Broadband Frequency-Independent Beamforming

by

Thomas Clayton Chou

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Bachelor of Science in Electrical Engineering

and

Bachelor of Science in Computer Science

and

Master of Science in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

January 1995

Author .....
Department of Electrical Engineering and Computer Science
January 10, 1995

Certified by ....
William M. Rabinowitz
Principal Research Scientist
Thesis Supervisor

Certified by ....
Gary W. Elko
Member Technical Staff, AT&T Bell Laboratories
Thesis Supervisor

Accepted by ....
Professor Frederic R. Morgenthaler
Chairman, Departmental Committee on Graduate Students

Broadband Frequency-Independent Beamforming

by

Thomas Clayton Chou

Submitted to the Department of Electrical Engineering and Computer Science
on January 20, 1995, in partial fulfillment of the
requirements for the degrees of
Bachelor of Science in Electrical Engineering
and
Bachelor of Science in Computer Science
and
Master of Science in Electrical Engineering

Abstract

A transducer array processing technique known as *beamforming* is often used to distinguish waves propagating in different directions. Applications include sonar, radar, teleconferencing, imaging, and mobile communications. Many of these applications involve broadband signals (such as speech, acoustic impulses, and spread-spectrum signals) and require beamformers with frequency-independent behavior. In particular, lobes (directions of high sensitivity) and nulls (directions of low sensitivity) should be invariant with frequency. An existing class of solutions achieves some frequency-independent behavior for a single lobe, but not for other lobes or nulls.

This work designs and builds a highly frequency-independent beamformer using a technique called filter-and-sum beamforming. In particular, frequency-independent nulls and sidelobes are demonstrated. The hardware implementation employs a log-periodic transducer geometry which is hardware efficient.

Thesis Supervisor: William M. Rabinowitz
Title: Principal Research Scientist

Thesis Supervisor: Gary W. Elko
Title: Member Technical Staff, AT&T Bell Laboratories

Acknowledgments

I would like to thank all who made this work possible: my mentor Gary Elko, for his ears and for everything in between; my advisor William Rabinowitz, for thousands of suggestions and overall high standards; my housemates, for being so crazy; many fellow students - you know who you are - who taught me how to live life, particularly Christina Park (1975 - 1993) who gave me more than she will ever know; Gary again, for assorted heroics; Bob Kubli, for infinitely reliable assistance and patience with us lesser beings; Robert Lustberg, on whose thesis I built my own; Charlie Guin, for discussions on all things imaginable, and some things not; and finally to David Berkley, Larry Rabiner, the Acoustics Research Department, Laboratory 1122 and AT&T Bell Laboratories for fostering a uniquely stimulating working environment.

Contents

List of Figures

Chapter 1

Introduction

In media such as air, water, free space, etc., propagating waves are often distinguishable by their directions of propagation. This fact is significant in applications which detect, resolve, or receive propagating waves in the presence of interfering waves. Acoustic applications include teleconferencing systems, speakerphones, hearing aids, sonar, seismic equipment, and acoustic imaging systems. Analogous electromagnetic uses include radar, mobile communications, and radio astronomy.

Transducer arrays are thus useful for their ability to: 1) focus *beams*[1] in the directions of desired sources, and 2) aim *nulls* in the directions of interfering sources. Often, these sources emit broadband signals, such as speech, music, or spread spectrum radio transmissions and thus require array designs having frequency-invariant behavior. Unfortunately, simple beamformers often exhibit beamshapes and null directions which are highly frequency-dependent.

Several schemes exist for constructing beamformers having constant mainlobe widths [5, 8, 11, 12, 23, 26]. These designs, which are outlined in Appendix B, do not generally achieve frequency-independent null directions.

The shortcomings of these so-called constant-beamwidth designs motivate the following goals of this thesis:

- Design a highly frequency-independent beamformer using the flexibility offered

[1] In this chapter, italicized words are standard beamforming terms which are reviewed in Appendix A.

by a technique known as filter-and-sum beamforming.

- Explore the performance limitations of the filter-and-sum approach.
- Implement the design in a hardware-efficient manner.

These goals are pursued in the following stages of this work:

- Theory - filter-and-sum beamformer design is shown to be related to two-dimensional filter design.
- Filter design - a simple two-dimensional filter design algorithm is developed and carried out. The performance limitations of the filter-design algorithm are examined.
- Simulation - the beamformer's theoretical plane-wave response is computed.
- Implementation - a real-time digital beamformer is constructed. Plane-wave responses are measured and compared with simulations.

The remainder of this introduction specifies the basic problems encountered in broadband beamforming, and outlines the proposed solution.

1.1 Broadband beamforming challenges

Broadband beamformer design involves several issues not encountered in narrowband beamformer design. Some of these issues are now illustrated with respect to a specific design example.

Consider a microphone array with 11 identical omnidirectional elements and 4 cm inter-element spacing.[2] Suppose further that the array output is formed by directly summing the elements' outputs. Figure 1-1 shows this array's *array pattern* at selected frequencies, while Figure 1-2 shows its frequency response at several selected angles.

This simple beamformer exhibits the following shortcomings:

[2]Not coincidentally, this is the configuration of one of the subarrays used in the real-time implementation, to be described in Figure 3-2.

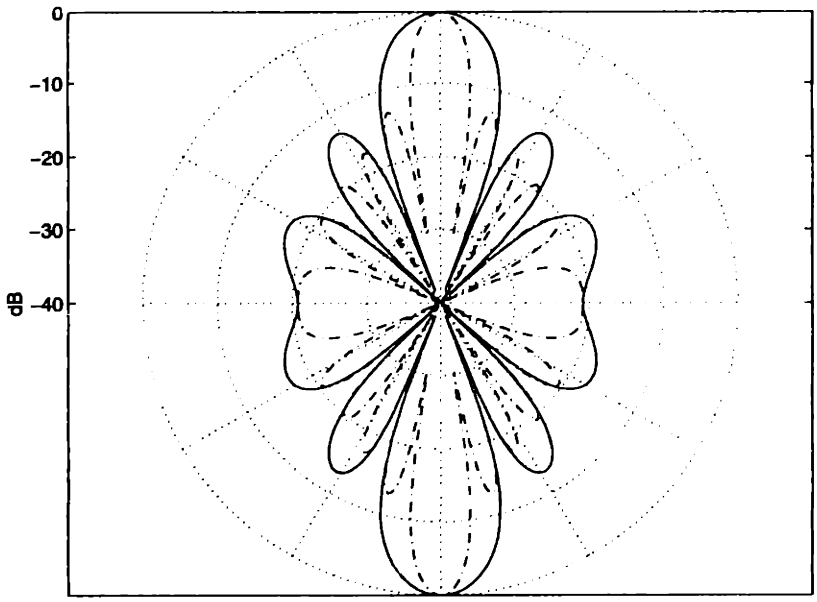

Figure 1-1: Far-field response of simple linear microphone array at 2 kHz (solid line) and 4 kHz (dotted line). In this polar and subsequent polar plots, *broadside* beams point up and down, while *endfire* beams point left and right.

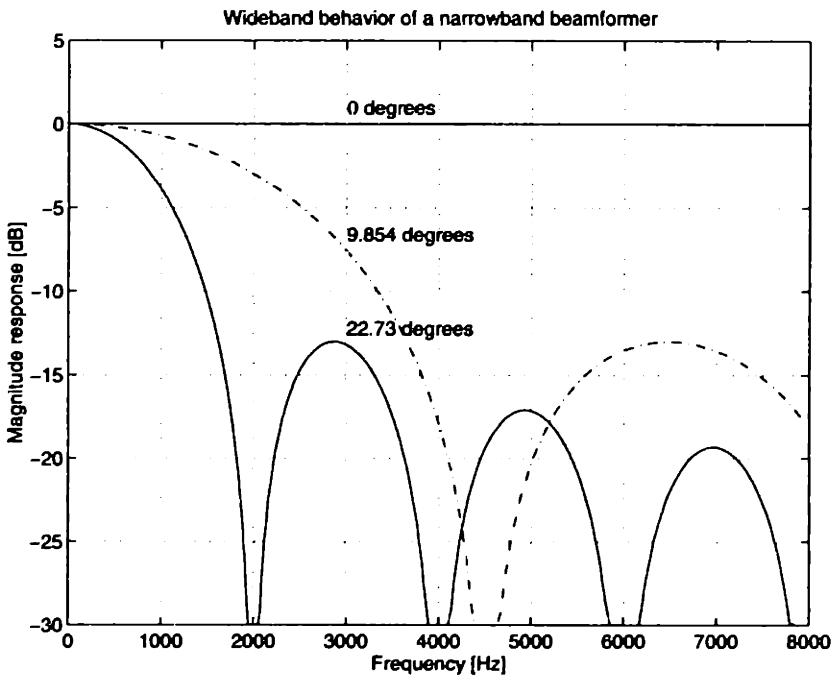

Figure 1-2: Distortion in frequency response of simple array. The frequency response curves are labeled with the incidence angle relative to broadside

.

1. Undesirable sensitivity to *steering* errors:

   Signals which impinge on the array at *broadside* are received with a flat frequency response. However, as signals move away from broadside, a low-pass characteristic emerges since the mainlobe width decreases as frequency increases. Not only is this spectral distortion generally undesirable, it is difficult to correct by highpass filtering since the required correction depends on incidence angle, which is often unknown a priori.

   Figure 1-2 clearly shows the rejection of high frequencies, as well as the trend towards greater rejection as incidence angle increases.

2. Inability to cancel broadband interferers:

   In this example, there is no incidence angle at which a broadband source will not be received at some (indeed, most) frequencies. For example, a source located 23 degrees from broadside is canceled completely at 2, 4, 6, and 8 kHz, but is substantially received at intermediate frequencies (see Figure 1-2).

3. High element count:

   The number of transducers is determined by the need to meet *spatial sampling* requirements at high frequencies, and the need for large *apertures* at low frequencies. In a uniformly spaced array, as exemplified here, the element count is thus proportional to the ratio of highest to lowest frequencies of operation. Covering the audible frequency range (typically two to three orders of magnitude) would be costly with this type of array.

1.2 Description of solution

We address these three problems with a design that combines the concepts of harmonic nesting with filter-and-sum beamforming. Harmonic nesting decomposes the broadband problem into a set of octave problems, leading to a log-periodic structure that reduces element count. The filter-and-sum approach then addresses frequency variations within each octave.

1.2.1 Harmonic nesting

A harmonically nested array is a set of superimposed subarrays, each designed for a particular frequency band. The subarray outputs are summed via bandpass filters such that each subarray only contributes to the total output in the frequency band for which it was designed. This nesting procedure has two features: it allows narrowband beamformer designs to be extended to broader frequency bands with little modification, and it allows array element counts to be *logarithmically* related to the ratio of highest to lowest operating frequencies.

Figure 1-3 displays frequency responses of such an array composed of four nested subarrays covering one octave each. The element configuration is obtained by superimposing four scaled versions of the array already described, as shown in Figure 3-2. Thus, the performance between 2 kHz and 4 kHz is as before, and the frequency responses in the other three octaves (500 Hz - 1 kHz, 1 kHz - 2 kHz, 4 kHz - 8 kHz) exactly replicate that of the 2 kHz - 4 kHz band. For simplicity, the computed responses assume that the octave bandpass filters associated with each subarray have ideally flat passbands and sharp cutoffs.

1.2.2 Filter-and-sum beamforming

The filter-and-sum approach is applied within each of the subarrays described above. Each transducer output is passed through an *elemental filter* whose outputs are in turn summed to produce the total subarray output. These filters are used to make the array's effective aperture (i.e. the actively contributing elements) frequency dependent. Typically, the effective aperture width should be inversely proportional to frequency, making the effective array length a constant multiple of the wavelength. This would lead to elemental filters which are lowpass for the endmost elements, and highpass for the central elements.

Unfortunately, a discrete set of transducers cannot exactly implement a continuous range of aperture sizes. Thus, designing the elemental filters is a challenging problem addressed in Chapter 4 of this thesis.

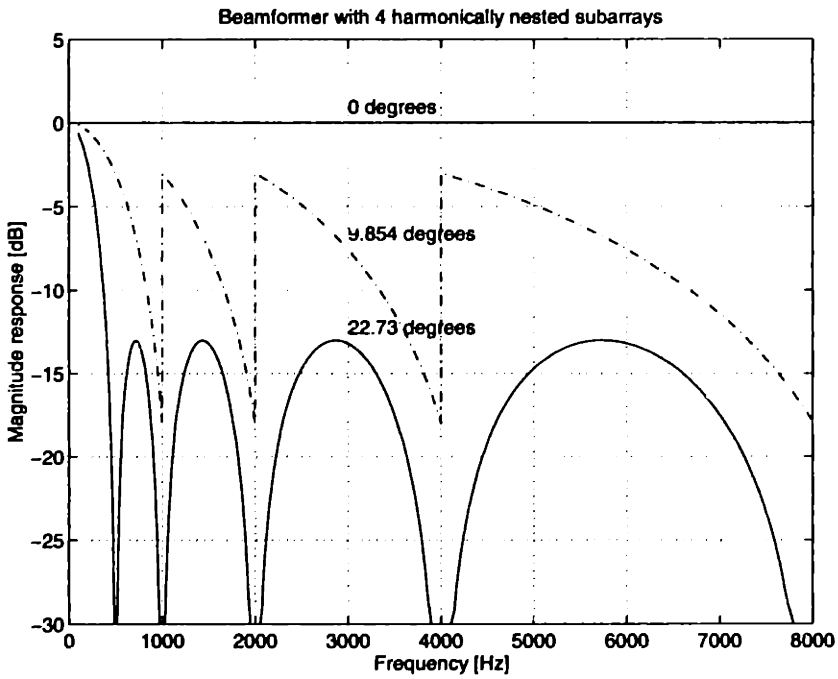

Figure 1-3: Frequency response of four-octave harmonically nested array. The broadband problem has been reduced to a set of octave problems.

1.3 Organization of thesis

Chapter 2 establishes the equivalence between filter-and-sum beamformer design and FIR filter design.

Chapter 3 describes a real-time filter-and-sum beamformer hardware and software platform. Arbitrary beamforming tasks may be performed by downloading a table of elemental filter coefficients into this platform.

Chapter 4 develops and carries out a method for designing the elemental filters.

Chapter 5 presents computer simulations of the beamformer design.

Chapter 6 presents experimental data for the same design.

Chapter 7 summarizes major accomplishments and identifies future research directions.

Appendix A reviews basic beamforming concepts for the novice reader.

Appendix B reviews existing constant-width solutions.

Appendix C describes the frequency sampling method of FIR filter design, which is an essential part of the procedure outlined in Chapter 4.

Appendix D repeats the steps described in Chapters 4, 5, and 6 using different performance goals. This process leads to a beamformer with minimax sidelobes at any particular frequency.

Chapter 2

Sampled aperture and beamforming theory

This chapter establishes the mathematical equivalence between filter-and-sum beamformer design and two-dimensional filter design. Since the theory assumes ideal system components, a brief discussion is also provided on the application of the theory to actual systems composed of non-ideal components (particularly transducers).

2.1 Sampled apertures

An aperture is a region over which energy is received. Apertures can either be continuous, as in parabolic dishes, or discrete (also called sampled) as in transducer arrays. Although sampled aperture systems generally must perform more signal processing, they offer several advantages over continuous aperture systems. Sampled apertures can be steered by purely electronic means, unlike continuous apertures which are often steered mechanically, if at all. Sampled apertures also offer unique capabilities such as multiple simultaneous beams, adaptive processing and superdirectivity [7]. Finally, sampled apertures can generate or compensate for frequency-dependent behavior, using a technique called filter-and-sum beamforming.

2.2 Delay-and-sum beamforming

Before exploring filter-and-sum beamforming, an old and simple technique called *delay-and-sum* beamforming is examined. This technique is used in *narrowband* operation to focus arrays onto a particular point or direction.[1]

The delay-and-sum beamformer output is a weighted sum of time-delayed element outputs:

$$r(t) = \sum_{i=1}^{N} a_n f(\vec{x}_n, t - \tau_n), \tag{2.1}$$

where $\vec{x}_n$, a_n, and τ_n are the element position, element weight, and time-delay associated with the ith sensor. The delays are chosen to compensate for differences in the relative propagation delays from the point of interest to the individual elements. Thus, signals originating from the desired location are summed in phase, while other signals undergo some destructive cancellation. By electrically manipulating the weights and time delays, the lobes and nulls of this array can be *steered* in desired directions.

2.3 Uniformly-spaced linear arrays

While array geometries are in principle arbitrary, certain configurations are especially useful and amenable to mathematical description. For example, consider an array of uniformly spaced collinear elements. Assume an odd number of elements, numbered from $-N$ to N, and positioned on the x-axis such that element n has coordinate nd, d being the inter-element spacing. Thus, endfire directions lie on the x-axis, while broadside directions lie in the y-z plane. Much as linear time-domain systems are analyzed by their response to complex exponential time functions, this space-time system will be considered in terms of its response to complex exponential plane waves.

Consider a delay-and-sum beamformer steered towards a direction making an

[1] A direction is merely a special case of a point which is located at infinite distance.

angle θ_0 measured clockwise from the y-axis. The time delays τ_i would be:

$$\tau_n = n\frac{d}{c}\sin\theta_0, \tag{2.2}$$

where θ_0 is termed the *look* or *steering* direction, and represents the direction from which waves are received with maximum response.

Now, consider the plane-wave arriving from a direction θ, (not necessarily equal to θ_0). This wave is represented by the expression:

$$f(\vec{x}, t) = e^{j(\omega t - \vec{k}\cdot\vec{x})}, \tag{2.3}$$

where the wavevector $\vec{k}$ is:

$$\vec{k} = -\frac{\omega}{c}\begin{bmatrix} \sin\theta \\ \cos\theta \\ 0 \end{bmatrix}. \tag{2.4}$$

Since the far-field pattern is rotationally symmetric about the array axis, it was assumed without loss of generality that the wavevector lies in the $x - y$ plane.

Substituting the wavefield and time delay expressions into the beamformer output expression (Equation 2.1) yields, after adjusting the indices of summation:

$$r(t) = e^{j\omega t}\sum_{n=-N}^{N} a_n e^{j\frac{\omega}{c}nd(\sin\theta - \sin\theta_0)}. \tag{2.5}$$

Removing the time-varying $e^{j\omega t}$ term yields an expression which we term the *array response*:

$$R(\theta, \omega) = \sum_{n=-N}^{N} a_n e^{j\frac{\omega}{c}nd(\sin\theta - \sin\theta_0)}. \tag{2.6}$$

The substitutions $u = \sin\theta$ and $u_0 = \sin\theta_0$ are often made:

$$R(u, \omega) = \sum_{n=-N}^{N} a_n e^{j\frac{\omega}{c}nd(u - u_0)}. \tag{2.7}$$

In the common special case when the weights are symmetric ($a_n = a_{-n}$), the response

becomes an entirely real expression:

$$R(u,\omega) = w_0 + 2\sum_{n=1}^{N} a_n \cos(\frac{\omega}{c} nd(u-u_0)). \tag{2.8}$$

2.4 Array response structure

The array response as expressed in Equation 2.7 is a function in a two dimensional $u-\omega$ space. There is a definite structure to this function, in that the following identity holds for any real constant K:

$$R((\frac{u}{K}+u_0),\ K\omega) = R(u+u_0,\omega). \tag{2.9}$$

This expression means the spatial response of the array contracts towards the steering direction as frequency increases. Thus, the nature of the array's frequency-dependence cannot be specified independently of its spatial dependence. The delay-and-sum beamformer does not have sufficient flexibility to design a frequency-independent system (except for the trivial case of an omnidirectional "beamformer").

Another important aspect of the array function is that it is periodic in u for any given ω:

$$R(u,\omega) = R(u+\frac{2\pi c}{\omega d},\omega). \tag{2.10}$$

Thus, the mainlobe has an infinite set of identical copies which are called *grating lobes.* Aliasing occurs when two or more lobes both appear in the visible region defined by $|u| \leq 1$. This can be avoided at all steering angles if adjacent lobes are separated by a distance greater than 2 in u-space:

$$2 < \frac{2\pi c}{\omega d}, \tag{2.11}$$

or

$$d < \frac{\pi c}{\omega} = \frac{\pi c}{2\pi f} = \frac{\lambda}{2}, \tag{2.12}$$

where λ is the wavelength. This is a *spatial sampling* requirement which is analogous

to the Nyquist sampling rate for time-domain signals.

2.5 Analogy with FIR filter design

The discussion thus far expresses the beamformer behavior in terms of the array weights. The opposite process, finding the weights for a given desired behavior, is also extremely important. This process is now shown to be related to finite impulse response (FIR) filter design:

Consider the Fourier transform [17] of the array weights:

$$X(e^{j\omega_t}) = \sum_{n=-N}^{N} a_n e^{-j\omega_t n}. \tag{2.13}$$

Comparing this with Equation 2.7 yields the identity:

$$R(u_0 + \frac{\omega_t c}{\omega_0 d},\ \omega_0) = X(e^{j\omega_t}), \tag{2.14}$$

for any ω_0.

Thus, finding the optimal weights a_n for a given $R(u, \omega_0)$ is equivalent to an FIR filter design problem for a given desired frequency response $X(e^{j\omega_t})$, where $\omega_t = \frac{(u-u_0)\omega_0 d}{c}$.

2.6 Filter-and-sum beamforming

The analysis of the delay-and-sum beamformer is now extended to a filter-and-sum beamformer operating in *discrete-time*. Thus, the elemental filters are assumed to be FIR filters of length M. The beamformer output sequence, for integer κ, is:

$$r[\kappa] = \sum_{n=-N}^{N} \sum_{m=0}^{M-1} f(\vec{x}, (\kappa - m)T)\ h_n[m], \tag{2.15}$$

where κ is the time sample, $h_n[m]$ is the impulse response of the filter associated with the nth element (i.e., the nth elemental filter), T is the sampling period, and $f(\vec{x}, t)$

is the incident wavefield. As before, we consider the complex exponential plane wave:

$$f(\vec{x}, \kappa T) = e^{j(\omega_0 \kappa T - \vec{k}_0 \cdot \vec{x})}, \tag{2.16}$$

for some frequency ω_0 and wavenumber $\vec{k}_0$. The beamformer output due to this wave is:

$$r[\kappa] = e^{j\omega_0 \kappa T} \sum_{n=-N}^{N} \sum_{m=0}^{M} h_n[m] e^{j(-\omega_0 mT + \frac{\omega_0}{c} u_0 nd)}, \tag{2.17}$$

where $u_0 = \sin\theta$ as before.

Once again, we remove the time-dependent portion to give the array response:

$$R(u, \omega) = \sum_{n=-N}^{N} \sum_{m=0}^{M-1} h_n[m] e^{-j(\omega mT - \frac{\omega}{c} und)}. \tag{2.18}$$

2.7 Two-dimensional filter design

Just as delay-and-sum beamformer design is related to FIR filter design, we now show that filter-and-sum beamformer design is related to two-dimensional (2-D) FIR filter design.

Consider the following function, which is an $N \times M$ set of 2-D dirac delta functions, weighted according to the elemental filter coefficients:

$$f(x, y) = \sum_{n=1}^{N} \sum_{m=1}^{M} \delta(x - n, y - m) h_n[m] \tag{2.19}$$

Now, the discrete 2-D Fourier transform of this function is:

$$X(\omega_1, \omega_2) = \sum_{n=-N}^{N} \sum_{m=0}^{M-1} h_n[m] e^{-j(\omega_1 n + \omega_2 m)}. \tag{2.20}$$

Note that this function is periodic with period 2π in both ω_1 and ω_2 dimensions. Combining this with Equation 2.18 yields:

$$R(\frac{\omega_1 cT}{\omega_2 d}, \frac{\omega_2}{T}) = X(\omega_1, \omega_2). \tag{2.21}$$

This relation is quite important. It implies that the filter-and-sum beamformer can be conceptualized as a 2-D filter whose Fourier transform is closely related to the desired array response. This idea is expanded into a design algorithm in Chapter 4.

2.8 $\omega_1 - \omega_2$ space

Note that $X(\omega_1, \omega_2)$ is exactly the familiar wavenumber-frequency response [10], after converting the wavenumber and frequency variables k_x, ω to the quantities $\omega_1 = k_x d$, $\omega_2 = \omega T$, which have units of pure radians. This representation is used in this chapter and Chapter 4 instead of the conventional wavenumber-frequency representation since it allows the filter design process to be specified independently of the spatial and temporal sampling rates d and T.

2.9 U-space

Throughout most of this chapter, the incidence angle was represented by the variable $u = \sin\theta$ rather than by the angle itself. This convention has several advantages beyond mere compactness. First, a change of steering direction amounts to a single linear translation in u-space. Thus, beamwidth measured in the u domain does not change with phase steering, while beams measured in θ space tend to broaden near endfire. Second, one can consider the response in the so-called *invisible region* corresponding to $|u| > 1$. These features are important to conceptualize and consider because they affect the array's noise sensitivity and also the beamshape when these features are brought into the visible region by phase steering. Finally, an array performance characteristic called the *directivity index* has a simpler expression in terms of u than θ.

2.10 Modeling assumptions

The above formulations contain some implicit assumptions about the physical and computational processes they represent. For example, they assume zero noise, zero quantization error, and omnidirectional transducers with flat frequency responses. The assumptions regarding the transducers are clearly not true of the microphones used in this thesis. But, if the assumption is made that all the transducers, while not ideal, are at least *identical*, then the equations will still apply to real systems after some simple corrections.

In particular, a non-flat microphone frequency response can be effectively corrected with a single inverse filter placed at the output of the beamformer. Non-omnidirectional directivity patterns cannot generally be compensated for, but one can mathematically model their behavior using the principle of *pattern multiplication*. This principle states that the total far-field response can be obtained from the far-field response of an equivalent array constructed using omnidirectional elements, and the far-field response of any of the individual microphones. Thus, the total array response is:

$$R(u,\omega) = R_m(u,\omega) \times R_a(u,\omega), \tag{2.22}$$

where $R_m(u,\omega)$ is the far-field response of a single microphone, and $R_a(u,\omega)$ is the far-field response of the equivalent array composed of omnidirectional elements.

Chapter 3

Firmware

This chapter describes the hardware and software which embody the mathematical descriptions of filter-and-sum beamforming described in Chapter 2.

3.1 Hardware

The beamformer hardware consists of analog and digital circuitry, much of which was custom built at Bell Laboratories' Acoustics Research Department [14]. A block diagram of the hardware is shown in Figure 3-1. Its major components are now summarized:

3.1.1 Transducer array

The transducer array contains four subarrays of eleven elements each, (see Figure 3-2). The microphones [18] are first-order gradient elements having a dipole response proportional to $\cos\theta$. These elements are mounted in circular holes in a flat metal bar and aligned such that the peaks of the dipole patterns lie in a direction normal to the bar.

The interelement spacings for the four subarrays are, from largest to smallest, 16, 8, 4, and 2 cm. For the intended frequencies of operation (as labeled in Figure 3-2), the interelement spacing never exceeds 0.47 wavelengths, thus slightly oversatisfying

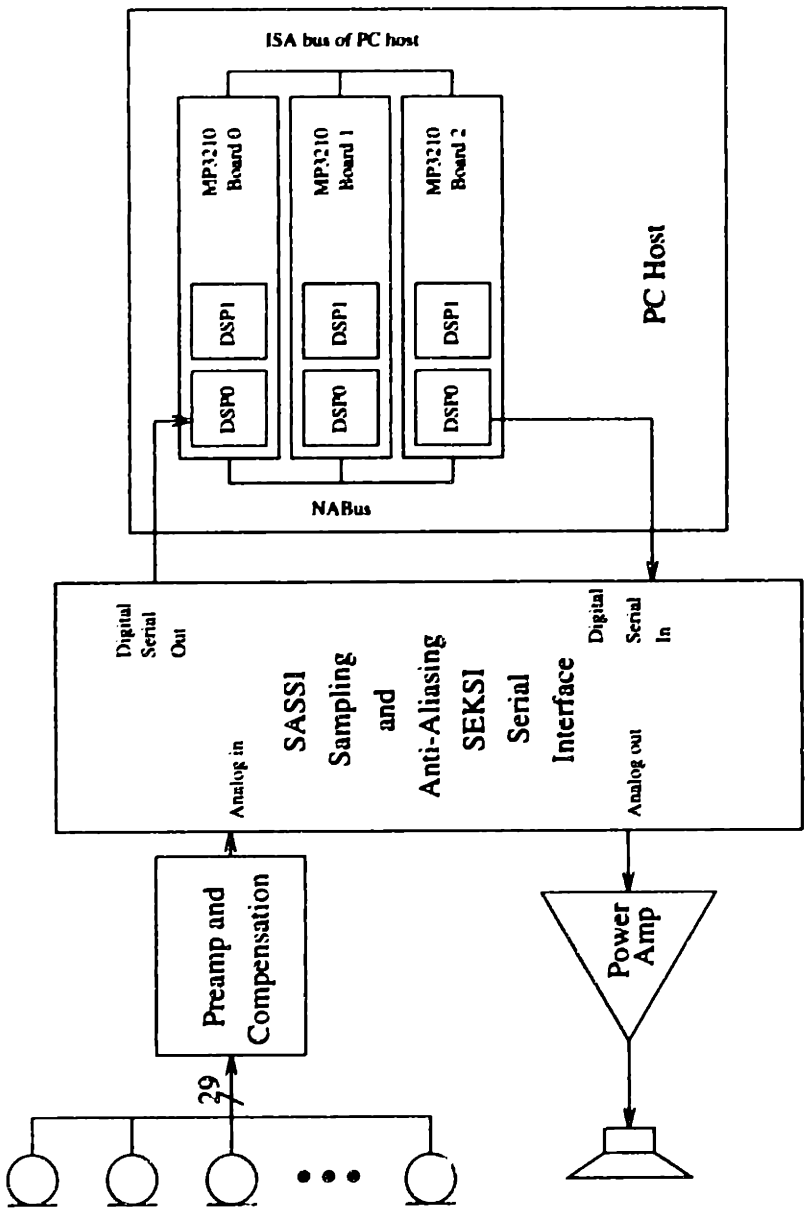

Figure 3-1: Block diagram of hardware platform

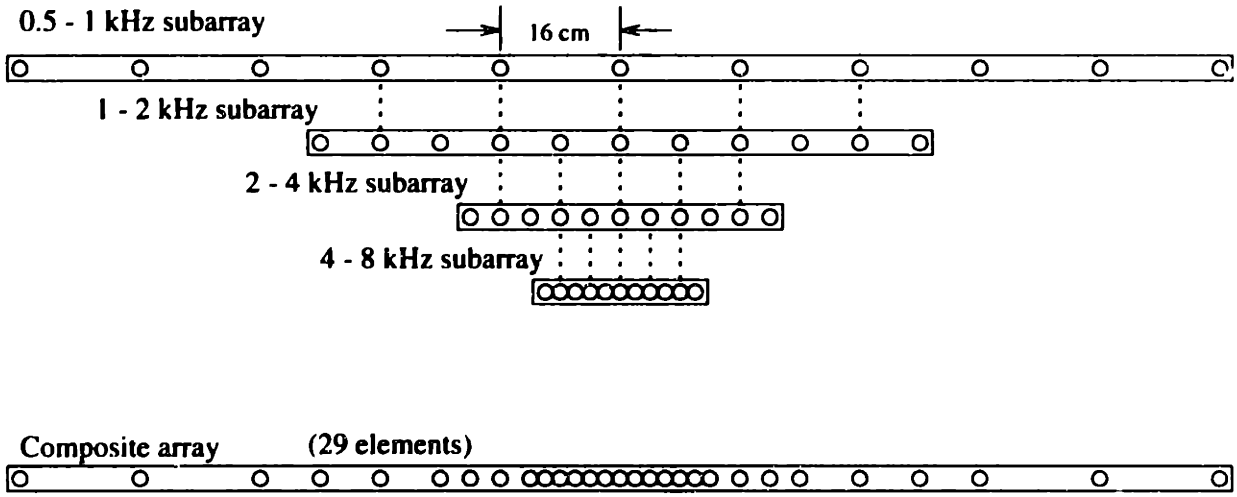

Figure 3-2: Microphone configuration in four-octave harmonically nested array (shown to 1/10 scale). Dashed vertical lines signify that elements in the composite array may belong to several subarrays.

the half-wavelength spatial sampling requirement.

3.1.2 Preamplifier and compensation

This unit compensates for the high-pass frequency response characteristics of the gradient microphones. The signal level is also increased to achieve an output of one volt for a 100 dB sound pressure level at the microphone diaphragm [14].

3.1.3 Sampling and Anti-Aliasing SEKSI Serial Interface

This unit, abbreviated *SASSI*, interfaces the analog and digital environments. The preamplifier's analog outputs are anti-alias filtered and then converted to 16-bit signed integers at a 16 kHz sampling rate. A full set of 29 microphone samples are obtained and digitized each sample period, and then converted into a serial bitstream which is fed to one of the DSP boards. Within the same time period, a single 16-bit integer is read from another DSP board and converted to an analog output.

3.1.4 DSP boards

The signal processing workload is handled by Ariel Corporation's MP3210 boards designed around AT&T DSP3210 digital signal processors. Support for multiple-board systems is provided by the so-called *NABus*, which provides a direct link between board address spaces, avoiding the bus of the host PC [2].

The choice of DSP platform was based on availability, a substantial amount of existing code, and a belief that a previously encountered hardware bug was being fixed. This bug manifested itself in code which performed simultaneous NABus accesses from different boards. Under such circumstances, a processor would occasionally halt, split open, and release noxious odors.

The hardware bug still exists, although we have learned how to program around it using a technique discussed in Section 3.2.1.

3.2 Software (DSP3210)

The software component consists of a real-time DSP3210 implementation of the filter structure shown in Figure 3-3. This figure shows the four sets of FIR filters, the IIR bandpass filters, and some preprocessing sections. The banks of elemental FIR filters implement a frequency-dependent array shading in a manner which was analyzed in Chapter 2. The functions of the other blocks in the diagram are now summarized:

The section labeled "redirection" sorts the microphone outputs by subarray. This section has more outputs than inputs because some microphones belong to multiple subarrays.

The section labeled "optional folding", when used, sums the outputs of elements which are symmetric with respect to the center element. This is useful in the common case of symmetric shading, since it reduces the number of required filters by almost two without affecting performance.

The section labeled "phase steering" is enclosed in dotted lines since phase steering is not investigated in this thesis. Note that phase steering is not incompatible with folding, provided that phase steering is implemented before folding, as indicated in

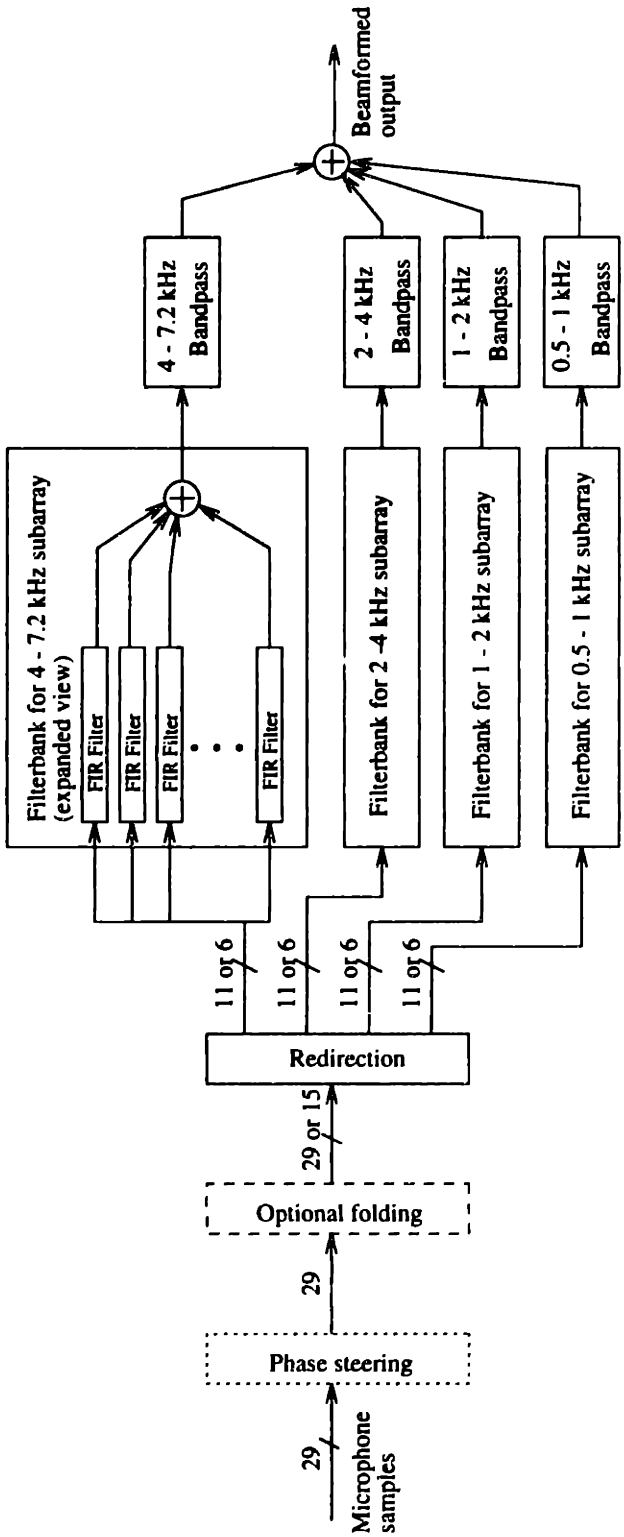

Figure 3-3: Block diagram of the DSP real-time signal processing routines. The four filterbanks are identical in structure, although only the lowest one is shown in expanded form. When folding is used, the smaller bus sizes should be assumed; otherwise the large values are assumed.

the block diagram.

Finally, the bandpass filters are used to suppress inter-band interference. These filters could be implemented as FIR filters as well. However, the overall system phase response was not considered crucial in this application. Thus, the bandpass filters were implemented as infinite length impulse response (IIR) filters, which generally require fewer coefficients than FIR filters to meet a given magnitude response requirement.

3.2.1 Timing of run-time code

In addition to the actual filter implementations, the run-time software contains mechanisms to synchronize activity among the processors. This avoids NABus conflicts and therefore avoids the hardware bug described in Subsection 3.1.4.

The DSP code is synchronized by a *token*, which can only belong to one DSP chip at any particular time, and which grants its owner the right to access the NABus. During each 1/16000 second sampling period, the token makes a complete circuit around the six processors. The token is implemented with a set of registers, one in the address space of each processor, and with a protocol in which a value of "one" is written into a register to indicate that the corresponding DSP owns the token.

The processing responsibilities are distributed among the six DSPs as described below, which are listed in the same sequence as visited by the token:

1. Board 0, Processor 0: Handles direct memory access (DMA) acquisition of microphone outputs. (DMA allows data transfers to occur without explicit processor intervention.)

   This processor also implements the folding and redirection sections shown in Figure 3-3. The redirected results are written via the NABus to one of the following four processors.

2. Board 0, Processor 1: Implements FIR filters for 500 Hz - 1 kHz subarray.

3. Board 1, Processor 0: Implements FIR filters for the 1 - 2 kHz subarray.

4. Board 1, Processor 1: Implements FIR filters for the 2 - 4 kHz subarray.

5. Board 2, Processor 1: Implements FIR filters for the 4 - 8 kHz subarray.

6. Board 2, Processor 0: Implements IIR bandpass filters for all subarrays. Since this processor never accesses the NABus, it uses the token not for the NABus privileges it confers, but for synchronization purposes. The single value obtained by summing the IIR filter outputs is sent to the SASSI box, also using the DMA subsystem of the DSP3210.

3.3 Platform capabilities

3.3.1 Hardware capabilities

This platform can achieve a beamwidth of 18 degrees (measured between half-power directions) over the 500 Hz - 7.2 kHz frequency range. The lower limit is set not by the analog/digital circuitry, which can process lower frequencies, but by the array geometry, which yields increasing beamwidths as frequency drops below 500 Hz. The upper limit of 7.2 kHz is set by the upper cutoff frequency of the bandpass filter associated with the highest frequency subarray. The net effect of these bandpass filters is shown in Figure 3-4. In addition to setting frequency limits, these filters introduce a small amount of ripple (generally less than 2 dB) into the overall system transfer function. This level of ripple is probably insignificant in typical audio applications.

3.3.2 Digital processing limitations

An important statistic is the maximum length FIR filters which can be implemented. The lengths act as constraints on the FIR filter design process and limit the precision to which desired array responses can be realized. The following discussion estimates the maximum allowable filter lengths based on empirical measurements.

Each DSP3210 has a 55 MHz clock, and can execute a single multiply-accumulate operation (or any other instruction) in four clock cycles. Given the 16 kHz sampling

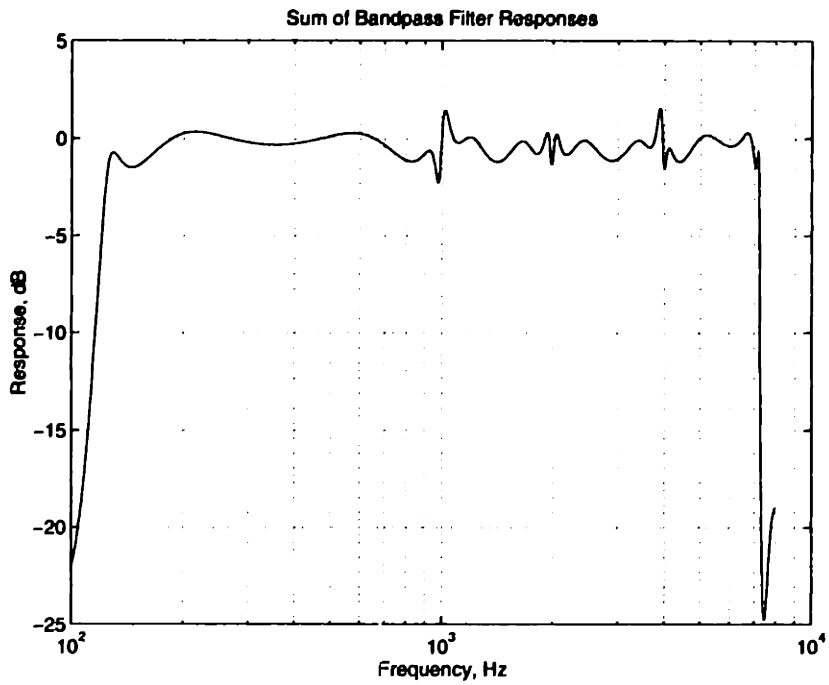

Figure 3-4: Sum of four IIR bandpass filters. The beamformer frequency response at broadside should be very close to this plot. Note irregular behavior around band transitions (at 1, 2, and 4 kHz).

rate, this leaves:

$$\frac{5.5 \times 10^7}{4 \times 1.6 \times 10^4} = 859.4 \tag{3.1}$$

instructions per sample period. From this limit, the maximum filter length will be calculated, using assumptions drawn from runtime examination of the MP3210 board's profiling counter registers. It has been observed that only about 700 to 750 of these cycles are actually available for executing the loop which implements the FIR filters for each microphone element. In our code, one DSP chip handles all the FIR filters for each subarray. Executing these filters requires, as a function of FIR filter length,

$$11 \times (FIRfilterlength) + X \tag{3.2}$$

cycles, where X is an overhead factor that should be constant across processors, since they all run identical code (with only the table of filter coefficients and some constants changed). However, empirical evidence shows that the overhead varies from about 72 to about 118 for reasons not fully understood. In the worst case, this means a maximum filter length of about:

$$[\frac{700 - 118}{11}] = 52 \tag{3.3}$$

where the brackets indicate values should be rounded downward. For symmetric arrays, there are only six filters, since opposite sensor outputs are added together. The folding occurs on a processor not running FIR filter code, and thus does not affect the filter computations. Here, the overhead factor X was not actually measured. Assuming that it does not increase, an upper bound on the FIR filter length is:

$$[\frac{700 - 118}{6}] = 97. \tag{3.4}$$

On-chip SRAM is also a precious resource, since it operates with no wait states, unlike off-chip memory. There are 8 kilobytes of on-chip SRAM, most of it available for storing coefficients that each occupy four bytes. However, for fast execution, between

0.5 and 1.0 kilobytes of the DSP code are copied into this RAM. Also required in the on-chip SRAM is a set of state variables which is almost the size of the coefficient table itself (although one element shorter for each microphone). Thus, the on-chip memory can accommodate filter lengths of up to about:

$$\left[\frac{7000}{2 \times 11 \times 4}\right] = 79 \tag{3.5}$$

for non-symmetric arrays and about:

$$\left[\frac{7000}{2 \times 6 \times 4}\right] = 145 \tag{3.6}$$

for symmetrically weighted arrays.

Clearly, processor time is the more precious computational resource, as it limits filter sizes to about 52 elements (97 if folding is used).

Chapter 4

Elemental filter design

This chapter discusses a process for designing the elemental filters required by the filter-and-sum technique. The filter coefficients can be downloaded onto the real-time beamforming platform described in Chapter 3, or used to simulate system behavior, as in Chapter 5.

The design process is first outlined in abstract form, and then carried out using the 1 - 2 kHz subarray as a concrete example.

4.1 Process outline

Section 2.7 showed that a filter-and-sum beamformer can be viewed as a 2-dimensional filter whose frequency response is exactly the wavenumber-frequency response expressed in normalized frequency units. Thus, the elemental filters can be designed using 2-D filter design techniques. However, the desired beamformer response must first be converted to a desired 2-D filter response, as follows:

1. Determine the desired beampattern, $\hat{R}(u)$. This determination should be based on the intended beamformer application, as well as on knowledge of what desired beampatterns can actually be realized. This selection is discussed further in the design example in Section 4.2.

2. The desired 2-D filter response $\hat{X}(\omega_1, \omega_2)$ is constructed by mapping $\hat{R}(u)$:

$$\hat{X}(\omega_1, \omega_2) = \begin{cases} \hat{R}(\frac{\omega_1 cT}{\omega_2 d}) & \text{for in-band temporal frequencies} \\ 0 & \text{for out-of-band temporal frequencies} \end{cases}, \quad (4.1)$$

where $\omega_1 = k_x d = |\vec{k}|\, d \sin\theta$, and $\omega_2 = \omega T$ are the radian spatial and temporal frequencies discussed in Section 2.8.

The beamforming problem is thus amenable to existing 2-D filter design techniques. However, it was found that many of these techniques were unnecessarily complex. Excellent results were achieved using a simple 2-D design by decomposition into two sets of 1-D design problems. Such a design algorithm is carried out by the following steps:

1. A finite rectangular grid in $\omega_1 - \omega_2$ space is defined.

   This so-called *sampling grid* is important in that this algorithm exactly achieves the desired response at these points. Also, the number of grid points exactly determines the 2-D filter size, and the grid point density is strongly correlated with the accuracy to which the desired response is approached.

   The grid points are defined by a set of samples in the ω_1 domain, and another in the ω_2 domain.

   The ω_1 domain is sampled at the following $2N + 1$ values:

$$\frac{2\pi n}{2N+1}, -N \leq n \leq N. \quad (4.2)$$

   Thus, there are exactly as many samples as transducers, and the samples are uniformly spaced in the range $[-\pi, \pi]$, which is significant since the function $\hat{X}(\omega_1, \omega_2)$ has periodicity 2π.

   Frequency samples can be chosen with more flexibility, and this selection is an important step in the overall filter design. The distribution of samples affects filter performance in that interpolation errors are generally reduced by placing

frequency samples closer together. Thus, a rule of thumb is that frequency sample density should be higher in frequency bands where greater accuracy is required.

A typical selection scheme might be to select six uniformly spaced samples within the octave of interest, and to select out-of-band samples at half the density of the in-band samples, since the out-of-band interpolation errors are strongly suppressed by the IIR bandpass filters.

2. The single-frequency beamforming problem associated with each constant-frequency grid line is solved.

   Each line of constant ω_2 value in the grids just defined contains $2N + 1$ sample points. The inverse discrete Fourier transform of the desired function values at these points yields transducer weights which achieve the desired beampattern at the sample points.

   Single-frequency beamforming is a very well-understood problem, and many variations on this algorithm can be constructed by varying the manner in which this stage of the design is carried out. One variation using the Dolph-Chebyshev approach is discussed in Appendix D.

3. Elemental filters are designed.

   Each element's frequency response has now been determined at the chosen set of frequencies. For each element an FIR filter is then designed that has the computed response at the corresponding frequency samples. Appendix C discusses an FIR filter design technique suitable for addressing this type of problem. This filter design technique reduces to an inverse Fourier transform if the frequency samples are uniformly spaced.

4. Interpolation errors are examined.

   Since the filter response is only explicitly constrained at a finite set of points, interpolation errors may occur in between these points. If such errors cause design criteria to be violated, the filter design steps must be repeated after

increasing the ω_2 sample density in the offending frequency regions. Likewise, economy may be achieved by lowering the sample density in regions where design criteria are oversatisfied.

It should be noted that since the ω_1 sample density is determined by the array hardware, interpolation errors in that dimension tend to predominate in the designs discussed in this thesis.

4.1.1 Discussion

A significant special case of the filter-design procedure occurs when the sampling grid points are uniformly spaced in both dimensions. The procedure then reduces simply to a 2-D Fourier transform.

The hardware used in this thesis has ample computational power for our purposes, and thus it was not essential that filter length be minimized. This may not be the case in lower-cost systems. In these cases, some of the more powerful 2-D filter design tools may achieve the same performance with a lower filter order than possible with the procedure outlined above. A basic review of 2-D filtering is given by Dudgeon and Mersereau [4]. Minimax approaches are discussed by Charlambous [1] and by Harris and Mersereau [6]. An l_p-norm optimization approach is described by Lodge and Fahmy [13]. Unfortunately, many optimal 2-D design approaches suffer from one or more of the following defects: high computational complexity, susceptibility to numerical instability, and restriction to square filters.

4.2 Design example

The design steps are now applied to the 1 - 2 kHz subarray depicted in Figure 3-2. This subarray contains 11 elements, and has an inter-element spacing of 8 cm (about a quarter of a wavelength at 1 kHz).

4.2.1 Selection of desired beampattern

Selection of the desired beampattern is crucial and nontrivial. In general, a beampattern $\hat{R}(u)$ cannot be exactly realized at all frequencies of interest. Thus, a pattern must be found which can be approximated with good accuracy at these frequencies.

Two desired patterns are suggested. The first is a sinc function:

$$\hat{R}(u) = \frac{\sin(\frac{11\pi u}{4})}{\frac{11\pi u}{4}}, \tag{4.3}$$

where $u = \sin\theta$, θ being the incidence angle. This function is exactly the far-field pattern of a uniformly shaded *continuous* aperture having length equal to $11/4 = 2.75$ wavelengths. Since this function cannot be precisely achieved at *any* frequency with a discrete array, the following "periodic sinc" function is proposed:

$$\hat{R}(u) = \frac{\sin(\frac{5}{2}\pi u)}{\sin(\frac{1}{2}\pi u)}. \tag{4.4}$$

This function has the form of the Fourier transform of five uniformly spaced impulses. This beampattern can be exactly realized at 2 kHz (by turning on only the elements numbered $0, \pm 1, \pm 2$) and at 1 kHz (by turning on only the elements numbered $0, \pm 2, \pm 4$). At intermediate frequencies, all 11 elements may be active, and the beampattern is only approximately achieved.

The two possible beampatterns (Equations 4.3 and 4.4) are shown in Figure 4-1. The patterns are generally similar within the visible region ($|u| < 1$). Grating lobes exist for the periodic sinc, but they are confined to the invisible region ($|u| > 1$).

The above design steps are carried out for each of the two patterns. As will be shown in Section 4.3, the periodic sinc leads to a more highly frequency-independent beamformer. This is achieved at the expense of a 10% wider mainlobe, somewhat higher sidelobes, and a greater sensitivity to uncorrelated noise because of the grating lobes in the invisible region.

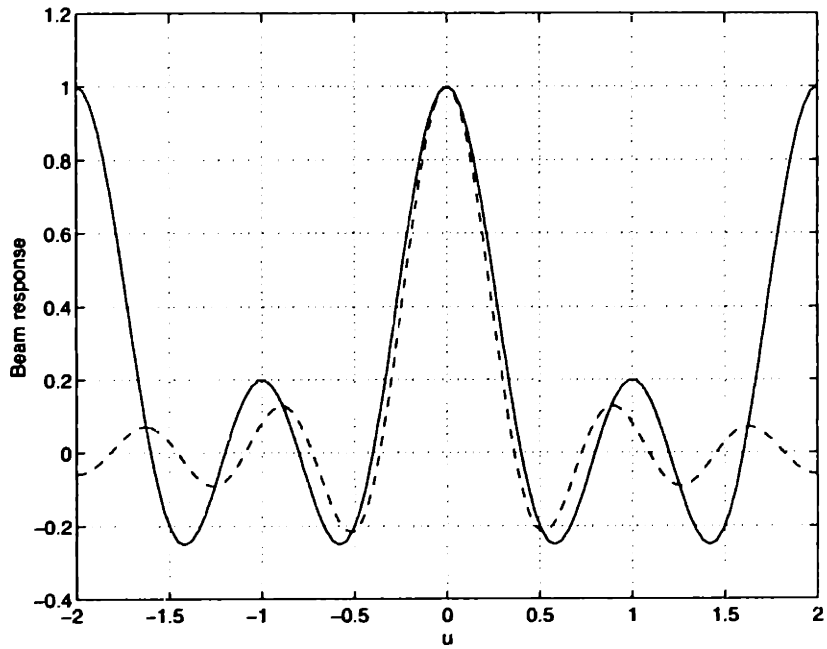

Figure 4-1: Possible desired array patterns. Dotted line is pattern #1, the sinc function; solid line is pattern #2, the "periodic sinc".

4.2.2 Mapping into desired filter response

Figure 4-2 shows the desired 2-D filter response function $\hat{X}(\omega_1, \omega_2)$ obtained from Equation 4.3. (The higher frequencies, which have zero response, are removed from the plot to emphasize the structure within the 1 kHz - 2 kHz band.) To help engage the reader's physical intuitions, this and subsequent filter response functions are labeled in the familiar units of wavenumber and frequency (in units of radians/meter and kilohertz), not the pure radian quantities ω_1 and ω_2.

The structure of the resulting wavenumber-frequency response is such that, within the 1-2 kHz range, the function has constant value along all straight lines passing through the origin.

4.2.3 2-D Filter design

The sampling grids for all four subarrays of our beamformer are shown in Figure 4-3. Note that each sampling grid exhibits six in-band frequency samples, and an out-of-

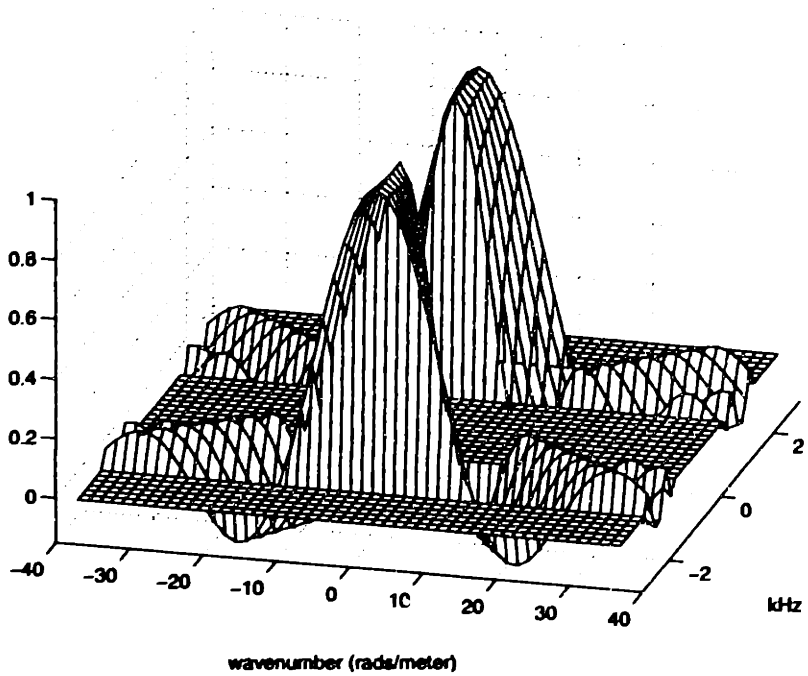

Figure 4-2: Desired wavenumber-frequency (*planewave*) response. Negative frequencies are shown to emphasize the quarter-plane symmetry. Frequencies above 3 kHz are not shown, in order to magnify in-band details.

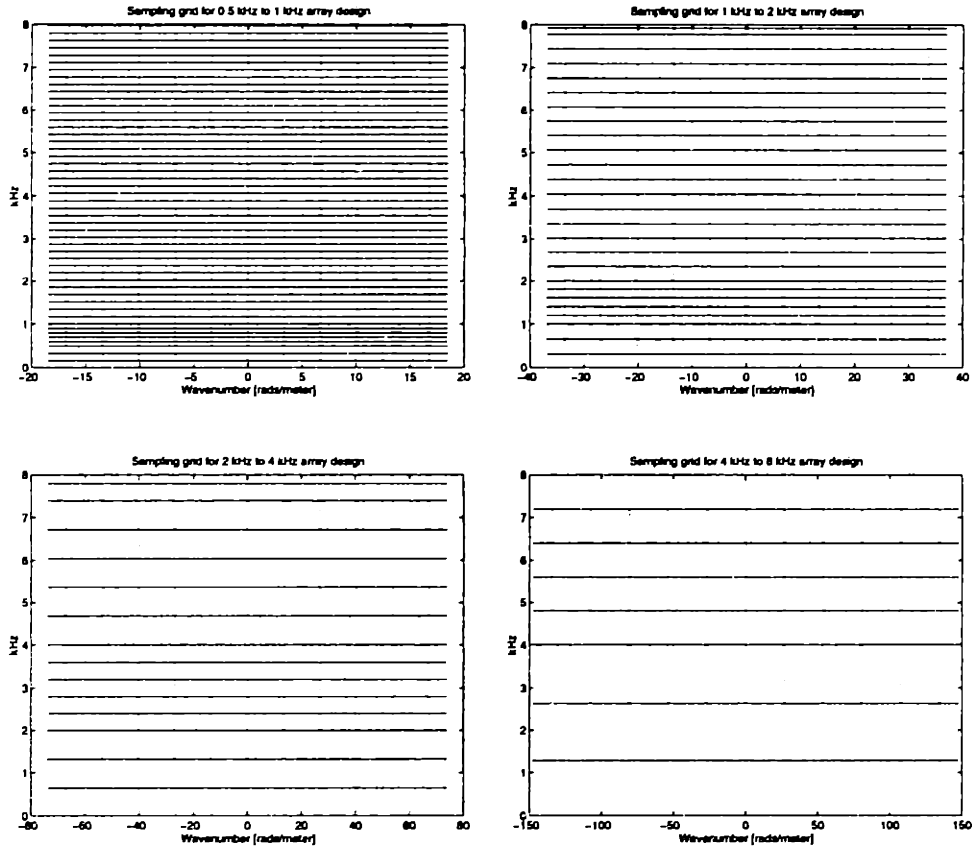

Figure 4-3: Plots of lines whose intersections define the sampling grid points. The desired wavenumber-frequency response is exactly achieved at each grid point.

band sample spacing which is 1.7 times larger than the in-band spacing. Performance was not found to be very sensitive to changes in either of these values, suggesting that the number of grid points could probably be substantially reduced without adversely affecting performance.

The axes are labeled in units of radians/meter and kilohertz, and correspond to values of ω_1 in the range $[-\pi, \pi]$ and values of ω_2 in the range $[0, \pi]$. Thus, negative temporal frequencies are not shown, thus omitting about one half of a symmetric set of points.

The numerical computations were carried out as prescribed, yielding the coefficients shown in Figures 4-4 through 4-7. Coefficients are only shown for the design based on the regular (as opposed to periodic) sinc function.

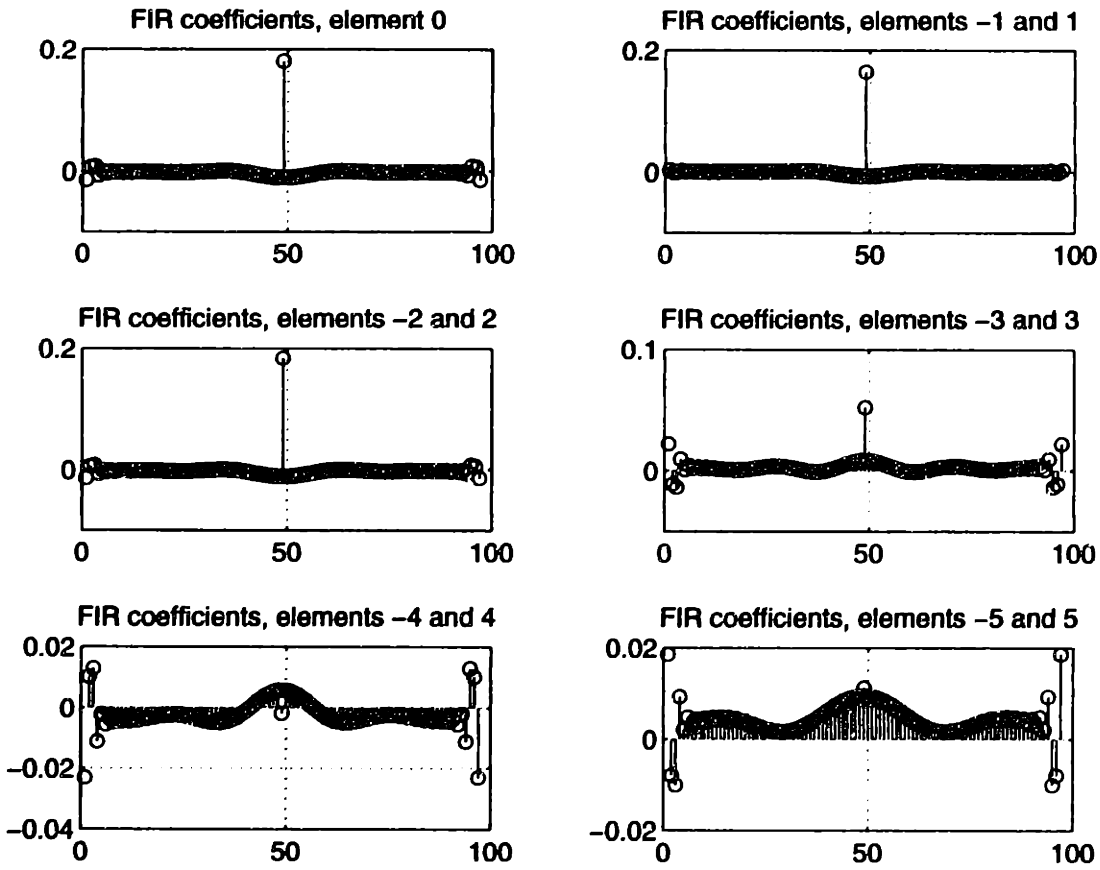

Figure 4-4: Impulse responses for 0.5 kHz - 1 kHz subarray. Filters are of length 97.

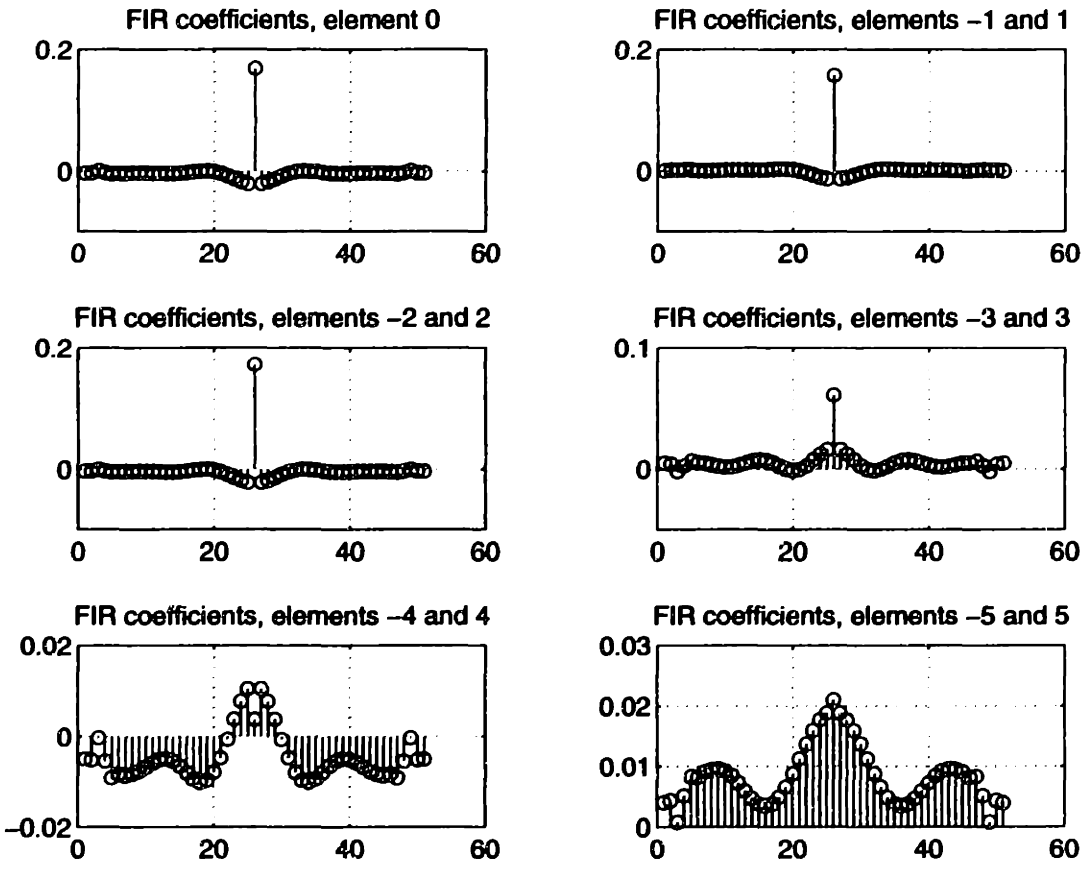

Figure 4-5: Impulse responses for 1 kHz - 2 kHz subarray. Filters are of length 51.

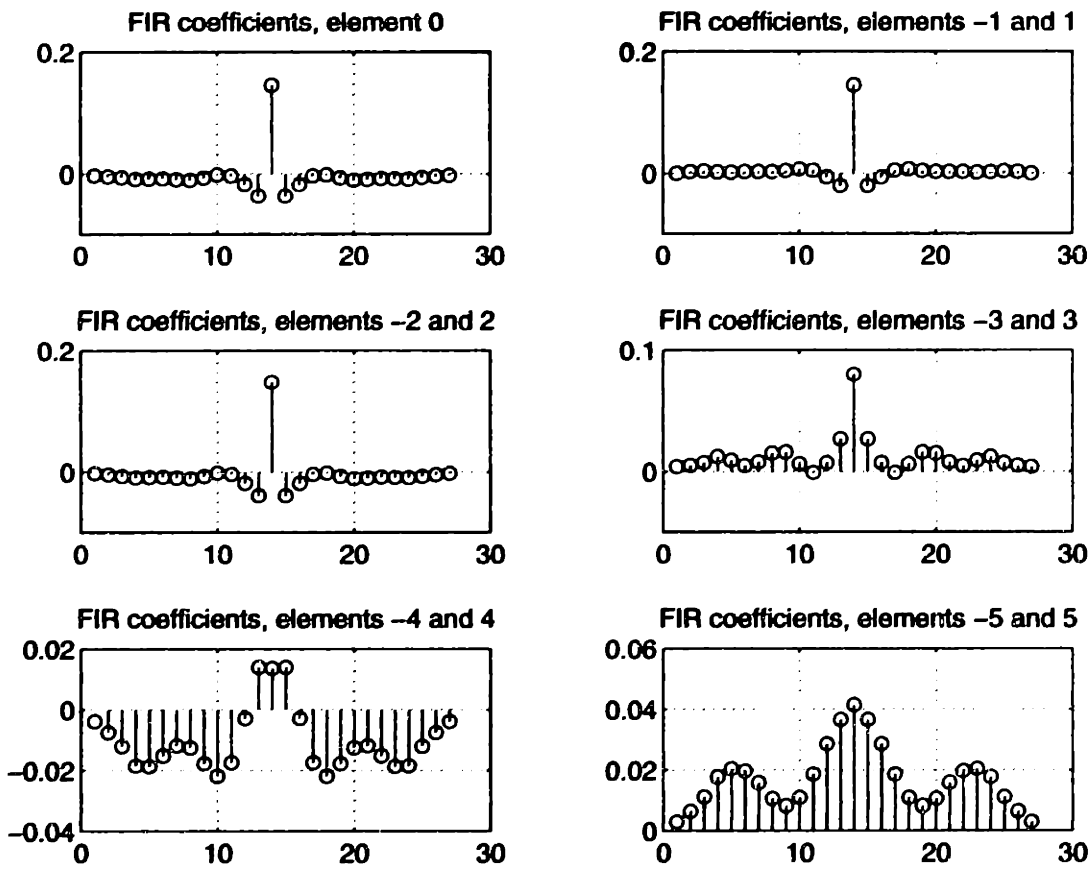

Figure 4-6: Impulse responses for 2 kHz - 4 kHz subarray. Filters are of length 27.

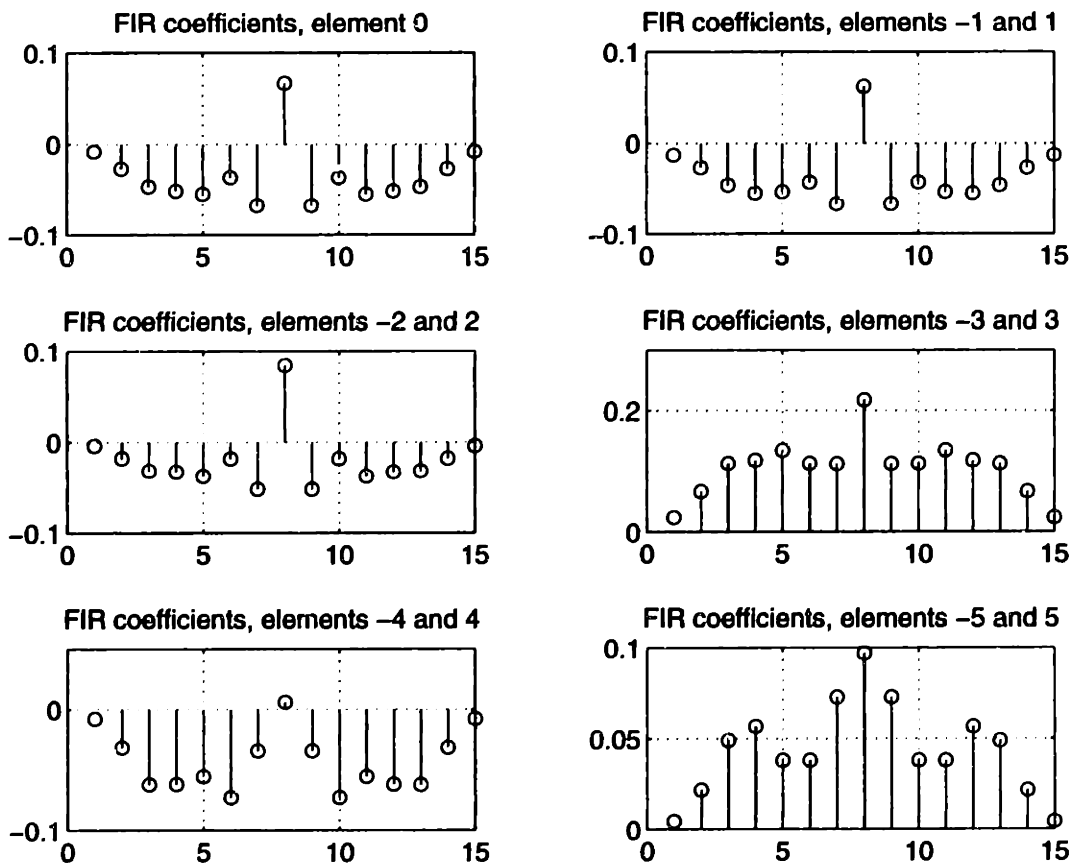

Figure 4-7: Impulse responses for 4 kHz - 8 kHz subarray. Filters are of length 15.

4.3 Performance limits

In general, a desired array response is asymptotically approached as the sampling grid density approaches infinity in both dimensions. However, the wavenumber sample density is proportional to the number of transducers, which may be an expensive number to increase. Even if the frequency sample density approaches infinity, interpolation errors in the wavenumber dimension generally remain, establishing a fundamental performance limit.

These limits were estimated by designing beamformers having a high number of frequency samples (26 in-band, 99 out-of-band). Figures 4-8 and 4-9 show the difference between actual and desired behavior, $X(\omega_1, \omega_2) - \hat{X}(\omega_1, \omega_2)$ across all incidence angles and in-band frequencies.

The two error surfaces are similar except at the endfire directions (where $|u| = 1$). There, the periodic sinc design has a peak error half as large as that of the ordinary sinc design. Thus, the periodic sinc pattern is used in the simulations of Chapter 5. However, the experimental measurements shown in Chapter 6 are based on the ordinary sinc since the periodic sinc idea was conceived too late for inclusion in anechoic chamber tests.

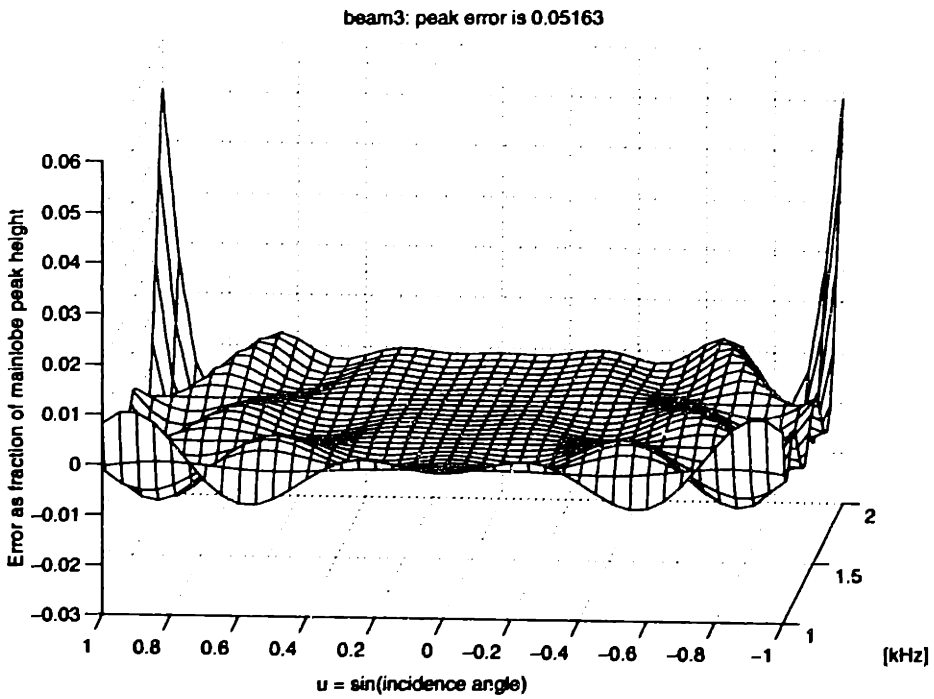

Figure 4-8: Deviations from desired pattern (pattern #1, sinc).

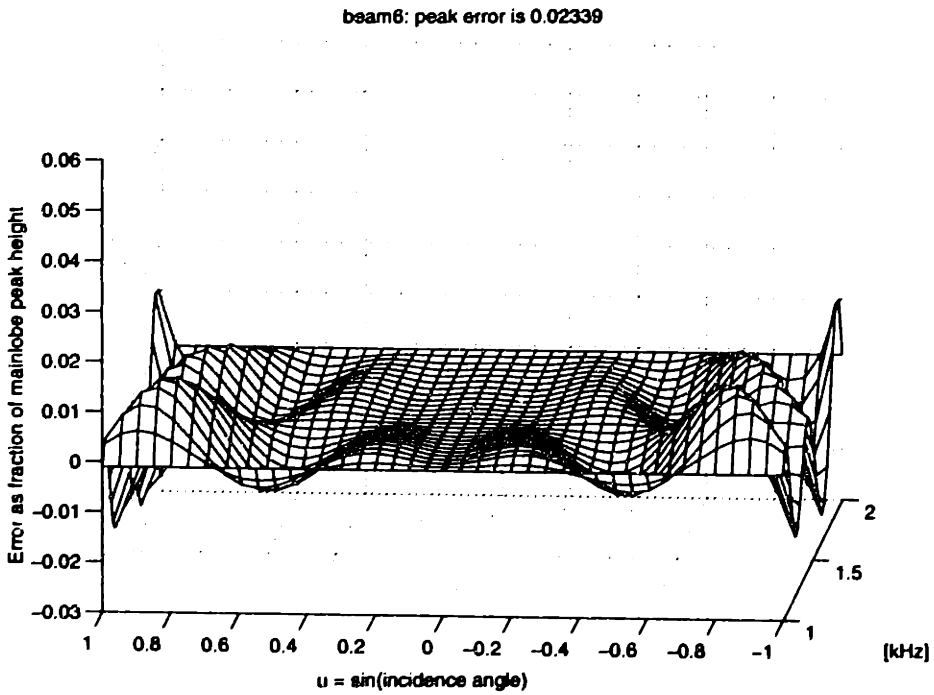

Figure 4-9: Deviations from desired pattern (pattern #2, periodic sinc). The peak error has been reduced by almost a factor of two - this reduction may not be obvious because the vertical axis is significantly magnified compared to the previous figure.

Chapter 5

Performance simulations

Differences between design objectives and actual performance are due mainly to: 1) approximations made by FIR and IIR filter design processes, and 2) non-ideal behavior of physical system components. The first type of error is examined in this chapter. The second type is examined in Chapter 6.

The effects of the filter design processes are assessed by numerically simulating the beamformer behavior, assuming ideal system components. Simulations were performed on individually operating subarrays, as well as the complete four-octave system. The individual subarray simulations were used to evaluate our 2-D FIR filter design process, while the four-octave simulations also assess IIR filter performance and interference between subarrays.

5.1 Individual subarray simulations

Simulations for the 1 - 2 kHz subarray show highly frequency-independent behavior according to several measures of array performance. Results for the other three subarrays are very similar.

Simulated data were obtained by using Equations 2.20 and 2.21 to compute the planewave response for a large number of wavenumber-frequency values. For frequencies within 1 - 2 kHz, Figure 5-1 shows the difference between these computed values and the desired response. Note that this difference plot is almost indistinguishable

from Figure 4-9, i.e., the performance is almost the best achievable using this design method in conjunction with the specified array geometry and desired beampattern.

Out-of-band deviations from desired responses (i.e. deviations from zero) also exist. Since a subarray's out-of-band response is suppressed at least 30 dB by the associated bandpass filter, the out-of-band deviations are not as serious and their discussion is delayed to Section 5.3.

The in-band data are further interpreted in terms of the following measures of array performance: 1) frequency response, 2) array pattern, and 3) beamwidth variations. These representations are displayed in Figure 5-2. The upper left subplot shows frequency responses at selected incidence angles; the upper right subplot shows angular responses at selected frequencies; the lower left subplot is a mesh plot showing all available in-band data; finally, the lower right subplot shows the beamwidths at various thresholds of beamwidth.

The subplots all show behavior which is extremely frequency-independent. The frequency response variations around the mainlobe are well under 1 dB, and probably inaudible. The array patterns are nearly coincident. The mesh plot shows few frequency variations, except for a slight ripple (under 2 dB) in the height of the endfire sidelobe, as well as some larger ripples (about 5 dB) in the frequency responses at angles near the nulls. These latter ripples are to be expected, since the decibel scale tends to magnify absolute deviations of numbers close to zero. Finally, the fourth subplot shows that two-sided beamwidths vary by less than a degree over the octave.

Finally, Figure 5-3 shows that the nulls are extremely constant with frequency, varying by less than a degree over the frequency band of interest.

5.2 Four-octave simulations

Figures 5-4 and 5-5 show the behavior of the entire system of filters (which was shown in block form in Figure 3-3). The first figure shows the magnitude responses at broadside (0 degrees), the half-power direction (9.8 degrees), the direction of the first null (21 degrees), and at an intermediate angle of 15 degrees. The behavior is

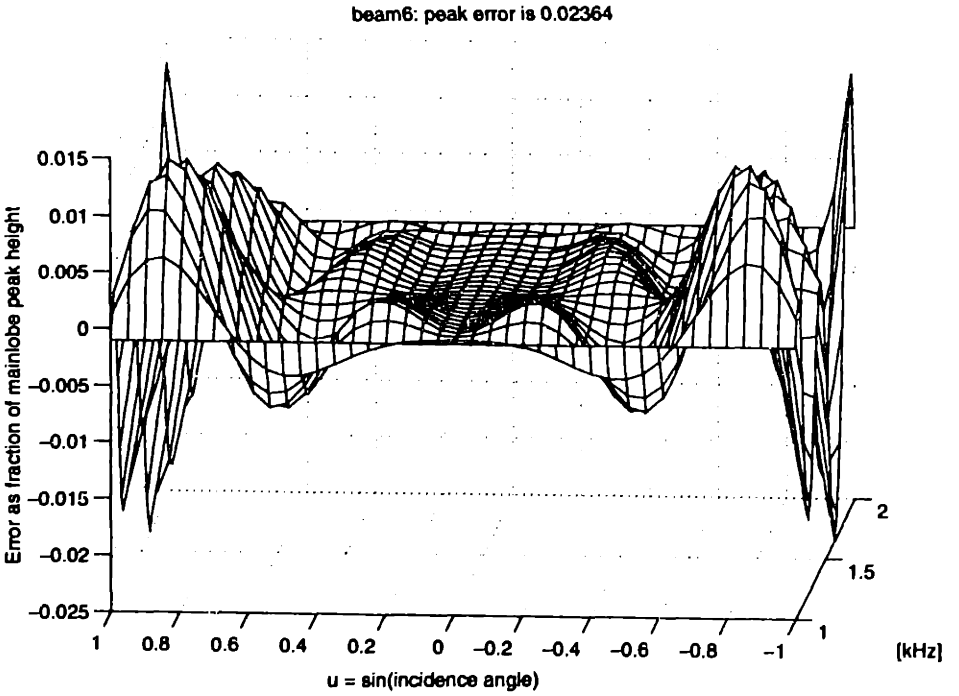

Figure 5-1: Simulated deviations from desired pattern. Note that these values are very close to the fundamental performance limits.

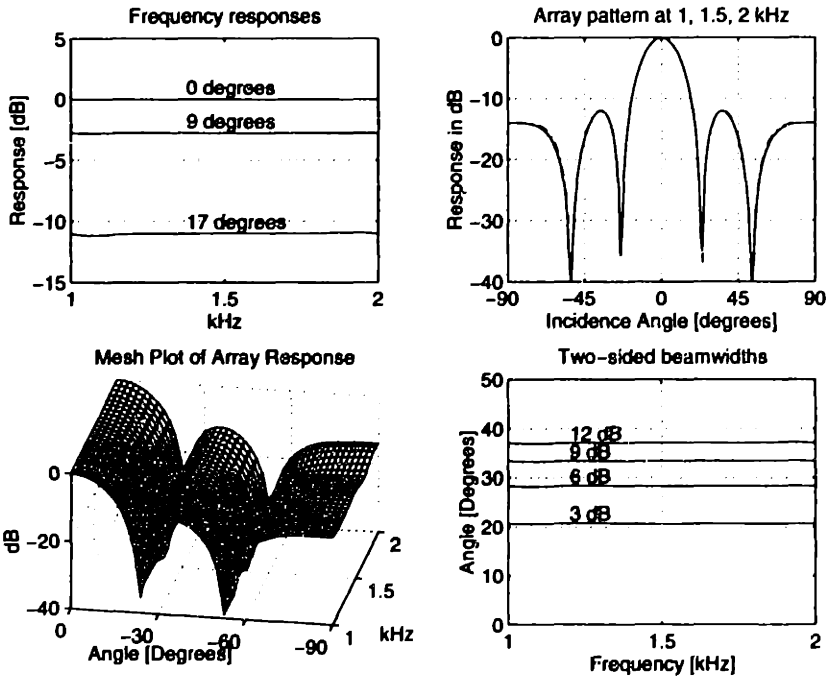

Figure 5-2: Simulated single-octave performance measurements.

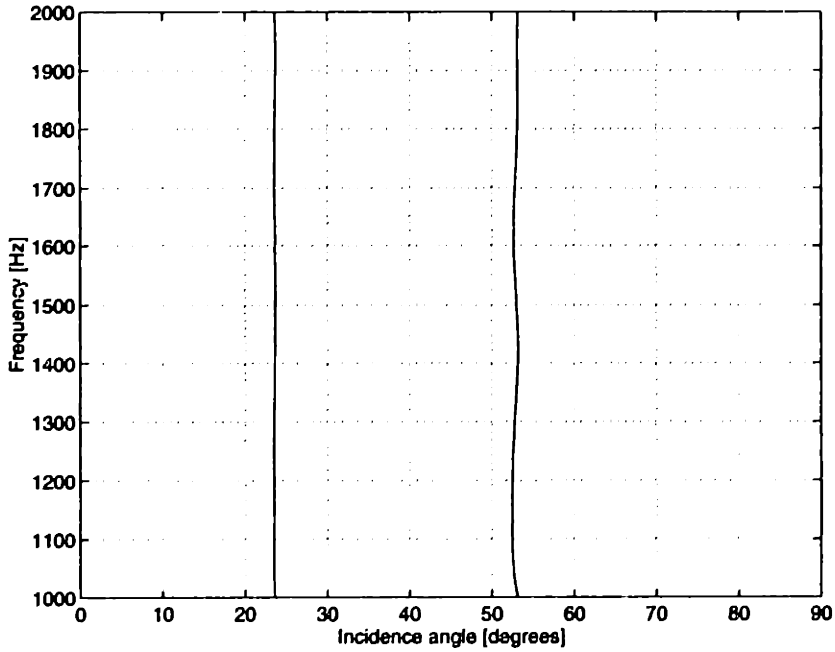

Figure 5-3: Simulated null locations of the single-octave beamformer.

close to ideal with two notable exceptions: the ripples introduced by the combined effect of four bandpass filters (essentially the result of multiplication by the curve shown in Figure 3-4, which is the sum of the four bandpass filter responses), and the tendency toward omnidirectionality below 500 Hz. The former effect, as discussed in Section 3.3.1 is not severe, and the latter effect for this design approach is a limitation imposed by the array length.

In contrast to the single-octave performance, the four-octave performance exhibits a null having somewhat greater variations with frequency (see Figure 5-5), as well as some erratic locations of zero-response above 7.2 kHz. However, the null location still only varies by about 3 degrees. Also, the apparent irregularities above 7.2 kHz can be ignored since they occur at frequencies where all four subarray responses are greatly attenuated (see Figure 5-6).

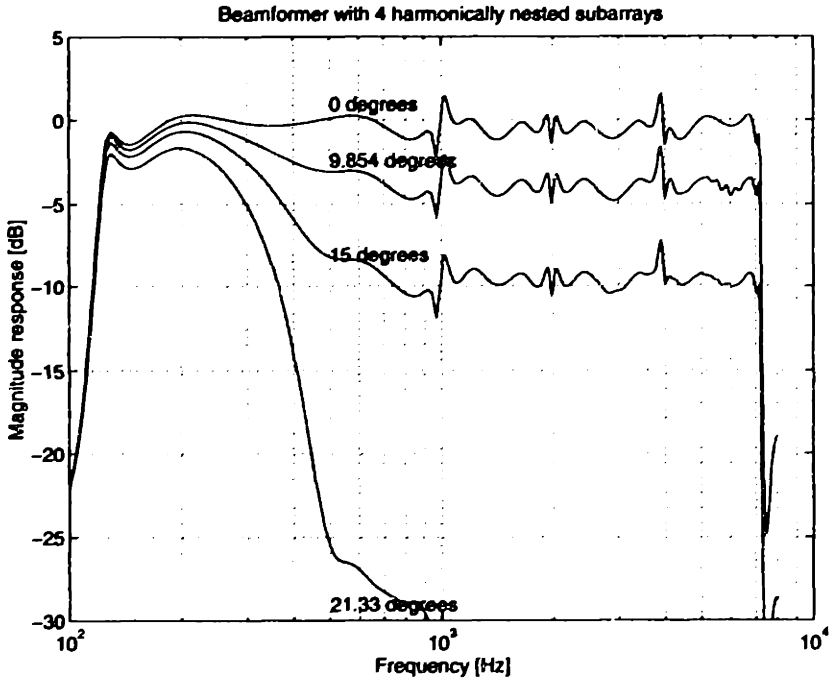

Figure 5-4: Simulated frequency responses of four-subarray beamforming system

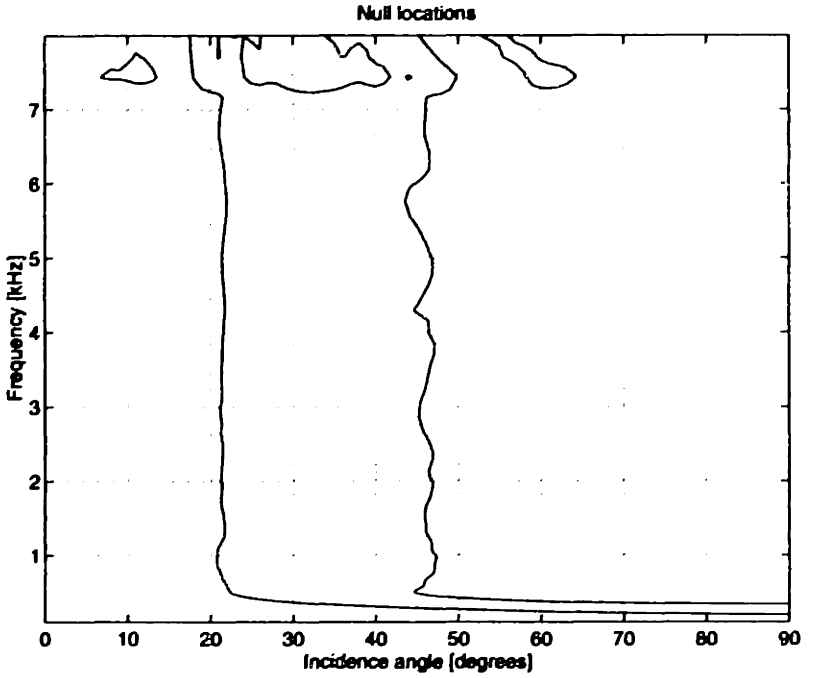

Figure 5-5: Simulated null locations for complete four-subarray beamforming system

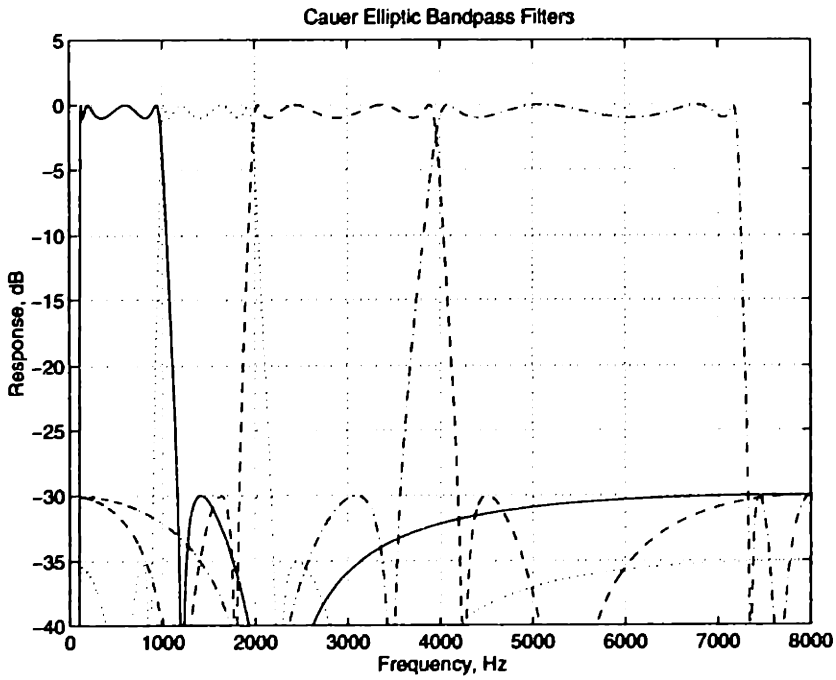

Figure 5-6: Responses of IIR bandpass filters

5.3 Elemental filter magnitude responses

The elemental filter magnitude responses are displayed and discussed in this section. Examining these responses is important since the out-of-band frequencies were not tightly controlled by the design process and may exhibit unexpected behavior.

The data for all 44 elemental filters are shown in Figures 5-7 through 5-10. The elements are numbered from -5 to 5, in accordance with the convention adopted in Chapter 2.

The filter magnitude responses are given in two representations; on a per element basis as a set of frequency responses (top six subplots of each of Figures 5-7 through 5-10), and on a per frequency basis as a set of array shading patterns (bottom six subgraphs of each of the figures).

The first representation illustrates the desired and achieved magnitude responses. The desired responses are the constraints located by the circles on the graphs. The in-band desired response constraints are marked additionally by x's. The achieved

response functions are then superimposed, and pass through all the constraints.

Since the in-band errors have already been closely analyzed, the elemental filter responses are now examined in terms of out-of-band errors. Most of the elemental filters have out-of-band responses with the same or lower magnitude than the peak in-band response. These deviations are probably adequately suppressed by the 30 dB stopband rejection of the bandpass filters. The notable exception is the 4 - 8 kHz subarray, in which the DC responses of several elemental filters are between 10 dB and 20 dB above the peak in-band response. Since there were no performance guarantees below 500 Hz, the low frequency behavior is not a serious issue. However, the elevated out-of-band responses extend up to 1 kHz for two of the filters (corresponding to elements 3 and 5), and might be expected to interfere with the lowest subarray. However, no unexpected disturbances are apparent in Figure 5-5. Thus, the out-of-band errors apparently do not degrade system performance.

Finally, the magnitude responses are also represented in the form of the array shading at selected frequencies within the 1 - 2 kHz octave. These pictures provide intuitive insight into the operations of the beamforming algorithm. The array shading is seen to approximate a boxcar whose width shrinks and height grows in proportion to frequency. Indeed, this simple "shrinking-aperture" model has been explicitly and elegantly used in at least one other frequency-invariant beamforming scheme [25].

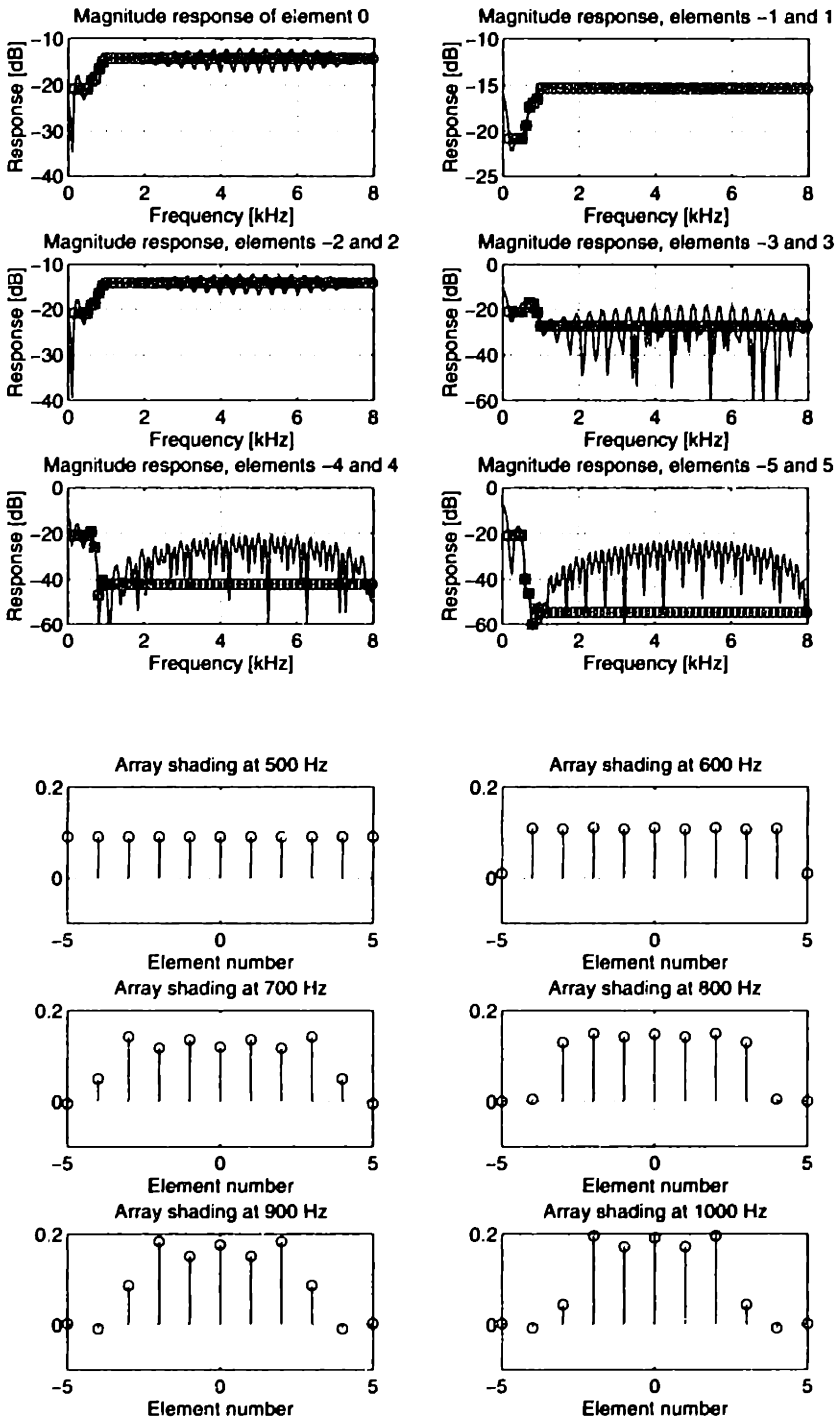

Figure 5-7: Magnitude responses and array shadings for 0.5 kHz - 1 kHz subarray.

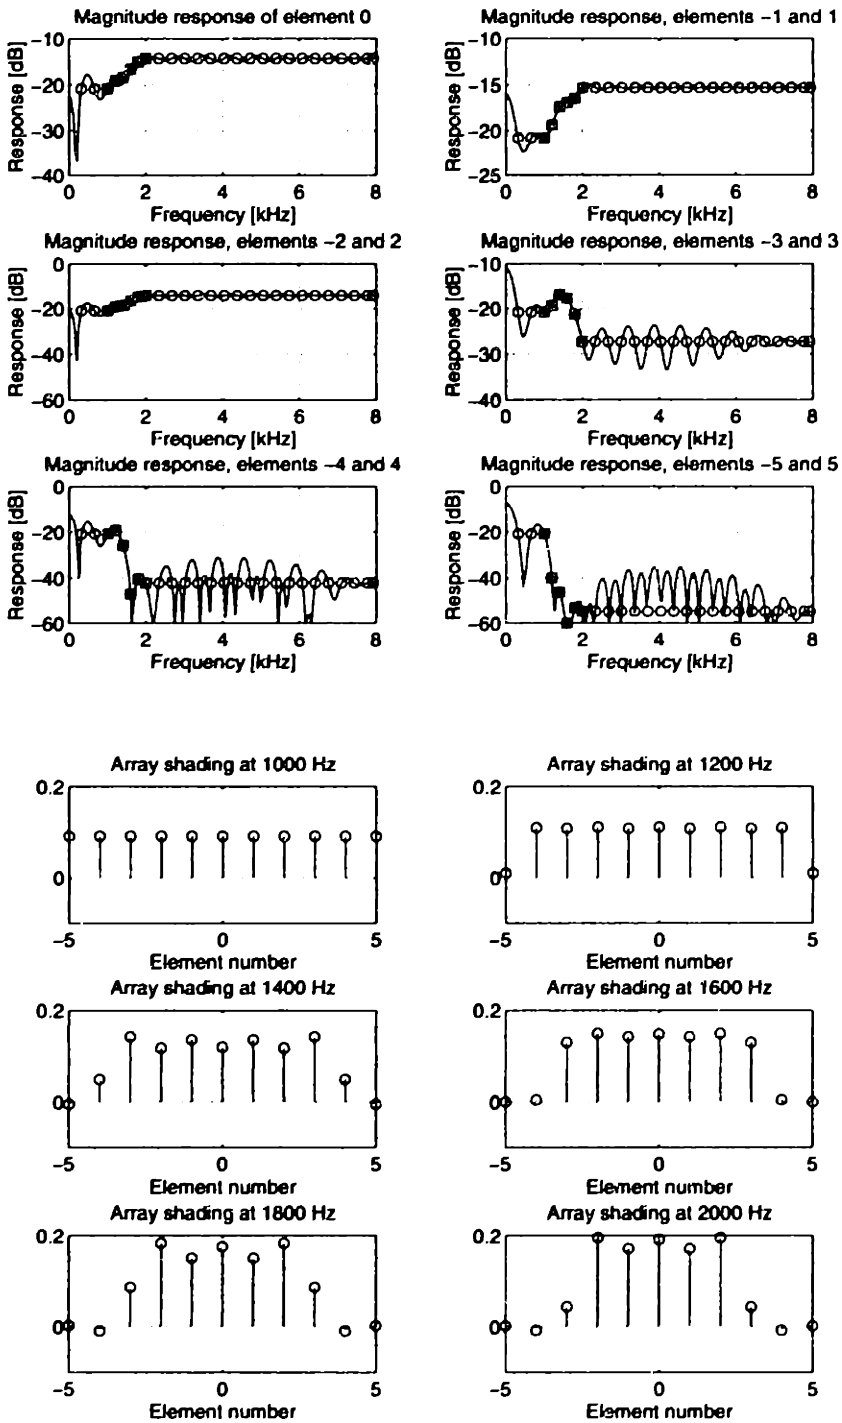

Figure 5-8: Magnitude responses and array shadings for 1 kHz - 2 kHz subarray.

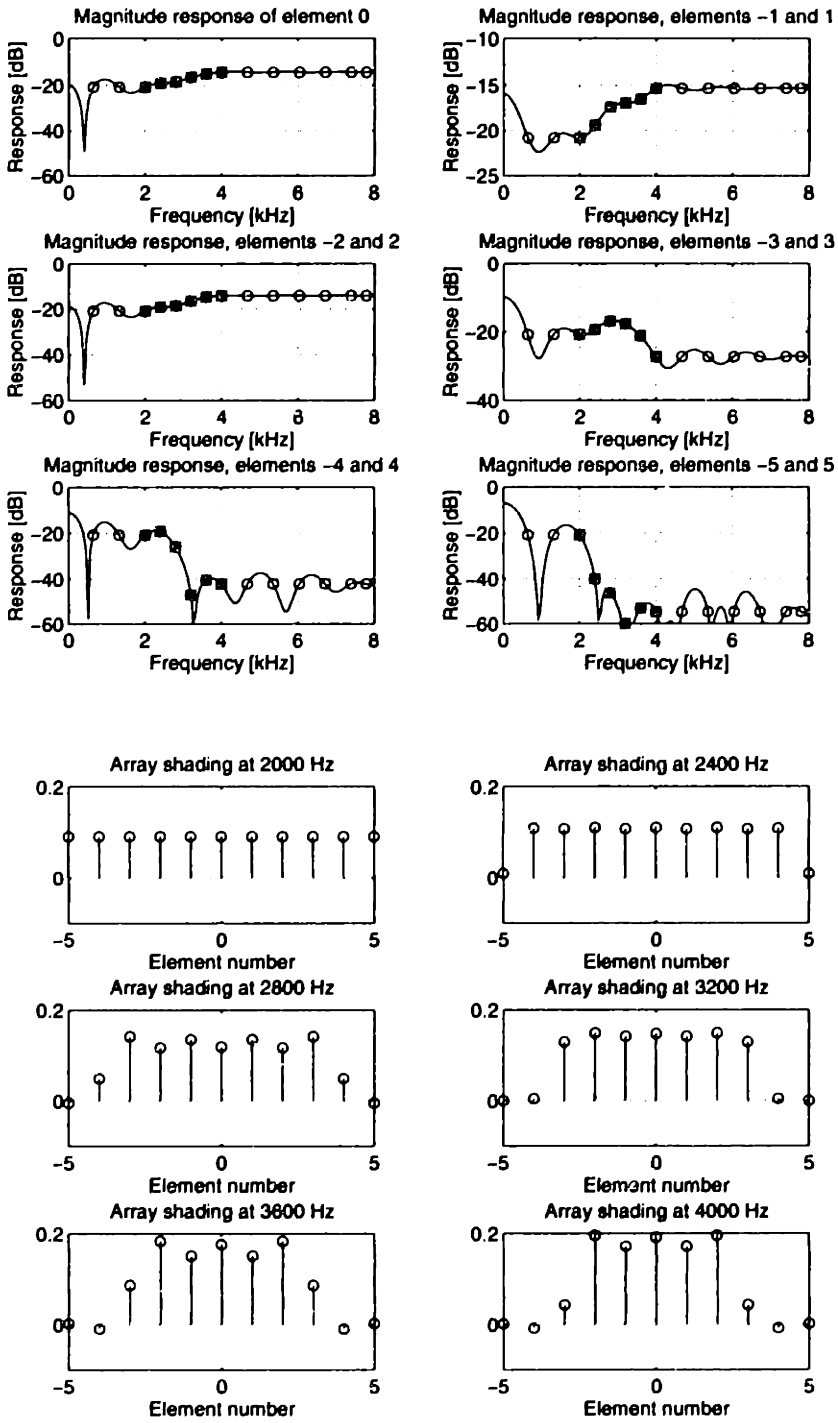

Figure 5-9: Magnitude responses and array shadings for 2 kHz - 4 kHz subarray.

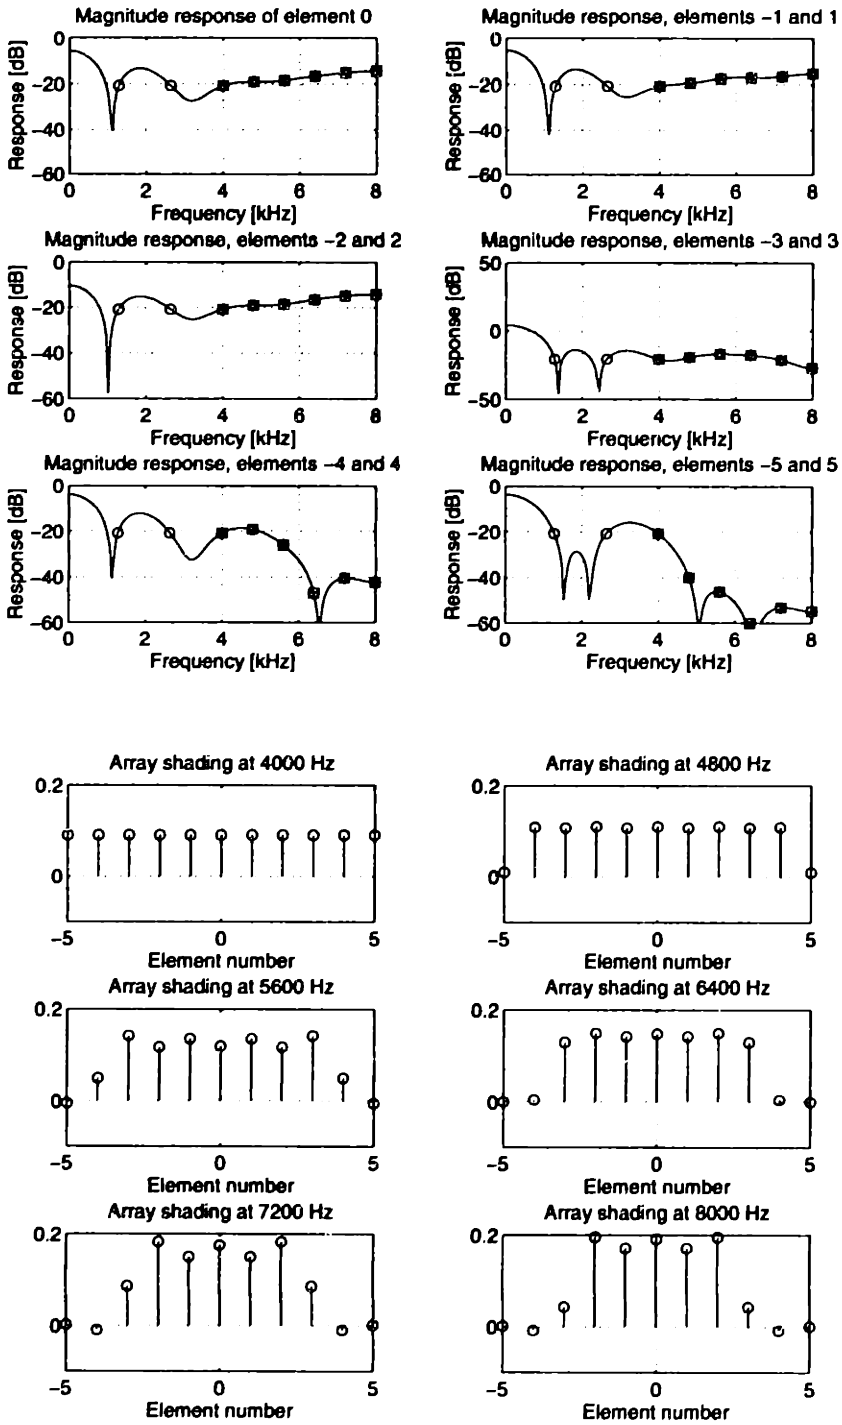

Figure 5-10: Magnitude responses and array shadings for 4 kHz - 8 kHz subarray.

Chapter 6

Performance measurements

The real-time beamformer far-field responses were physically measured, and then compared with the simulations computed in Chapter 5. These measurements were used to identify the effects of system components (particularly transducers) with non-ideal behavior. As with the simulations, measurements were made on individual octaves, as well as on the system as a whole.

6.1 Experimental setup

All measurements were made in the anechoic chamber at the Murray Hill location of AT&T Bell Laboratories. A loudspeaker was placed in one corner of the chamber, while the array was mounted horizontally on a rotating shaft suspended from the ceiling at the other corner of the chamber. The array was positioned such that the axes of symmetry of the dipole microphones were all lying in a horizontal plane, mimicking the intended real-world orientation. The array's magnitude responses were measured while rotating the array through a full circle in the presence of sinusoidal signals emitted by the loudspeaker.

6.2 Single-octave measurements

Figure 6-1 compares anechoic chamber measurements (solid lines) with computed directivity patterns (dashed lines) at six frequencies selected within the octave of interest (1 - 2 kHz). Each polar plot is scaled such that the peak is unity, thus compensating for the frequency responses of the array and the measuring system itself.

In general, there is good correspondence between the calculated and measured values. In particular, none of the sidelobes are more than 3 dB higher than predicted. Also, the nulls are all deeper than -20 dB. This level of performance represents an improvement over previous results obtained with this hardware [14], and was achieved by compensating the filter coefficients for relative differences in microphone sensitivities. Figure 6-2 shows the compensation factors, which are scale factors obtained by taking reciprocals of the magnitude responses of the microphones at 1 kHz. Clearly, the assumption is made that inter-microphone differences can largely be explained by a single scale factor. This assumption appears to be validated by the measured results.

6.3 Four-octave measurements

Performance measurements are also made over the entire 8 kHz operating range. Figure 6-3 shows the beamformer's frequency response to signals arriving at broadside (zero incidence angle). Figure 6-4 shows responses, relative to the broadside response, for signals with incidence angles of 10 and 20 degrees. The latter figure shows that the relative responses are quite frequency independent, with attenuations near 4 dB ($\theta = 10$ degrees), and 18 dB ($\theta = 20$ degrees). Such frequency-independence was one of the main design goals.

The broadside response is not as well behaved, showing a marked low-pass characteristic, in addition to other irregularities. The low-pass effect is principally due to the frequency response of the test equipment (particularly the loudspeaker used

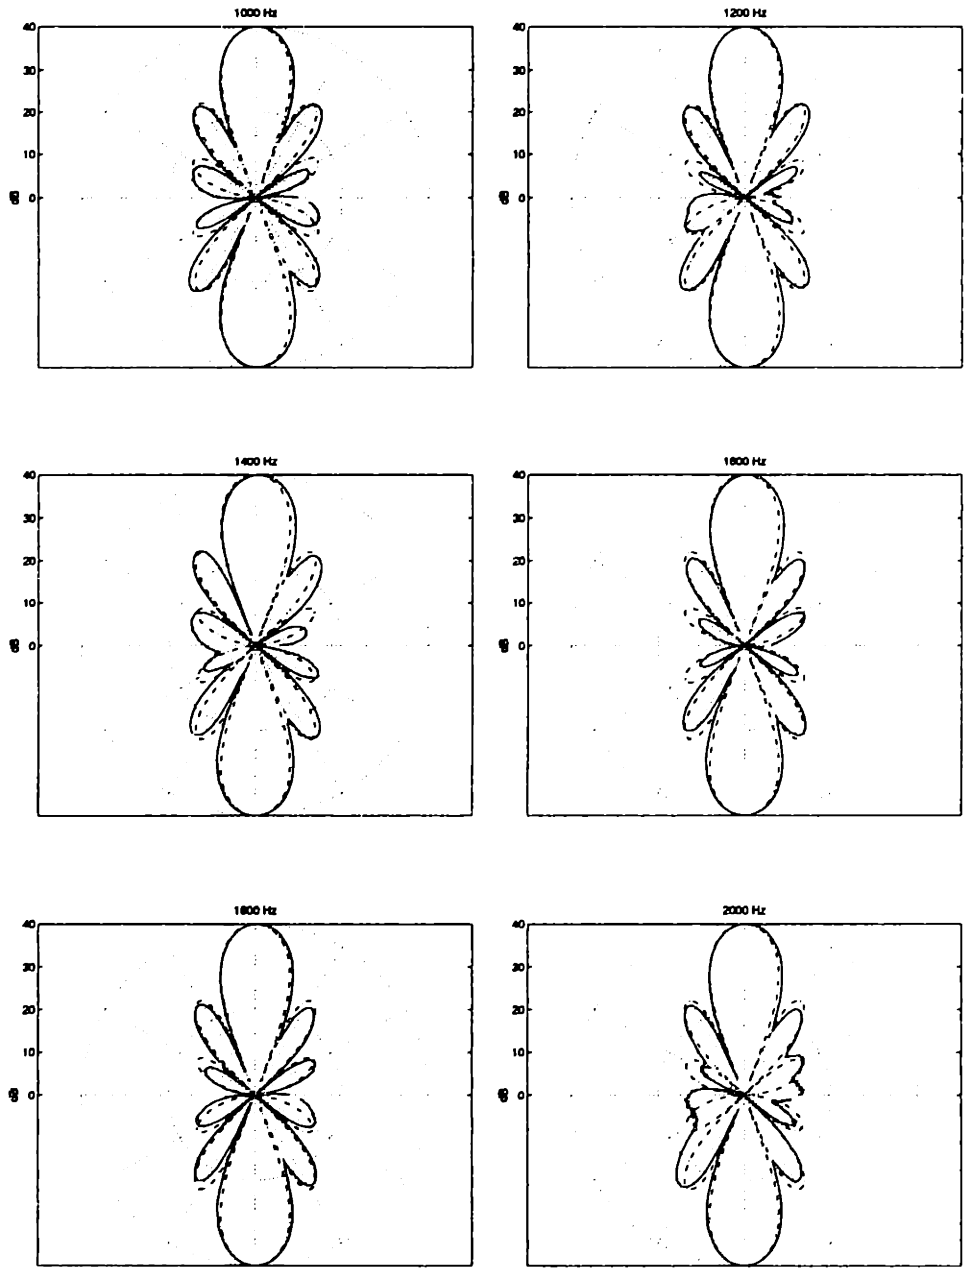

Figure 6-1: Measured responses after simple compensation for microphone variations. Dashed lines again are theoretical patterns.

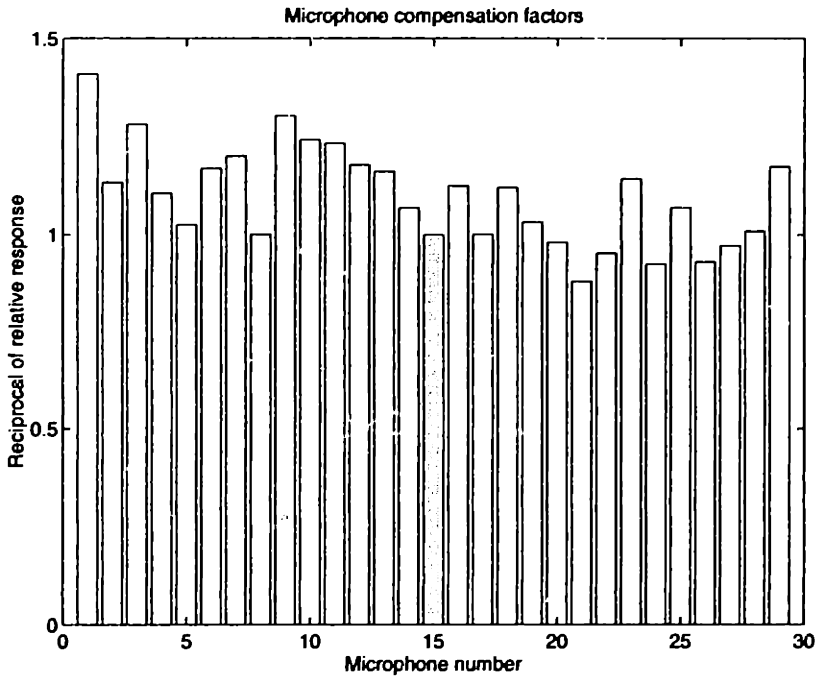

Figure 6-2: Correction factors for inter-microphone variations.

to generate test signals), as well as the microphones in the array. The broadside response also drops significantly (nearly 20 dB in one case) at the transition frequencies between adjacent octaves. These notches arise from phase differences between the outputs of adjacent subarrays at these transition frequencies. The perceptual impact of these notches is unclear; they can probably be removed using a technique called *phase centering* [14], in which the bandpass filters for each subarray are designed to have the same relative phases.

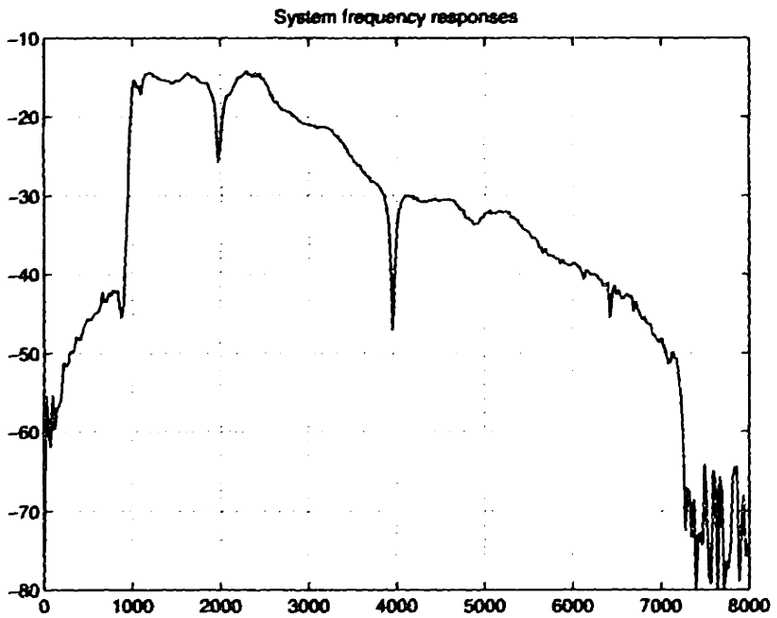

Figure 6-3: Frequency response of four-octave system to broadside signals. The lowest subarray was turned off during these tests, due to hardware malfunction.

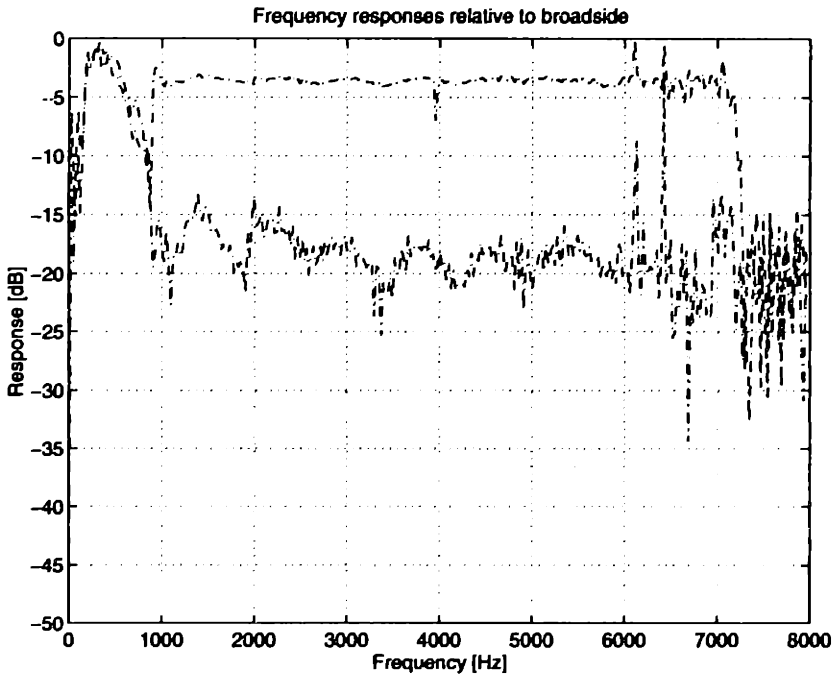

Figure 6-4: Frequency response to signals arriving at 10 and 20 degrees from broadside. Decibel values are shown relative to broadside response. Responses are valid above 1 kHz - see previous figure.

Chapter 7

Conclusions and future directions

A highly frequency-independent beamformer was designed, simulated, implemented, and experimentally evaluated.

This design operates over nearly a four-octave range, from 500 Hz to 7.2 kHz, and features a harmonically nested array structure combined with a filter-and-sum signal processing structure. The array structure allows a hardware-efficient implementation, while the filter-and-sum approach provides the flexibility needed to control the beamformer's frequency-variations.

Simulated results show very small differences from ideal frequency-independent behavior (most plane-wave responses differ from ideal by under 1% of peak plane-wave response). Differences for experimental responses are somewhat larger, but are still under 5% of the peak plane-wave response. In particular, experimentally measured nulls are all more than 20 dB deep, which is probably adequate for operation in typical room environments. Finally, the null directions are extremely frequency-independent in both simulated and measured data, thereby enabling broadband cancellation along specific directions.

The following sections describe major directions for future work.

7.1 Higher dimensional broadband arrays

Conceptually, our approach easily extends to two-dimensional (planar) arrays. Indeed, the design problem then becomes analogous to three-dimensional filter design. However, implementational problems are great due to the large array structures and high signal processing demands. One active area of investigation involves a parallel architecture called a *systolic* array, which can efficiently execute common beamforming tasks using a network of processors in which only adjacent processors are actually required to exchange information.

7.2 Computational efficiency via subband processing

Our harmonically nested structure is amenable to subband processing, significantly reducing the lengths of the elemental FIR filters. This technique relies on the fact that the lower three of the four subarrays do not need to be sampled at 16 kHz, since the upper limit of the signals they process is much less than 8 kHz. For example, the 500 Hz - 1 kHz subarray signals could be downsampled by a factor of eight (after suitable anti-alias filtering), likewise shortening the lengths of the associated FIR filters. Treating all the subarrays in this manner reduces the total number of FIR coefficients by a factor of almost four. The tradeoffs are the addition of a relatively short IIR filter before each elemental filter, and the addition of a single interpolation (upsampling) filter at each subarray output.

7.3 Adaptive frequency-independent beamforming

Real-world situations are often dynamic. Signal sources may move (e.g. if they are people), while receivers may also be mobile, as in cellular communications. Such environments may require adaptive beamformers.

A fully adaptive filter-and-sum beamformer could have a large number of adaptive parameters, since each coefficient of each elemental filter is in principle adjustable. Full adaptivity is computationally expensive, and probably unnecessary. We describe a much more economical *partially adaptive* system:

Consider a *basis* set of beams having the property that the mainlobe of each beam coincides with a null of each of the other beams (so-called Rayleigh-spaced beams). Our algorithm can then be carried out for each of the beamshapes in this basis set, yielding a set of pre-formed frequency-independent beams. An arbitrary beamshape can be formed by an appropriate linear combination of the pre-formed beams. Thus, broadband adaptive behavior is achieved with no more adaptive parameters than in a typical narrowband adaptive beamformer.

Appendix A

Review of basic beamforming terminology and concepts

A beamformer is a transducer array that transmits or receives energy in a directionally selective manner. The following discussion introduces basic beamforming jargon in the context of a simple beamforming receiver. The concepts apply equally well to beamforming transmitters.

A.1 Array pattern

We consider the following simple beamformer: 11 identical, uniformly spaced, collinear microphones whose outputs are directly summed to form the beamformer output. Assume further that this beamformer exists in a narrowband world in which all propagating waves have a wavelength of twice the inter-element spacing.

A convenient way of characterizing this beamformer is by considering its response to sinusoidal plane waves. Figure A-1 plots the response versus plane-wave propagation direction. This figure can be called the array pattern, beampattern, or far-field response. This figure assumes that the elements lie on the x-axis, which is shown by the line piercing the center of the structure.

This array is said to be steered to *broadside*, since it is most sensitive to plane waves propagating in directions perpendicular to the array axis. Directions which are

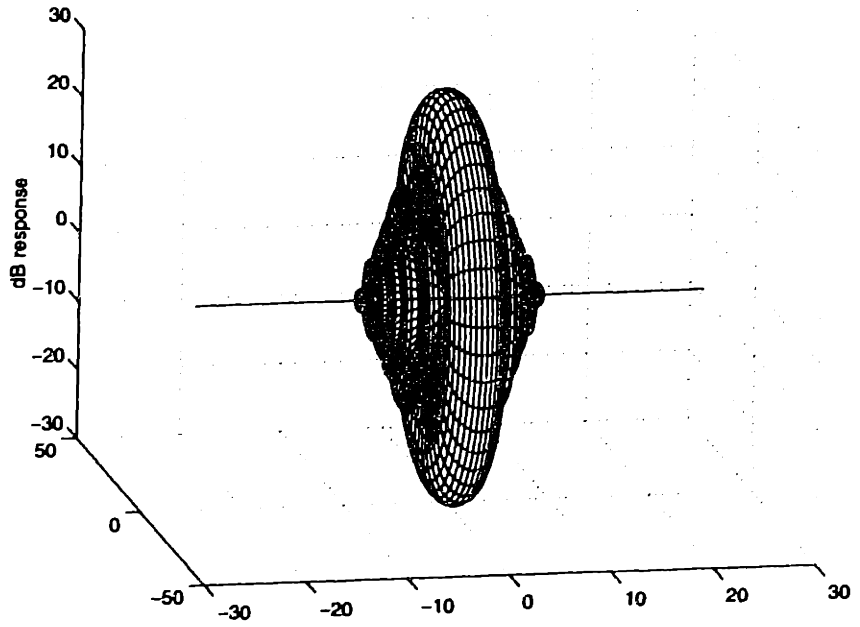

Figure A-1: Array pattern in three dimensions.

parallel to the array axis are also important, and are referred to as *endfire.*

Several features of the array pattern are noteworthy. The central bulge which is bisected by the y-z plane is called the *mainlobe*, while the smaller conical ridges on either side of it are called *sidelobes.* The regions of zero response between adjacent lobes are called *nulls.*

Since this pattern, and indeed any pattern corresponding to an array of collinear omnidirectional elements, is circularly symmetric about the array axis, the array response can be viewed in cross section without any loss of information. This is illustrated in Figure A-2 which is the intersection of the previous figure and the x-y plane. All the features identified before still exist in the polar response.

A.2 Shading

By varying the weights of the elements before summing them, one can produce different array patterns. This process can be termed *weighting*, *shading*, or sometimes

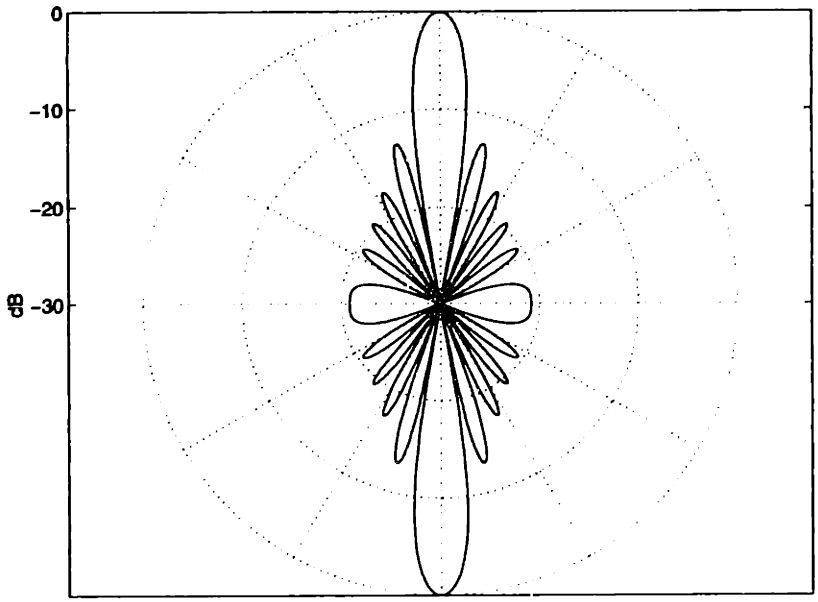

Figure A-2: Cross section of array pattern containing the array axis.

tapering. Appropriately chosen weights can lower the sidelobe level, usually at the cost of broadening the mainlobe. A well-known technique for achieving optimally low (in a minimax sense) sidelobes is described by Dolph [3]. Figure A-3 shows an application of Dolph's technique to produce a set of –25 dB sidelobes.

A.3 Steering

Steering is the adjustment of time or phase delays to electronically aim array lobes. Figure A-4 shows the array pattern for our example array when the mainlobe is steered to endfire. Note that the mainlobe has a larger angular extent at endfire than it did at broadside. This is typical behavior.

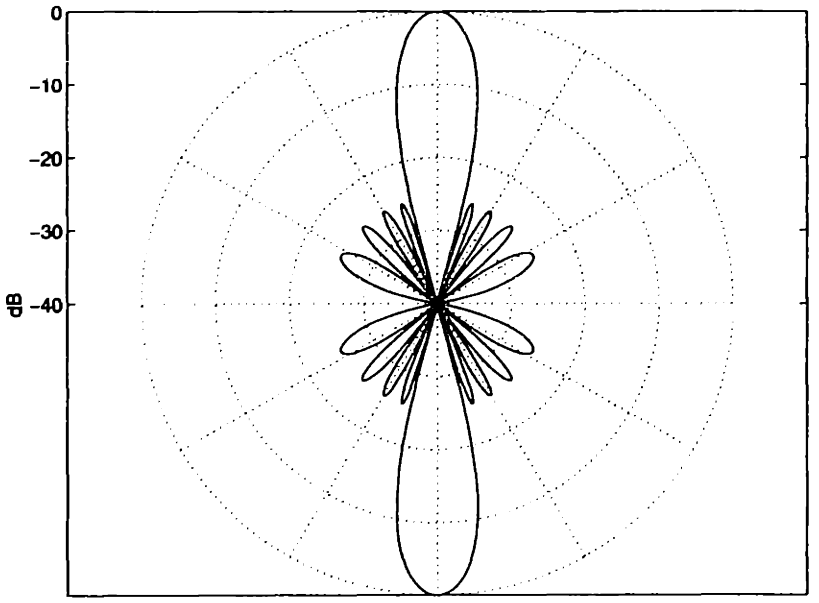

Figure A-3: Example array pattern due to dolph-chebyshev weighting.

A.4 Aliasing and spatial sampling

In our initial example, the interelement spacing was carefully chosen to be one-half the signal wavelength. If one attempts to economize using a larger interelement spacing, a phenomenon called *aliasing* may occur.

For example, suppose the interelement spacing were exactly one wavelength, and the wavefield consisted of a single plane-wave having an unknown direction of propagation. To this array, waves arriving at broadside are indistinguishable from those arriving at endfire, since the element outputs are precisely in phase in both cases. The term aliasing refers to this ambiguity with respect to propagation direction.

It is well known that an array which meets *spatial sampling* requirements (i.e. interelement spacing less than $\lambda/2$) does not suffer from aliasing at any steering angle. The spatial sampling requirement can be relaxed in some cases, i.e. when arrival directions are known a priori.

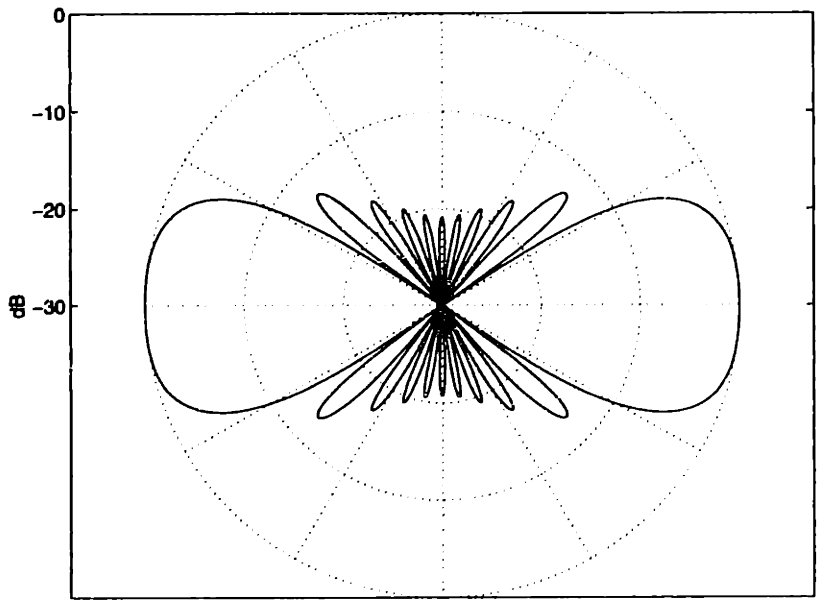

Figure A-4: Example array steered to endfire (180 degrees). Array pattern illustrates grating lobe at opposite direction.

A.5 Far-field vs. near-field

So far, the discussion has only considered the array's response to plane waves. In real environments, acoustic waves often emanate from a finite distance from the array, and therefore appear to be spherical waves. Still, the plane-wave response is useful when the distance from the array to the source is much greater than the span of the array. In these cases, the spherical wave's curvature is low enough that the wave appears planar in the vicinity of the array. Such sources are said to lie in the *far-field* of the array. Conversely, the *near-field* consists of the locations for which the plane-wave approximation is inaccurate.

A.6 Pattern multiplication

Another basic idea concerns the nature of the individual transducers in the arrays. The previous discussion has assumed omnidirectional transducers. However, directional transducers are commonly used to augment the directional behavior of the overall array. For a given direction, the response of an array of identical directional transducers is the product of the transducer response at that direction with the response of the array if it were made using omnidirectional elements.

For example, the hardware array used in this thesis contains microphones whose elements have a dipole response. The array has a spatial response similar to that shown in Figure A-5. This has the advantage of reducing the sensitivity to sounds arriving from endfire directions (as well as from the floor or ceiling). If these sounds are a priori unlikely to be the desired signal, some improvement in signal to noise ratio may be achieved.

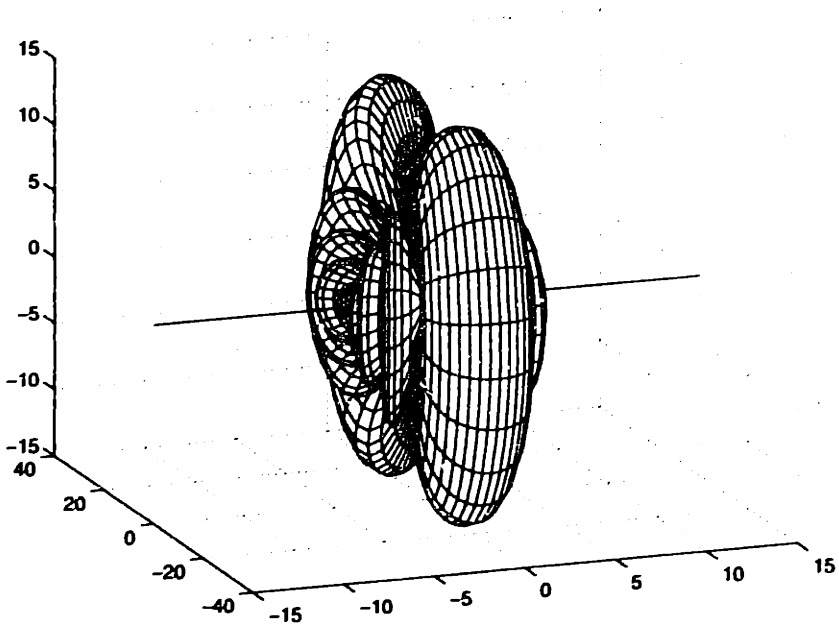

Figure A-5: Array pattern of array with dipole elements.

Appendix B

Prior solutions

Prior attempts to reduce the frequency-dependent behavior of transducer arrays generally aimed at the simpler problem of achieving a frequency-independent mainlobe width. These techniques are called constant-beamwidth solutions. Other techniques rely on transducers with novel geometric properties. A representative sample of such techniques is given in this appendix.

B.1 SHA technique

An especially simple approach was independently proposed by Smith [23] and Hixson and Au [8], hence the acronym. This technique has relatively low hardware requirements, and is applicable to a wide variety of array geometries.

B.1.1 Method

The underlying observation is that two beams of different widths can be averaged to form a beam of intermediate width. A beamformer which operates over a desired frequency range is thus constructed by building two beamformers: one has the desired beamshape at the lowest frequency, and the other has the desired beamshape at the highest frequency of interest. At intermediate frequencies, some linear combination of the two subarrays can often produce a satisfactory beamshape. Thus, a pair of

complementary filters is required, which controls the relative contributions of each subarray.

B.1.2 Example

An example is considered, and analyzed. In this example, the array transducers are as shown in Figure B-1. The complementary filters are generated by satisfying two constraints: 1) the mainlobe peak response must be frequency invariant, and 2) the half-power beamwidth must be frequency invariant. The example array is designed for operation over 1 - 2 kHz.

B.1.3 Performance.

The example SHA beamformer behaves well at directions between the half-power directions. However, performance degrades as sources move farther from the mainlobe. Figure B-2 shows variations of over 20 dB in the frequency response to sources located 17 degrees from broadside. Even further away from the mainlobe are the nulls, which are highly frequency-dependent.

B.1.4 Anomalous sidelobe

This SHA example also exemplifies an artifact termed the *anomalous sidelobe* [26]. This extra lobe is caused by the fact that the longer of the two constituent subarrays was designed for 1 kHz operation, and does not meet spatial sampling requirements at other frequencies in the 1 - 2 kHz octave. Thus, at some steering angles, an extra lobe appears which is essentially a grating lobe of the longer subarray, attenuated by the filter associated with that subarray. This seemingly clear observation was first noted in the literature by Webster [26] 20 years after the original papers by Smith, Hixson and Au.

This anomalous lobe can be removed by decreasing the inter-element spacing in the longer subarray. Interestingly, this subarray does not need to meet the half-wavelength spatial sampling requirement throughout the *entire* octave. This is due to

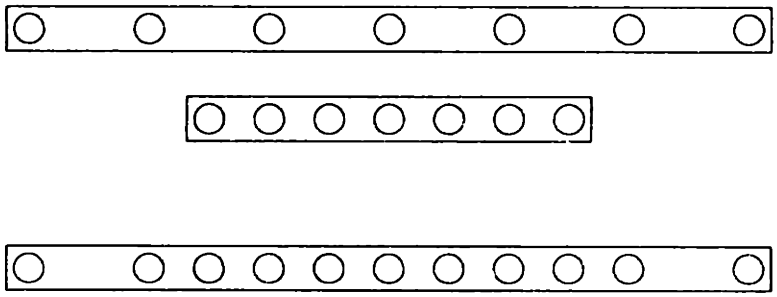

Figure B-1: Subarrays used to form the SHA array.

the fact that the longer subarray's filter tends to have lowpass characteristics, thereby attenuating its output at exactly the frequencies at which the aliasing problem should be most acute. Thus, some hardware economization is possible since one can violate the half-wavelength requirement at some higher frequencies in the octave without adverse effect.

B.2 SHA with hexagonal array

An interesting configuration reported by Lardies [11] consists of two concentric rings each having six elements spaced uniformly along the circumference. The outer ring has exactly twice the radius of the inner ring.

When properly weighted, this array has constant beamwidth, as in the linear arrays described above, *and* remarkably constant null placement. This configuration appears to deserve further examination, to determine, for example, if the frequency-independent null property is retained for rings with more then six elements.

B.3 Multiple beamforming

An important advantage of discrete transducer arrays is their ability to form multiple simultaneous beams. This capability leads to another scheme for constant-beamwidth beamforming.

Specifically, a number of beams is formed such that at some lower frequency

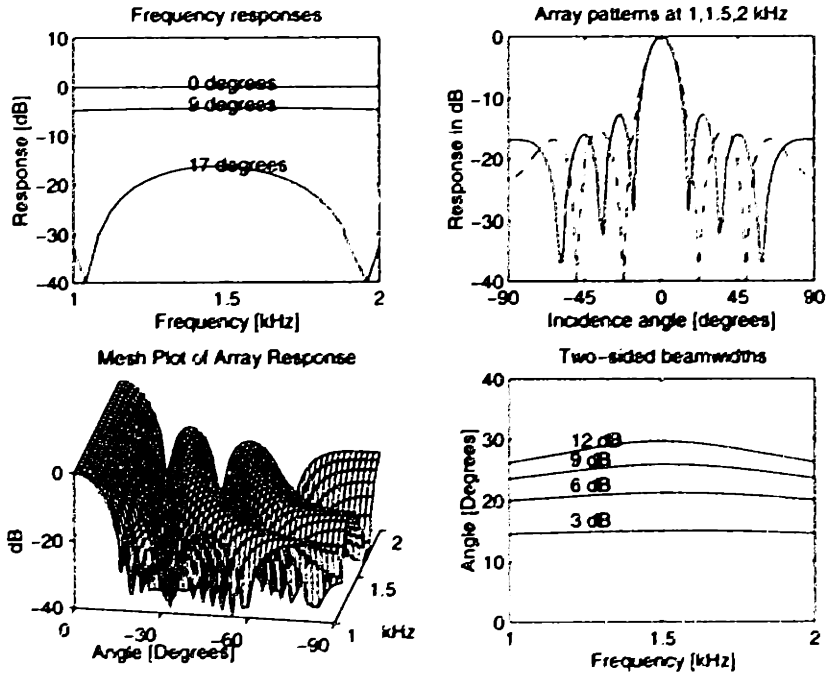

Figure B-2: Performance of SHA array

limit the beams coincide, and that as frequency increases the beams are gradually steered outward in such a way that the shoulder of the outermost pair of beams stays in the same spot. Such a scheme was proposed in 1957 by Tucker [24]. An analog implementation was described in 1961 by Morris [16], and a digital solution by Goodwin [5] arrived much later. When analyzed, this is shown to require a filter having affine phase, i.e. phase is a linear function of frequency, with a possible constant offset [5, 24].

This approach is attractive partly because of an elegant analog delay line implementation due to Smith [23]. However, when implemented digitally, this scheme requires the construction of a set of arbitrary elemental filters, and thus has no obvious computational advantage over the filter-and-sum beamformer.

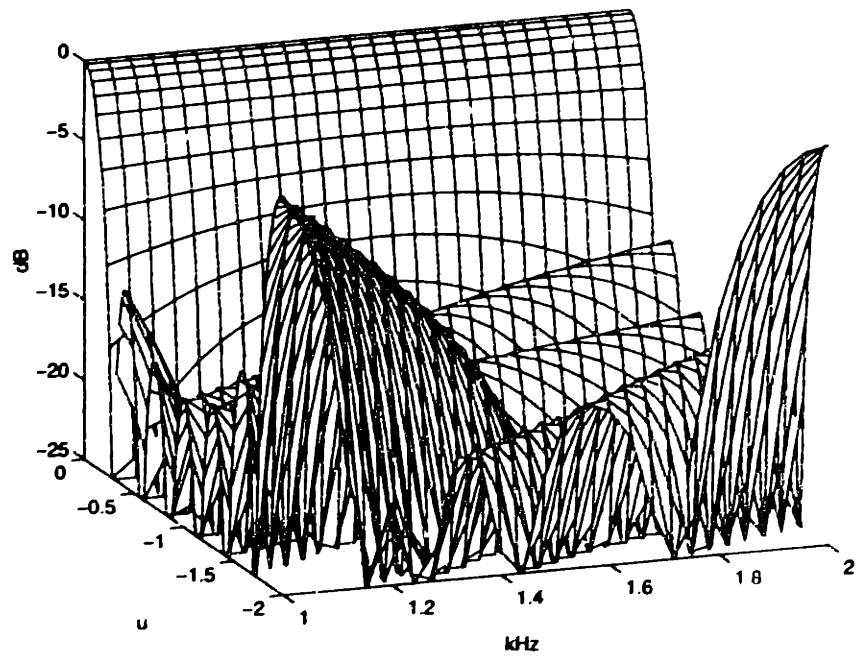

Figure B-3: SHA beamformer response outside of visible region shows anomalous sidelobe which peaks around $|u| = 2$, $f = 1$ kHz.

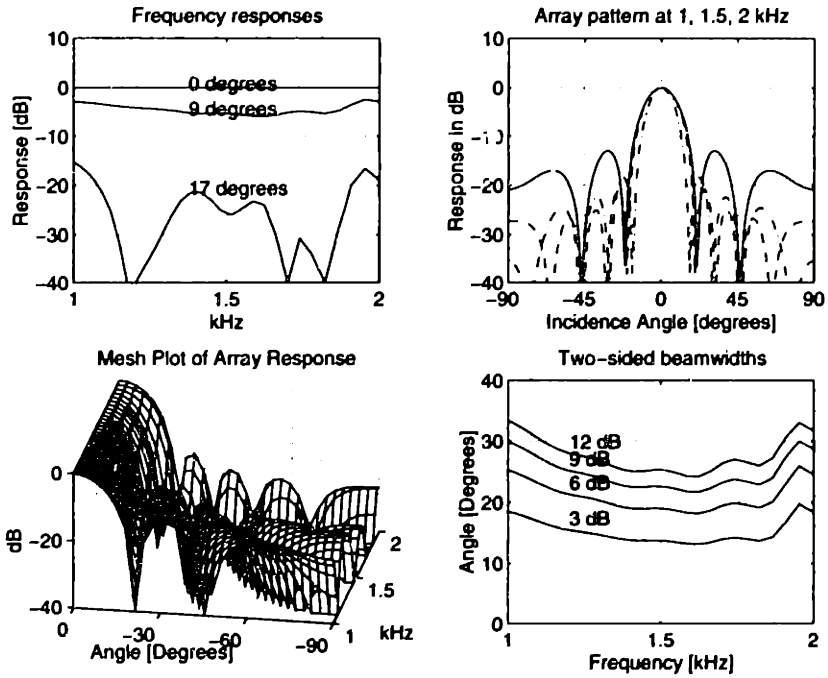

Figure B-4: Constant width beams approximated by overlapping several narrower beams

B.3.1 Performance

A five-beam, one-octave system is designed and simulated. The results are displayed in Figure B-4.

This approach actually achieves some constancy in the location of the null pair immediately flanking the mainlobe. The frequency responses are acceptably flat near the mainlobe, but exhibit wide variations away from the mainlobe. However, other nulls are not so well behaved.

B.4 Curved transducer surfaces

Several studies have involved transducers with special surface geometries that naturally lead to frequency-independent beamwidths above a certain threshold frequency which is determined by the transducer dimensions. These techniques are most useful in ultrasonic applications where wavelengths are conveniently small. Transducers

having spherical surfaces are described by a number of researchers [9, 15, 22]. Morris [15] also describes a transducer having discrete elements placed on a hyperbolic-paraboloid. These configurations do achieve fairly constant mainlobe widths, but do not control null location.

Appendix C

FIR filter design using non-uniform frequency sampling

This appendix describes a filter design technique which is well suited for use in the 2-D filter design algorithm of Chapter 4.

C.1 Review of uniform frequency sampling

Frequency sampling is a simple, noniterative method of FIR filter design which can be applied to design problems requiring arbitrary frequency responses. Briefly, filter design occurs by sampling the desired response at N equally spaced points in the frequency range $[-\frac{1}{T}, \frac{1}{T}]$, where T is the sampling period. The desired coefficients are then obtained by an inverse discrete Fourier transform, yielding a filter having exactly the desired response at the N chosen points (the *frequency samples*). Some filter design problems may leave the desired response unspecified in certain regions (e.g. between the passband and stopbands). Sample points which happen to fall in unconstrained regions may be adjusted via optimization techniques such as linear programming to reduce the error between the N frequency samples. A detailed discussion is given by Rabiner and Gold [21].

C.2 Non-uniform spacing

Linear programming is not the only method of trading off response accuracy in some frequency bands versus others. A simpler method is that of non-uniform frequency sampling, in which the frequency samples are concentrated more densely in the regions where accurate response is desired.

The approach is now mathematically developed for linear-phase filters. Consider a list of frequencies $\omega_0, \omega_1, \ldots \omega_N$, $0 \leq \omega_i \leq \pi$, and a corresponding list of real desired responses, one for each of the frequencies listed: $H_0, H_1, \ldots H_N$. Now, an FIR filter with taps $[a_n, a_{n-1}, \cdots, a_0, \cdots, a_{n-1}, a_n]$ has frequency response:

$$H(e^{j\omega}) = a_0 + 2\sum_{n=1}^{N} a_n \cos(n\omega), \tag{C.1}$$

where ω is a normalized frequency variable. Substituting all $N+1$ constraints into Equation C.1 yields a set of simultaneous equations which can be expressed in matrix form:

$$\begin{bmatrix} H_0 \\ H_1 \\ \vdots \\ H_N \end{bmatrix} = \begin{bmatrix} 1 & \cos\omega_0 & \cdots & \cos N\omega_0 \\ 1 & \cos\omega_1 & \cdots & \cos N\omega_1 \\ & & \vdots & \\ 1 & \cos\omega_N & \cdots & \cos N\omega_N \end{bmatrix} \begin{bmatrix} a_0 \\ 2a_1 \\ \vdots \\ 2a_N \end{bmatrix}. \tag{C.2}$$

All the elements in the vector on the left-hand-side and the matrix on the right-hand-side of the equation are known. Thus, the coefficients a_i can be computed by premultiplying each side of the equation by the inverse of the matrix. Note that when the frequency samples are equally spaced, this equation computes a Fourier transform.

C.3 Example

The following example illustrates the application of this method, as well as the practical aspects involved when selecting frequency samples. This example uses a filter which is representative of those encountered in the designed efforts of this thesis.

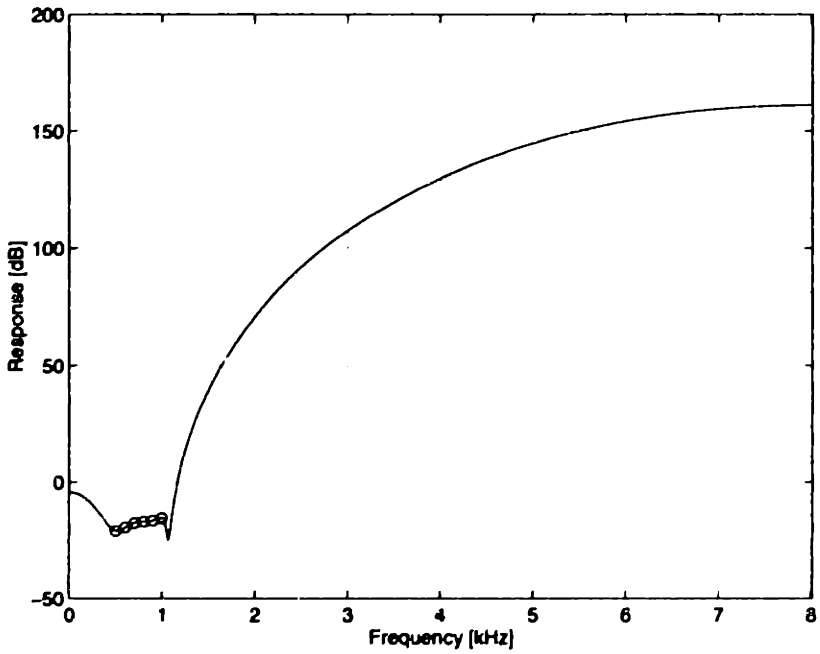

Figure C-1: Magnitude response of a naively designed FIR filter. The circles represent samples of the desired response of an actual filter encountered when computing the coefficients that approximate the desired response depicted earlier in this chapter.

Figure C-1 shows a naive design attempt, in which six frequency samples are selected within the octave of interest (0.5 - 1 kHz), and no out-of-band samples are selected. While the filter order of six is admirably low and the response is indeed correct at the six frequencies, the out-of-band response reaches almost 200 dB above the in-band response. While it is mathematically possible to later suppress this out-of-band response with a bandpass filter, considerable roundoff error would result in a practical system. For example, a typical DSP has 24-bit mantissas, corresponding to a decibel value of $20\log(2^{24}) = 144.5$.

A practical solution is to add extra constraints outside of the octave band, resulting in a filter of somewhat higher order, but with much lower out-of-band responses. If a certain increase in the out-of-band response, say 20 dB, is deemed tolerable, then the out-of-band constraints can be chosen to have a wider spacing than the in-band constraints. This is shown in Figure C-2, in which 37 additional out-of-band

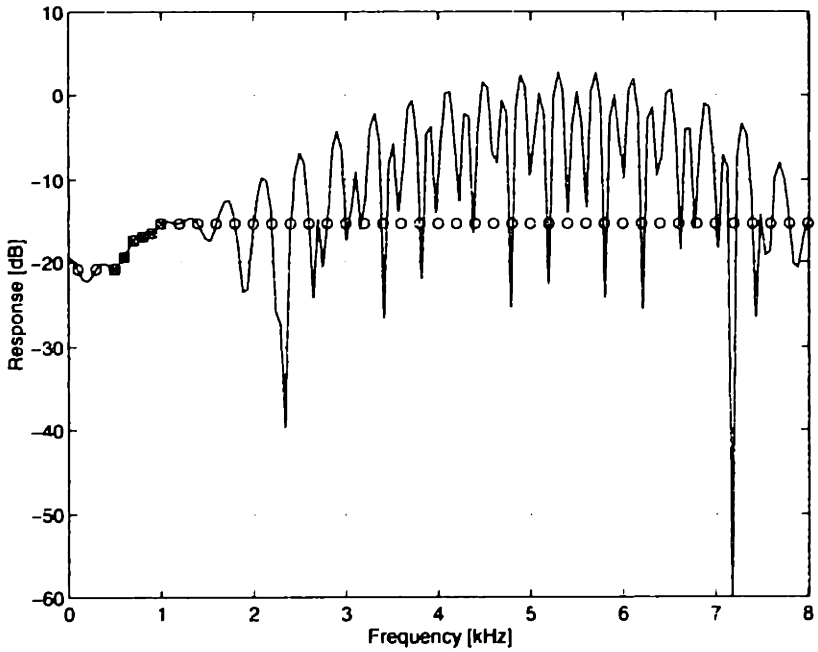

Figure C-2: Magnitude response of FIR filter designed to the same specifications as in the previous figure, but with additional constraints outside of the region of interest.

constraints are chosen to smoothly continue the responses at the edges of the octave region.

The exact number of additional out-of-band constraints depends on the desired performance. If the out-of-band constraints had been chosen closer together, for example with the same spacing as the in-band constraints, then the out-of-band responses would be very well controlled. However, this would have required 75 extra constraints, instead of 37.

Appendix D

Broadband dolph chebyshev beamformer

In Chapter 4, a sampling grid was defined, containing a set of horizontal lines which defined a set of narrowband beamforming problems into which the broadband problem was decomposed. Reexamination of the narrowband solutions leads to new variations on our broadband beamforming design.

This appendix discusses a design variation in which the narrowband problems are solved using the dolph-chebyshev approach, yielding sidelobes which are in some sense optimally low for a given beamwidth.

D.1 Dolph-Chebyshev design

The Dolph-Chebyshev algorithm presented a moderate challenge in itself. The basic problem was to design an array with a given number of elements and a specific 3 dB beamwidth. The mathematical software package MATLAB contains a function which designs Dolph-Chebyshev arrays given the desired ripple level, not the desired beamwidth. Since this did not suit our objectives, the problem had to be solved from more basic principles. These principles were:

1. Uniform linear array design problems can be cast as *equivalent FIR filter* design problems, as was shown in Section 2.5.

2. Such an FIR filter design problem can be recast as an *equivalent polynomial approximation* problem, since any FIR filter magnitude response can be expressed exactly as a polynomial in $x = \cos\omega$, ω being the frequency in radians per sample period. This equivalence is a key step in the Parks-McClellan FIR filter design algorithm [19, 20], and is discussed in many textbook treatments such as that of Oppenheim and Schafer [17, pp. 466-7].

3. The Chebyshev polynomials are equiripple in the interval $[-1\ 1]$. A linear transformation in the polynomial variable can stretch or compress the equiripple region to occupy any desired interval.

The design constraints described at the beginning of this appendix can thus be transformed into a set of constraints on an equivalent polynomial approximation problem. The transformations are described as follows:

1. Constraining the half power points to occur at $\theta = \theta_0$ implies that the equivalent FIR lowpass filter has its 3 dB cutoff at $\omega = 2\pi kd\sin\theta_0$, where k is the planewave's spatial frequency (radians/meter), and ω is the frequency parameter (radians/sample) of the FIR filter. Thus, the equivalent polynomial passes through the point $(x_0, \sqrt{2}/2)$, where $x_0 = \cos\omega_0$.

2. The equiripple-sidelobe requirement implies that the corresponding FIR filter is lowpass with equiripple sidelobes. The equivalent polynomial then must be equiripple in a region whose left edge is at $\cos\pi/2 = -1$.

One more constraint is needed, and is provided by an arbitrary scale factor which makes the broadside gain equal to unity. This implies that the equivalent polynomial passes through the point $(\cos 0,1)$. Now, if $T_n(x)$ is the nth degree Chebyshev polynomial, then the required polynomial is of the form:

$$P(x) = a\ T_n(b(x+1) - 1). \tag{D.1}$$

Since $P(1) = 1$, we can find a in terms of b:

$$a = \frac{1}{T_n(2b-1)}. \tag{D.2}$$

Combining the previous two equations with the half power constraint, $P(x_0) = \sqrt{2}$ yields the following equation in the unknown b:

$$\sqrt{2} = \frac{T_n(b(x_0+1)-1)}{T_n(2b-1)}, \tag{D.3}$$

where $x_0 = \cos\omega$ as defined above. This is solved by expressing the previous equation as a polynomial in b, and solving for the largest root, since smaller roots will satisfy the 3 dB constraint using one of the "ripples" of the Chebyshev polynomial, rather than the "tail" portion outside of the equi-amplitude ripples.

Knowing b determines a, and in turn $P(x)$. Finally, the filter coefficients can be determined from the polynomial coefficients using the equivalence already described.

D.2 Performance

As in the main body of this thesis, the FIR filter design process was evaluated both by computer simulation and by measurements performed on a real-time implementation. The simulation results are shown in Figure D-1. The significant features of the simulation are the extremely low sidelobes towards the upper end of the octave, and the small change in mainlobe shape with frequency. The null performance is relatively unimportant at the higher frequencies, since the sidelobe levels are low enough that the entire region outside the mainlobe is effectively nulled.

The measured performance is shown in Figure D-2 for one of the four octaves. (The other three are similar, except for a lower signal to noise ratio at the higher frequencies.) The patterns above 1400 Hz do show extraneous sidelobes; however, most of these lobes are over 30 dB below the response at the center of the mainlobe.

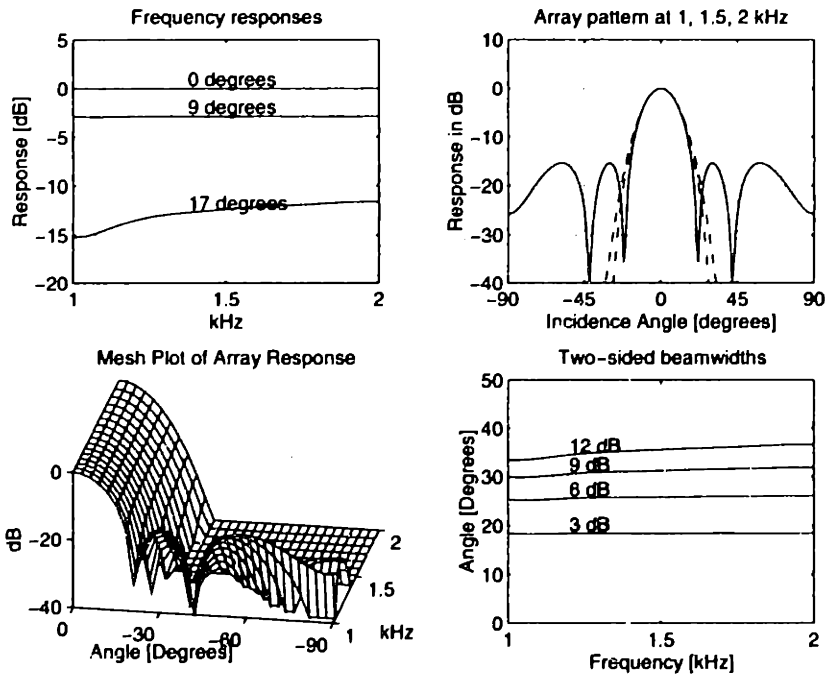

Figure D-1: Simulated performance of minimax design.

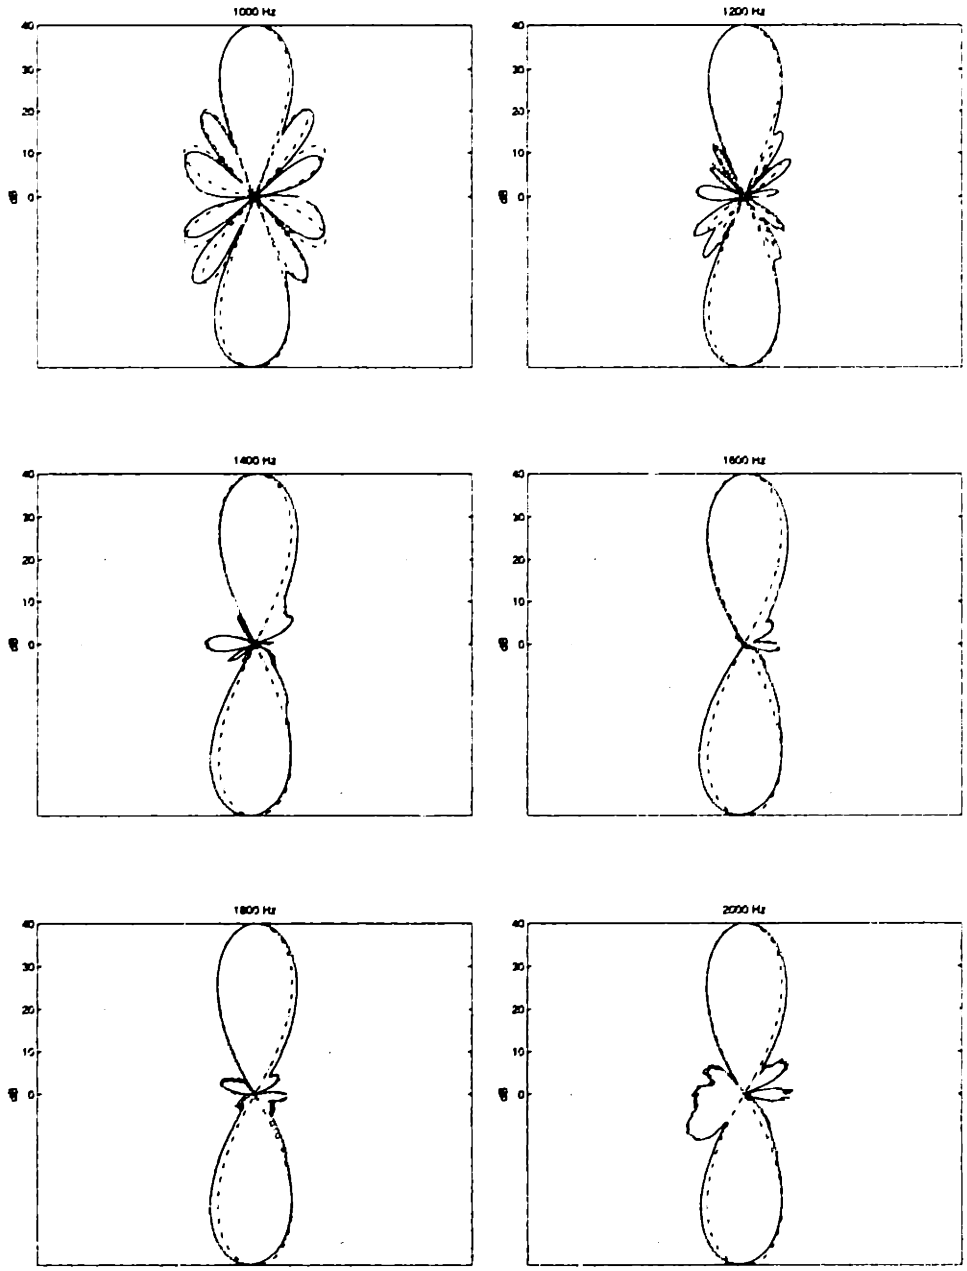

Figure D-2: Array patterns of Chebyshev constant-beamwidth beamformer. Measured patterns are solid lines, simulated patterns are dotted lines.

Bibliography

[1] Christakis Charlambous. The performance of an algorithm for minimax design of two-dimensional linear phase fir digital filters. *IEEE Transactions on Circuits and Systems*, 32(10):1016–1028, October 1985.

[2] Ariel Corporation. *Preliminary User's Manual for the MP3210.* Ariel Corporation, second edition, 1992.

[3] C. L. Dolph. A current distribution for broadside arrays which optimizes the relationship between beamwidth and sidelobe level. *Proceedings IRE*, 34:335–348, June 1946.

[4] Dan E. Dudgeon and Russell M. Mersereau. *Multidimensional Digital Signal Processing.* Prentice-Hall, 1984.

[5] Michael M. Goodwin and Gary W. Elko. Constant beamwidth beamforming. *International Conference on Acoustics, Speech, and Signal Processing*, pages 169–172, 1993.

[6] D. B. Harris and R. M. Mersereau. A comparison of algorithms for minimax design of two-dimensional linear-phase fir digital filters. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 25:492–500, 1977.

[7] Simon S. Haykin. *Adaptive Filter Theory.* Prentice Hall, 1991.

[8] E. L. Hixson and K. T. Au. Wide-bandwidth constant beamwidth acoustic array. *Journal of the Acoustical Society of America*, 48(1):117, July 1970.

[9] J. Jarzynski and W. James Trott. Array shading for a broadband constant directivity transducer. *Journal of the Acoustical Society of America*, 64(5):1266–9, November 1978.

[10] Don H. Johnson and Dan E. Dudgeon. *Array Signal Processing, Concepts and Technique.* Signal Processing. Prentice Hall, 1993.

[11] J. Lardies. An optimal broadband constant beamwidth concentric ring array. *Acoustics Letters*, 11(1):5–8, 1987.

[12] J. Lardies and J. P. Guilhot. Realization of a broadband constant beamwidth end-fire line array. *Acoustics Letters*, 11(1):5–8, 1987.

[13] John H. Lodge and Moustafa M. Fahmy. An efficient l_p optimization technique for the design of two-dimensional linear-phase fir digital filters. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, ASSP-28:308–312, June 1980.

[14] Robert J. Lustberg. Acoustic beamforming using microphone arrays. *MIT Master's Thesis*, 1993.

[15] J. C. Morris. Broad-band constant beam-width transducers. *Journal of Sound and Vibration*, 1:28–40, 1964.

[16] J. C. Morris and E. Hands. Constant-beamwidth arrays for wide frequency bands. *Acustica*, 11:341–7, 1961.

[17] Alan V. Oppenheim and Robert W. Schafer. *Discrete-Time Signal Processing.* Signal Processing. Prentice-Hall, 1989.

[18] Ltd. Panasonic/Matsushita Communication Industrial Co. Noise-cancelling back electret condenser microphone cartridge - wm-55d103.

[19] T. W. Parks and J. H. McClellan. Chebyshev approximation for nonrecursive digital filters with linear phase. *IEEE Trans. Circuit Theory*, CT-19:189–194, March 1972.

[20] T. W. Parks and J. H. McClellan. A program for the design of linear phase finite impulse response filters. *IEEE Trans. Audio Electroacoustics*, AU-20(3):195–199, August 1972.

[21] L. R. Rabiner and Bernard Gold. *Theory and Application of Digital Signal Processing.* Prentice-Hall, Inc., 1975.

[22] Peter H. Rogers and A. L. Van Buren. New approach to a constant beamwidth transducer. *Journal of the Acoustical Society of America*, 64(1):38–43, July 1978.

[23] R. P. Smith. Constant beamwidth receiving arrays for broad band sonar systems. *Acustica*, 23:21–26, 1970.

[24] D. G. Tucker. Arrays with constant beam-width over a wide frequency-range. *Nature*, 180:496–7, September 1957.

[25] Darren B. Ward. Design of frequency-invariant broadband far-field sensor arrays. *Unpublished article*, 1994.

[26] R. J. Webster and T. N. Lang. Prescribed sidelobes for the constant beamwidth array. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(4):727–30, April 1990.