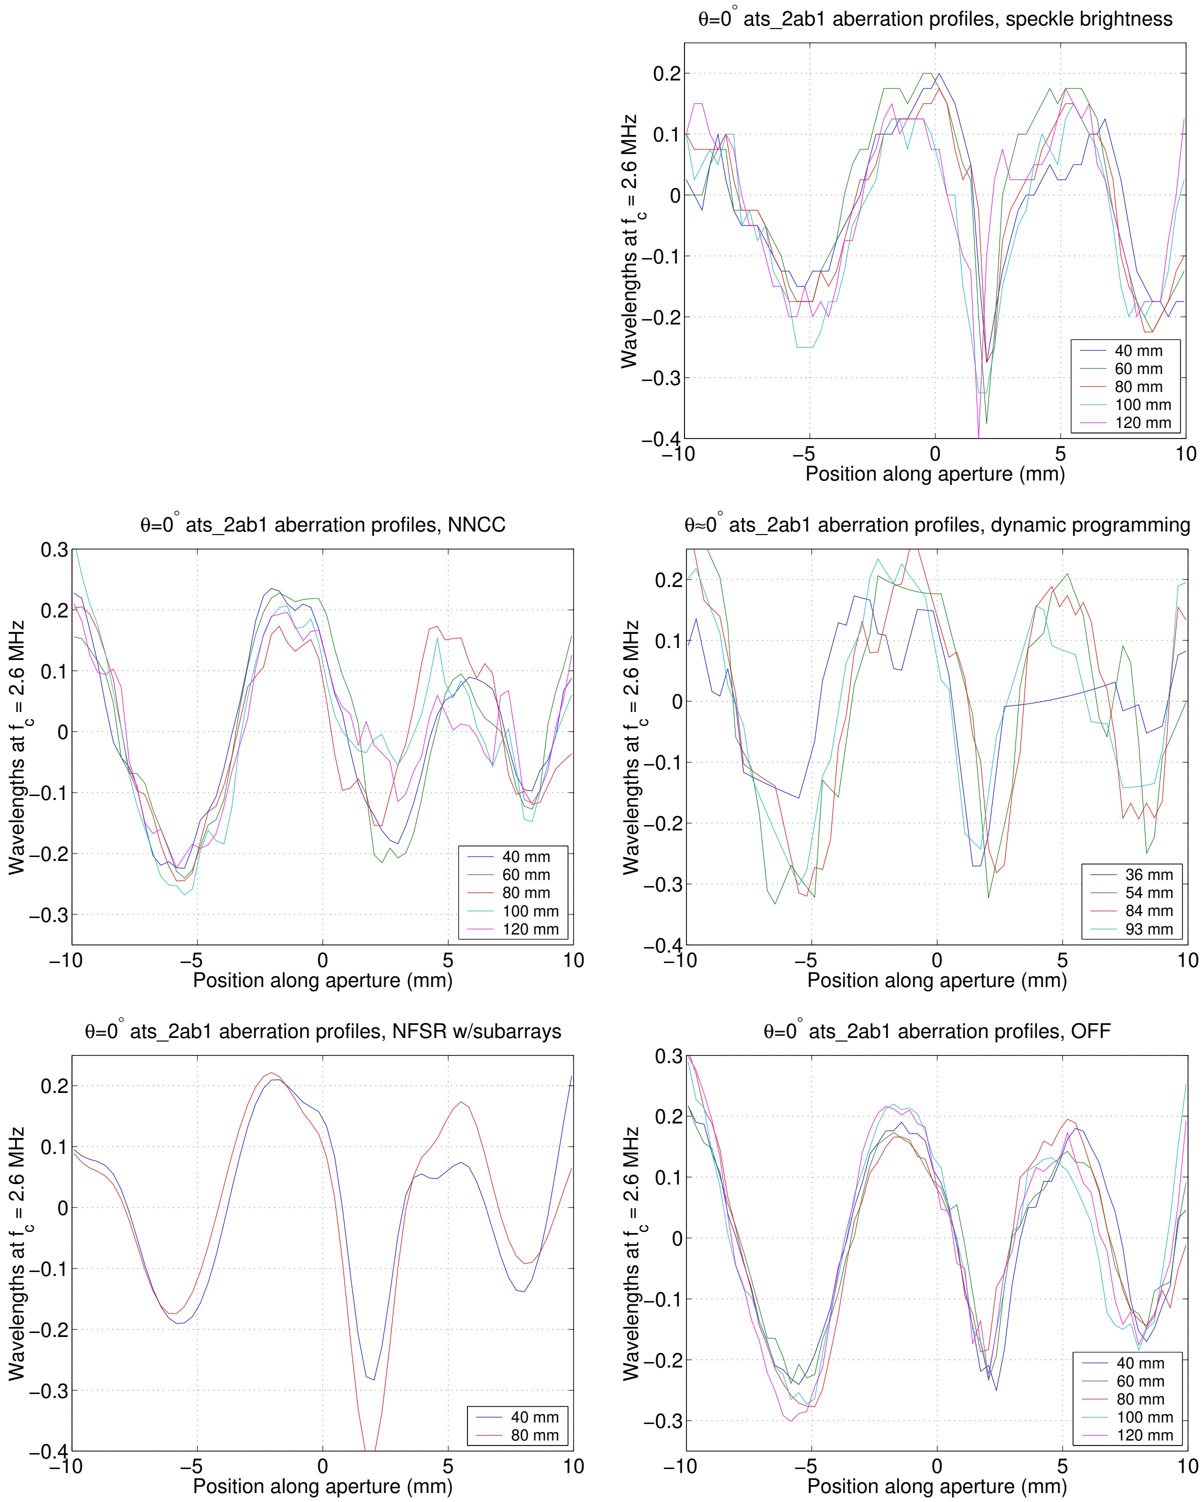

Figure 6.17 Aberration-correction performance generally deteriorates with increasing depth, as illustrated by these estimated profiles for a series of depths in `ats_2ab1`, broadside to the array. Missing curves indicate a grossly incorrect profile at that depth. Note the excellent agreement of the profiles for the overdetermined, fan-filtering (OFF) algorithm.

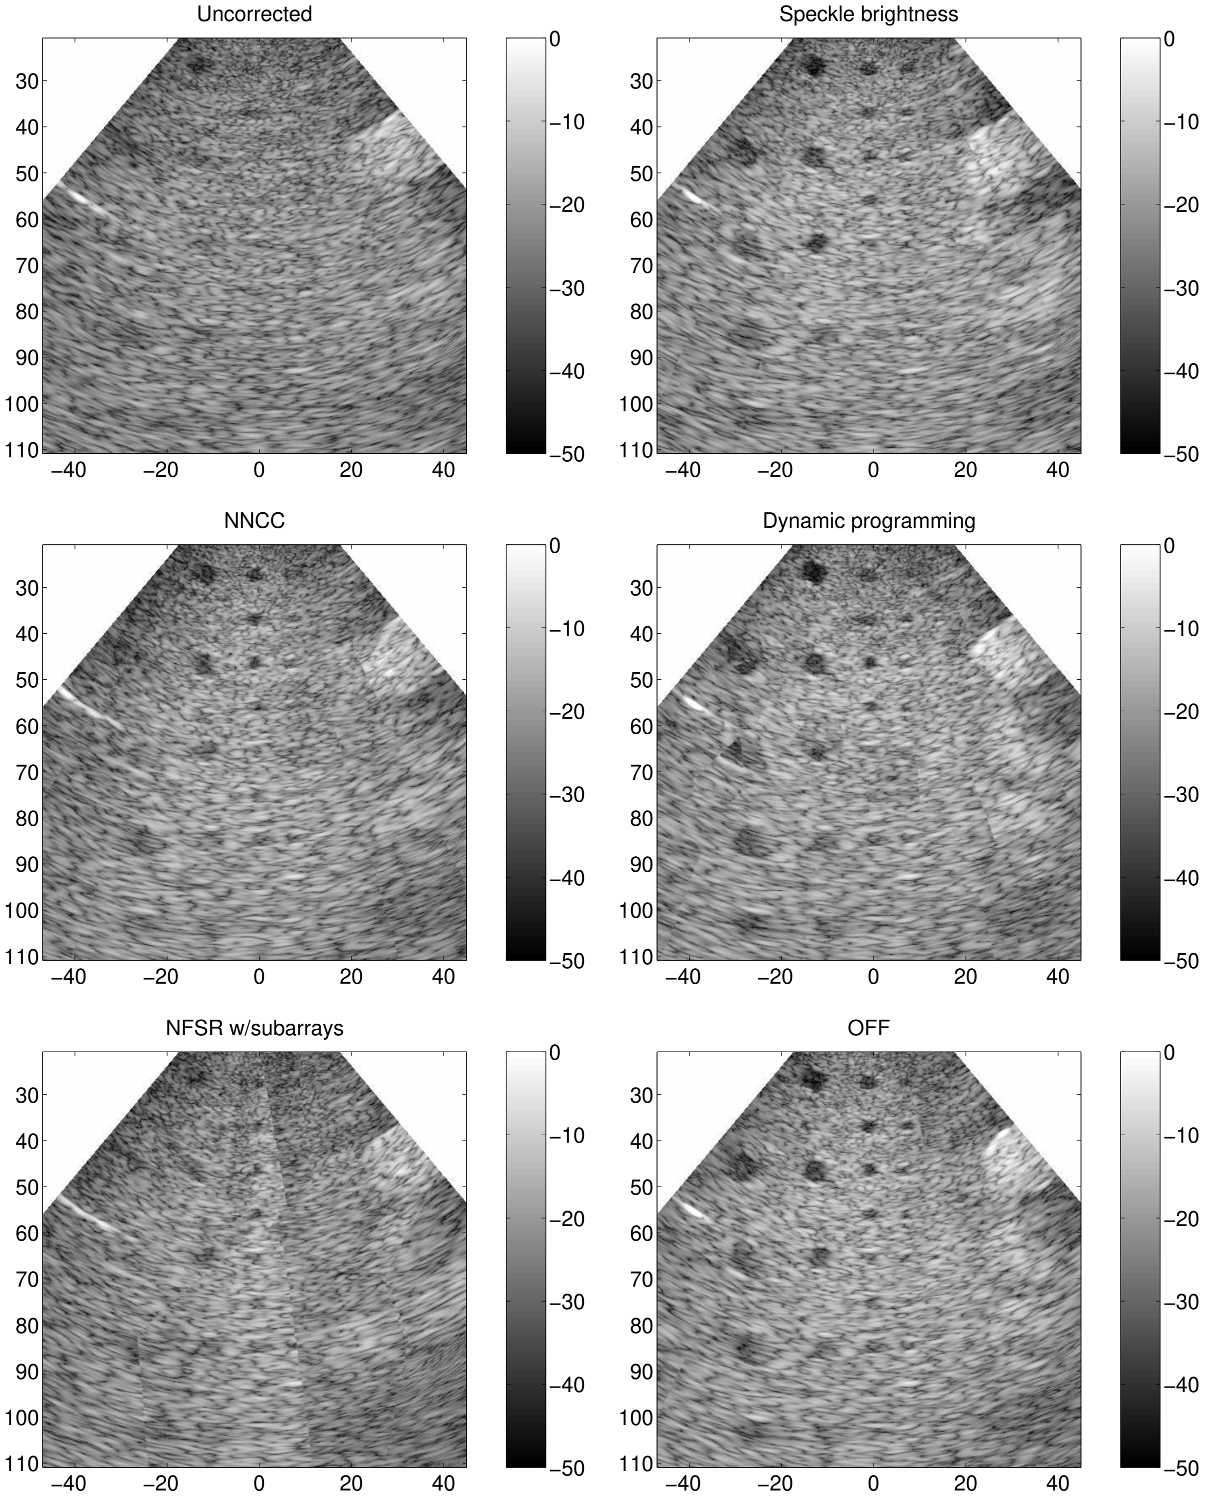

Figure 6.18 Images of data set `ats_2ab1`—uncorrected, and corrected with aberration profiles supplied by five different algorithms. (Axis labels are in millimeters.)

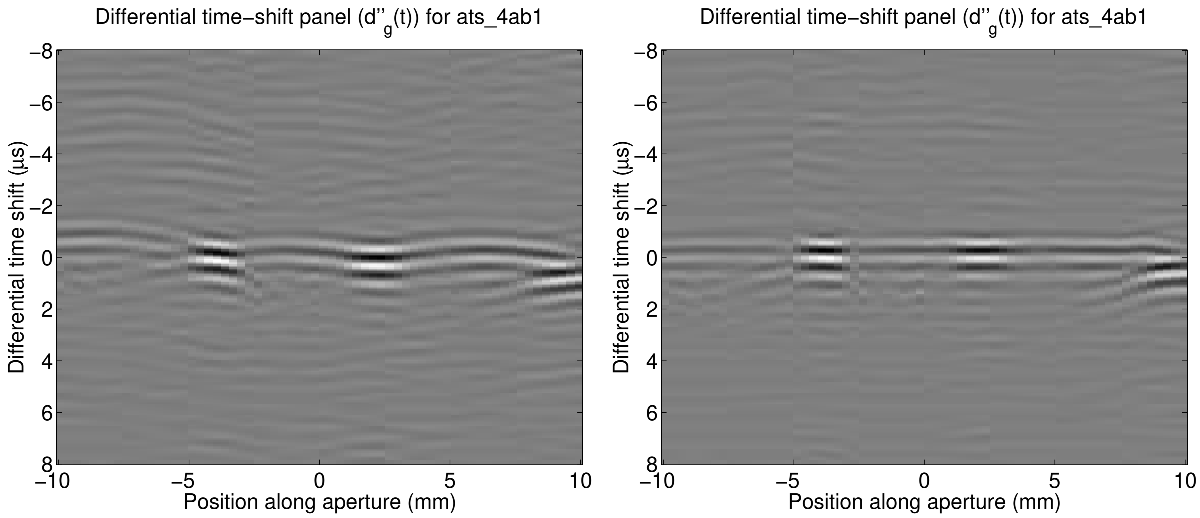

Figure 6.19 The reflection from a point target in data set `ats_4ab1` exhibits strong amplitude fluctuations and diffraction artifacts in this differential time-shift (DTS) panel.

brightness image, 25 for the NNCC image, 23 for the NFSR image, and 21 for the OFF image. The dynamic programming algorithm found 57 profiles.

Although the speckle brightness and NNCC algorithms yield marginal image improvement, only the OFF and dynamic programming algorithms are successful in revealing numerous anechoic cysts. The cause of the bright arcs in the upper right-hand corner of the images is unknown.

Experimentation with the dynamic programming algorithm suggested that a "bootstrap" method would allow it to converge with higher probability on a greater number of targets in the image. An experiment was performed using the aberration profiles estimated by the OFF algorithm to construct initial focusing operators for the dynamic programming algorithm. The result is shown in Figure 6.21. Although not every part of the image has improved, the overall image is better than either the OFF-only or the dynamic-programming-only results. This combination of one algorithm using common-midpoint signals and another algorithm based on a transmit focus deserves further investigation. They appear complementary, since the former can yield rough aberration estimates even in severe aberration, while the latter is subject to misconvergence if not carefully initialized.

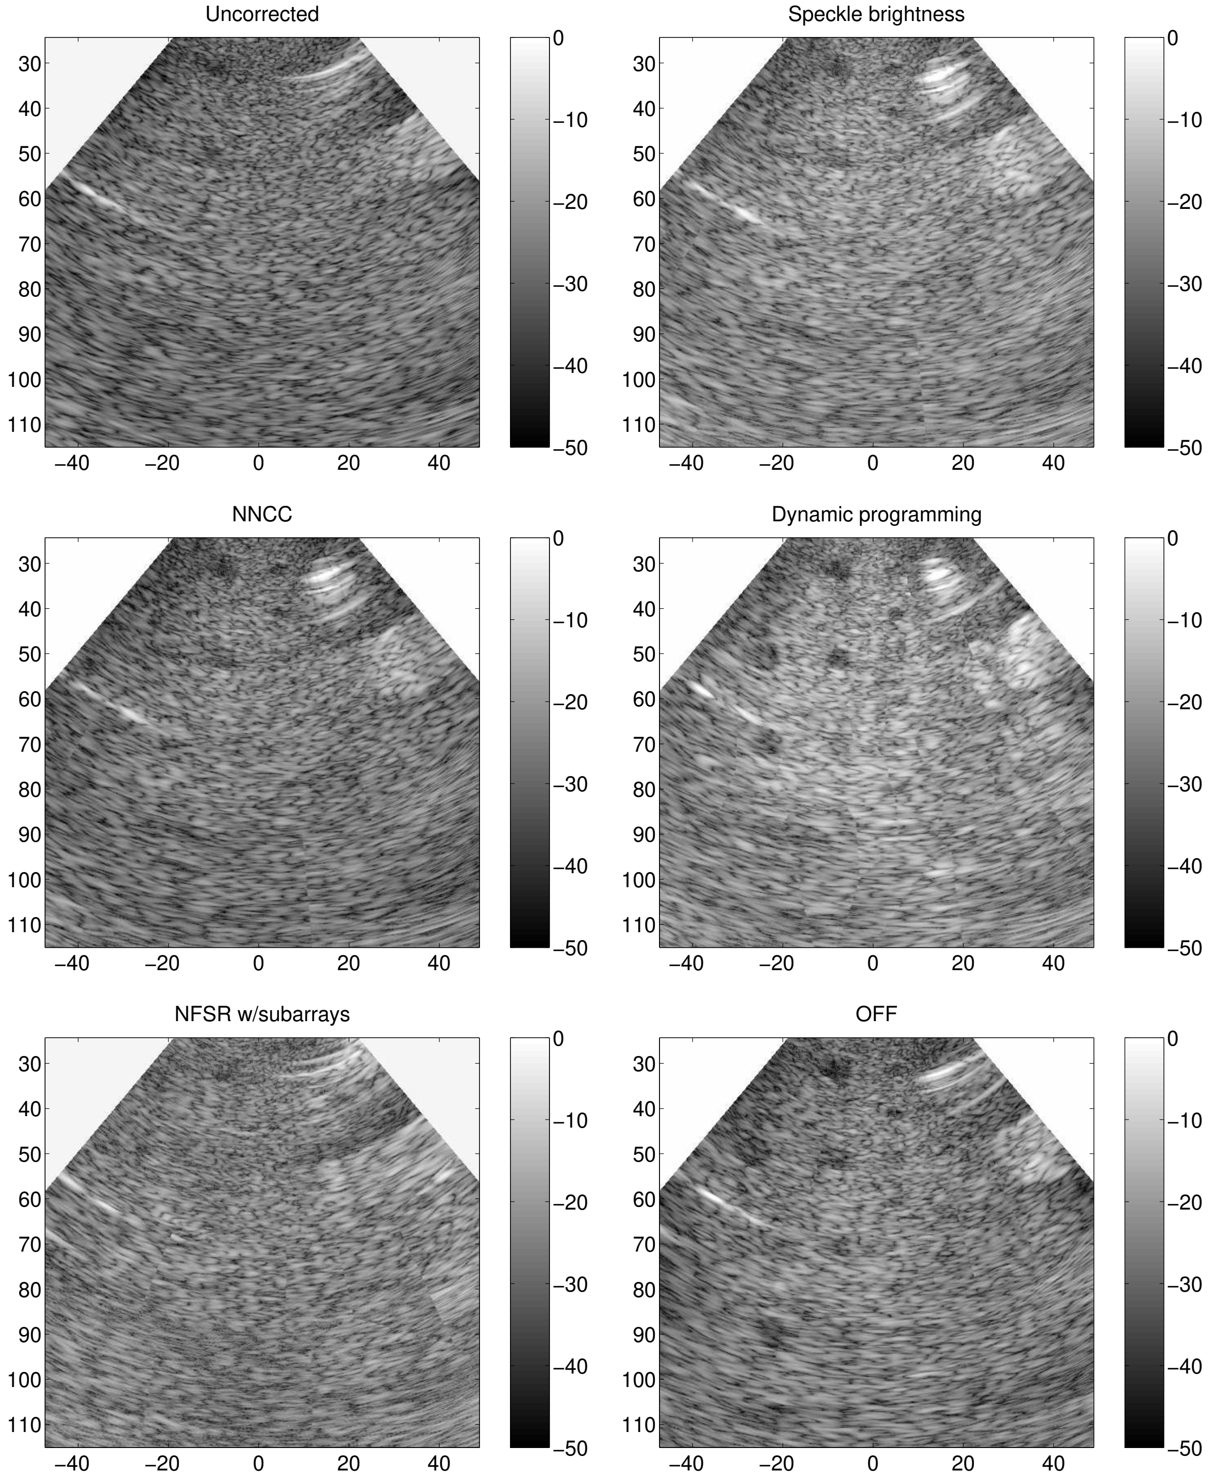

Figure 6.20 Images of data set `ats_4ab1`—uncorrected, and corrected with aberration profiles supplied by five different algorithms. (Axis labels are in millimeters.)

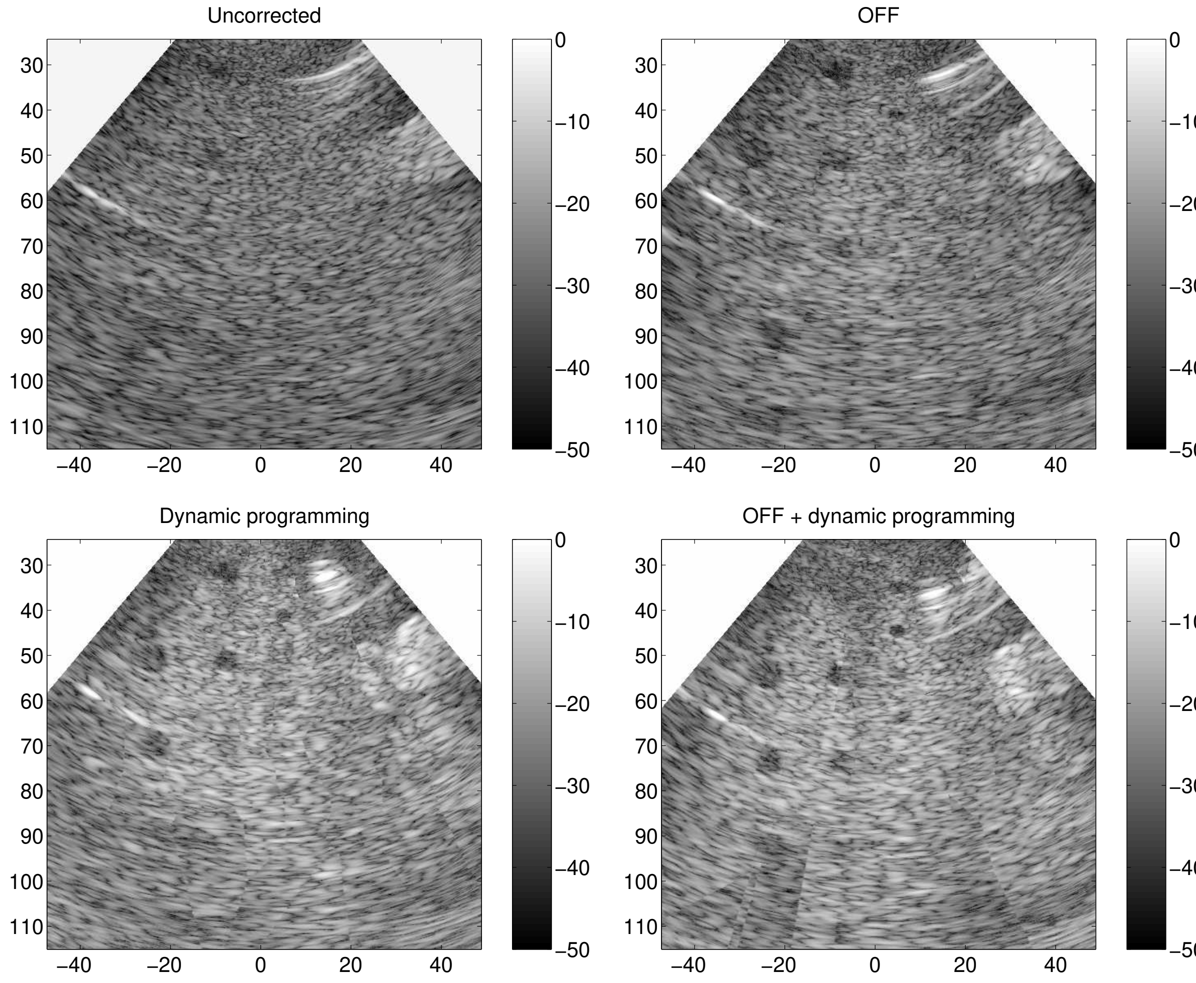

Figure 6.21 Deriving initial aberration profiles using OFF, then refining them with the dynamic-programming, operator-updating algorithm results in a better image of data set `ats_4ab1` than those obtained with either algorithm alone. (Axis labels are in millimeters.)

CHAPTER 7

EFFICIENT, THREE-DIMENSIONAL, CYLINDRICAL-APERTURE IMAGING

An ongoing research project is developing very small ultrasound transducers that can be fabricated on the side of a needle and operated *in vivo* at high frequencies. Such probes could eventually provide a minimally invasive alternative to biopsy and speed the diagnosis of tumors. Crucial to the success of this effort are imaging algorithms adapted to the geometry imposed by these ultrasonic microprobes [4].

This image formation problem is challenging due to the microprobes' shape. With one probe, only two transducer motions are possible: inward and outward travel, and rotation about the needle axis. The imaging aperture is therefore highly curved and its spatial extent is severely limited. Because the microprobe transducer will be surrounded by a scattering volume, good resolution in three dimensions is desirable so that quality 2-D image slices may be obtained. This can be accomplished using synthetic aperture techniques, but a significant 2-D aperture is required. It is imperative, then, that the available probe diameter be used efficiently.

Many of these constraints also apply to other imaging systems. Some intravascular ultrasound (IVUS) systems, for example, use a circular array of transducers on a catheter to image the interior of blood vessels [3]. As new ultrasound imaging modalities are developed, it is anticipated that cylindrical apertures will become more common and will benefit from ongoing work in this area.

7.1 Synthetic Aperture Imaging in a Cylindrical Geometry

Circular apertures have previously been used for intravascular ultrasound (IVUS) imaging systems, where synthetic aperture or array focusing has usually been carried out in the time domain [3, 72]. Frequency-domain algorithms, however, have a large speed advantage over traditional beamforming methods due to the computational efficiency of the fast Fourier transform [73]. (The computational requirements of conventional beamforming and the proposed frequency-domain algorithm are compared at the end of this section.)

A frequency-domain imaging method was recently proposed for use with IVUS systems [74]. The authors start with a geometrically derived, two-dimensional point spread function (PSF) and obtain the Fourier transform of the imaging kernel for monostatic (colocated transmitter and receiver) and bistatic cases using the method of stationary phase.

In the following derivation, the PSF for three-dimensional, monostatic (colocated source and receiver) imaging from a cylindrical aperture is obtained using the Rayleigh-Sommerfeld formulation of scalar diffraction theory. This PSF is then compared with the PSF for wave propagation in Cartesian coordinates, which has a well-known Fourier transform. This approach makes clear the narrow-beamwidth approximation necessary to put the PSF for cylindrical coordinates into the same form and obtain its Fourier transform. Because of the monostatic imaging assumption, the resulting algorithm is best suited to applications where it is only practical to have a single-element transducer.

We start with the Rayleigh-Sommerfeld formula [75],

$$U(P_0) = \iint\limits_{\Sigma} \frac{1}{j\lambda} U(P_1) \frac{e^{jkr_{01}}}{r_{01}} \cos(\vec{n}, \vec{r}_{01}) ds, \tag{7.1}$$

which expresses the field at P_0, located on the imaging aperture, in terms of a source distribution on the surface Σ shown in Figure 7.1. Given monostatic data acquisition in a constant-sound-speed, weakly scattering medium, no generality is lost by considering the scatterers on Σ to be the original source of ultrasonic waves traveling at speed $c/2$; this is the "exploding reflectors" model sometimes used in exploration seismology [9]. Rewriting (7.1) using cylindrical coordinates, we have

$$\begin{aligned} U(R, \phi, z) = {} & \int_{-\infty}^{\infty} \int_{-\pi}^{\pi} \frac{1}{j\lambda} U(r_i, \phi_i, z_i) \\ & \cdot \frac{\exp(jk\sqrt{R^2 + r_i^2 - 2Rr_i \cos(\phi - \phi_i) + (z - z_i)^2})}{R^2 + r_i^2 - 2Rr_i \cos(\phi - \phi_i) + (z - z_i)^2} \\ & \cdot (r_i - R\cos(\phi - \phi_i)) r_i d\phi_i dz_i \ . \end{aligned} \tag{7.2}$$

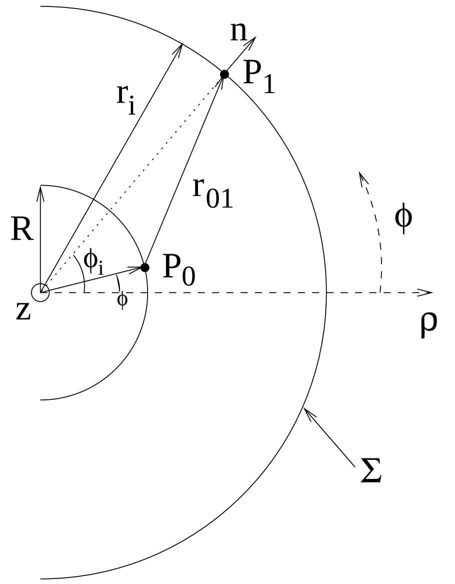

Figure 7.1 Cylindrical geometry for derivation of the point spread function (PSF).

This is a convolution integral; the point spread function is

$$h(\phi, z) = \frac{1}{j\lambda}(r_i - R\cos\phi)r_i \frac{\exp(jk\sqrt{R^2 + r_i^2 - 2Rr_i\cos\phi + z^2})}{R^2 + r_i^2 - 2Rr_i\cos\phi + z^2} \; . \tag{7.3}$$

If the transducer beamwidth is not too wide, the cosine terms in the exponential and the denominator may be approximated to second order as $1 - \phi^2/2$. If the other cosine term in the numerator is approximated to first order, we have

$$h(\phi, z) \approx \frac{1}{j\lambda}(r_i - R)r_i \frac{\exp(jk\sqrt{(R - r_i)^2 + Rr_i\phi^2 + z^2})}{(R - r_i)^2 + Rr_i\phi^2 + z^2} \; . \tag{7.4}$$

When is the approximation $\cos\phi \approx 1 - \phi^2/2$ valid? The square-root quantity being approximated in (7.2) is r_{01}, the distance from the transducer to a reflector in the scattering volume, expressed in cylindrical coordinates. Set $z = z_i = 0$ and consider the angle range, $\pm\phi$, over which echoes may be collected from a point reflector at radius r_i. From Figure 7.2, ϕ is related to the transducer beamwidth $\gamma_{r\phi}$ by

$$\frac{r_i \sin\phi}{\sqrt{R^2 + r_i^2 - 2Rr_i\cos\phi}} = \sin\frac{\gamma_{r\phi}}{2} \; . \tag{7.5}$$

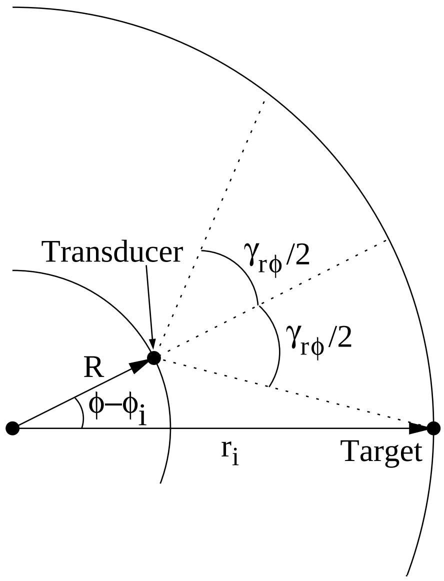

Figure 7.2 Transducer beamwidth in relation to the positions of the transducer and target in cylindrical coordinates.

This is a quadratic equation in $\cos\phi$; the solution (assuming all angles are in the first quadrant) is

$$\phi = \arccos\left(\frac{R\sin^2\frac{\gamma_{r\phi}}{2} + \cos\frac{\gamma_{r\phi}}{2}\sqrt{r_i^2 - R^2\sin^2\frac{\gamma_{r\phi}}{2}}}{r_i}\right) . \tag{7.6}$$

A simple (though computationally demanding) reconstruction method from synthetic aperture data is to back-project the time series recorded at each transducer position, smearing the data back onto the loci of constant travel time in the target space. Even though the imaging algorithm proposed next operates in the Fourier domain, simple back-projection is a useful concept for understanding the effects of the distance approximation (see Figure 7.3).

The validity of the narrow-beamwidth approximation may be tested for a given cylindrical imaging problem by imposing a limit on the distance error

$$\sqrt{(R-r_i)^2 + Rr_i\phi^2} - \sqrt{R^2 + r_i^2 - 2Rr_i\cos\phi} \le K \tag{7.7}$$

with $K = \lambda/2$, for example. (For simplicity, it is assumed that $z = z_i = 0$.) Sample plots of the error (in wavelengths) versus the ratio r_i/R are given in Figure 7.4 for three different transducer beamwidths when $R = 50\lambda$, a typical value for the high-frequency ultrasonic microprobes being developed. Although the exact errors will be different for every cylindrical imaging scenario, the

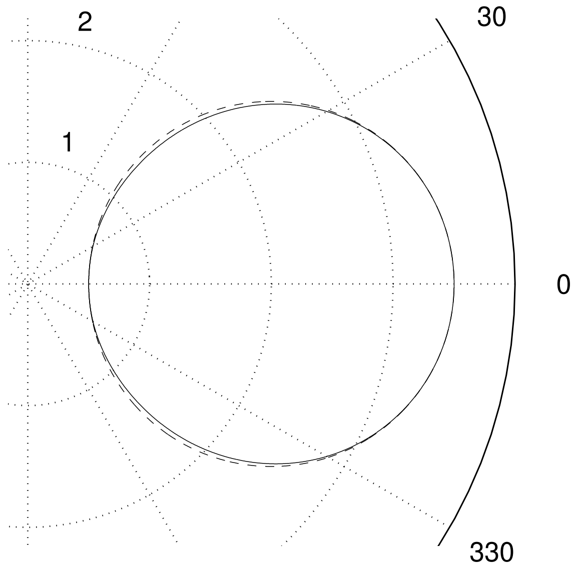

Figure 7.3 The effects of the second-order cosine approximation may be visualized by considering circles of constant travel time. Here, a circular wavefront departing $R = 2$, $\phi = 0$ at $t = 0$ is shown at $ct = 1.5$ using the exact (dashed line) and approximate (solid line) forms of the distance function. The close match at small angles is evident.

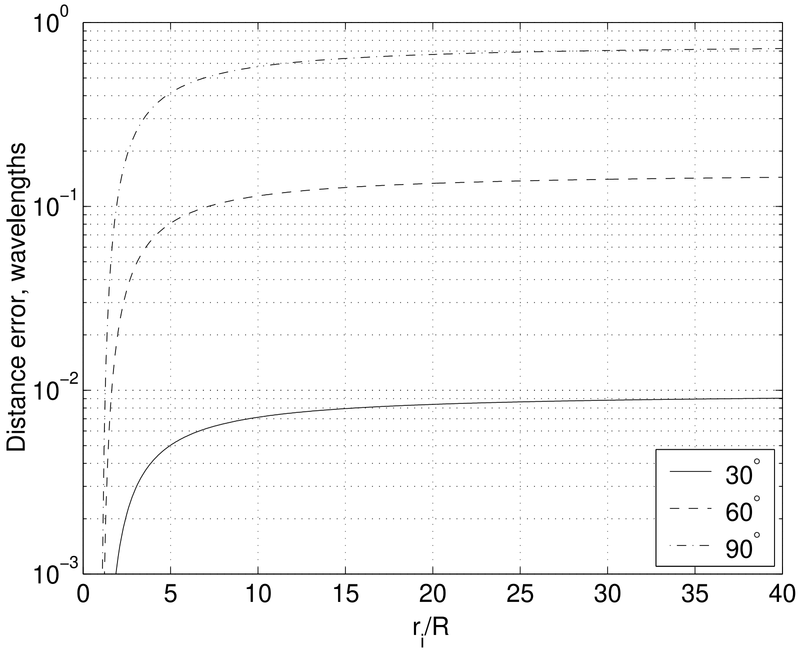

Figure 7.4 Error analysis for the second-order cosine approximation. These curves are plotted for transducer beamwidths of 30°, 60°, and 90° when $R = 50\lambda$.

narrow-beamwidth approximation will usually break down only when the transducer beamwidth becomes extremely wide.

Returning to Equation (7.4), compare this approximate PSF with the PSF obtained for a rectangular aperture in Cartesian coordinates:

$$h(x,y) = \frac{1}{j\lambda} \cdot \frac{d\exp(jk\sqrt{d^2+x^2+y^2})}{d^2+x^2+y^2} \ . \tag{7.8}$$

Clearly they have the same form, except that the angular variable is scaled by the geometric mean of the transducer and reflector radii. The Fourier transform of $h(x,y)$ is well known [75], and the transform of $h(\phi,z)$ follows easily using the scaling property:

$$H(f_\phi, f_z) \approx \sqrt{\frac{r_i}{R}}\ e^{j2\pi(r_i-R)\sqrt{\frac{1}{\lambda^2}-\frac{f_\phi^2}{Rr_i}-f_z^2}}, \tag{7.9}$$

where the fact that ϕ is an angular variable (and $h(\phi,z)$ is thus periodic in ϕ) has been ignored, and $\lambda = c/2f$ is in accordance with the exploding reflectors model.

In the spatial-frequency (f_ϕ, f_z) domain, multiplication by H is equivalent to propagating the wave field from one concentric cylindrical surface to another. This is the principle behind what are known as *Fourier migration* algorithms in the seismic exploration community [13] and *wavenumber* or $\omega - k$ algorithms in the radar community [14]. The most significant difference here is that now H is a function not only of the relative distance between the target and recording surfaces, but also of the absolute radii of those surfaces.

Following [76], a simple way to obtain a two-dimensional image at one desired depth is to first compute the 3-D FFT of the raw data, taking it into the (f_ϕ, f_z, f) domain. Then, for each temporal frequency f, the (f_ϕ, f_z) spatial frequency planes are multiplied by the appropriate $H(f_\phi, f_z)$ for the target reconstruction depth. Finally, the resulting data cube is summed over temporal frequency f, and an inverse 2-D FFT yields the focused image. The required discretization of this process is straightforward, but the spatial sampling in ϕ and z must be dense enough to prevent artifacts; an in-depth discussion of this issue may be found in [55].

In practice, if the full three-dimensional image is desired, the H operator will be applied repeatedly to back-propagate the wavefield in increments of the axial resolution. At each iteration, the data are summed over f and inverse transformed to yield the next cylindrical image slice. This is analogous to the seismic migration technique known as *phase-shift migration* [12].

Table 7.1 Computational cost of the proposed frequency-domain imaging algorithm compared to conventional focusing.

Operation	Real multiplications	Real additions
(Conventional beamforming)		
Delay and sum N_s^2 elements, $N_s^2 N_t$ times	$N_s^4 N_t$	$N_s^4 N_t$
(Proposed algorithm)		
3-D FFT	$2N_s^2 N_t(2\log_2 N_s + \log_2 N_t)$	$3N_s^2 N_t(2\log_2 N_s + \log_2 N_t)$
Mult. by H, N_t times	$4N_s^2 N_t^2$	$2N_s^2 N_t^2$
Sum over f, N_t times	0	$2N_s^2 N_t^2$
2-D FFT, N_t times	$4N_s^2 N_t \log_2 N_s$	$6N_s^2 N_t \log_2 N_s$
Total	$2N_s^2 N_t(2N_t + 4\log_2 N_s + \log_2 N_t)$	$3N_s^2 N_t(\frac{4}{3}N_t + 4\log_2 N_s + \log_2 N_t)$

To evaluate the computational efficiency of this approach, consider an experiment in which echoes are recorded at $N_s \times N_s$ positions on a curved aperture. Let N_t time samples be recorded for each signal. The imaging algorithm will be tasked with computing $N_s^2 N_t$ voxels in the target volume. Conventional beamforming is here viewed as a series of delay-and-sum operations on the raw (ϕ, z, t) data set. Operation counts for the frequency-domain algorithm follow directly from the steps described above. The results are listed in Table 7.1.

For the parameters used in the simulations and experiments of the following sections ($N_s = 128$, $N_t = 512$), conventional beamforming would require about 1.4×10^{11} real multiplies and adds. The proposed frequency-domain algorithm requires about 1.8×10^{10} multiplies and adds, a savings of 87%. While a host of implementation-dependent factors make it impossible to categorically declare one algorithm as "best," the potential gains of working in the frequency domain are evident. Other factors which might favor a frequency-domain implementation are the ease of parallelizing the FFTs and the elimination of oversampling requirements.

7.2 Options for Effective Use of the Available Aperture

When a single-element, unfocused transducer is significantly larger than a wavelength and has an effective aperture D, the best lateral resolution available with synthetic aperture processing is approximately $D/2$ [77]. For high-frequency ultrasound applications, this would seem to suggest using a transducer that is extremely small, on the order of one wavelength! Such a small transducer would place severe limits on the available pulse energy in an application where the low signal-to-noise ratio is already a major concern.

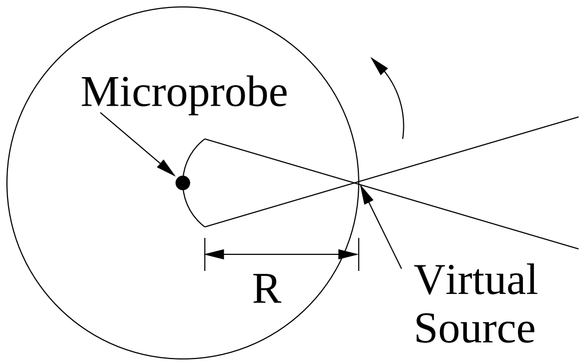

Figure 7.5 Using a focused transducer to create a virtual source element.

One possible solution is to use the full probe diameter for a focused transducer and treat the focus as a virtual source [78]. The virtual source traces out a cylindrical surface, as shown in Figure 7.5. The imaging algorithm proceeds as if there were an unfocused transducer located at the focus, treating that point as a source of diverging waves. It can then form images at depths past the real focus. A recent study exploring this technique found that the resolution achievable beyond the focus is comparable to the resolution at the focus, this fundamental limit being set by the focal ratio of the transducer [79].

Using a focused transducer to create a virtual source has the advantage that more transducer area, and hence more energy, is available for transmitting. It is important to realize, however, that we are not getting something for nothing: The usable aperture is still the same, even though the virtual source may be traveling on a much larger surface. For a transducer of some constant diameter, as the focal length is increased, the width of the focal region increases and the beamwidth past the focus decreases. These factors conspire to limit the attainable resolution in the virtual source case.

7.3 Resolution

7.3.1 Axial resolution

As in all pulse-echo imaging systems, the axial resolution is determined by the length of the pulse, although this may be shorter than the actual transmitted waveform if pulse compression is used. If the transmitted energy is uniform over a bandwidth W, the axial resolution (using the Rayleigh criterion) will be approximately

$$\delta_r = \frac{c}{2W} \tag{7.10}$$

corresponding to $c/2$ times the interval from peak to first zero-crossing of the sinc function $\sin(\pi W t)/\pi t$.

7.3.2 Lateral resolution

The lateral resolution attainable in this imaging geometry is easiest to calculate when the transducer size is small compared to the wavelength. As the probe moves, the echoes from any given scattering point are phase modulated by the changing scatterer-transducer distance. The amount of phase modulation governs the extent of the data collected in the spatial frequency domain and hence the resolution. Larger transducers cause a smearing of the raw data that imposes a window on the data in the spatial frequency domain and degrades the resolution.

From Equation (7.2), the phase of the received signal at (R, ϕ, z) due to a point reflector at (r_i, ϕ_i, z_i) is, using the exact form of the distance function,

$$\psi(\phi, z) = 2\pi f \frac{2}{c} \sqrt{R^2 + r_i^2 - 2Rr_i \cos(\phi - \phi_i) + (z - z_i)^2}\ . \tag{7.11}$$

The spatial derivatives of this function yield the instantaneous spatial frequencies of the unprocessed data:

$$\begin{aligned} f_\phi &= \frac{1}{2\pi} \frac{d}{d\phi} \psi(\phi, z) \\ &= \frac{2f}{c} \cdot \frac{Rr_i \sin(\phi - \phi_i)}{\sqrt{R^2 + r_i^2 - 2Rr_i \cos(\phi - \phi_i) + (z - z_i)^2}} \\ |f_\phi|_{max} &= \frac{2R}{\lambda} \sin \frac{\gamma_{r\phi}}{2}, \end{aligned} \tag{7.12}$$

where $\gamma_{r\phi}/2$ is the half-beamwidth of the ultrasound transducer in the $r\phi$ plane (see Figure 7.2). We also have

$$\begin{aligned} f_z &= \frac{1}{2\pi} \frac{d}{dz} \psi(\phi, z) \\ &= \frac{2f}{c} \cdot \frac{z - z_i}{\sqrt{R^2 + r_i^2 - 2Rr_i \cos(\phi - \phi_i) + (z - z_i)^2}} \end{aligned}$$

$$|f_z|_{max} = \frac{2}{\lambda}\sin\frac{\gamma_{rz}}{2}, \tag{7.13}$$

where $\gamma_{rz}/2$ is the half-beamwidth of the transducer in the rz plane.

We can now estimate the lateral resolution in the z and ϕ directions for the case of a small transducer. The variation of f_z and f_ϕ defines an approximate rectangle in the spatial frequency domain, symmetric about $(f_\phi, f_z) = (0,0)$. Following [80], the resolution under the Rayleigh criterion is determined from the extent of the spatial frequency data as

$$\delta_\phi = \frac{\lambda}{4R\sin\frac{\gamma_{r\phi}}{2}} \text{ rad} = \frac{\lambda r_i}{4R\sin\frac{\gamma_{r\phi}}{2}} \text{ m at depth } r_i \tag{7.14}$$

$$\delta_z = \frac{\lambda}{4\sin\frac{\gamma_{rz}}{2}} \text{ m.} \tag{7.15}$$

What if the transducer is not small with respect to the wavelength? In this case, the aperture impulse response must be convolved with the transmitted pulse, which causes the spatial frequency data to be windowed. See [77] for a detailed derivation of this effect. If the transducer has an effective aperture of D, then the lateral resolution will be approximately $D/2$ (in the z-direction; the resolution in the ϕ-direction is depth-dependent).

If the virtual source technique is used, calculation of the expected lateral resolution is more difficult because of the complex field pattern beyond the focus. The above formulas are not applicable, for they assume a relatively flat response within the transducer beamwidth and sharp cutoffs at angles of $\pm\gamma_{r\phi}/2$ and $\pm\gamma_{rz}/2$, leading to a sinc-like response for point targets. For focused transducers, the natural apodization provided by the past-focus beam pattern will broaden the main lobe and reduce the sidelobes. Simulations and experiments have shown the achievable resolution to be about the same as the size of the focal region: approximately $F\lambda$, where F is the focused transducer's focal ratio, or "f-number."

7.4 Simulations

Two types of computer simulations were performed to verify the performance of this imaging method. In the first, an infinitesimally small, unfocused transducer was simulated to test the algorithm and its inherent approximations, independent of the virtual source technique. Five point reflectors in an "X" pattern at radius $r_i = 15$ mm were imaged by a transducer at radius $R = 8$ mm having a beamwidth of 20°. A Gaussian-weighted 5-MHz sinusoid with a 3-MHz bandwidth was used for the transmit pulse. Simulated echoes were collected on a 128x128 grid with ϕ ranging from $-55°$ to $+55°$ and z ranging from -6.4 to $+6.4$ mm. These data were processed with the algorithm described in Section 7.1; Figure 7.6 shows the results. The lateral resolution is comparable to the values predicted by Equations (7.14) and (7.15): 3.1° in ϕ and 0.43 mm in z. Due to the lack of any apodization in the simulated beam pattern, high sidelobes are to be expected, and in fact the first sidelobes are only 22 dB down.

In the second type of simulation, the *Field II* program [81] was used to simulate a spherically focused transducer and thus validate the virtual source technique used in the experiments. Point reflectors at radius $r_i = 57$ mm were imaged by an f/1.33, 19.1-mm diameter transducer at radius 22.3 mm, yielding a virtual source radius of $R = 47.7$ mm. A Gaussian-weighted 2.25-MHz sinusoid with a half-power bandwidth of 1.1 MHz was used as the transmit pulse. Simulated echoes were collected on a 128x128 grid with a ϕ step size of 0.14° and a z step size of 120 μm. All of these parameters were selected to match as closely as possible the parameters of the first experiment described in the following section.

As can be seen in Figure 7.7, in this simulation the main lobe has broadened and the side lobes are lower, due to the tapering of the transducer beam pattern beyond the focus. The half-power resolution is about 600 μm; as expected, this is comparable to the width of the focal region.

7.5 Experimental Results

A number of experiments were performed in a water tank with a precision positioning system and conventional focused ultrasound transducers in pulse-echo mode. In the first experiment, a 19.1-mm diameter, 2.25-MHz transducer with a measured focal length of 26.5 mm was mounted to a vertical support arm in the water tank and used to scan a target consisting of three 100-μm wires crossing at the center of a plastic holder (Figure 7.8). The central part of the target was approximately 10 mm beyond the focus of the transducer. The transducer was scanned up and down and rotated about the axis of its support arm, covering an area of 15.4 mm by 17.9° in 128 by 128 steps.

After processing, "fly-through" movies were generated showing the imaged wires on a series of cylindrical shells. Reflections are present over a range of depths because the imaging cylinder cuts through the planar target at different places depending on the depth chosen for focusing. To create the right-hand panel in Figure 7.9, these images were summed over depth, creating a 2-D projection with a complete view of the target. Compare the detail visible in this image with

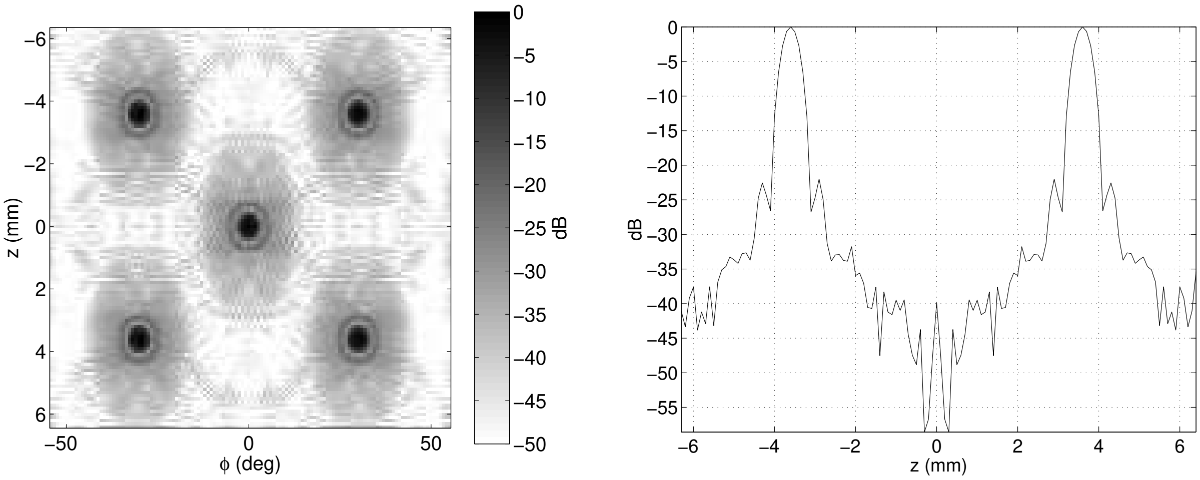

(a) (b)

Figure 7.6 Five point-reflectors at $r_i = 15$ mm imaged with an infinitesimally small transducer at $R = 8$ mm. (a) Focused image and (b) slice at constant ϕ.

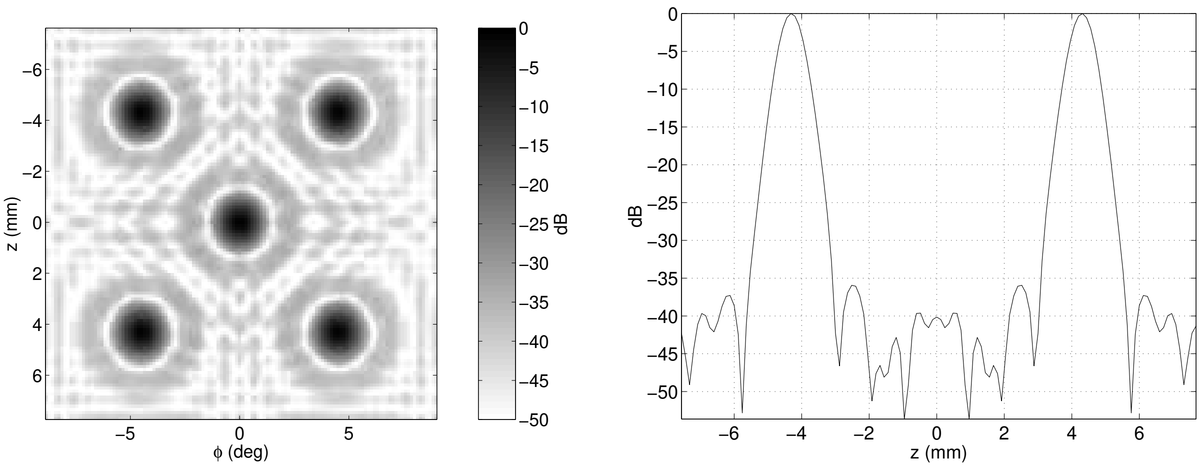

(a) (b)

Figure 7.7 Five point-reflectors at $r_i = 57$ mm imaged using the virtual source technique with a simulated spherically focused transducer at $R = 22.3$ mm. (a) Focused image and (b) slice at constant ϕ.

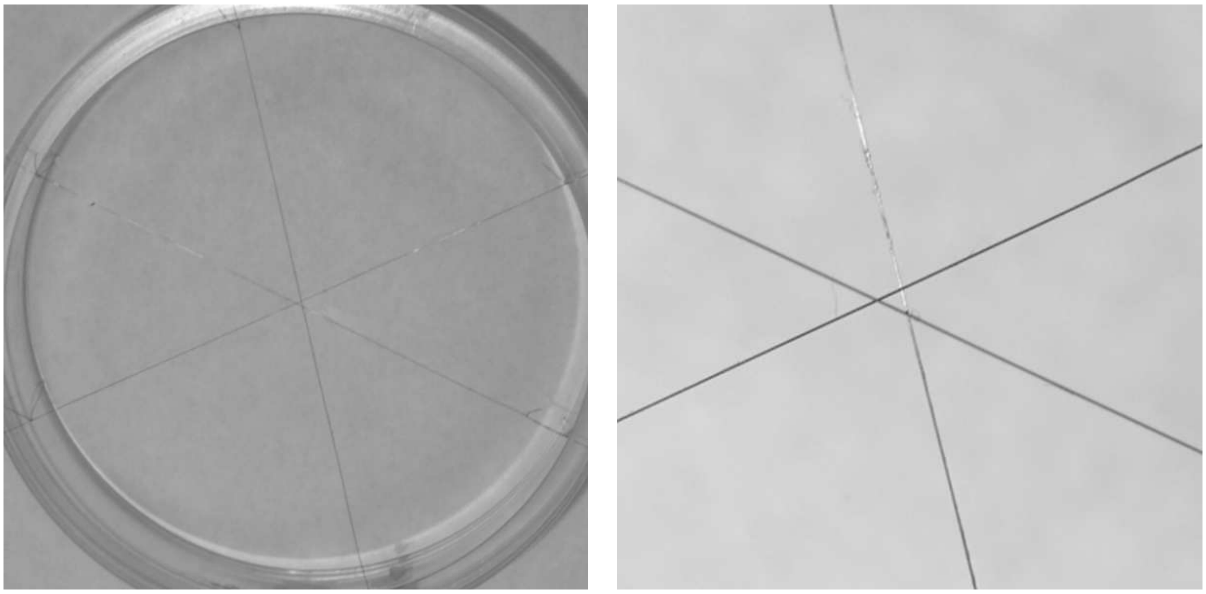

(a) (b)

Figure 7.8 (a) 100-μm wire target and (b) close-up.

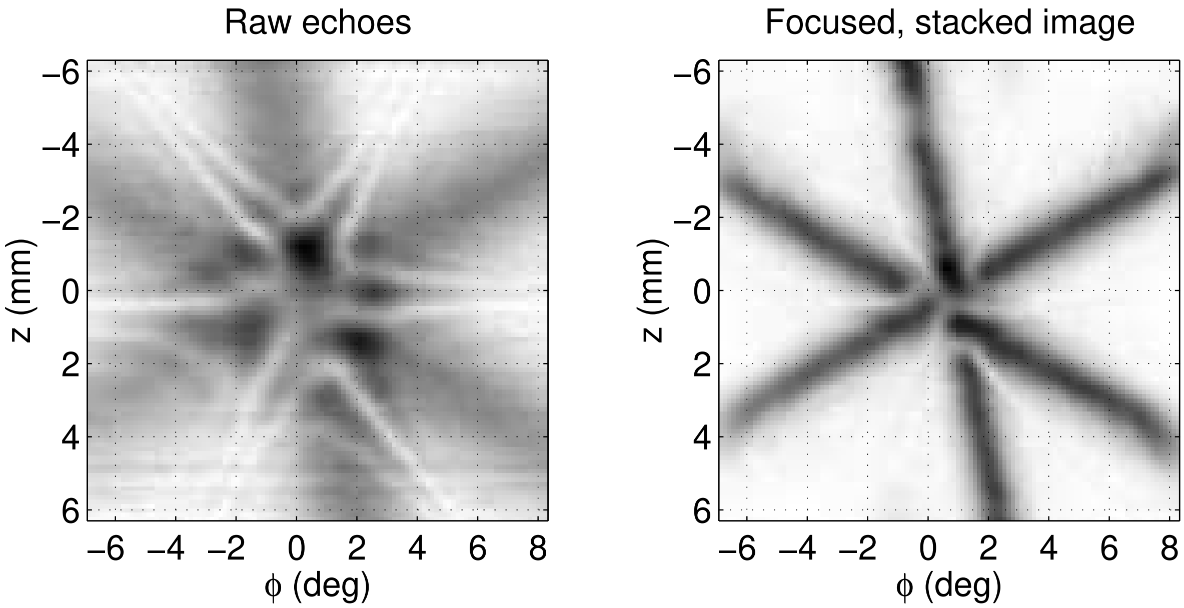

(a) (b)

Figure 7.9 Results of imaging a 100-μm wire target located beyond the transducer focus. (a) Raw echoes summed absolutely over time. (b) Stack of focused images summed absolutely over depth.

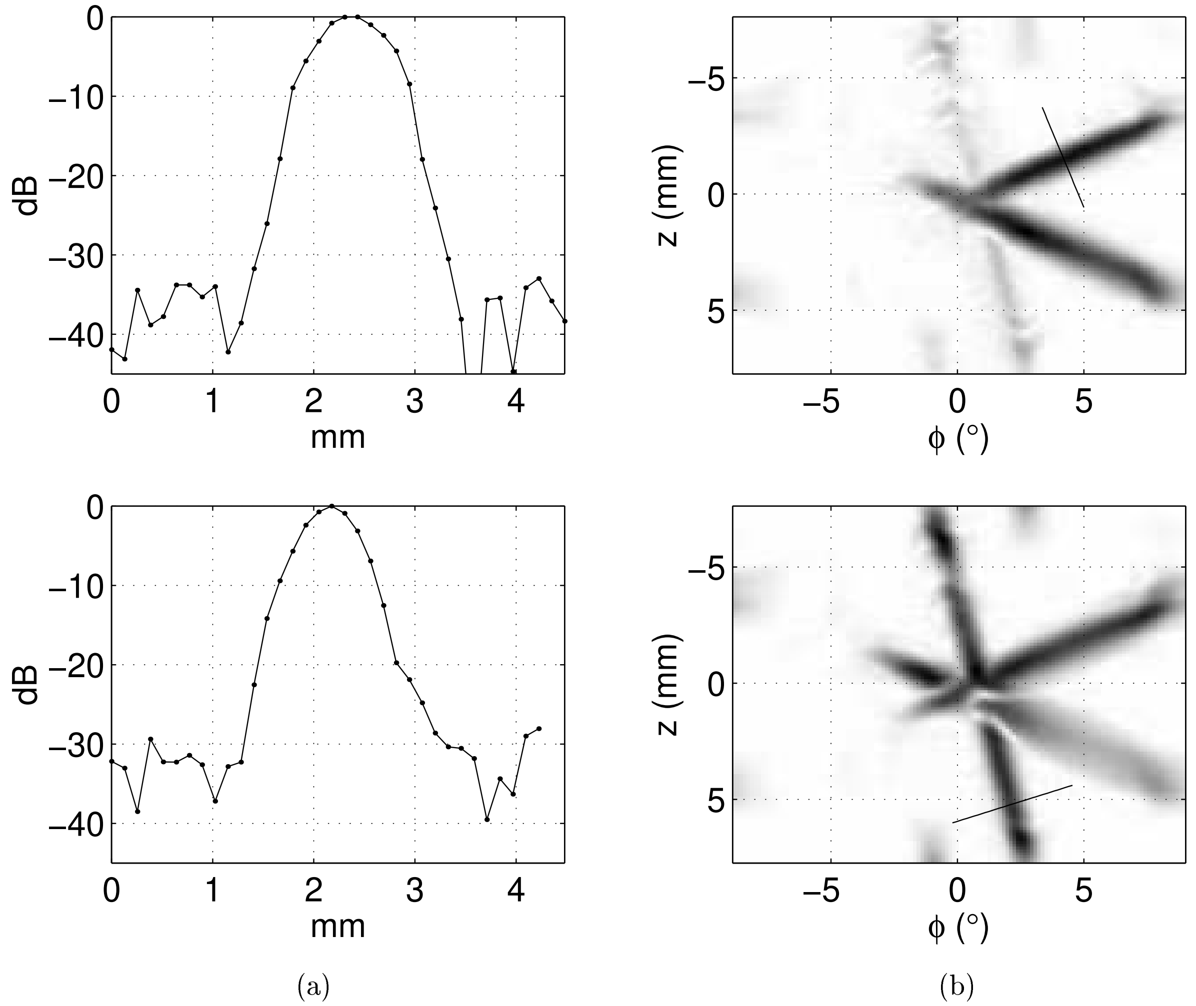

Figure 7.10 (a) One-dimensional profiles of (b) cuts across the in-focus wires taken at different reconstruction depths and wire orientations. At the depths being considered, the linear distance scale is approximately the same on both the z and ϕ axes.

that in the left-hand panel, which was obtained from the raw echoes by an absolute summation of the received waveform at each scan location. Figure 7.10 plots the one-dimensional profiles across the in-focus wires at two different reconstruction depths. The measured half-power [synthesized] beamwidth at the target distance is about 600 μm in either the z or ϕ directions, due to the choice of step sizes and distance to the target. This compares favorably to the transducer's resolution at focus, and is in excellent agreement with the *Field II* simulation of Section 7.4.

In another experiment, a 15-MHz transducer having a 12.7-mm diameter, a 19.1-mm focal length, and a theoretical resolution at focus of $\lambda D/f = 150$ μm was used. The target was a piece of ordinary 1.6-mm aluminum screen held at constant radius from the transducer support arm and about 5 mm beyond the focus of the transducer (see Figure 7.11(a)). The area scanned for this experiment was 6.4 mm by 8.32° in 128 by 128 steps.

Figures 7.11(b) and 7.11(c) show two log-scaled images of the wire screen, at distances of 5.17 mm and 5.30 mm beyond the focus of the transducer. The mesh consists of vertical wires that are nearly straight, parallel, and normal to the view direction, and horizontal wires which weave through them. The upper image shows some of the vertical wires and also the over-crossings of the horizontal wires. In the lower image, the deeper focus reveals other vertical wires, indicating that the wire mesh was held at a not-quite-constant radius. Most of the wire over-crossings are still visible, but the rest of the horizontal wires remain invisible due to their angles with respect to the view direction (note that the wire diameter, at 280 μm, is much larger than the acoustic wavelength).

7.6 Summary

An efficient and accurate three-dimensional image formation algorithm has been obtained directly through simple approximations to the point spread function for 3-D wave propagation in cylindrical coordinates. Computer simulations and experimental results verify its good performance. When combined with the virtual source technique, this algorithm should allow high quality, near-diffraction-limited imaging from small cylindrical platforms, whether they be needles, catheters, or others yet to be developed.

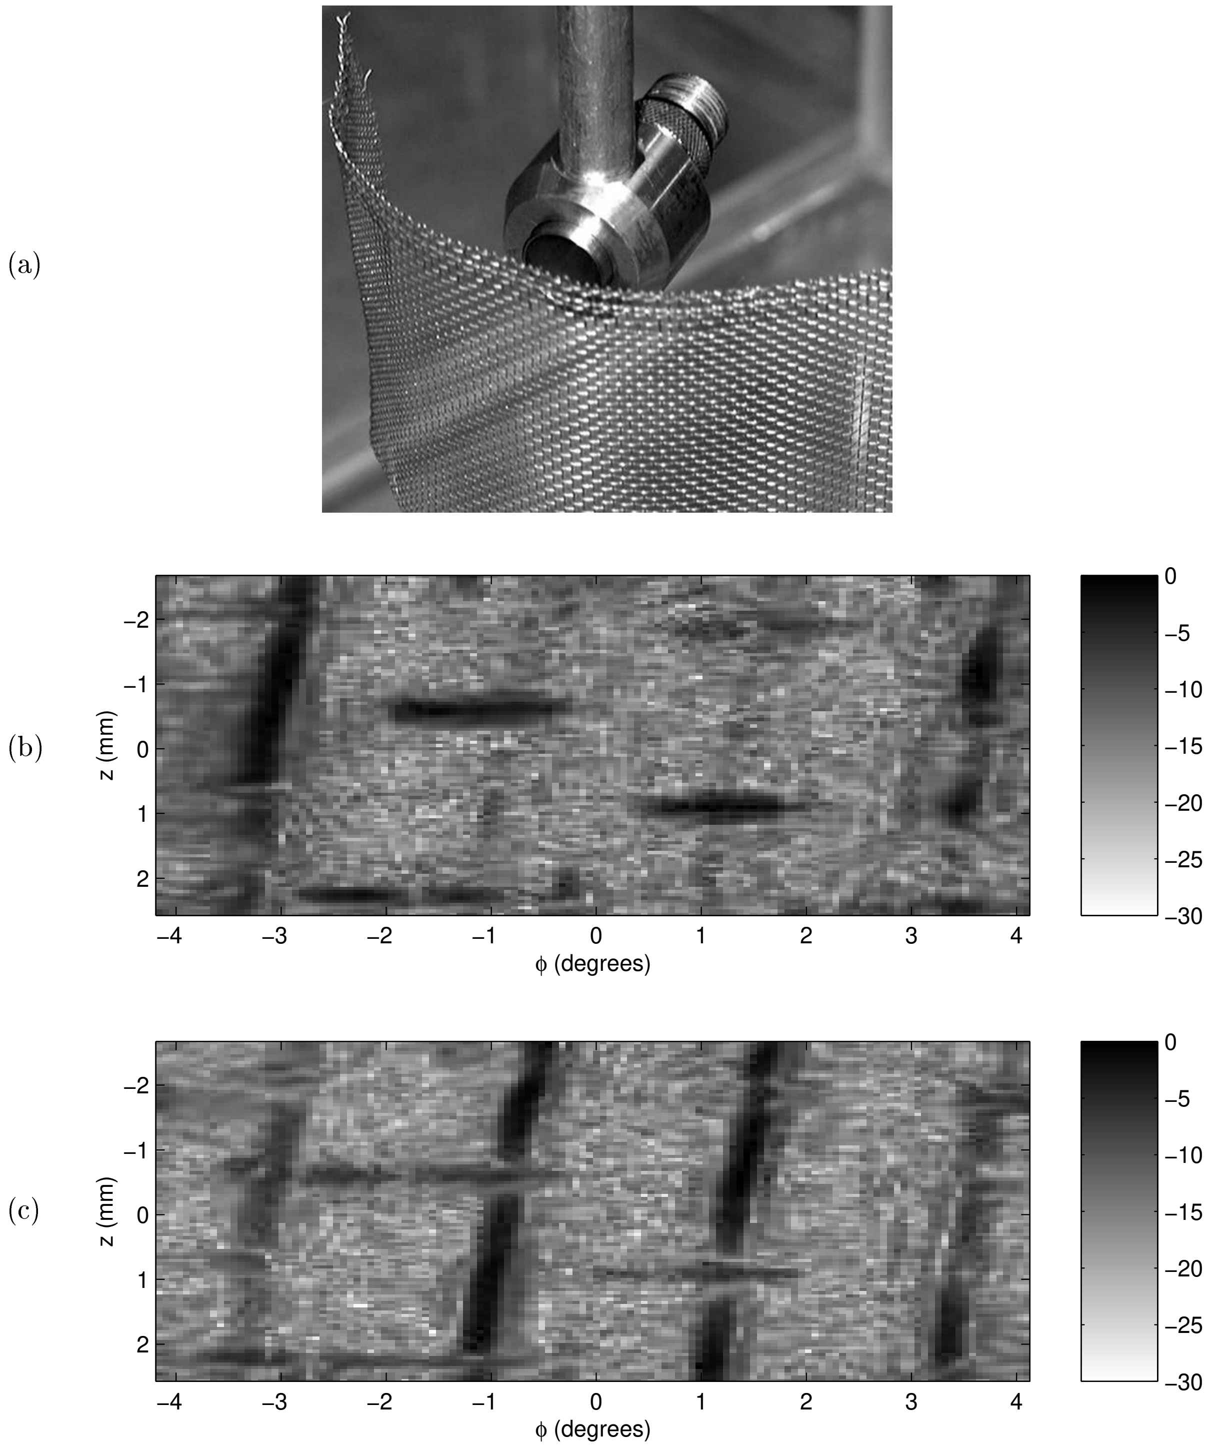

Figure 7.11 Wire mesh imaging experiment: (a) Experimental set-up. (b) and (c) Reconstructed images of the wire mesh at 5.17 mm and 5.30 mm beyond the transducer's focus.

CHAPTER 8

CONCLUSIONS

Tissue-induced aberration is still a problem plaguing medical ultrasound imaging. Solutions based on single-valued focusing operators may not be sufficient to restore diffraction-limited resolution in all cases, but the results in Chapter 6 clearly demonstrate that significant improvement over published algorithms is possible. Concepts borrowed from other imaging disciplines can provide new insights into the aberration problem in medical ultrasound; in particular, some of the same issues have been studied for many years to improve acoustic imaging of the nonhomogeneous Earth. Invariably, such insights also improve the understanding of existing approaches in the medical ultrasound literature. Three of these concepts have been considered in this dissertation.

The tremendous redundancy of a complete data set can be exploited for aberration correction by analyzing the time shifts on common-midpoint gathers. Until now, however, the wide-angle, random-scattering nature of medical ultrasound targets has limited the accuracy and robustness of this approach, particularly when estimating azimuth-dependent aberration profiles. Prefiltering the data with two-dimensional fan filters largely solves this problem and produces an aberration-correction algorithm (OFF) that outperforms the most popular existing algorithms in almost all cases.

The concept of focusing-operator updating, recently popular in seismic imaging, has provided insight into aberration-correction algorithms based on an initial transmit focus. We have developed a new updating procedure based on dynamic programming. When coupled with a careful selection of initial focus points, this results in an algorithm with improved performance. It outperforms existing algorithms in some experiments.

The seismic imaging community has understood for some time that layered media may be approximated by constant-sound-speed media for beamforming purposes, leading to so-called time-migration algorithms. This raises the possibility of medical ultrasound applications—in particular, brain imaging through the adult human skull. While our simulation results have been encouraging, experiments with animal skulls have been inconclusive due to the high attenuation of ultrasound in skull bone. Further research may validate the time-migration concept for brain imaging.

This work has at least several implications for future aberration-correction research. First, imaging with single-valued focusing operators may be able to correct for most of the aberration encountered in soft tissues. Our results demonstrate obvious improvement under the screen model, and there is every reason to believe that further improvement is possible without moving to a more complicated model. Second, increasing aperture should not be viewed merely as a source of aberration, but as an opportunity to do a better job of correcting it. Even if a reduced aperture is used for imaging, the extra aperture can provide more information for aberration correction. In the OFF algorithm, for example, having a full spread of offsets at every midpoint would improve the estimation accuracy at the "ends" of the imaging array. Third, the noise penalty for using complete data sets may not be as serious a problem as conventionally assumed. The extra information contained in the signals partially compensates for this problem. Finally, the performance of hybrid algorithms should be investigated, where aberration estimates from common-midpoint signals in complete data sets are used to obtain a good initial transmit focus, followed by iterative focusing-operator updates.

APPENDIX A

LOW-NOISE PREAMPLIFIERS

The array imaging system described in Section 6.1.1 routes all signals except the 64 same-element pulse-echo signals through four low-noise preamplifiers. The schematic diagram for these custom preamps is shown in Figure A.1. Figure A.2 is a photograph of one of the finished preamps. The frequency response (Figure A.3) was chosen to maximize the gain at the array transducer's center frequency of 2.6 MHz.

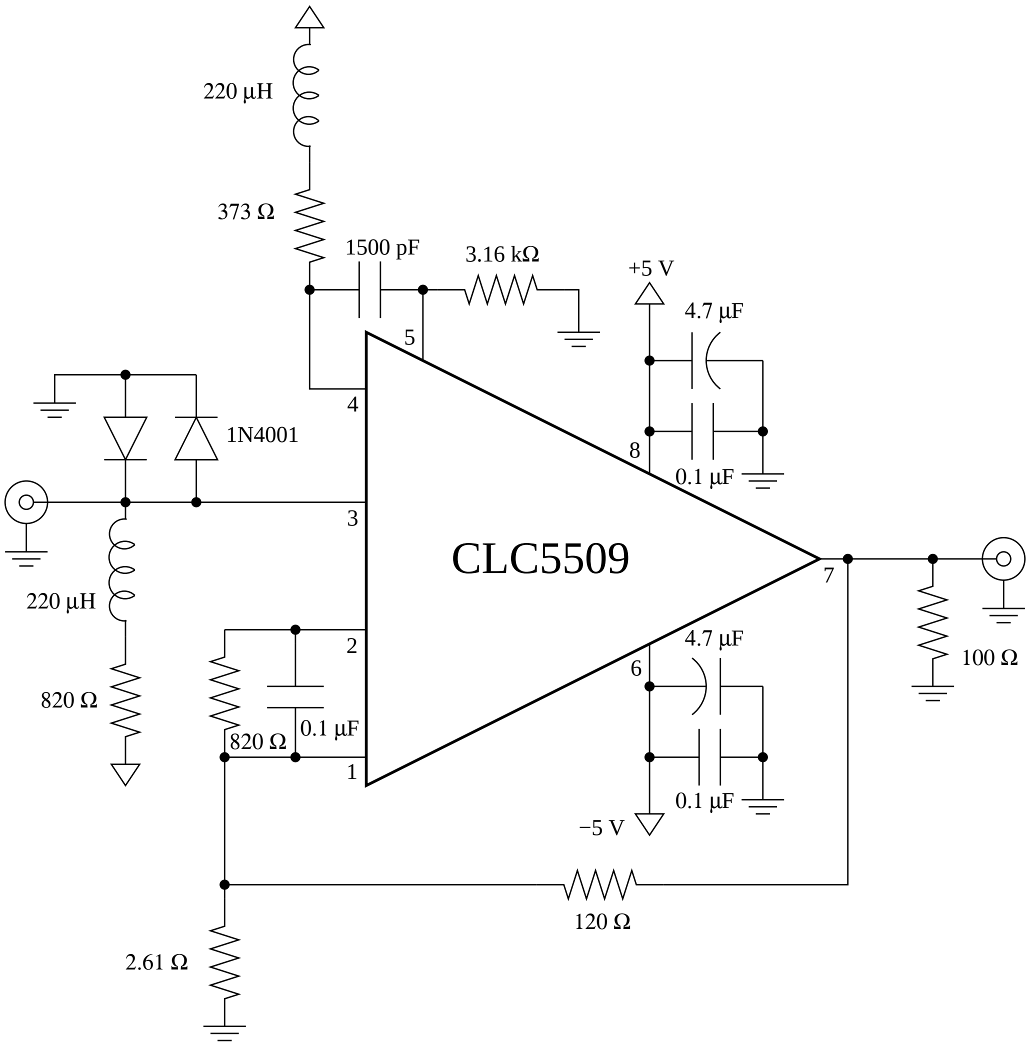

Figure A.1 Schematic diagram for the low-noise preamps.

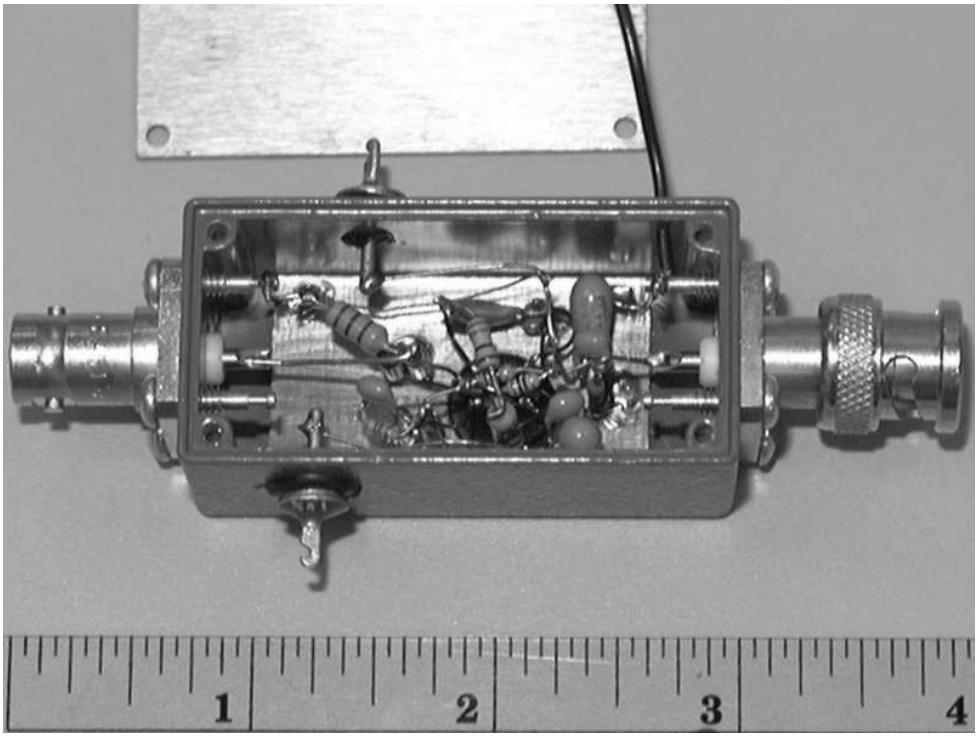

Figure A.2 Close-up of a low-noise preamp.

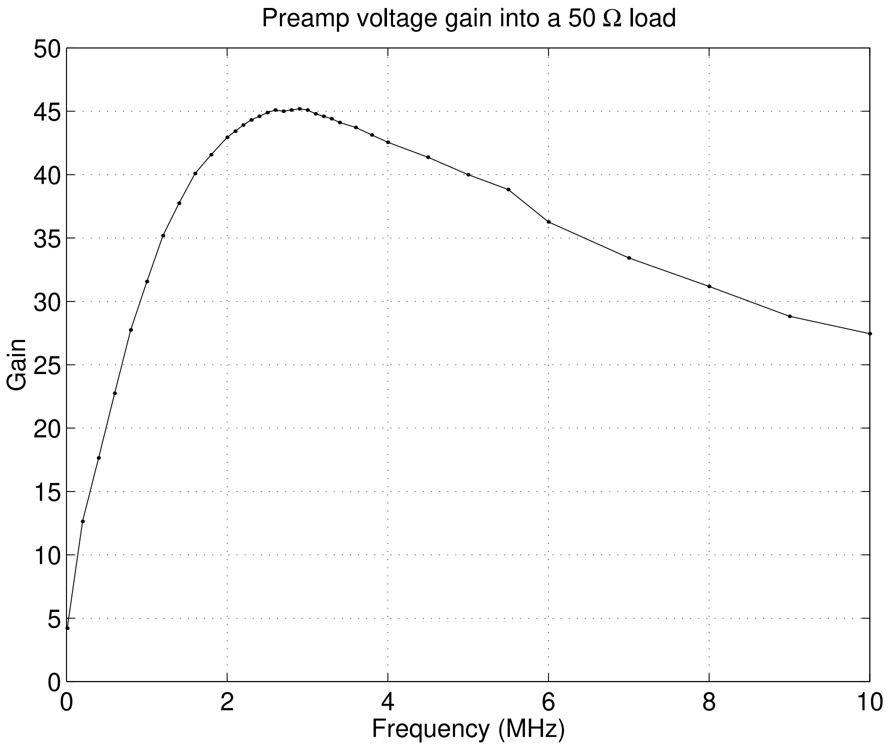

Figure A.3 Preamp voltage gain vs. frequency.

REFERENCES

[1] K. S. Deffeyes, *Hubbert's Peak: The Impending World Oil Shortage*, Princeton: Princeton University Press, 2001.

[2] K. Thomenius, "Beamformation in medical ultrasound," Short course handouts at the 2000 IEEE Ultrasonics Symposium.

[3] M. O'Donnell, M. J. Eberle, D. N. Stephens, J. L. Litzza, K. San Vicente, and B. M. Shapo, "Synthetic phased arrays for intraluminal imaging of coronary arteries," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 3, pp. 714–721, May 1997.

[4] M. A. Haun, D. L. Jones, and W. D. O'Brien, Jr., "Efficient three-dimensional imaging from a small cylindrical aperture," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 49, no. 7, pp. 861–870, July 2002.

[5] S. A. Goss, R. L. Johnston, and F. Dunn, "Comprehensive compilation of empirical ultrasonic properties of mammalian tissues," *Journal of the Acoustical Society of America*, vol. 64, no. 2, pp. 423–457, August 1978.

[6] S. A. Goss, R. L. Johnston, and F. Dunn, "Comprehensive compilation of empirical ultrasonic properties of mammalian tissues, II," *Journal of the Acoustical Society of America*, vol. 68, no. 1, pp. 93–108, July 1980.

[7] J. W. Goodman, "Statistical properties of laser speckle patterns," in *Laser Speckle and Related Phenomena*, 2nd ed., J. C. Dainty, Ed., New York: Springer-Verlag, 1984.

[8] R. Brown, "William Peter Haseman and the birth of his idea," in *Timelines*, the newsletter of the Geophysical Society of Tulsa, February 2002, `http://gst.seg.org/TL/2002/02/TidBits.shtml`.

[9] J. F. Claerbout, *Imaging the Earth's Interior*, Palo Alto: Blackwell Scientific Publications, 1985.

[10] M. J. O'Brien, N. D. Whitmore, S. Brandsberg-Dahl, and G. E. Murphy, "Multicomponent modeling of the Valhall field," in *61st. Conf. European Assoc. of Geoscientists and Engineers*, Helsinki, Finland, 1999.

[11] S. H. Gray, J. Etgen, J. Dellinger, and D. Whitmore, "Seismic migration problems and solutions," *Geophysics*, vol. 66, no. 5, pp. 1622–1640, Sept.–Oct. 2001.

[12] J. Gazdag, "Wave equation migration with the phase-shift method," *Geophysics*, vol. 43, no. 7, pp. 1342–1351, December 1978.

[13] R. H. Stolt, "Migration by Fourier transform," *Geophysics*, vol. 43, no. 1, pp. 23–48, February 1978.

[14] C. Cafforio, C. Prati, and E. Rocca, "SAR data focusing using seismic migration techniques," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 27, no. 2, pp. 194–206, March 1991.

[15] M. E. Anderson, M. S. McKeag, and G. E. Trahey, "The impact of sound speed errors on medical ultrasound imaging," *Journal of the Acoustical Society of America*, vol. 107, no. 6, pp. 3540–3548, June 2000.

[16] Y. Sumino and R. C. Waag, "Measurements of ultrasonic pulse arrival time differences produced by abdominal wall specimens," *Journal of the Acoustical Society of America*, vol. 90, no. 6, pp. 2924–2930, December 1991.

[17] L. M. Hinkelman, D.-L. Liu, L. A. Metlay, and R. C. Waag, "Measurements of ultrasonic pulse arrival time and energy level variations produced by propagation through abdominal wall," *Journal of the Acoustical Society of America*, vol. 95, no. 1, pp. 530–541, January 1994.

[18] M. Moshfeghi and R. C. Waag, "*In vivo* and *in vitro* ultrasound beam distortion measurements of a large aperture and a conventional aperture focussed transducer," *Ultrasound in Med. and Biol.*, vol. 14, no. 5, pp. 415–428, 1988.

[19] Q. Zhu and B. D. Steinberg, "Large-transducer measurements of wavefront distortion in the female breast," *Ultrasonic Imaging*, vol. 14, pp. 276–299, 1992.

[20] D. N. White, J. M. Clark, J. N. Chesebrough, M. N. White, and J. K. Campbell, "Effect of the skull in degrading the display of echoencephalographic B and C scans," *Journal of the Acoustical Society of America*, vol. 44, no. 5, pp. 1339–1346, 1968.

[21] J. A. Hileman, P. Embree, and J. C. Pflueger, "Automated static corrections," *Geophysical Prospecting*, vol. 16, pp. 326–358, 1968.

[22] M. T. Taner, F. Koehler, and K. A. Alhilali, "Estimation and correction of near-surface time anomalies," *Geophysics*, vol. 39, no. 4, pp. 441–463, August 1974.

[23] S. W. Flax and M. O'Donnell, "Phase-aberration correction using signals from point reflectors and diffuse scatterers: Basic principles," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 35, no. 6, pp. 758–767, November 1988.

[24] M. O'Donnell and S. W. Flax, "Phase-aberration correction using signals from point reflectors and diffuse scatterers: Measurements," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 35, no. 6, pp. 768–774, November 1988.

[25] R. Mallart and M. Fink, "The van Cittert-Zernike theorem in pulse echo measurements," *Journal of the Acoustical Society of America*, vol. 90, no. 5, pp. 2718–2727, November 1991.

[26] R. Mallart and M. Fink, "Adaptive focusing in scattering media through sound-speed inhomogeneities: The van Cittert-Zernike approach and focusing criterion," *Journal of the Acoustical Society of America*, vol. 96, no. 6, pp. 3721–3732, December 1994.

[27] D. Rachlin, "Direct estimation of aberrating delays in pulse-echo imaging systems," *Journal of the Acoustical Society of America*, vol. 88, no. 1, pp. 191–198, July 1990.

[28] Y. Li, "Phase aberration correction using near-field signal redundancy—Part I: Principles," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 2, pp. 355–371, March 1997.

[29] Y. Li, D. Robinson, and D. Carpenter, "Phase aberration correction using near-field signal redundancy—Part II: Experimental results," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 44, no. 2, pp. 372–379, March 1997.

[30] W. F. Walker and G. E. Trahey, "Speckle coherence and implications for adaptive imaging," *Journal of the Acoustical Society of America*, vol. 101, no. 4, pp. 1847–1858, April 1997.

[31] L. Nock, G. E. Trahey, and S. W. Smith, "Phase aberration correction in medical ultrasound using speckle brightness as a quality factor," *Journal of the Acoustical Society of America*, vol. 85, no. 5, pp. 1819–1833, May 1989.

[32] G. Trahey, D. Zhao, J. A. Miglin, and S. W. Smith, "Experimental results with a real-time adaptive ultrasonic imaging system for viewing through distorting media," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 37, no. 5, pp. 418–427, September 1990.

[33] D. Zhao and G. E. Trahey, "Comparisons of image quality factors for phase aberration correction with diffuse and point targets: Theory and experiments," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 38, no. 2, pp. 125–132, March 1991.

[34] R. A. Muller and A. Buffington, "Real-time correction of atmospherically degraded telescope images through image sharpening," *Journal of the Optical Society of America*, vol. 64, no. 9, pp. 1200–1209, 1974.

[35] D.-L. Liu and R. C. Waag, "Correction of ultrasonic wavefront distortion using backpropagation and a reference waveform method for time-shift compensation," *Journal of the Acoustical Society of America*, vol. 96, no. 2, pp. 649–660, August 1994.

[36] F. Audebert, D. Nichols, T. Rekdal, B. Biondi, D. Lumley, and H. Urdaneta, "Imaging complex geologic structure with single-arrival Kirchhoff prestack depth migration," *Geophysics*, vol. 62, no. 5, pp. 1533–1543, Sept.–Oct. 1997.

[37] C. Dorme and M. Fink, "Focusing in transmit-receive mode through inhomogeneous media: The time reversal matched filter approach," *Journal of the Acoustical Society of America*, vol. 98, no. 2, pp. 1155–1162, August 1995.

[38] P. Blomgren, G. Papanicolaou, and H. Zhao, "Super-resolution in time-reversal acoustics," *Journal of the Acoustical Society of America*, vol. 111, no. 1, pp. 230–248, January 2002.

[39] C. Dorme and M. Fink, "Ultrasonic beam steering through inhomogeneous layers with a time reversal mirror," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, no. 1, pp. 167–175, January 1996.

[40] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*, New York: IEEE Press, 1988.

[41] A. J. Devaney, "A filtered backpropagation algorithm for diffraction tomography," *Ultrasonic Imaging*, vol. 4, pp. 336–350, 1982.

[42] N. Sponheim and I. Johansen, "Experimental results in ultrasonic tomography using a filtered backpropagation algorithm," *Ultrasonic Imaging*, vol. 13, pp. 56–70, 1991.

[43] A. B. Weglein, K. H. Matson, D. J. Foster, P. M. Carvalho, D. Corrigan, and S. A. Shaw, "Imaging and inversion at depth without a velocity model," in *Proc. 70th Annual Meeting of the Society of Exploration Geophysicists*, Calgary, Canada, August 2000, pp. 1016–1019.

[44] A. B. Weglein (private communication), May 2001.

[45] A. J. Berkhout, "Pushing the limits of seismic imaging, Part I: Prestack migration in terms of double dynamic focusing," *Geophysics*, vol. 62, no. 3, pp. 937–953, May–June 1997.

[46] A. J. Berkhout, "Pushing the limits of seismic imaging, Part II: Integration of prestack migration, velocity estimation, and AVO analysis," *Geophysics*, vol. 62, no. 3, pp. 954–968, May–June 1997.

[47] A. J. Berkhout and D. J. Verschuur, "Seismic imaging beyond depth migration," *Geophysics*, vol. 66, no. 6, pp. 1895–1912, Nov.–Dec. 2001.

[48] O. Yilmaz and R. Chambers, "Migration velocity analysis by wave-field extrapolation," *Geophysics*, vol. 49, no. 10, pp. 1664–1674, October 1984.

[49] M. T. Taner and F. Koehler, "Velocity spectra—digital computer derivation and applications of velocity functions," *Geophysics*, vol. 34, no. 6, pp. 859–881, December 1969.

[50] R. J. Castle, "A theory of normal moveout," *Geophysics*, vol. 59, no. 6, pp. 983–999, June 1994.

[51] C. H. Dix, "Seismic velocities from surface measurements," *Geophysics*, vol. 20, no. 1, pp. 68–86, January 1955.

[52] S. W. Smith, G. E. Trahey, and O. T. von Ramm, "Phased array ultrasound imaging through planar tissue layers," *Ultrasound in Medicine and Biology*, vol. 12, no. 3, pp. 229–243, 1986.

[53] S. MacKay and R. Abma, "Imaging and velocity estimation with depth-focusing analysis," *Geophysics*, vol. 57, no. 12, pp. 1608–1622, December 1992.

[54] S. MacKay and R. Abma, "Depth-focusing analysis using a wavefront-curvature criterion," *Geophysics*, vol. 58, no. 8, pp. 1148–1156, August 1993.

[55] P. T. Gough and D. W. Hawkins, "Imaging algorithms for a strip-map synthetic aperture sonar: Minimizing the effects of aperture errors and aperture undersampling," *IEEE Journal of Oceanic Engineering*, vol. 22, no. 1, pp. 27–39, January 1997.

[56] H. Choi, "Signal processing issues in synthetic aperture radar and computer tomography," Ph.D. dissertation, University of Illinois at Urbana-Champaign, 1998.

[57] Center for Wave Phenomena, Colorado School of Mines, Golden, CO, "Seismic Unix," 2001, `http://www.cwp.mines.edu/cwpcodes/`.

[58] Y. Li, "Small element array algorithm for correcting phase aberration using near-field signal redundancy—Part I: Principles," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 47, no. 1, pp. 29–48, January 2000.

[59] Y. Li and B. Robinson, "Small element array algorithm for correcting phase aberration using near-field signal redundancy—Part II: Experimental results," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 47, no. 1, pp. 49–57, January 2000.

[60] T. S. Huang, "Two-dimensional windows," *IEEE Transactions on Audio and Electroacoustics*, vol. AU-20, pp. 88–89, March 1972.

[61] J. G. Proakis and M. Salehi, *Communication Systems Engineering*, Englewood Cliffs: Prentice Hall, 1994.

[62] D. H. Rothman, "Large near-surface anomalies, seismic reflection data, and simulated annealing," Ph.D. dissertation, Stanford University, 1985.

[63] D. P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Englewood Cliffs: Prentice-Hall, 1987.

[64] A. J. Viterbi, "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm," *IEEE Transactions on Information Theory*, vol. IT-13, no. 2, pp. 260–269, April 1967.

[65] G. D. Forney, Jr., "The Viterbi algorithm," *Proceedings of the IEEE*, vol. 61, no. 3, pp. 268–278, March 1973.

[66] J. A. Dellinger, N. R. Goins, and T. A. Reilly, "Matching P wave and shear wave sections in seismic exploration," U.S. patent 4,686,657, August 1987.

[67] K. W. Hollman, K. W. Rigby, and M. O'Donnell, "Coherence factor of speckle from a multi-row probe," in *Proc. IEEE Ultrasonics Symposium*, Lake Tahoe, Nevada, October 1999, pp. 1257–1260.

[68] O. Yilmaz, *Seismic Data Analysis*, Tulsa: Society of Exploration Geophysicists, 2001.

[69] J. C. Lacefield and R. C. Waag, "Time-shift estimation and focusing through distributed aberration using multirow arrays," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 48, no. 6, pp. 1606–1624, November 2001.

[70] D.-L. Liu, "Element size effect on phase aberration correction," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 49, no. 9, pp. 1212–1221, September 2002.

[71] Biomedical Ultrasonics Laboratory, University of Michigan, Ann Arbor, MI, "Online ultrasound data," 2001, `http://bul.eecs.umich.edu/ultrasound_data.html`.

[72] M. O'Donnell, "Efficient synthetic aperture imaging from a circular aperture with possible application to catheter-based imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 3, pp. 366–380, May 1992.

[73] M. Soumekh, *Synthetic Aperture Radar Signal Processing*, New York: Wiley-Intersciences, 1999.

[74] D. Vray, C. Haas, T. Rastello, M. Krueger, E. Brusseau, K. Schroeder, G. Gimenez, and H. Ermert, "Synthetic aperture-based beam compression for intravascular ultrasound imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 48, no. 1, pp. 189–201, January 2001.

[75] J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed., New York: McGraw-Hill, 1996.

[76] L. J. Busse, "Three-dimensional imaging using a frequency-domain synthetic aperture focusing technique," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 2, pp. 174–179, March 1992.

[77] P. T. Gough and D. W. Hawkins, "Unified framework for modern synthetic aperture imaging algorithms," *International Journal of Imaging Systems and Technology*, vol. 8, no. 4, pp. 343–358, 1997.

[78] C. Passman and H. Ermert, "A 100-MHz ultrasound imaging system for dermatologic and opthalmologic diagnostics," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 43, no. 4, pp. 545–552, July 1996.

[79] C. H. Frazier and W. D. O'Brien, Jr., "Synthetic aperture techniques with a virtual source element," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 45, no. 1, pp. 196–207, January 1998.

[80] D. C. Munson, Jr., J. D. O'Brien, and W. K. Jenkins, "A tomographic formulation of spotlight-mode synthetic aperture radar," *Proceedings of the IEEE*, vol. 71, no. 8, pp. 917–925, August 1983.

[81] 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*, vol. 39, no. 2, pp. 262–267, March 1992.

VITA

Mark Alden Haun was born in Carmichael, California, in 1972. He received the B.S.E. degree in electrical engineering from Walla Walla College in 1996 and the M.S. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 1999.

While studying at the University of Illinois, Haun worked as a teaching assistant and research assistant. During the summer of 2001, he was the instructor for ECE 310, Digital Signal Processing.

Haun has published a paper in the *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* and several conference papers on topics related to acoustic imaging.