CENTER FOR COMPUTER RESEARCH IN MUSIC AND ACOUSTICS
DECEMBER 1995

DEPARTMENT OF MUSIC
REPORT NO. STAN-M-96

SOME ASPECTS OF APPLYING PSYCHOACOUSTIC PRINCIPL TO SOUNDFIELD RECONSTRUCTION

STEVEN D. TRAUTMANN

CCRMA
DEPARTMENT OF MUSIC
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305-8180

SOME ASPECTS OF APPLYING PSYCHOACOUSTIC PRINCIPLES TO SOUNDFIELD RECONSTRUCTION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MUSIC
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Steven D. Trautmann
December 1995

where the real source location and time lag are denoted with the 1 subscript and the virtual source information with the 2 subscript. If for convenience the real and virtual sources are placed at $(b, 0)$ and $(-b, 0)$ respectively, then

$$\sqrt{(x-b)^2+y^2} - \sqrt{(x+b)^2+y^2} = c\,(t_{+b} - t_{-b}) \tag{B.11}$$

squaring and collecting terms gives

$$bx + \left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)^2 = -\left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)\sqrt{(x+b)^2+y^2} \tag{B.12}$$

squaring and collecting terms again gives

$$x^2\left(b^2 - \left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)^2\right) - \left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)^2 y^2 = \tag{B.13}$$

$$\left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)^2\left(b^2 - \left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)^2\right) \tag{B.14}$$

which can be put into the form

$$\frac{x^2}{\left(\frac{c(t_{+b}-t_{-b})}{2}\right)^2} - \frac{y^2}{b^2 - \left(\frac{c(t_{+b}-t_{-b})}{2}\right)^2} = 1 \tag{B.15}$$

This is the classic form of the hyperbola. It should be noted that only one side of the hyperbola is correct, the mirror solution can be thought of as what would happen if time was run backwards and the spherical wave extrapolated though the source prior to its actual existence. Also the term

$$b^2 - \left(\frac{c\,(t_{+b} - t_{-b})}{2}\right)^2 \tag{B.16}$$

is of interest. It shows that if the distance between the two sources is shorter than the distance the sound covers in the time delay, then no (real) solution exists.[1] Since the speed of sound is constant, a large enough time lag to allow one source's impulse to get past another source before that source can create an impulse, then the second impulse will never catch up to the first to intersect it.

Now consider the set of spheres around each monopole source in a steady state as being surfaces of equal amplitude, but the initial amplitude can be adjusted by a constant. The points where the amplitudes from each source are equal are the points where $A_2 d_1 = A_1 d_2$ with A_1, A_2 constants applied to the initial amplitude and d_1, d_2 the distances to sources 1 and 2 respectively. If the real and virtual sources are placed at $(b, 0)$ and $(-b, 0)$ respectively, then

$$A_2\sqrt{(x-b)^2+y^2} = A_1\sqrt{(x+b)^2+y^2} \tag{B.17}$$

squaring both sides

$$A_2^2\left(\left(x^2 - 2bx + b^2\right) + y^2\right) = A_1^2\left(\left(x^2 + 2bx + b^2\right) + y^2\right) \tag{B.18}$$

[1] As a side note, in the case where the original source, the loudspeaker and one point of time equality not between the two are co-linear, then all the colinear points beyond either the original source or the loudspeaker will also have time equality, (the side is the same as where the one point of time equality lies) and the error will be of amplitude only.

collecting terms gives the general form of a circle

$$(A_1^2 - A_2^2)x^2 + (A_1^2 - A_2^2)y^2 - (A_1^2 + A_2^2)2cx + (A_1^2 - A_2^2)c^2 = 0 \tag{B.19}$$

so if $A_1 \neq A_2$ then dividing by $(A_1^2 - A_2^2)$ and completing the square puts the circle in standard form

$$\left(x - \frac{c(A_1^2 + A_2^2)}{(A_1^2 - A_2^2)}\right)^2 + y^2 = c^2\left(\frac{(A_1^2 + A_2^2)^2}{(A_1^2 - A_2^2)^2} - 1\right) \tag{B.20}$$

In three dimensions, the surface of revolution defines a sphere. The points where time and amplitude are both synchronized is the intersection of these surfaces of revolution, the hyperboloid and sphere. This intersection can be either no points, a single point, one or two circles around the line defined by the two sources, where the hyperbola and sphere intersect. In one case the hyperbola and sphere can coincide, since the limit of each is a plane halfway between the two sources, which happens if both sources have identical amplitudes and timing. By requiring $\frac{A_1}{dist(p,p_1)} = \frac{A_2}{dist(p,p_2)}$ and $dist\,(p,p_1) = c\,(t - t_1)\,, dist\,(p,p_2) = c\,(t - t_2)$ the locus of points satisfying the constraints is found to be those satisfying $A_1 dist\,(p,p_1) - A_2 dist\,(p,p_2) = 0$ and $dist\,(p,p_1) - dist\,(p,p_2) = c\,(t_1 - t_2)$. This implies that $A_1 = A_2 = 1, t_1 - t_2 = 0$ and $dist\,(p,p_1) = dist\,(p,p_2)$ which defines a plane equal distant between the two sources with identical time lags and identical amplitude adjustments. Adjusting the time difference and the amplitude of the real source, the signal at any given point produced by the virtual source can be exactly mimicked by the real source.

By way of returning to comparing plane waves and spherical waves, here is a simple illustration. Consider a single source at the origin viewed along one dimension. Assume that there is a point of concern some distance from a loudspeaker, and these two points determine the x-axis. Since N sources can produce exactly the correct signal at N points, in this case the correct signal can be reproduced at one point x_0. For simplicity the virtual source to be mimicked at this point is also on the x-axis at $-\infty$, (which eliminates the $\frac{1}{d}$ dependence).[2] If x_0 is chosen to be +1.5 units distance from the loudspeaker, and the signal from the virtual source has unit amplitude, then all that needs to be done is to synchronize signals and set the gain on the loudspeaker to 1.5 to compensate for the loss due to distance.

Thanks to the sampling theorem, the number of samples in an interval determines the highest cutoff frequency beyond which aliasing will occur. Even though the one source can be used to get fluctuations for only one sample point exactly right for all time, this is enough to determine a signal everywhere, if and only if that signal happens to be bandlimited to D.C. Since the virtual source at $-\infty$ has no $1/d$ component, it truly is bandlimited the same way any of its input signals are bandlimited. All that is needed is conversion from time to spatial location. Giving the virtual source a D.C. signal, having a source at the origin mimic at the point x_0 and taking the difference on [1,2] gives the plot B.6. Although the amplitude at x_0 is right, the error becomes progressively larger the farther away from x_0 one goes.

B.0.2 Comparison of Plane and Spherical Fields

Since there will always be some error (except in the trivial case where the virtual source and the real source happen to coincide and have identical properties), understanding the

[2] Although impossible in reality since $\frac{1}{d}$ goes to 0, such a non-dependence on distance can be achieved through the use of acoustic lenses, similar to the use of lenses in optics

nature of the error becomes important so that questions about its perceptual effects can be answered. In order to put a bound on the error in any given situation, a comparison must be made between the desired sound field and the reconstructed field. If the desired field consists of a plane wave and there is only one ideal source for reproduction, the error can be put in terms of distance, which determines the nature of both the amplitude error and the time (or phase) error. Consider a 2-dimensional plane wave with amplitude A at some frequency. At some distance R from the 2-dimensional point source, the reconstruction signal has amplitude A at the same frequency. Assume that the edge of this circle represents time t_0 and the corresponding time in the plane wave intersects this sphere. (see Figure B.7) Then at least at this intersection both time and phase are "in sync." Everywhere else on this intersecting plane, there is an error which is a function of the distance of that point from the source minus R. This in turn is a function of the angle ϕ determined by the line connecting the point and source and the shortest line connecting the source to the plane (with length k). The difference in distance is

$$\frac{k}{\cos\theta} - R \tag{B.21}$$

and the ratio of distance between the correct points and other points along the plane is

$$\frac{R\cos\phi}{k} \tag{B.22}$$

Thus, the amplitude error at any point along the intersecting line is

$$error_{amp} = A - A \cdot \frac{R\cos\phi}{k} \tag{B.23}$$

and the time difference along the line is

$$error_t = \frac{1}{c}(\frac{k}{\cos\phi} - R) \tag{B.24}$$

where c = the speed of propagation.

Thus, between any two sample points where there is no error, the maximum amplitude error is

$$A - A \cdot \frac{R}{k} \tag{B.25}$$

and the time error is

$$\frac{1}{c}(k - R) \tag{B.26}$$

which can be related to phase error at the midpoint as

$$\frac{2\pi(k - R)}{w} \tag{B.27}$$

where w is the wavelength. Notice that the maximum phase error may not occur at the midpoint, but until 180 degrees are reached, the midpoint will have the maximum phase error. After that point, the maximum phase error between the two points will still be at 180 degrees, but the points (and there can be any number) where this maximum phase error occurs will not necessarily be at the midpoint.

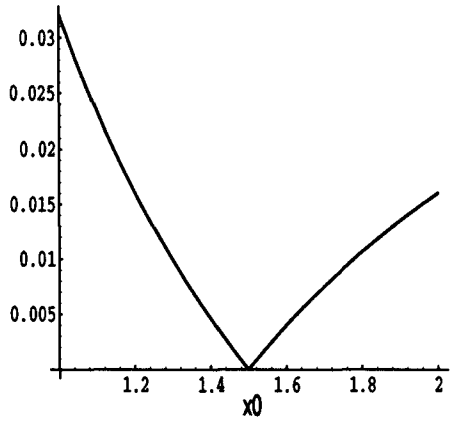

Figure B.6: Difference between the amplitude of a sound from $-\infty$ and a sound from the origin made to have equal amplitude at x0.

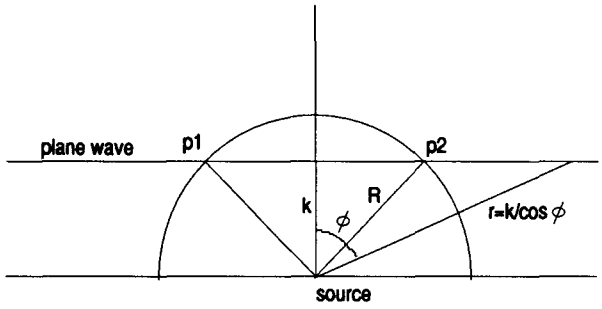

Figure B.7: Diagram of line intersecting a circle.

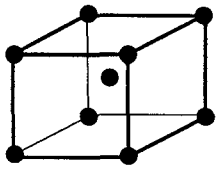

Figure B.8: Diagram of a point within region bounded by sample locations.

B.0.3 Bounding the Error

In order for the sampling theorem to be of any use, the error between spatial samples must be bounded. Consider the region inside eight samples creating a cube (see Figure B.8) There must be eight or more speakers in order to make sure the signal at these sample points is correct. The signal at these points due the speakers is

$$\begin{bmatrix} P_1(z) \\ \vdots \\ P_8(z) \end{bmatrix} = \begin{bmatrix} F_{1,1} & \cdots & F_{N,1} \\ \vdots & \cdots & \vdots \\ F_{1,8} & \cdots & F_{N,8} \end{bmatrix} \cdot \begin{bmatrix} SP_1(z) \\ \vdots \\ SP_N(z) \end{bmatrix} \tag{B.28}$$

Where $N \geq 8$. To find the maximum amplitude error at D.C. within this region, it is useful to find the points within the region which are closest and furthest from each source. These points may or may not be distinct, but there will always be at least two. Two new formulations can be made for the direct signal

$$\begin{bmatrix} maxP_1(z) \\ \vdots \\ maxP_N(z) \end{bmatrix} = \begin{bmatrix} \frac{1}{dist_{1,1}} & \cdots & \frac{1}{dist_{N,1}} \\ \vdots & \cdots & \vdots \\ \frac{1}{dist_{1,N}} & \cdots & \frac{1}{dist_{N,N}} \end{bmatrix} \cdot \begin{bmatrix} SP_1(z) \\ \vdots \\ SP_N(z) \end{bmatrix} \tag{B.29}$$

and

$$\begin{bmatrix} minP_1(z) \\ \vdots \\ minP_N(z) \end{bmatrix} = \begin{bmatrix} \frac{1}{dist_{1,1}} & \cdots & \frac{1}{dist_{N,1}} \\ \vdots & \cdots & \vdots \\ \frac{1}{dist_{1,N}} & \cdots & \frac{1}{dist_{N,N}} \end{bmatrix} \cdot \begin{bmatrix} SP_1(z) \\ \vdots \\ SP_N(z) \end{bmatrix} \tag{B.30}$$

Now, by summing the amplitude of these points of maximum amplitude for each speaker, an upper bound on the amplitude possible within the region will be found. This is the same as summing the maximum distance constant from each column of the transfer matrix.

$$P_{upperbound} = \sum_{i=1}^{N} \max\left(\frac{1}{dist_{i,1}}, \frac{1}{dist_{i,2}}, \cdots, \frac{1}{dist_{i,N}}\right) \cdot SP_i \tag{B.31}$$

This upperbound will only be achieved if the same point is closest to all the sources, otherwise all points within the region will be below the upperbound. Similarly, the sum of the smallest amplitudes can be defined from the second matrix

$$P_{lowerbound} = \sum_{i=1}^{N} \min\left(\frac{1}{dist_{i,1}}, \frac{1}{dist_{i,2}}, \cdots, \frac{1}{dist_{i,N}}\right) \cdot SP_i \tag{B.32}$$

Again, this lower bound will only be achieve if the same point in the region is farthest from all the speakers. Since the D.C. case was being considered, the signal at each of the sample points will be the same constant C. Thus, the maximum error with in the region is $\max(P_{upperbound} - C, C - P_{lowerbound})$. Since at D.C. all the vectors can be thought to have zero phase, the upperbound can be applied to all frequencies. The lowerbound at frequencies other than D.C. can be affected by cancelations caused by differing phases in the sum, however.

A bound can also be put on the phase error for a given frequency. Consider the same argument above applied to a given frequency. The same set of points nearest and furthest from each speaker is used

$$\begin{bmatrix} nearP_1(\omega) \\ \vdots \\ nearP_N(\omega) \end{bmatrix} = \begin{bmatrix} F_{1,1} & \cdots & F_{N,1} \\ \vdots & \cdots & \vdots \\ F_{1,N} & \cdots & F_{N,N} \end{bmatrix} \cdot \begin{bmatrix} SP_1(\omega) \\ \vdots \\ SP_N(\omega) \end{bmatrix} \tag{B.33}$$

and

$$\begin{bmatrix} farP_1(\omega) \\ \vdots \\ farP_N(\omega) \end{bmatrix} = \begin{bmatrix} F_{1,1} & \cdots & F_{N,1} \\ \vdots & \cdots & \vdots \\ F_{1,N} & \cdots & F_{N,N} \end{bmatrix} \cdot \begin{bmatrix} SP_1(\omega) \\ \vdots \\ SP_N(\omega) \end{bmatrix} \tag{B.34}$$

Now, the maximum and minimum relative phases for each source over both matrices are taken to be the maximum and minimum phase bounds since they cannot be exceeded through vector addition. The minimum phase for each source can be adjusted to zero, and the maximum phase over all the sources determines the region in which the phase resulting from vector addition will lie in. (see Figure B.9) It should be noted that bounding the phase error in this way can be applied to bound cancelation at a given frequency and can be used in conjunction with the D.C. upperbound to bound amplitude error at a given frequency. Unfortunately this method of bounding error is useful only at wavelengths much longer than the distances within the sampled region. The estimates for both the upper and lower bounds of both amplitude and phase can be improved if the actual signals, locations and properties of both the real and virtual sources is known.

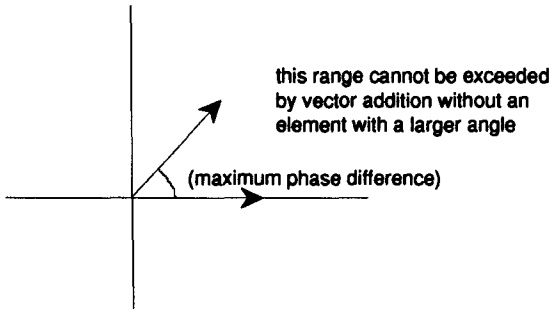

Figure B.9: Diagram of maximum phase deviation within region.

There are several possibilities for lessening the effect of the non-bandlimited attenuation with distance. First, the distance between the speakers and the listening area can be increased. This decreases the relative change in distance, and thus in effect flattens the spatial decay term regardless of amplitude. This requires of course, a large room, but a large room also helps with reducing unwanted reverberation. Another possibility is to focus the radiation from the speakers such that the spatial decay is lessened.

Appendix C

Minimax and Algorithm 1

The algorithm for trading phase for amplitude can be used in a novel way in conjunction with the quest of finding a minimax error solution with respect to amplitude along a line within the sound field. As discussed in the last chapter, one way to find the minimax solution for certain types of problems is to search for an extremal basis. Part of this process can be described with matrix notation. Indeed, the algorithm of de la Vallée Poussin can be described as follows [37].

1. Given an initial basis of N+1 points, compute the transfer matrix with N parameters and alternating signs for the error.

2. Solve for the error and coefficients.

3. Compute the largest error on the interval.

4. Insert the location of this error in the basis without changing the order nor alternating sign of the error property.

5. Repeat with new basis until the extremal basis is found, and the error no longer increases.

For example, given five sources located at (-2,0), (-1,0), (0,0), (1,0), (2,0) with 1/d decay, find the minimax solution on the interval from (-2,-1) to (2,-1) relative to a source at (0,1). (see Figure C.1)

Solution: Choose an initial equally spaced basis along the line y=1 at x=-2, -1.2, -0.4, 0.4, 1.2, 2. Since there are one to many points, an extra parameter is required, namely the error h which will be made to alternate signs. This amounts to solving

$$\begin{bmatrix} 1 & f_0(y_0) & f_1(y_0) & \cdots & f_n(y_0) \\ -1 & f_0(y_1) & f_1(y_1) & \cdots & f_n(y_1) \\ \vdots & & & \ddots & \vdots \\ (-1)^{n+1} & f_0(y_{n+1}) & f_1(y_{n+1}) & \cdots & f_n(y_{n+1}) \end{bmatrix} \begin{bmatrix} h \\ a_0 \\ a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} f(y_0) \\ f(y_1) \\ \vdots \\ f(y_n) \\ f(y_{n+1}) \end{bmatrix} \tag{C.1}$$

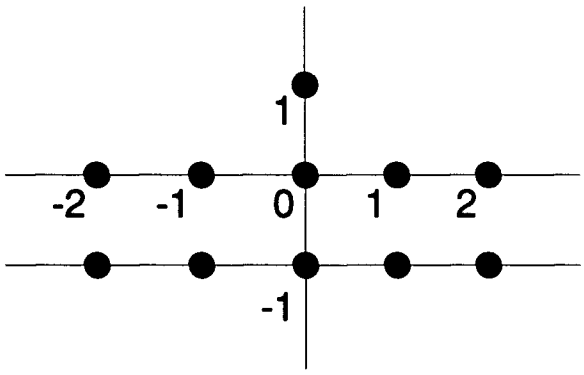

Figure C.1: Initial setup with 5 speakers

This method was actually used to get the result shown in Figure C.2 and Figure C.3. The results of this method can be compared with a simple inversion to get the correct answer at 5 equally spaced points (Figure C.5) and a least squares solution with 250 samples over the region (Figure C.5). Note however, that this minimax approach will only work for real numbers. If the values are allowed to be complex, then in general the error will also be complex, still alternating in sign. In Figure C.6 the circles represent the target amplitude, and the actual amplitude vector is labeled.

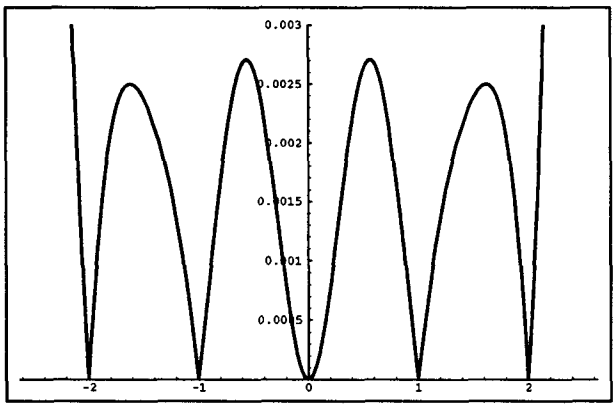

Figure C.2: Absolute error of the five point inversion solution.

Since the error is one vector (complex number) it will always point in some direction or 180 degrees opposite (when the sign reverse) regardless of the complex value at some point. Since we are concerned with only the magnitude, it would be nice if the error pointed along the same line containing the complex value at a given point. This can be achieved in two stages. First the error vector can be made to form the same angles at every point by a simple manipulation.

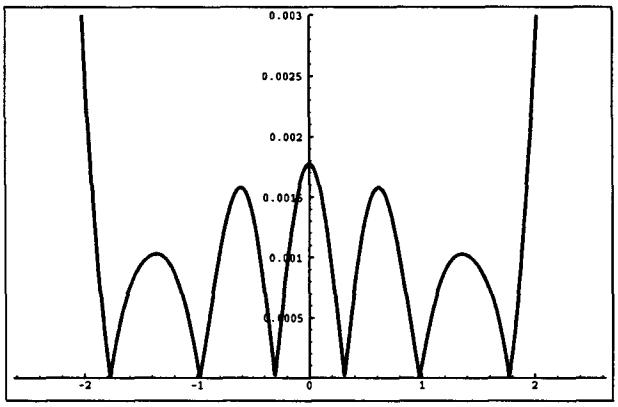

Figure C.3: Absolute error of the least squares solution over 250 equally spaced samples.

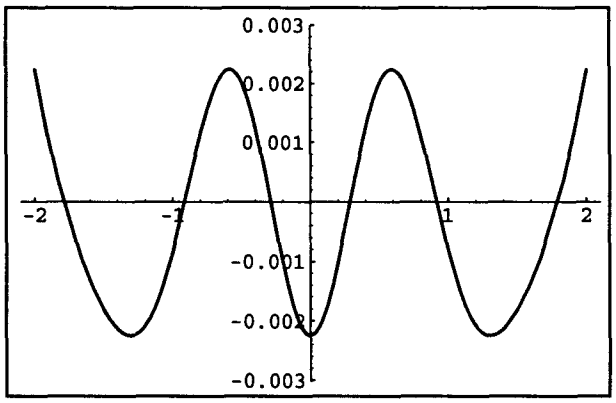

Figure C.4: Error of the minimax solution.

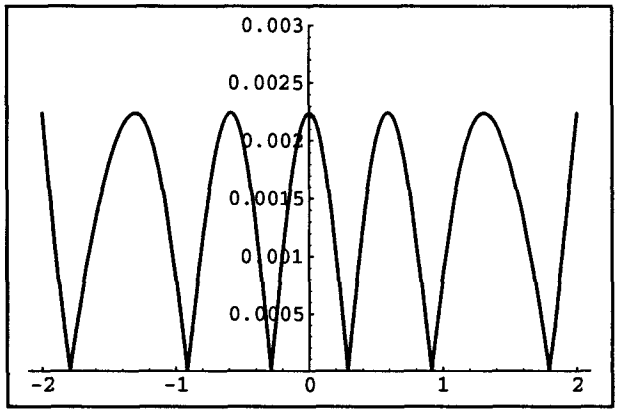

Figure C.5: Absolute error of the minimax solution.

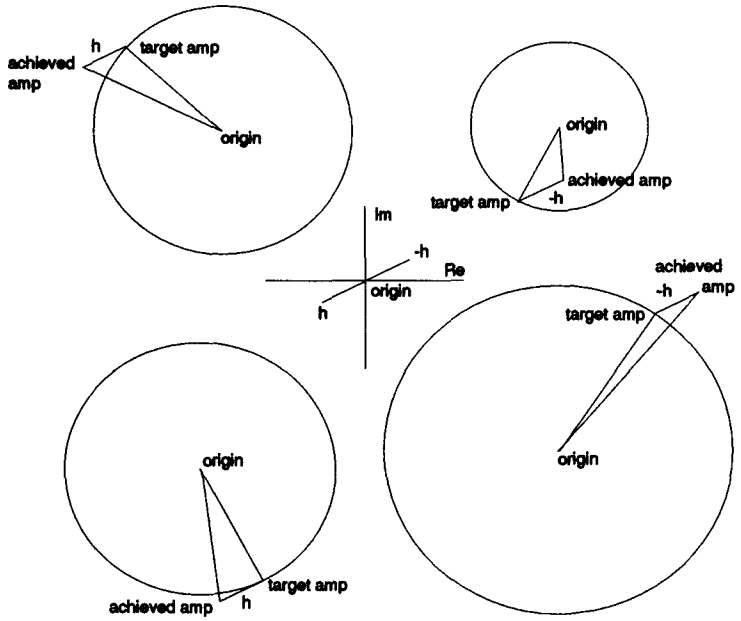

Figure C.6: Possible sample values with reconstruction error vectors h, -h.

$$\begin{bmatrix} e^{j\arg(f(y_0))} & f_0(y_0) & f_1(y_0) & \cdots & f_n(y_0) \\ -e^{j\arg(f(y_1))} & f_0(y_1) & f_1(y_1) & \cdots & f_n(y_1) \\ \vdots & & & \ddots & \vdots \\ (-1)^{n+1}e^{j\arg(f(y_{n+1}))} & f_0(y_{n+1}) & f_1(y_{n+1}) & \cdots & f_n(y_{n+1}) \end{bmatrix} \begin{bmatrix} h \\ a_0 \\ a_1 \\ \vdots \\ a_n \end{bmatrix} = \begin{bmatrix} f(y_0) \\ f(y_1) \\ \vdots \\ f(y_n) \\ f(y_{n+1}) \end{bmatrix} \tag{C.2}$$

An example of this adjustment to keep the relative phase of the error vector with the target vector constant is shown in Figure C.7. The errors can be further weighted to correspond with the magnitude of the target.

Next the algorithm trading phase accuracy for amplitude accuracy can be applied. Since the error vector has the same magnitude at each point and alternates direction, the act of replacing the phase with the phase of what the speakers achieve at a certain point will change the error vector in two possible ways depending on whether it was facing inward or outward. If facing outward with magnitude $|h|$, then the new magnitude will be $(h^2 - (r\sin\theta)^2)^{0.5} - (r - r\cos\theta)$ while if facing inward the new value will be $(h^2 - (r\sin\theta)^2)^{0.5} + (r - r\cos\theta)$. These formulas can be seen by studying the geometry of Figure C.8 and Figure C.9 respectively. A third case which is much rarer can also occur in which the error vector begins inside the constant amplitude circle, but then goes out. Two examples of this are given in Figure C.10.

Thus the amount by which the error vector is reduced depends on both the amplitude at that point as well as the amount of phase adjustment. However, since the error is reduced each time, this process will converge. (If the adjustment to make the direction of the vector h the same except for sign relative to field was not done, then the process could converge only to vectors pointing in the same directions as the error vectors)

I certify that I have read this dissertation and that in my opinion it is fully adequat in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Julius O. Smith III (Principal Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequat in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Chris Chafe

I certify that I have read this dissertation and that in my opinion it is fully adequat in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Perry R. Cook

Abstract

Too few loudspeakers are applied with current technology to reproduce a soundfield exactly to high audio frequencies over a large region. However, it is feasible to use reasonable numbers of speakers to reproduce a soundfield in a lower frequency range. The prior work of this nature has primarily grown out of transaural techniques, noise control approaches and applications of the Kirchhoff-Helmholtz integral theorem, either explicitly or implicitly. These approaches to soundfield reconstruction are essentially similar in that they all use some form of a linear transfer matrix to relate the speaker array output to the soundfield. Knowing this allows seemingly disparate techniques to be combined. Possible perceptual tradeoffs are discussed in the context of optimizing with a clear perceptual goal depending on carefully weighed considerations. The necessary steps to achieve this 'perceptual' reconstruction are discussed on a higher level before details of possible implementations are mentioned. Emphasis is placed on a class of iterative algorithms wherein the target soundfield is replaced by a perceptually similar target soundfield which is easier to obtain using loudspeakers and simple methods such as least squares approximation. To explore various techniques and ideas, simple models are used consisting of monopole monochromatic sources in a simulated anechoic environment. The iterative algorithms are discussed in the context of implications for a simple psychoacoustic model. By taking the difference between the target soundfield and the reconstructed soundfield, it was found that modest improvement is possible, reducing amplitude error at a set of sample locations, and relative phase error between members of such a set. The cost of this improvement is a worsened absolute phase error, which is much less perceptually relevant. A direct implementation based on some of these discrete frequency results using the short time Fourier transform is examined. Also, a simple method of approximating a least squares solution to an ill-conditioned problem is discussed informally.

Acknowledgments

I would like to take this opportunity to thank the many people who have been supportive and helpful, enabling me to do this work. Julius Smith, through his energy, enthusiasm and ability to convey to me part of his vast knowledge and experience, is particularly responsible for inspiring within me the desire to do this sort of work. His support, along with the support of John Chowning, Chris Chafe and Perry Cook not only made my continued stay at CCRMA possible, but also very productive and rewarding. They, along with many others have created an environment that I'm glad to have been a part. Heidi Kugler, Patty Wood and Fernando Lopez-Lezcano are also to be thanked for keeping CCRMA running smoothly on a day-to-day basis as well as their many other contributions that made life and work here enjoyable. I would also like to thank all the other people at CCRMA, fellow students, scholars and faculty, composers and researchers who are too numerous to name, for making it such a wonderful environment. Thanks also to the CCRMA affiliates, who not only provided funding for part of my stay, but inspired me yearly to present at the annual meeting. Finally I wish to thank my mom, dad and sister who have always been supportive of me and my unusual choice to stay at Stanford to pursue this work.

Contents

List of Tables

List of Figures

Chapter 1

Background

1.1 Introductory Comments

Although, in principle, the goal of loudspeakers is to reproduce arbitrary audible signals, special emphasis has been placed on sound localization, since in nature most sounds come from localized sources and people are often able to perceive roughly both the direction and, to a lesser extent, distance. The historical development of the use of speakers to localize sounds was and is dependent on several factors, including the technology of transducers, recording media, signal processing equipment as well as our understanding of how humans localize sound. Generally, attempts to localize sound through speakers have fallen into two overlapping categories, qualitative or quantitative. The qualitative approach emphasizes psychoacoustic clues to localization, while the quantitative approach emphasizes getting the correct physical air pressure fluctuations at the ears of the listener. Needless to say, this has never been a strict dichotomy, and every effective attempt has ramifications for both aspects. Likewise, this dissertation will emphasize the quantitative approach at first but quickly place even greater weight on qualitative aspects. Historically, by necessity, more attempts have focused on the qualitative strategy since there are many more difficulties in getting the signal right quantitatively at the ears of the listener(s) through loudspeakers.

1.2 Qualitative Approaches

In the case where no calculation is made of the actual soundfield produced by the reproduction system, any attempt to make the reproduction 'sound good' must be qualitative in nature. By this it is meant that the signal driving the speaker has been made to give a certain impression to the listener without any attempt at reconstructing some target soundfield. If the output is from a recording, certain aspect of the original soundfield are contained in the reproduced field, but the differences are large with respect to localization, reverberation (especially if the listening room adds much to this), background noise and coloration. Much has been written about techniques of recording to create a pleasant listening situation, and there are various methods for setting microphones and adding effects and so on [98] [75]. These techniques are sometimes somewhat quantitative in nature, such as making sure there is no unacceptable cancelation in certain places at certain frequencies. The setup of listening environments can also be somewhat quantitative in nature as when special schemes of speaker placement are used, or when materials are applied to deaden the room or diffuse the sound [89]. Although some common methods used in mainstream audio

enter the grey area between qualitative and quantitative, all are based on the perceived outcome without the careful control of the soundfield typical of a quantitative approach.

1.2.1 Affecting the Source

Certain types of processing are applied to the source signal in many cases. These can be effects like chorusing, flanging, distortion, etc., or can be the relationship between sources through mixing and editing [75]. The goal is to enhance the recorded signal to sound as good as possible. Then, any degradation caused by less than ideal listening situations must begin at a higher level. Also by having an interesting source, other problems such as localization or room effect may receive less attention.

1.2.2 Room or Environment

Generally, artificial reverberation is added to most recordings which creates certain impressions of the current or previous sound environment's reveberant qualities. The rational for using it is not always to create a room effect but may be as simple as the source sounds better with it [75]. Nonetheless, reverberation does create an impression on the listener of the dimensions and quality of the supposed recording or reproduction environment which seem to be the reverberation's source. Generally two approaches to reverberation have been used, ray tracing and modal analysis [20] [81] [59].

Ray tracing attempts to track how a sound reflects off walls or obstacles by tracing the wavefront from the source to the listener in all possible paths. It is most useful for early reflections, before the sound has had a chance to diffuse. Since refections off flat surfaces are at the same angle, a popular method to compute these paths is by the image method, where the source is reflected off the boundaries to create virtual sources, allowing the calculation of the signal at the listening point through a free field calculation. This is most easily done in rectilinear rooms, but has been extended by Jeff Barish [20] to be applicable to any environment consisting of straight boundaries, for example polygonal rooms.

Modal analysis of rooms is more useful for modelling the spectral content of latter reverberation through finding the modes of a room, which reinforce certain frequencies [59]. Although modal analysis can be applied to irregularly shaped rooms using approximation methods such as finite element analysis, it is most useful for rectangular rooms where the analysis is more straightforward. These methods of modeling reverberation are often used to help create localization effects as discussed in the next section.

1.2.3 Localization

Peoples' ability to figure out the distance and direction a sound was coming from continues to play a key role in their survival, since it allows danger to be located when in the peripheral vision or outside of the visual range. The ability to use sound localization together with visual localization and kinetic localization allows people to form perceptions of a three dimensional world. Large discrepancies between a visual and and corresponding auditory localization will cause some confusion, thus many researchers are concerned about localization when combining sound with visual stimulus [41].

It also may be the case that a clearly localized sound source is easier to separate mentally from other competing sources. This may play a part in the "cocktail party effect", where a listener can zero in on a speaker, while excluding competing speakers [23].

Finally, sound localization has been important to composers for a long time. From antiphonal singing to Henry Brant's extensive use of position as a compositional component (for instance [40] [22]), composers have sought to influence our perceptions of localization. Often this component has been under utilized by composers in that no specific instructions are made by them. Nonetheless every performance has localized sources, and musicians attempt to arrange themselves to create a pleasant soundfield. In electronic music Stockhausen has been vocal about the use of localization, and many computer pieces by many composers use localization as a parameter. In some pieces, such as Chowning's *Turenas* [31], it is a key factor. In order to control localization qualitatively it is important to understand how people localize sound.

Distance Cues

Since the intensity of a sound is inversely proportional to its distance from the listener, a sound known to be at a certain intensity at a certain usual distance, (for example, thunder is loud, a fly is quiet) but perceived differently can provide cues to distance (the storm is far away, the fly is right by my ear). If the sound's usual approximate intensity is not known, a situation uncommon in the real world, then other cues must be found.

Another very important cue is the ratio of reverberant energy to direct energy in enclosed environments [19]. Due to the nature of reverberation, its intensity decreases much less with distance than the intensity of the direct path. Since sound coming by the direct path arrives at the listener first, direct energy and reverberant energy are distinguishable and can be compared consciously or unconsciously.

Finally there is an effect of air absorption of higher frequencies over large distances. This effect is not noticeable for normal distances, but for sounds occurring far away the cumulative effect of air on the propagation is like a lowpass filter [83].

Intraural Time Difference

In the steady state of lower frequencies, the phase difference between the two ears provides the most important cue to angular localization. Because of the larger wavelengths there is little intensity difference at the two ears caused by the head. Above about 1500Hz, the size of the head is such that phase difference is no longer a reliable cue [95].

Intraural Intensity Difference

At frequencies between 3000Hz and 6000Hz the head is a better barrier and intensity difference between the two ears becomes the dominant cue for localization [95]. Essentially the head casts a shadow at these frequencies so the ear nearer the source gets a stronger signal and the relative amplitude of signal between the two ears is dependent on the direction of the head to the source.

Pinna Cues

The pinna provides a distinctive filter, which helps separate known sounds originating in the front and back of the head [95] [11]. This filtering of sound provided by the pinna, and to a lesser extent reflections off the rest of the body, especially the head and shoulders, provides spectral cues for discerning direction. Essentially the spectrum of a source will vary based on its location in three dimensions, and certain key frequencies are especially

important. Often time delay and intensity difference are not enough to determine whether a source is in front or back, nor how high or low, and this additional filtering provides extra clues. Pinna cues have been studied [19] [11] and vary from individual to individual. Some pinna filterings are more effective than others, and this combined with the ability of people to recognize direction based on other people's pinna filterings has lead to extensive experimentation involving such filters delivered over headphones [12]. Actual use of pinna filterings over loudspeakers is wrought with difficulty since the sound will be refiltered by the listeners own pinna. However awareness of the role of the pinna is crucial to designing localization systems for loudspeakers, if only to see why they don't work well.[1] Combining the pinna effects with the rest of the head and shoulders allows the creation of head related transfer functions, which will be discussed in section 1.3.1 under quantitative techniques since these transfer functions are carefully designed for exact soundfield reproduction in the ear canal.

Minimal Audible Angle

The ability to localize sounds is not always the same for an individual depending on where a sound is in relation to the head, and depending on the characteristics of the sound. Knowing the minimum audible angle for most people for a sound at a certain frequency and direction can be useful. People may be made more susceptible to cues hinting at one localization over another when their minimal audible angle is taken into account [95].

Doppler Shift

Doppler shift refers to the change in frequency due to the relative motion between a sound source, a sound receiver and the sound transmitter (air). Since people have heard these pitch shifts and associate them with changing relative distance, Doppler shift is a useful clue to aid in the localization of moving sounds in the real world. Likewise a simulated Doppler shift combined with other localization clues can help create a perception of a moving localized sound. The general formula for the changed frequency due to Doppler shift is

$$f_{new} = f\left(\frac{c \pm v_o}{c \pm v_s}\right) \tag{1.1}$$

where c is the speed of sound, v_o is the component of the observers velocity directly toward or away from the source and v_s is the component of the source's velocity toward or away from the observer, with both velocities being measured relative to the medium (air) [67].

1.2.4 Implementations

CCRMA has a long history of being involved in sound localization through speakers using psychoacoustic clues. Pioneering work done by John Chowning [30] has been followed by numerous researchers and composers including Marnia Bosi [21], William Schottstaedt, Fernando Lopez-Lezcano and others [60] [61]. Also important work has been done elsewhere by F. R. Moore [69], Jean-Marc Jot [49] [50] and many other people.

[1] for example, why it is hard to get localizations in back from speakers in front etc.

Heightening the Cues

It has been found that in order to overcome the undesired but natural localization and room effect caused by the listening environment, it is often helpful to make the cues larger than psychoacoustic research might indicate. The reason for this might be similar to the hypothesis given earlier that so much attention is given to making the source sound 'better than natural' in part to distract from more apparent problems. It may be that by increasing certain cues beyond their 'normal' values attracts attention to them and away from other phenomenon in addition to the hypothesis that the augmented cues help weight the perception in the case where attention is constantly applied in the same proportion to the cue and other phenomenon regardless of the cue's strength. An example of this increase in the cue is given in the summary of some of Lopez-Lezcano's methods [60] [61].

Chowning

The early work of John Chowning is summarized in [30], which described a program which simulated doppler shift, adjusted the ratio of direct to reverberant sound, and used a non-linear function when adjusting each channels amplitude to help "fill the hole" in between speakers. This was done for a quadraphonic system though mention was made of application to stereo. This system was controlled by a graphical user interface. The use of these clues in a musical work such as *Turenas* demonstrates the effectiveness of this approach for moving sounds as well as the limitations.

Bosi

A generalization of Chowning's approach was demonstrated in [21], expanding the equations to allow for N co-planer speakers, and implementing the system on four speakers in real time using a graphical interface on an Apple MacII computer to send a source (from a Yamaha DX-7 synthesizer) through two Yamaha DMP-7s (digital mixers) with two reverb modules (all controlled via MIDI) creating the appropriate output.

Lopez-Lezcano

More recently Fernando Lopez-Lezcano has added functions to William Schottstaedt's Common Lisp Music (CLM) which allow simulation of moving sounds on a quadraphonic system controlled by a NeXT computer. His approach, described in [60] and [61] implements distance using a varying length delay line, which has the advantages of 1) giving approximately the correct arrival time for two sources with simultaneous attacks, but different distances from the listener, 2) producing the correct doppler-shift as a side effect, 3) allowing the addition of a filter at the end of the delay line to model air absorption and 4) allowing adjustable attenuation at the end (the delay line is lossless). This adjustable attenuation is desirable because while the physical attenuation would be 1/distance, C. Sheeline found that raising $\left(\frac{1}{distance}\right)$ to a higher exponent (like 1.5) was more effective, especially for unfamiliar sounds [87]. Furthermore, the system for controlling the apparent motion of sources makes use of user determined points through which the virtual source moves, following an user controlled interpolation of points in both time and space by Bezier curve fitting.

Moore

F. R. Moore has taken the approach of thinking of the listening space as a room within a room, where the speakers approximate small apertures through which sound from the desired environment enters.[2] Attempts are made to figure out what signals would be arriving at the speaker positions from the imaginary environment. These signals are then broadcast over the speakers to create a room within a room effect [69]. Although this line of thought attempts to do what is possible with limited resources, very few people would go to a concert hall only to sit in a box with a few small holes in it.

Reverberation has been seen to play a key role in the localization of sounds as well as in characterizing the environment in which sounds are located. Some researchers have attempted to capitalize on this by producing multi-speakers systems, where reverberation plays a key role. In 'the first reverberation system that recreates the full spatiotemporal soundfield of a natural environment' [54] the spatial reverberator uses N output channels to model direct sound as well as M reverberation signals. The reverberation technique used attempts to physically model what the actual reverberation would be in a given environment.

Jot et al.

Much work has been done at ICRAM (Institut de Recherche et Coordination Acoustique/Musique) on sound spatialization. The result has been a system called the Spatialisateur [50], which runs in real time processing dry sources to create spatial and room effects in a carefully controlled way. The controls which the user has access to are perceptually based parameters with perceptual names, such as brilliance and warmth. Each of the eleven parameters, two of which determine the direction of the source, while nine control perceptual parameters based on room source interaction, can be individually controlled. The output is formatted for various listening situations, such as plain stereo, quadraphonic as well as quantitative formats discussed in the next section, which are head related transfer functions over headphones or transaural stereo [49] [50]. Despite employing the possibility of quantitative techniques in some modes of reproduction, the philosophy behind the Spatialisateur is clearly qualitative in nature.

In general, qualitative systems use the same localization cues to simulate localized sources. The differences in systems are in the implementation of these cues. For instance reverberation can be added using an artificial reverberation filter such as the Schroeder reverberator [91], or modelled explicitly using one or more of several techniques including ray tracing, modal analysis or convolution with actual impulse responses of rooms. Another difference between qualitative systems is the amount and type of control over the perceptual parameters given to the user.

Surround Sound and Ambisonics

Ever since stereo replaced mono, thought has been given to how to add more channels. Original work with stereo actually involved more than two channels, however LPs could only record two channels and two channel stereo became the norm around 1958. Other systems involving more channels have been available (such as quadraphonic), but never reached a wide commercial market. Some early work on multichannel audio is due to

[2] Actually, the idea of a small aperture is just an analogy as the radiation from most speakers will be quite different from a real "hole in the wall" as far as directivity and frequency response.

Harvey Fletcher, and a system devised by him was almost adopted by the movie industry in the 1940's [9] [2]. However, only recently have multichannel techniques become standard in the movie industry. The driving force behind this has been Dolby Inc., which has long been famous for its noise reduction processes [2].

Ideas which can be applied to noise reduction are often also useful for compressing audio information, which was important to allow encoding of additional channels. Dolby Stereo introduced in film encoded left, center, right and surround channels on two tracks through a MP (Motion Picture) Matrix encoder [2]. This was made into consumer products with some adjustments and called Dolby Surround in 1982. In 1987 Dolby Pro Logic Surround was introduced in the consumer market with more complicated processing [1]. A more recent system called Dolby Surround AC-3 uses compression to put 5.1 channels of audio in less space than required by one channel on a CD [3]. The 5.1 channels are left, center, right, left surround, right surround with ranges from 3Hz to 20,000Hz and a low frequency effects channel with a limited range from 3Hz to 120Hz, which is considered the 0.1 channel [3].

Ambisonic technique finds the main directional components of a soundfield [5]. An omnidirectional (monopole) component (called W) is combined with bi-directional (dipole) components along the X, Y, and Z axis, (called X,Y, and Z) . Further components corresponding to quadripoles etc. can also be used within the framework but are not in practice. The W and Y components are the same as the M and S components when using M-S stereo recording. The Z signal is sometime omitted since this requires speakers above and below the listening region. Transformations are also used to encode these channels on two channels.

In [62] comparisons some of the various surround sound formats are taken on the basis of interaural time difference and interaural intensity difference. Only these two cues as measured through a dummy head were used to create a simple model of the imaging capability of the systems in a "real room" (nonanechoic). The frequency dependent interaural intensity difference (IID) was calculated as the ratio of the Fourier transforms of the signals at the two 'ears'. The interaural time difference (ITD) was measured using the cross correlation between the two signals. The formats compared were spectral stereo, KEMAR reconstruction, and ambisonics. Spectral stereo refers to the attempt to place sound images by accounting for diffraction on a simple model of the head (a sphere). KEMAR reconstruction goes a step further by using data from a KEMAR acoustic mannequin rather than a spherical model. The main result was the usefulness of these techniques for measuring how well a system creates sound images. By way of demonstrating the technique it was found that that both the spectral and KEMAR approaches did reasonably well although the KEMAR approach was more sensitive to listener placement. The Ambisonic approach lagged behind, but only certain angles were chosen so the results were not conclusive.

1.3 Quantitative Approaches

Quantitative approaches include those which seek to produce sounds accurately in the physical sense, without regards to perceptual effects. Work toward the design of accurate (non-distorting) speakers and other sound equipment lays the basis for this category. Without fine control and knowledge of speaker outputs, any attempt to control pressure fluctuations away from the speaker are doomed. Probably the simplest approach under this heading is to place the speaker where the imitated source would be. An example of this is placing a speaker in the mouth of or just behind the head of a mannequin to simulate talking. In

the case of multiple speakers, the direct signal can degenerate into this case when a speaker happens to be coincident with the source's intended position. There are at least two essential problems with this approach. Firstly, the directivity of a speaker is not often exactly the same as the original source, so that only within a small region will the sound waves be close to identical in amplitude and phase with the original source. Secondly it is almost always impractical to have speakers for each intended source location. This is especially difficult if the source is moving (which human performers often can do, subtly or not so subtly) as the speaker would have to move the same way[3].

1.3.1 Head-Related Transfer Functions (HRTF)

Within the large grey area between qualitative and quantitative approaches lies HRTF approaches to sound localization. These approaches use filters to simulate the lost processing of the pinna, and sometimes also the head and shoulders, to create a stereo sound heard through headphones [12].

The process of filtering with a HRTF ideally produces the effect of localization. Although psychoacoustics plays a major role regarding the processing of the outer ear, ideally the soundfield reaching the eardrum is the same as one from the original virtual source. In this strict sense, HRTFs are just an example of soundfield reconstruction, where the reconstructed soundfield has already interacted with the outer ear.

1.3.2 Reconstruction of Soundfields at Specific Points

There are at least two distinct approaches to soundfield reconstruction, although both fit together neatly in a broader catagory. The first approach attempts to get the correct signal at some set of points in space directly, hereafter referred to as the matrix approach, since despite many possible variations, matrices play a key role explicitly or implicitly. The second approach attempts to get the correct signal over some continuum by applying the Kirchhoff-Helmholtz integral theorem and will be referred to as the integral approach. Both names are for convenience only, based on observation, and do not imply the necessity of integrals or matrices.[4] One goal of this dissertation is to relate these seemingly different approaches.

Schroeder and Atal

The first application of what might be referred to as 'two point' soundfield reconstruction (through loudspeakers) is probably due to Schroeder and Atal who in 1963 devised a system using two speakers to create two points near a person's ears which simulated localized sound sources as well as associated reverberant sound [88]. This approach was used by them in 1970 for research in auditorium acoustics and proved to be an effective tool [90]. This basic approach has been generically called "transaural stereo" [34] and has continued to be developed. This approach can be generalized to larger numbers of speakers and points quite easily and forms the basis of the matrix approach.

[3]To be very strict, the microphone recording the original would have to move as well, or some signal processing solution would have to be found to avoid doubling the apparent movements as far as intensity, doppler-shift, direct sound to reverb ratio etc.

[4]The HRTF approach can be thought of as falling in the first category even though matrices are not needed to compensate for the two speakers since they are isolated in headphones.

The approach taken in [88] is to assume two loudspeakers placed symmetrically in front of a symmetrical person. The $h_1(t)$ and $h_2(t)$ are the impulse responses produced at the left and right ears respectively by the first speaker, while $h_2(t)$ and $h_1(t)$ are the impulse responses produced at the left and right ears respectively by the second speaker. This makes the strong assumption that there is equivalence between the impulse response from the first speaker to the left ear and from the second speaker to the right ear, as well as equivalence between the response from the first speaker to the right ear and second speaker to the left ear. Assuming this the signals $y_1(t)$ and $y_2(t)$ at the left and right ears are then

$$\begin{aligned} y_1(t) &= x_1(t) * h_1(t) + x_2(t) * h_2(t) \\ y_2(t) &= x_1(t) * h_2(t) + x_2(t) * h_1(t) \end{aligned} \tag{1.2}$$

taking the Fourier transforms

$$\begin{aligned} Y_1(\omega) &= X_1(\omega) \cdot H_1(\omega) + X_2(\omega) \cdot H_2(\omega) \\ Y_2(\omega) &= X_1(\omega) \cdot H_2(\omega) + X_2(\omega) \cdot H_1(\omega) \end{aligned} \tag{1.3}$$

or in matrix form

$$\begin{bmatrix} Y_1(\omega) \\ Y_2(\omega) \end{bmatrix} = \begin{bmatrix} H_1(\omega) & H_2(\omega) \\ H_2(\omega) & H_1(\omega) \end{bmatrix} \begin{bmatrix} X_1(\omega) \\ X_2(\omega) \end{bmatrix} \tag{1.4}$$

It is now easy to solve for $X_1(\omega)$ and $X_2(\omega)$ and take the inverse Fourier transform to produce solutions for $x_1(t)$ and $x_2(t)$.

$$\begin{aligned} x_1(t) &= [y_1(t) * h_1(t) - y_2(t) * h_2(t)] * a(t) \\ x_2(t) &= [y_2(t) * h_1(t) - y_1(t) * h_2(t)] * a(t) \end{aligned} \tag{1.5}$$

where $a(t)$ is the inverse Fourier transform of

$$\frac{1}{H_1^2(\omega) - H_2^2(\omega)} \tag{1.6}$$

It should be noted that at some frequencies $H_1^2(\omega) - H_2^2(\omega)$ can get very close to zero causing eq. 1.6 to blow up. This is most likely to happen where the difference in direct path lengths from one speaker to the two ears is close to an integer multiple of the wavelength, so most of the linear independence of the direct signal is produced by slight differences in amplitude due to the different distances.

In order to graphically demonstrate what happens using various approaches, a simulation environment was created. This environment is described in detail in an appendix, but for now it will just be used to graphically represent a 100 by 100 unit section of a free field. To demonstrate this environment, and at the same time illustrate the method of Schroeder and Atal, consider a source located at (40,5) using arbitrary units of measure and some chosen origin. If there are two speakers at (25,25) and (75,25), and the points where the reproduction equals the original are chosen to be (47,50) and (53,50), possibly corresponding to a person's ear locations, then the accuracy of the reproduction can be graphically shown. For example, Figure 1.1 shows the result of subtracting the source from the reproduction in the field (complex values) and plotting magnitude of the difference. Here black indicates high magnitude (1 or above, depending on the resolution setting) and white indicates 0. Notice that the Y-axis increases going down the page instead of up as in the usual Cartesian coordinates. The original source is always given a magnitude of 1

at a distance of 1 unit, unless otherwise indicated. The resolution here is given a value of 2, which is just a multiplication factor to make the image more discernible. Thus black is achieved where ever the magnitude exceeds 0.5.

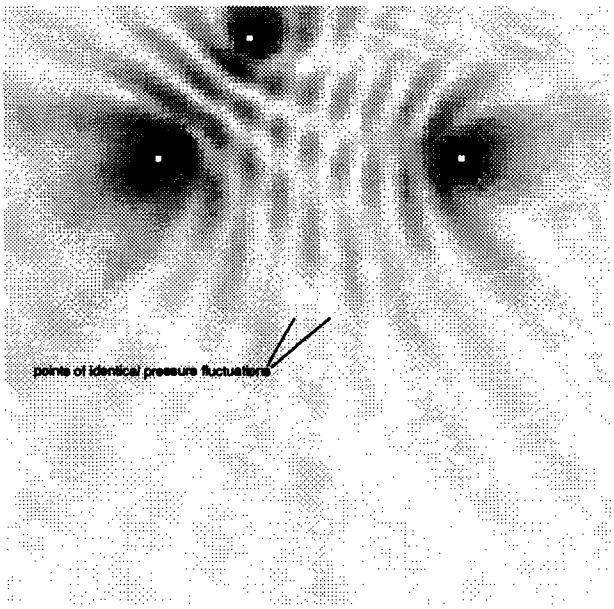

Figure 1.1: Amplitude of the difference using method of Schroeder and Atal

In addition to the simulation environment, more detailed analysis was done by simulating the response as specific points in Matlab, and moving the data into Mathematica.[5] Looking at 121 points located in a square grid pattern at the center of the room extending from (45,45) to (55,55) gives a clear picture of what happens in this region (Figure 1.2).

These is even more clear using a density plot Figure 1.3, in which black indicates the points of zero error. These points are exactly where they should occur. Another important measure is the difference in amplitudes between the target soundfield and the reproduced field. Fig.1.4 shows this measure for the center of the room.

Generalization

It is quite easy to generalize the approach for non-symmetrical speakers and people. In such a case the impulse responses at the left and right ears from the first speaker usually cannot be reversed to get the impulse responses from the second speaker. Thus four independent impulse responses are used and in the frequency domain the equation looks like

$$\begin{bmatrix} Y_1(\omega) \\ Y_2(\omega) \end{bmatrix} = \begin{bmatrix} H_{1,1}(\omega) & H_{2,1}(\omega) \\ H_{1,2}(\omega) & H_{2,2}(\omega) \end{bmatrix} \begin{bmatrix} X_1(\omega) \\ X_2(\omega) \end{bmatrix} \tag{1.7}$$

where $H_{2,1}(\omega)$ is the Fourier transform of the impulse response $h_{2,1}(t)$ which the subscripts indicate is taken from the second speaker to the first ear, and so on.

[5] Although both are just high level mathematical computer languages, Matlab was faster doing intensive caluculations while Mathemnatica was more straightforward for analysing data

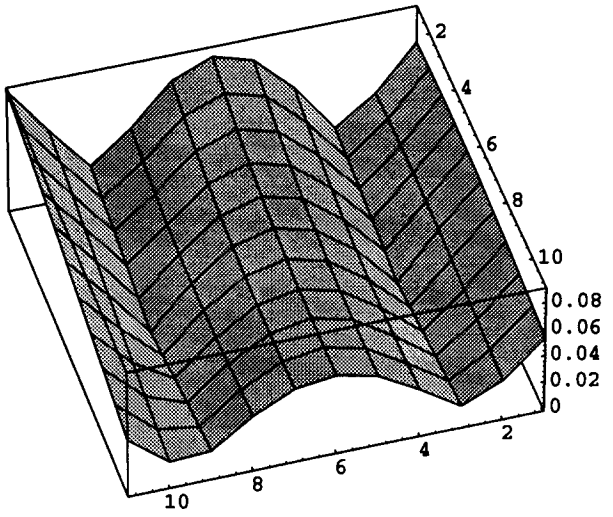

Figure 1.2: Contour plot of the difference in the center of the room

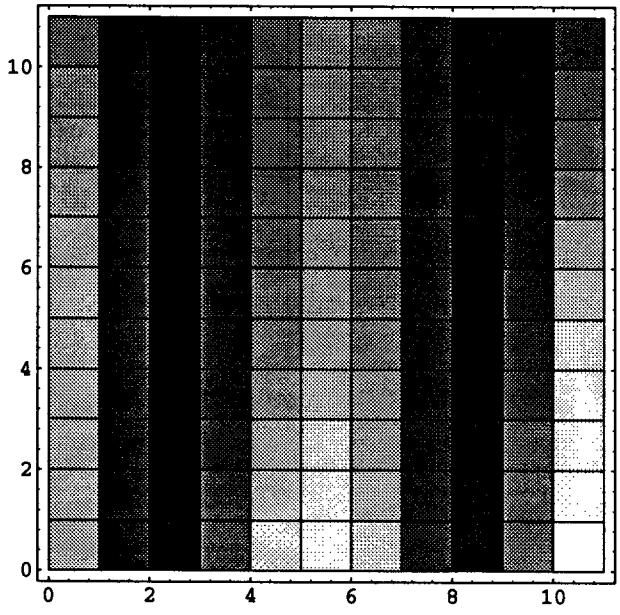

Figure 1.3: Density plot of the total difference

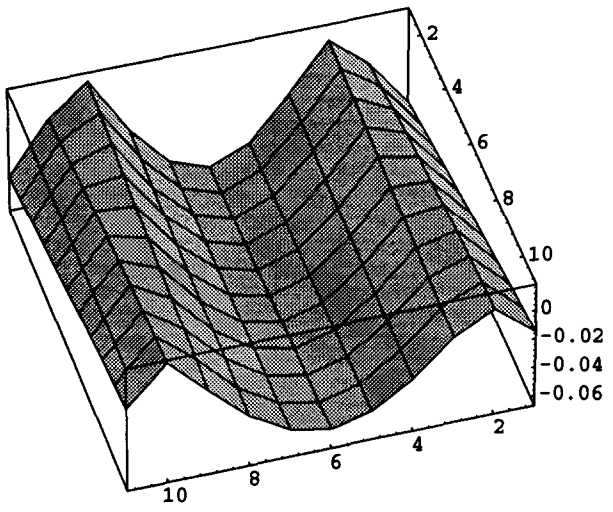

Figure 1.4: Contour plot of just the amplitude difference in the center of the room

It is also quite easy to generalize this approach to larger numbers of speakers and points in space. Experience shows that larger numbers of speakers produce better audio results. Whether the move from mono to stereo, stereo to quadraphonic, quadraphonic to eight speakers, the use of surround sound, as the number of speakers increases the better the audio quality. This seems intuitive in that, qualitatively, more speakers can decrease unwanted localization effects by diffusing the signal and if speakers are closer to a desired localization, the reproduction will be of higher quality. In the case of soundfield reconstruction this trend continues for similar reasons. In effect, each additional speaker introduces a new parameter which can be used to approximate a desired soundfield. A simple demonstration of this lays the foundation for finding the best approximation of a desired soundfield in any given situation.

It is well known that N linear equations with N unknowns can be solved provided that the N equations are linearly independent. Thus, if $Ax = b$ where A is an N by N matrix, x is an unknown length N and b is a known length N vector, then $x = A^{-1}b$ provided A^{-1} exists. This simple formula is the basis of many variations involved in soundfield reconstruction, including the above approach of Atal and Schroeder. As an example, consider the case with three speakers at locations (10,10), (50,90) and (90,10) trying to get the correct signal at the points (40,40), (50,50) and (65,60). A graph of the amplitude of the difference between the reproduced signal and a signal originating at (50,-10), both with wavelength 20, shows the identical nature of these signals at these points. Here the resolution has been increased to 4 for better contrast.

This matrix approach can be used to design filters which will work at all frequencies, given certain circumstances. In a non-reverberant environment, consider a finite number of point sources of radiation (in this case sound) within a bounded region. With no other influences, every point within this region will experience pressure fluctuations based on emanations from these point sources. In the case of sound, the pressure fluctuations at any point will be a linear combination based on the distance of the point from each source and

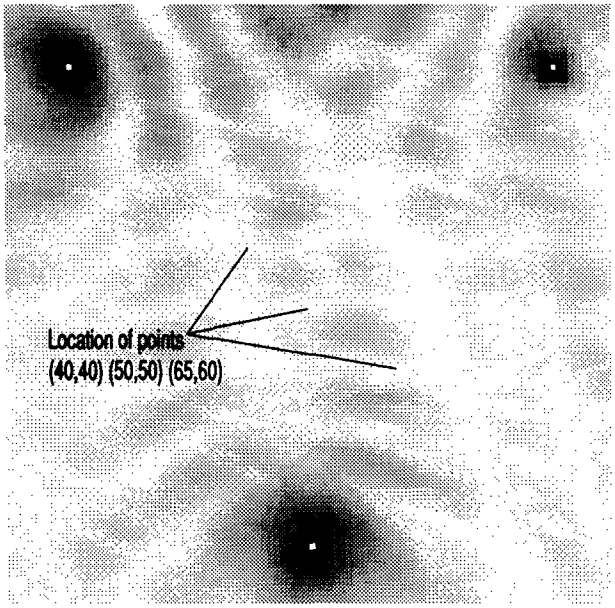

Figure 1.5: Amplitude of the difference

the speed of propagation. For any particular point p the pressure at p at time t is

$$f_p(t) = \sum_{j=1}^{N} \frac{1}{d_j} sp_j \left(t - \frac{d_j}{c}\right) \tag{1.8}$$

where c = the speed of sound
d_j = the distance between p and sp_j
N = the number of sources
$sp_j(t)$ = the disturbance at the point source (normalized)[6]

Now to simplify notation let $p(t) = f_p(t)$ then consider N speakers and N points. Then

$$\begin{array}{rcccccc} p_1(t) & = & \frac{1}{d_{1,1}}(sp_1(t-t_{1,1}) & + & \cdots & + & \frac{1}{d_{N,1}}(sp_N(t-t_{N,1}) \\ \vdots & & \vdots & & & & \vdots \\ p_N(t) & = & \frac{1}{d_{1,N}}(sp_1(t-t_{1,N}) & + & \cdots & + & \frac{1}{d_{N,N}}(sp_N(t-t_{N,N}) \end{array} \tag{1.9}$$

where $d_{i,j}$ =the distance between sp_j and p_i and $t_{i,j} = \frac{d_{i,j}}{c}$ =the time lag between sp_j and p_i. Now if the original signals are discrete with a sampling time of $\Delta t = \frac{\Delta d}{c}$ then the delay in samples between a source i and a point j is

$$\frac{t_{i,j}}{\Delta t} = \frac{\frac{d_{i,j}}{c}}{\frac{\Delta d}{c}} = \frac{d_{i,j}}{\Delta d} \tag{1.10}$$

[6]It is clear that as a p is taken closer to sp_j, $\frac{1}{d_j} \to \infty$. Thus the value of $sp_i(t)$ (for any value of i) is not the pressure fluctuation at the location of sp_i but rather at a point where $\frac{1}{d_i} = 1$, i.e. $d_i = 1$. However $sp_i(t)$ refers to this "normalized for distance" value at the location of sp_i not a point one unit away. This normalization simplifies the math later on a little bit.

The z-transform of $\frac{1}{d_{i,j}} sp\left(t - \frac{d_{i,j}}{c}\right)_{t=n\Delta t}$ is $\frac{1}{d_{i,j}} \cdot z^{-\left(\frac{d_{i,j}}{\Delta d}\right)} \cdot SP(z)$. The above system of equations can be z-transformed to become

$$\begin{bmatrix} P_1(z) \\ \vdots \\ P_N(z) \end{bmatrix} = \begin{bmatrix} F_{1,1} & \cdots & F_{M,1} \\ \vdots & \cdots & \vdots \\ F_{1,M} & \cdots & F_{M,M} \end{bmatrix} \cdot \begin{bmatrix} SP_1(z) \\ \vdots \\ SP_N(z) \end{bmatrix} \tag{1.11}$$

where $F_{i,j} = \frac{1}{d_{i,j}} \cdot z^{-\left(\frac{d_{i,j}}{\Delta d}\right)}$. In more condensed notation, $P = F \cdot SP$.

To be technically correct, we need to have $t_{i,j}$ divisible by Δt for all i and j. Although this is generally impossible, the z-transform is quite robust with fractional delays since this is equivalent to oversampling by a certain amount. Arbitrary delays can be approximated to any desired accuracy by a high enough sampling rate. In practice, arbitrary delays are approximated by some sort of interpolation function without upsampling. For now assume that the sampling rate is much higher than twice the highest temporal frequency, so that rounding $t_{i,j}$ to the nearest sample time will not produce large errors. (An analysis of the error cause by this rounding will be included later)

now, if $H = F^{-1}$, then $H \cdot P = SP$

Therefore there will be a solution if the system of equations are linearly independent. Stability will be achieved if each $H_{i,j}$ is stable. This will be the case if the roots of the determinant of F are within the unit circle since by Cramer's Rule

$$SP_i = \frac{\det |C_i|}{\det |F|} \tag{1.12}$$

where C_i is the matrix produced by substituting P for the i^{th} column and of H.

This approach can be applied to reverberant environments as well. In that case each term in the F matrix becomes more of a filter (as opposed to just a scaling and delay) while in the time domain equation each p_i is equal to the sum of the direct signal from each speaker, as before, along with the sum of every reflection reaching that point from each speaker. Furthermore, non-ideal directivity patterns of the speakers can be dealt with in a similar fashion, adjusting the F matrix for these new parameters.

Romano

In 1987, Anthony Romano arrived at some similar fundamental equations [81]. However he goes even further, suggesting the use of three sources per point to produce the the correct vector value at a single point in space, as opposed to using one source per point to get the correct scaler value.

A formulation for N points and N speakers in continuous time using the Fourier transform is given by Romano. Given N sources $x_i(t)$ the fluctuations at N points $y_j(t)$ is

$$\begin{array}{ccccccccc} x_1(t) * h_{1,1}(t) & + & x_2(t) * h_{2,1}(t) & + & \cdots & x_N(t) * h_{N,1}(t) & = & y_1(t) \\ \vdots & & & \cdots & & & = & \vdots \\ x_1(t) * h_{1,N}(t) & + & x_2(t) * h_{2,N}(t) & + & \cdots & x_N(t) * h_{N,N}(t) & = & y_N(t) \end{array} \tag{1.13}$$

where $h_{i,j}$ is the (room's) impulse response from source x_i to point y_j.
Taking the Fourier transform of this expression, and assuming non-singularity,

$$\begin{bmatrix} X_1(f) \\ X_2(f) \\ \vdots \\ X_N(f) \end{bmatrix} = \begin{bmatrix} H_{1,1}(f) & H_{2,1}(f) & \cdots & H_{N,1}(f) \\ H_{1,2}(f) & H_{2,2}(f) & \cdots & H_{N,2}(f) \\ \vdots & \vdots & \vdots & \vdots \\ H_{1,N}(f) & H_{2,N}(f) & \cdots & H_{N,N}(f) \end{bmatrix}^{-1} \begin{bmatrix} Y_1(f) \\ Y_2(f) \\ \vdots \\ Y_N(f) \end{bmatrix} \tag{1.14}$$

However, Romano argues that since each wave front has a direction, we should keep account of this by looking at a vector at each point rather than a scaler. However to get each point exactly right using a vector approach would require three speakers per point (following similar arguments as above). Thus, to get magnitude and direction of propagation correct at two points, $y_1(x, y, z|t) = (\hat{y}_1(x(t)), \hat{y}_1(y(t)), \hat{y}_1(z(t)))$ and $y_2(x, y, z|t) = (\hat{y}_2(x(t)), \hat{y}_2(y(t)), \hat{y}_2(z(t)))$ he uses the following formulation (note the distinction between the same names for vectors and the scalars of which they are composed)

$$\begin{array}{ccccccccc} x_1(t) * \hat{h}_{1,1}(x(t)) & + & x_2(t) * \hat{h}_{2,1}(x(t)) & + & \cdots & + & x_6(t) * h_{6,1}(x(t)) & = & \hat{y}_1(x(t)) \\ x_1(t) * \hat{h}_{1,1}(y(t)) & + & x_2(t) * \hat{h}_{2,1}(y(t)) & + & \cdots & + & x_6(t) * h_{6,1}(y(t)) & = & \hat{y}_1(y(t)) \\ x_1(t) * \hat{h}_{1,1}(z(t)) & + & x_2(t) * \hat{h}_{2,1}(z(t)) & + & \cdots & + & x_6(t) * h_{6,1}(z(t)) & = & \hat{y}_1(z(t)) \\ x_1(t) * \hat{h}_{1,2}(x(t)) & + & x_2(t) * \hat{h}_{2,2}(x(t)) & + & \cdots & + & x_6(t) * h_{6,2}(x(t)) & = & \hat{y}_2(x(t)) \\ x_1(t) * \hat{h}_{1,2}(y(t)) & + & x_2(t) * \hat{h}_{2,2}(y(t)) & + & \cdots & + & x_6(t) * h_{6,2}(y(t)) & = & \hat{y}_2(y(t)) \\ x_1(t) * \hat{h}_{1,2}(z(t)) & + & x_2(t) * \hat{h}_{2,2}(z(t)) & + & \cdots & + & x_6(t) * h_{6,2}(z(t)) & = & \hat{y}_2(z(t)) \end{array} \tag{1.15}$$

Taking the Fourier transform (still using x, y and z to denote orthogonal coordinates) the six sources can be solved

$$\begin{bmatrix} X_1(f) \\ X_2(f) \\ X_3(f) \\ X_4(f) \\ X_5(f) \\ X_6(f) \end{bmatrix} = \begin{bmatrix} H_{1,1}(x(f)) & H_{2,1}(x(f)) & \cdots & H_{N,1}(x(f)) \\ H_{1,1}(y(f)) & H_{2,1}(y(f)) & \cdots & H_{N,1}(y(f)) \\ H_{1,1}(z(f)) & H_{2,1}(z(f)) & \cdots & H_{N,1}(z(f)) \\ H_{1,2}(x(f)) & H_{2,2}(x(f)) & \cdots & H_{N,2}(x(f)) \\ H_{1,2}(y(f)) & H_{2,2}(y(f)) & \cdots & H_{N,2}(y(f)) \\ H_{1,2}(z(f)) & H_{2,2}(z(f)) & \cdots & H_{N,2}(z(f)) \end{bmatrix}^{-1} \begin{bmatrix} Y_1(f) \\ Y_2(f) \\ Y_3(f) \\ Y_4(f) \\ Y_5(f) \\ Y_6(f) \end{bmatrix} \tag{1.16}$$

Thus, six speakers can be used to produce the correct amplitude and direction of propagation at two points. The question arises of whether the gain is worth the cost, considering the good results produced by Schroeder and Atal using only only two speakers. If just pressure fluctuation is of concern, then only the magnitude of the vector must be known. It turns out that the apparent vector quality of a wavefront is a by-product of pressure fluctuations since uneven distributions in pressure causes the medium to oscillate in certain directions trying to achieve equilibrium creating wave propagation. Therefore, getting the pressure fluctuations right at enough points insures the correct wavefront direction without

resorting to three speakers per point in space. Romano argues that too many speakers are needed to get the correct vector quantities in this way. Still, using three speakers per point rather than one tends to exacerbate the problem and it is not at all clear that the benefits outweigh the costs in all but the most constrained circumstances. If there are not enough speakers, then the question becomes whether they should be used to get the right vector quantity at a few points, or the pressure right at three times as many points. Although the choice depends on the application, it seems that often getting the pressure fluctuations right at a larger number of locations will be preferred since there will be some tendency of the wavefront vectors to conform anyway because of the correct pressure distributions. Vector quantities can be computed by ray tracing models, but these are easily transformed into scalars since the magnitude of the sum of the vectors arriving at the same location at the same time gives the pressure fluctuation at that moment.

Romano uses a scheme to expand a discrete, finite convolution to create a serial product from which the input series can be obtained, creating a digital filter implementation. However, he seems to skirt the issue of stability entirely, and this turns out to be an important consideration. Another variation on the same theme is to design the filter directly from discrete impulse responses using a z-transform instead of a continuous time Fourier transform followed by discretization.

As a graphical example of Romano's method, consider a 2-dimensional case with an original source at (50,10) and speakers at (10,10) and (60,10) used to produce the correct directional values at the single point (50,50). Again both the entire room and a more detailed plot of the center of the room are shown, as well as the difference of just the amplitudes.

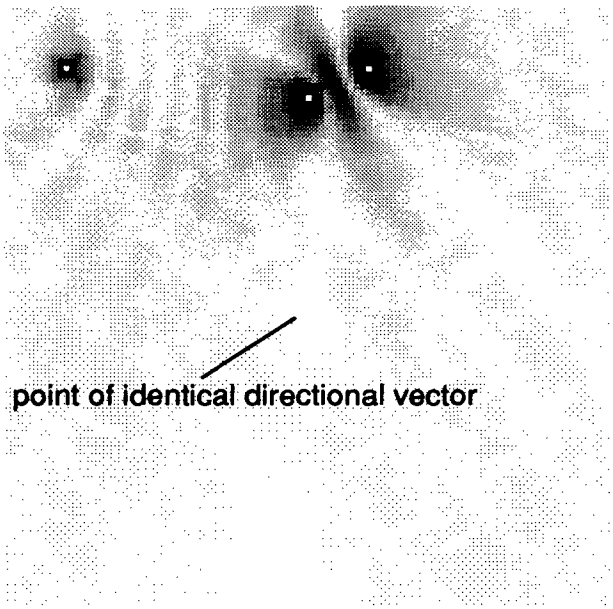

Figure 1.6: Amplitude of the difference using method of Romano

It appears that the region around the point at (50,50) where the amplitude of the difference is small covers more area than the Schroeder and Atal example. However that is to be expected since twice as many speakers are being spent to get aspects of this point

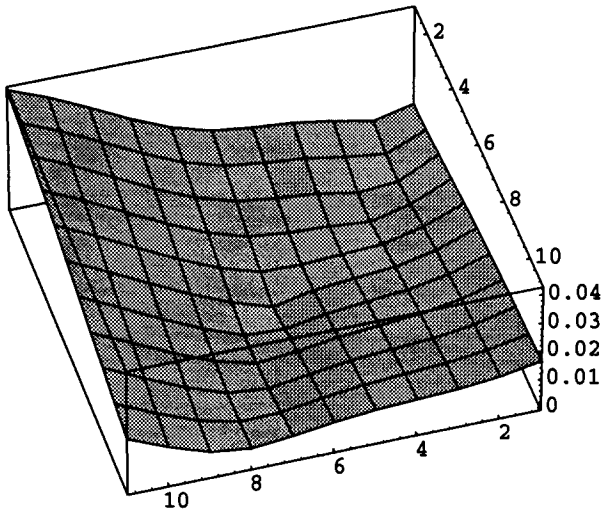

Figure 1.7: Amplitude of the difference at the center of the room

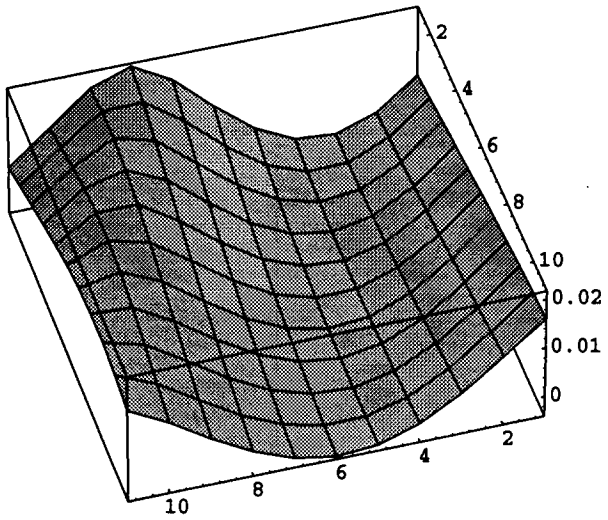

Figure 1.8: Difference in the amplitudes at the center of the room

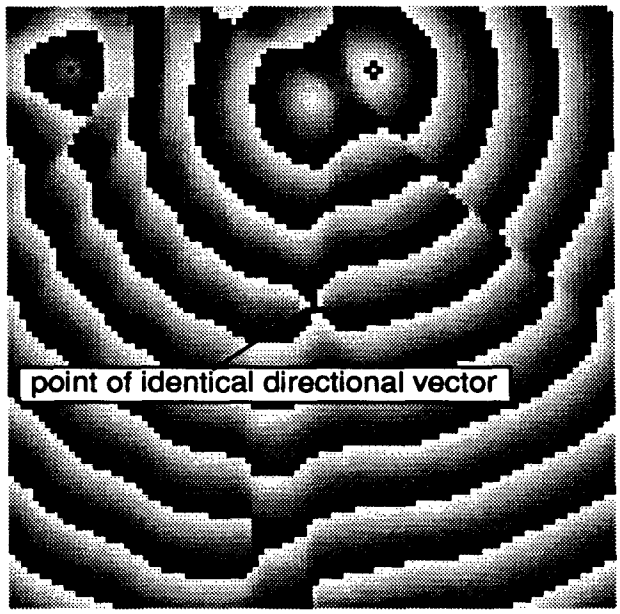

Figure 1.9: Phase of the difference using method of Romano

correct. The pinch in the phase diagram shows how small the point is where the directional wavefront is reproduced accurately. In this picture produced by the simulation environment, the shades from white to black indicate the range from $-\pi$ to π, and the principle branch of the phase is made to lie in that region. In fact, overlaying the phase portrait of the target field and the reconstruction for several successive 'times' indicates that the direction of propagation is similar only in a very small region around that point. Nearby, the directional vectors are very different. However there is a region of very good approximation to the desired phase and amplitude around the controlled point. Thus, forcing the directional vector to be correct at a given point seems to benefit the surrounding area with respect to amplitude and phase error, without making the region of approximately correct directional vectors extend very far.

A final point should be made, that in order for Romano's method to work, the speakers must be within 90 degrees of the source, since any greater angle would project the directional component of the speaker in the opposite direction as the source. This will correspond to a phase inversion of that directional component in the speaker which does not correspond to changing the direction of the wavefront. Thus, if any component is further than 90 degrees away, its inverted phase directional component will not contribute properly to the sum determining the directional component at the point. A simple example is this: given a speaker on the positive x-axis and a speaker on the positive y-axis, try to produce a wavefront direction at the origin going toward the speaker on the x-axis. Since the y-axis speaker is orthogonal, it can contribute nothing, and the x-axis speaker will always have a wavefront direction going the wrong way at the origin.

Intuitively it seems that the more points or parameters which are gotten exactly right in a listening environment, the better the result at the other locations since there is not very much room to go wrong. This does in fact turn out to be the case, but in order to know how much error is possible between these perfect points, further analysis is required. The

next section examines the case when the spatial frequency is band-limited. Unfortunately the spatial signal is usually not band-limited, and the following section on aliasing deals with these issues.

In order to produce the correct phase and amplitude only one speaker is needed for each point in space and this less strict approach is put forth as a basis for this dissertation. It can be shown that in the limit, three speakers per point in space are not needed to capture vector quantities. In fact, using multiple speakers to get additional properties correct at a point in space can be seen as a special limiting case of the one point per speaker approach.

Sakamoto, et al.

Sakamoto, et al., in [85] controlled localization by manipulating the filters into a form where both speakers have a common filter and only one speaker has an additional filter. This is described as controlling the sound pressure common to both ears by the first filter while the second controls the ratio of sound pressure between the ears by affecting one of the outputs. They showed that the solution to this is obtained by first solving for the filter which will be used on one channel to control the ratio of pressures between the ears. If the transfer matrix from an original source to the ears is

$$\begin{bmatrix} H_{\phi L} \\ H_{\phi R} \end{bmatrix} \cdot OS = \begin{bmatrix} P_L \\ P_R \end{bmatrix} \tag{1.17}$$

and the transfer matrix from the speakers to the ears is

$$\begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \begin{bmatrix} SP_L \\ SP_R \end{bmatrix} = \begin{bmatrix} P'_L \\ P'_R \end{bmatrix} \tag{1.18}$$

then setting $\frac{P'_L}{P'_R} = \frac{P_L}{P_R}$ causes

$$\frac{SP_L}{SP_R} = H = \frac{H_{11}H_{\phi R} - H_{12}H_{\phi L}}{H_{11}H_{\phi L} - H_{12}H_{\phi R}} \tag{1.19}$$

the common filter can be found after some algebraic manipulation as

$$G = SP_L = \left\{1 + \sum_{n=1}^{\infty} \left(\frac{H_{12}}{H_{11}}\right)^{2n}\right\} \left\{1 - \left(\frac{H_{12}}{H_{11}}\right)\left(\frac{H_{\phi 2}}{H_{\phi 1}}\right)\right\} \left(\frac{H_{\phi 2}}{H_{11}}\right) \tag{1.20}$$

The infinite sum is caused by a recursion, but because this filter feeds both ears it was found that the approximation $G \approx \frac{H_{\phi 1}}{H_{11}}$ does not harm localization. Application was made to conventional stereo inputs as well as certain quadraphonic schemes to improve sound image localization, but the technique seems better suited for its own coding scheme which would take full advantage of the separate control of a transfer function common to both ears and a ratio of the sound pressure levels at the ears.

M. Yanagida et al.

In [100], application of the least squares approach to soundfield reconstruction is described which, following in the approach of Schroeder and Atal [88], is formulated in terms of right and left ears but here the technique is to use a generalized inverse. Given a problem $Ax = b$ (A is a matrix, b and x are vectors, all of which can be complex) where b contains more

entries than x, (the overdetermined case) there may be no exact solution. However, if A^* represents the conjugate transpose, there almost always will be a solution to $A^*Ax = A^*b$ which is the generalized inverse corresponding to a least squares solution. i.e.

$$x = (A^*A)^{-1}A^*b \tag{1.21}$$

This solution gives the vector which minimizes the total squared error with respect to the original problem $Ax = b$ and can be called the generalized inverse in the overdetermined case. In the case with twice as many speakers as people, the least squares approximation would not be needed and each person could be made to perceive identical, or even arbitrary sensations. With more speakers than ears, the equations become underdetermined and this case is not treated here. The formulation used here is simply $\hat{C} = T^+H$ where C is a vector length n describing the output of the speakers and H of length $2m$ and is the transfer function vector from the original source to everyone's ears and T^+ is the generalized inverse of a matrix of transfer functions from the speakers to everyone's left and right ears and is arranged accordingly.

$$C = \begin{bmatrix} C_1 \\ \vdots \\ C_n \end{bmatrix} \qquad H = \begin{bmatrix} H_L \\ \vdots \\ H_L \\ H_R \\ \vdots \\ H_R \end{bmatrix} \qquad T = \begin{bmatrix} T_{11L} & T_{12L} & \cdots & T_{1nL} \\ \vdots & \vdots & & \vdots \\ T_{m1L} & T_{m2L} & \cdots & T_{mnL} \\ T_{11R} & T_{12R} & \cdots & T_{1nR} \\ \vdots & \vdots & & \vdots \\ T_{m1R} & T_{m2R} & \cdots & T_{mnR} \end{bmatrix} \tag{1.22}$$

Thus, the initial goal is for everyone to perceive the source as being in the same location relative to them.[7] Mention is made that the transfer function H from the original source to the ears can be modified for different sound localizations among individuals. In [100], no mention is made of the problems involved with finding the generalized inverse of a matrix of filters. In the overdetermined case where there are more ears than speakers, the generalized inverse will involve conjugates of s or z and not be analytic. Stability is another important issue. Nonetheless, finding the generalized inverse is not a problem when done at just one frequency. As an example of a least squares solution at one frequency (corresponding again to a wavelength of ten units) over the same 121 sample points in the middle of a room from the same two loudspeakers used in previous examples is shown in figures 1.10 and 1.11. Here, each sample in the target is determined by sampling the original source at the same points.

Bauck and Cooper

Transaural refers to the original method of Schroeder and Atal [88] and has continued to be developed. Recently, papers by Bauck and Cooper [34] [35] have generalized transaural techniques and developed accompanying theory. In [34] further improvements on the filter structure of Schroeder and Atal are made. Shuffler filters are designed to simplify the signal processing requirements of crosstalk cancelation. The Atal-Schroeder formulation is shown to be equivalent to a lattice filter, which can further be made equivalent to a sequence of shuffler filters.

[7] It is interesting to digress for a moment to note that to have each listener perceive a localized source with the same position relative to that listener is an example which could not occur in real life and shows the flexibility of soundfield reconstruction

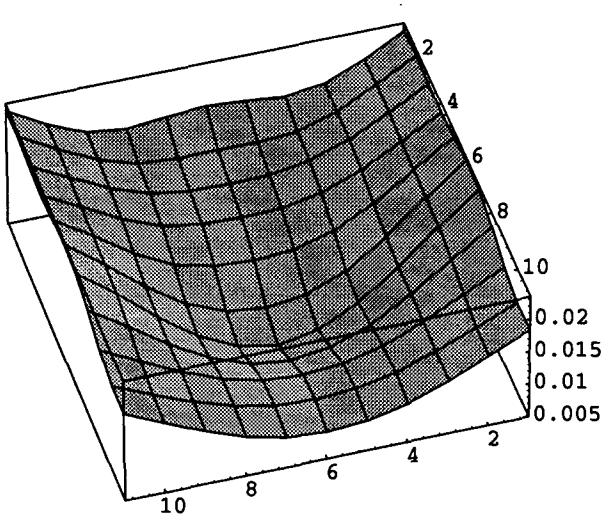

Figure 1.10: Amplitude of the difference at the center of the room

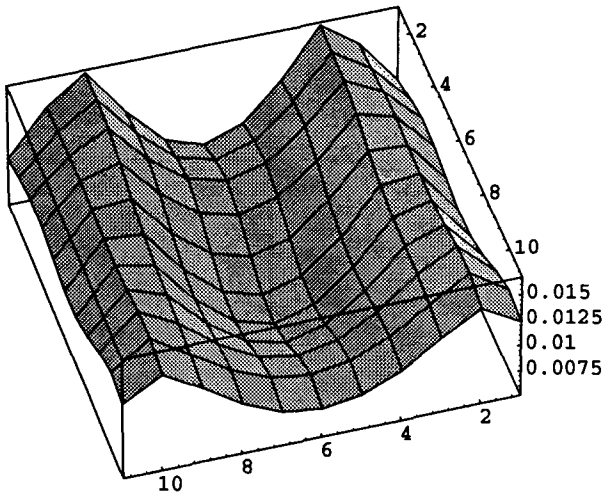

Figure 1.11: Difference in the amplitudes at the center of the room

$$A' = -\frac{H_{12}}{H_{11}^2 - H_{12}^2} \qquad S' = \frac{H_{11}}{H_{11}^2 - H_{12}^2} \qquad N' = \frac{S' - A'}{2} \qquad P' = \frac{S' + A'}{2} \tag{1.23}$$

These new filters can simplify implementations as can be seen in the equivalent matrix formulation below. The same forms can be obtained by taking the inverse of the transfer matrix from the speakers to the ears, inverting and finding a certain factorization.

$$H = \begin{bmatrix} H_{11} & H_{12} \\ H_{12} & H_{11} \end{bmatrix} \qquad H^{-1} = \frac{1}{H_{11}^2 - H_{12}^2} \begin{bmatrix} H_{11} & -H_{12} \\ -H_{12} & H_{11} \end{bmatrix} \tag{1.24}$$

$$H^{-1} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} \frac{1}{2(H_{11}+H_{12})} & 0 \\ 0 & \frac{1}{2(H_{11}-H_{12})} \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{1.25}$$

This formulation of the filter matrix was made more general and applied to several practical examples in [35] though such fortunate factorization do not occur for many situations. Here, the idea of using a generalized inverse is treated in some detail. Both overdetermined and underdetermined systems are analysed and use is made of the singular value decomposition for reducing the rank of singular matrices. The general formulation given includes a separate transfer matrix from the input channels to the speakers in addition to the transfer matrix describing the how the sound gets from the speakers to the ears. If the transfer from the speakers s to the ears e is $e = Xs$ and the transfer from the channels p to the speakers is $s = Yp$ then the transfer from the channels to the ears is $e = Zp$ where $Z = XY$. To solve these equations of the form $Ax = y$, use is made of the Moore-Penrose inverse or pseudo-inverse. This is the same as the generalized inverse used in [100] by Yanagida et al. In the case where the equations are overdetermined, where there are more equations than parameters, an exact solution may not exists and the band-limited, A^+ is

$$A^+ = (A^H A)^{-1} A^H \tag{1.26}$$

where A^H is the Hermitian transpose. If $e = Zp$ is underdetermined, the least squares solution is found by using the band-limited A^+ where

$$A^+ = A^H (A A^H)^{-1} A^H \tag{1.27}$$

which gives a minimum norm solution, that is, a solution where $||x||$ is minimum. When $rank(A) < \min(m, n)$ where A is $m \times n$, use can be made of the singular value decomposition, in which A is factored into three matrices $A = U \Sigma V^H$. If Σ^+ is defined by taking the replacing only the nonzero values of Σ by their reciprocals, then the pseudo-inverse of A can be defined as $A^+ = V \Sigma^+ U^+$. Mention is made that other approaches may also be beneficial in this case.

In the underdetermined case, further constraints on the solution can be added without degrading the transaural effect. Some other constraints mentioned are using minimum power to the speakers, and finding solutions, where the channels to the speakers can be coded easily. One example of the latter is the situation where there are three speakers, right, middle and left, and one listener. In this case, solutions can be found where the signals applied to the speakers are the sums and differences of two stereo channels. Another interesting approach is the use of dipole and monopole speakers. A solution is found by which one monopole speaker is used for an entire audience where each member has a small dipole speaker behind their head.

Nelson

In general, it is not the case that there will be enough speakers to adequately cover an area even with only one speaker per point in space. Philip Nelson, who approaches the problem from a background of noise control, takes more care about the problematic aspects of taking the Hermitian transpose of filter matrices which arise in this overdetermined case. The first part of [73] looks to the Kirchhoff-Helmholtz integral theorem as motivation toward soundfield reproduction by giving a solid theoretical basis for demonstrating that arbitrarily perfect reconstruction is possible in theory. However, the approach actually taken is similar to the generalization of Schroeder and Atal's approach to N speakers and M ears but concentrating on the situation where $M \geq N$ (whereas Bauck and Cooper [35] focused more attention on the situations where $M \leq N$ although they also deal with the overdetermined case). Nelson also separates the notion of sample points from the location of ears, which is an important idea not explicitly expressed in the previous 'transaural' sources, although the idea of sampling a soundfield is certainly not new. Thus, the issue of sampling interval becomes important, and Nelson states that sampling at 1/2 the highest wavelength of concern is sufficient, though higher spatial sampling is preferred. The treatment is put in terms of recreating the soundfield around a single listener in an anechoic environment using several speakers arranged in a circular pattern about the square region. He is also one of very few to consider the problem of soundfield reconstruction as one of optimizing certain criteria. His optimizations use a least squares criterion, and include a weighting for the amount of power used to drive individual loudspeakers. Furthermore, the problem of localized sound is reduced to approximating plane waves from the desired direction, which can be thought of as a perceptually based simplification. The cost function he used in [73] is

$$J(\omega) = e^H(\omega)e(\omega) + \beta v H(\omega)v(\omega) \tag{1.28}$$

where $e(\omega) = d(\omega) - \hat{d}(\omega)$ and $v(\omega)$ consists of the input voltages. β weighs the cost of the (squared) energy needed to drive the system while trying to minimize the squared errors.

The implementation of the filters described by the least squares solution in the frequency domain is problematic. The reason could be summarized as involving the non-analytic nature of the Hermitian transpose of a matrix of filters combined with non-conjugated matrices, but Nelson takes the approach of showing that this causes poles outside the unit circle. The issue of designing a stable causal filter matrix is treated by finding stable realizable filters which approximate the desired least squares response. Both adaptive FIR and IIR filters are looked at as possibilities. Issues raised about sampling and implementation of filters will be touched on later.

In addition, [74] seems to show that the configuration of the speaker array plays a role in least squares approximation of soundfields similar to what one might intuit, namely that an equally spaced circular array allows the best reconstruction of plane waves of arbitrary direction. This configuration has intuitive appeal as well, since an equally spaced circular arrangement has nice properties of minimizing the largest angle between speakers and making sure that no speaker is farther or closer to the listening environment center, assuming the center of the speaker circle coincides with the environment's center.

1.3.3 The Kirchhoff-Helmholtz Integral Theorem

For larger speaker arrays, use is often made of a powerful theorem which can be interpreted as saying that the soundfield in a region with no sources can be derived from knowledge about the soundfield on the boundary of the region. This makes sense intuitively since with no sources in the region, any sound in the region must have entered through the boundary. What is amazing is how relatively uncomplicated the relationship is, relying just on information about the pressure fluctuations at the boundary and the derivative of the pressure fluctuations in the direction normal to the boundary. Thus, if the number of speakers is large enough to approximate a sampled boundary surface, these relationships become very useful.

The development of the integral theorem from Green's theorem given here is based on the development given in [43]. Given a wavefield in space dependent on time $U(x,y,z,t)$ which satisfies the wave equation $\nabla^2 U - \frac{1}{c^2}\frac{\partial^2 U}{\partial t^2} = 0$, and the related Helmholtz equation $(\nabla^2 + k^2)U = 0$ the value at a particular point P_0 in the field can be calculated from values on a surface which contains that point in its enclosed volume. In particular, if U is composed of spherical and planar waves, these conditions are satisfied. To compute a value at a point inside S, use is made of Green's Theorem which states:

If U and G are functions of position and their first and second derivatives are single-valued and continuous on a surface S surrounding volume V, then if n is the outward normal vector from S, it is the case that

$$\iiint_V (G\nabla^2 U - U\nabla^2 G)dv = \iint_S \left(G\frac{\partial U}{\partial n} - U\frac{\partial G}{\partial n}\right) ds \tag{1.29}$$

To use this tool intelligently, an appropriate function G (called the Green's function) must be chosen. If r_{01} is the distance between P_0 and P_1, then by letting $G(P_1) = \frac{e^{jkr_{01}}}{r_{01}}$ the partial derivative is

$$\frac{\partial G(P_1)}{\partial n} = \cos(n, r_{01})\left(jk - \frac{1}{r_{01}}\right)\frac{e^{jkr_{01}}}{r_{01}} \tag{1.30}$$

Also, G satisfies the wave equation and substituting into the wave equation produces the Helmholtz equation

$$(\nabla^2 + k^2)G = 0 \tag{1.31}$$

Since G is undefined at P_0, it is necessary to define the additional surface S_ϵ which is a sphere with radius ϵ around the point P_0. In applying Green's Theorem, the volume (V') used will be that contained in S minus that contained in S_ϵ and the surface of integration is $S' = S + S_\epsilon$. Note that the unit normal on S_ϵ is pointing inward, away from V'. Substituting in to the right half of the Green's theorem formula produces

$$\iiint_{V'} (G\nabla^2 U - U\nabla^2 G)dv = -\iiint_{V'} (GUk^2 - UGk^2)dv = 0 \tag{1.32}$$

so

$$\iint_{S'} \left(G\frac{\partial U}{\partial n} - U\frac{\partial G}{\partial n}\right) ds = 0 \tag{1.33}$$

implying

$$-\iint_{S_\epsilon}\left(G\frac{\partial U}{\partial n}-U\frac{\partial G}{\partial n}\right)ds=\iint_{S}\left(G\frac{\partial U}{\partial n}-U\frac{\partial G}{\partial n}\right)ds \tag{1.34}$$

For P_1 on the surface S_ϵ, $\cos(n,r_{01})=-1$ and $G(P_1)=\frac{e^{jk\epsilon}}{\epsilon}$ and $\frac{\partial G(P_1)}{\partial n}=\frac{e^{jk\epsilon}}{\epsilon}\left(\frac{1}{\epsilon}-jk\right)$ so as ϵ goes to 0

$$\iint_{S_\epsilon}\left(G\frac{\partial U}{\partial n}-U\frac{\partial G}{\partial n}\right)ds=4\pi\epsilon^2\left[\frac{\partial U(P_0)}{\partial n}\frac{e^{jk\epsilon}}{\epsilon}-U(P_0)\frac{e^{jk\epsilon}}{\epsilon}\left(\frac{1}{\epsilon}-jk\right)\right]=-4\pi U(P_0) \tag{1.35}$$

Yielding the Kirchoff-Helmholtz Integral Theorem

$$U(P_0)=\frac{1}{4\pi}\iint_S\left[\frac{\partial U}{\partial n}\left(\frac{e^{jkr_{01}}}{r_{01}}\right)-U\frac{\partial}{\partial n}\left(\frac{e^{jkr_{01}}}{r_{01}}\right)\right]ds \tag{1.36}$$

Even if the waves of concern come through only a small part of the surface S, they will still have an effect by exiting the surface which would need to be included in the above formula. However, if the wave propagation satisfies the *Sommerfeld radiation condition*, namely that

$$\lim_{R\to\infty}R\left(\frac{\partial U}{\partial n}-jkU\right)=0 \tag{1.37}$$

then the contribution from the exiting wave can be made 0 by increasing the far surface distance from P_0 to infinity [43]. Thus, the surface integration can take place on a plane between the source and P_0. This integration can be further reduced by applying the Kirchhoff boundary conditions. Given a finite area Σ on the plane S between the source and P_0, U and $\frac{\partial U}{\partial n}$ are unaffected on σ and 0 elsewhere on S. Note that these conditions are mathematically inconsistent since it can be shown that if a function and its normal derivative are both 0 together then the entire function must be equal to 0. Still, these conditions can be arbitrarily close to true.

This mathematical inconsistency can be removed by a different choice of Green's function, such that either G or $\frac{\partial G}{\partial n}$ vanish everywhere, for instance the Rayleigh-Sommerfeld formulation uses $G\_(P_1)=\frac{e^{jkr_{01}}}{r_{01}}-\frac{e^{jk\tilde{r}_{01}}}{\tilde{r}_{01}}$ where $\tilde{r}_{01}=dist(\tilde{P}_0,P_1)$

$$U(P_0)=\frac{1}{4\pi}\iint_\Sigma\left(\frac{e^{jkr_{01}}}{r_{01}}\right)\left[\frac{\partial U}{\partial n}-U\cos(n,r_{01})\left(jk-\frac{1}{r_{01}}\right)\right]ds \tag{1.38}$$

Note that when $k\gg\frac{1}{r_{01}}$ then $\cos(n,r_{01})\left(jk-\frac{1}{r_{01}}\right)\approx jk\cos(n,r_{01})$. If $U(P_1)=\frac{Ae^{jkr_{21}}}{r_{21}}$ and $k\gg\frac{1}{r_{21}}$ then

$$U(P_0)\approx\frac{A}{j\lambda}\iint_\Sigma\frac{e^{jk(r_{21}+r_{01})}}{r_{21}r_{01}}\left[\frac{\cos(n,r_{01})-\cos(n,r_{21})}{2}\right]ds \tag{1.39}$$

which is the Fresnel-Kirchhoff diffraction formula for a single point source. Since in audio neither assumption is necessarily valid, more care must be taken with this formula than is done in optics.

Berkhout

The second approach, referred to as the integral approach, has its roots in the Kirchhoff-Helmholtz integral theorem. In [15], A.J. Berkhout expounds on a sound system called ACS (Acoustic Control System) which makes use of a holographic approach to control the direct soundfield as well as providing separate control over the reverberant field. Berkhout's formulation of the reconstruction function is taken from his background in seismic waves, where e^{-j} is used in contrast to other fields (such as electrical engineering) which tend to use e^{j} and vice versa. With this in mind, Berkhout's formulation is

$$P(x,y,z,\omega) = \iint_{S_1} P(x_1,n,q,\omega)(1+jkr)\cos\phi\frac{e^{jkr}}{r^2}dndq \tag{1.40}$$

where $x > x_1$ $k = \frac{\omega}{c}$ $r = dist[(x,y,z),(x_1,n,q)]$ and $\cos\phi = (x-x_1)/r$. This is made discrete by

$$P(x,y,z,\omega) = \sum\sum_{S_1} P(x_1,y_m,z_n,\omega)(1+jkr_{mn})\cos\phi_{mn}\frac{e^{jkr_{mn}}}{r_{mn}^2}\delta y\delta z \tag{1.41}$$

Berkhout requires $\delta x = \delta y \leq \lambda/2$ (actually exact reconstruction is possible as long as both δx and δy are less that $\lambda/2$) and states that sampling should be denser at higher frequencies.[8] In ACS, the wave is sampled by a planar microphone array, and extrapolated to a further plane to produce the output to the loudspeakers.

As a demonstration of Berkhout's approach, consider Figure 1.12. Here 25 speakers are used rather than 2. Unfortunately, the simulation environment currently only supports monopole sources, so the requirement by Berkhout that the sources be dipole is not fulfilled. The output of the speakers has been put $\pi/2$ radians ahead of the original source so that the resulting wavefront is in phase for comparison.[9]

Vogel

A similar approach seems to have been taken by Peter Vogel [94]. In the abstract, he states that the "Kirchhoff-Helmholtz representation theorem" is used toward "an amplification system for direct sound that can be used in theatres to enhance the actors' voices, without disturbing the localization cues in the soundfield" [94]. Use is made of simplifications based on human hearing to allow good localization even with wider speaker intervals (presumably wider than half the wavelength of the highest frequency) though some coloration is produced. The hall response is also taken into account.

Use of Time Reversal Mirrors

Cassereau, Chakroun, Wu and Datchi took an interesting approach to wavefront reconstruction by the use of time-reversal to get the inputs to the speaker array [27]. Using ultrasonic transducers, they were able to approximate a wavefront in the near field using both a full array of 1024 transducers and a sparse array the same size with only 128 randomly placed

[8] Exactly how dense the sampling needs to be is discussed in the next chapter but it turns out to be quite high in many cases.

[9] Note that an alternative Green's function can be choosen to eliminate the dipole sources and leaves only a monopole array. In this case the obliquity factor which was part of the directivity pattern of the dipoles is replaced by an adjustment to the monoples resulting from keeping the normal derivative to the array surface instead of the direct signal from the original source to the array.

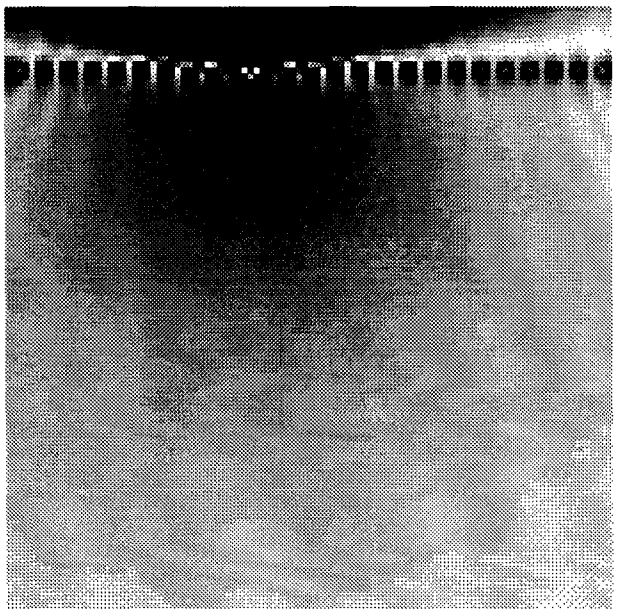

Figure 1.12: Amplitude of the difference using method of Berkhout (resolution 30)

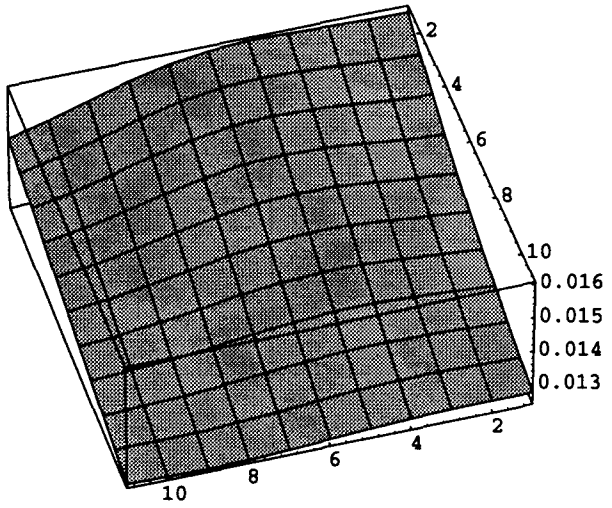

Figure 1.13: Amplitude of the difference using method of Berkhout (center of the room)

transducers. Their formulation of the problem in terms of Kirchhoff diffraction theory began with

$$p(\hat{r},t) = \int_{\Pi} \left[g(\hat{r},\hat{r}_s,t) \overset{*}{t} \partial_s p(\hat{r}_s,t) - \partial_s g(\hat{r},\hat{r}_s,t) \overset{*}{t} p(\hat{r}_s,t) \right] d^2\hat{r}_s \tag{1.42}$$

where $\overset{*}{t}$ represents time domain convolution. The Green's function is

$$g(\hat{r},\hat{r}_s,t) = g_f(\hat{r}-\hat{r}_s,t) - g_f(\hat{r}-\hat{r}_s,t) \tag{1.43}$$

where

$$g_f(\hat{r},t) = \frac{1}{4\pi|\hat{r}|}\delta\left(t - \frac{|\hat{r}|}{c}\right) \tag{1.44}$$

which was chosen to have the property that $g(\hat{r},\hat{r}_s,t) = 0$ and $\partial_s g(\hat{r},\hat{r}_s,t) = 2\partial_s g_f(\hat{r}-\hat{r}_s,t)$ for $\hat{r}_s$ on Π.

Solving for the transducer array gives

$$p(\hat{r},t) = K(z_0 - z)\left[\frac{1}{c^2 t^2} h(\hat{r},t) + \frac{1}{c}\frac{\partial}{\partial t}\left(\frac{1}{ct} h(\hat{r},t)\right)\right] \tag{1.45}$$

$$= \frac{K(z_0 - z)}{c^2 t}\frac{\partial}{\partial t} h(\hat{r},t) \tag{1.46}$$

In the case presented the impulse response function $h(\hat{r},t)$ corresponds to diffraction by a circular aperture (where the synthesized wavefront is desired) backward to the transducer array. The signals appearing at the array from such a diffraction can be time reversed and played back to approximate the desired wavefront, a circular impulse on the plane $z = z_0$.

In essence, what happens is that the desired wavefront is propagated back to the speakers from where it should originate, and this signal is recorded on a separate channel for each speaker. These recorded signals are then time reversed and played through the speaker to obtain the wavefront. In this case, the analytic solution of diffraction through a circular aperture was used to project the desired wavefront onto the speakers, but often this is achieved as the result of an echo off some object (such as a kidney stone). The process can then be repeated iteratively to provide a improved focus on the target. Good results were obtained when this was done for a full array of 1024 transducers and a sparse array of 128 transducers randomly distributed over the same area.

A more complete description of time-reversal is given in [39] and [28]. These describe time reversal as an iterative process which can be applied to non-homogeneous material, such as the bodily tissues to which ultrasound is applied in medical applications. The backward transmission causes signal arrivals at the speakers to be distorted in such a way that time reversal compensates for the distortion. This makes sense, since the slowest signals arrive at the speakers the latest, so that in the reproduction mode these signals come out the earliest to compensate for the slower path they must take. While the iterative use of time reversal has application for transmission through such material for the purpose of focusing sound, it is unclear how useful it might be in a more homogeneous medium such as the air at audio frequencies. Nonetheless, it provides a useful technique and perspective on the problem of creating a desired soundfield.

Use of Sound Focusing

In addition to these important papers, there have been others using the quantitative approach more loosely. One of these is [57], which proposes that focusing sound in front of a speaker array by using simple delays can create the impression of a sound coming from the focused region. The relationship of this approach to the others is a close one. Similar in principle to the time-reversal approach, a delay can be provided to speakers in an array so that the sounds arrive "focused" at a certain point, in the sense that at that point all the signals are time synchronized. For instance, to project a sound in front of the array, delays are used so that sound from the edge of the array leave first to compensate for the longer distance they have to travel to the point of focus. Although this paper is not mathematically sophisticated, the simple idea of using delay to "focus" the sound at a "real image" was subjectively evaluated, and was reasonably effective. One explanation offered was that the focused soundfield had similar characteristics when compared to a spherical field emanating from the point of focus with respect to equal phase and equal amplitude curves.

1.4 Additional Possibilities for the Kirchhoff-Helmholtz Integral

Another possible application of the Kirchhoff-Helmholtz integral theorem is to reconstruct the pressure field and the normal derivative field on a boundary away from the speakers. The approach of [15] was to use the speakers as sources on the boundary, and by using a particular Green's function the solution was formulated just in terms of dipole sources representing the correct normal derivative of the field. The biggest limitations of this approach are that the sources must be dipolar at all frequencies and must be in a linear arrangement, or possibly some other if a Green's function exists to eliminate the monopole component. By reconstructing the pressure and normal derivative fields away from the speakers, the directivity pattern of the speakers and their arrangement becomes flexible. However, since each speaker essentially buys one extra parameter, up to twice as many speakers may be required to get the same resolution on the boundary as Berkhout's approach.

Achieving the correct pressures at particular points in space has already been discussed. Achieving the correct normal derivative can be done in a way similar to Romano's approach, which produced the correct wavefront direction in some circumstances. What is needed in the Kirchhoff-Helmholtz integral formulation is the correct rate of change of the pressure on the surface in the direction normal to the surface. Thus, things are even more flexible, since only the derivative in one direction is required.

1.5 Comparison of Approaches

It should be noted that these approaches are not the only ones possible. Other people have also done work in similar veins, but these were chosen to represent a few of the possibilities. Although approaches based on matrices of transfer functions and those based on an integral formulation seem on the surface to be different, they are in fact related since the way sound propagation works is the same and by necessity any integral formulation must be made discrete for discrete speakers. In [16] a matrix formulation of the Kirchhoff-Helmholtz integral is given by Berkhout: Consider the 2-dimensional Kirchhoff-Helmholtz integral along a straight line boundary (here z is orthogonal to x)

$$P_A = \frac{1}{4\pi} \int \left[P \frac{\partial G}{\partial z} - G \frac{\partial P}{\partial z} \right] dx \tag{1.47}$$

The discrete version is

$$P(z_m) = \frac{1}{2} G_{01}(z_m, z_0) P(z_0) - \frac{1}{2} G_{00}(z_m, z_0) \partial_z P(z_0) \tag{1.48}$$

where $G_{00}(z_m, z_0)$ is the propagation matrix for (x_i, z_0) to (x_i, z_m) $G_{01}(z_m, z_0)$ is the derivative with respect to z $P(z_0)$ is the (monochromatic) pressure at (x_i, z_0) $\partial_z P(z_0)$ is the (monochromatic) z-derivative of pressure at (x_i, z_0) $P(z_m)$ is the (monochromatic) pressure at (x_i, z_m) taking the derivative with respect to z at z_m gives

$$\partial_z P(z_m) = \frac{1}{2} G_{11}(z_m, z_0) P(z_0) - \frac{1}{2} G_{10}(z_m, z_0) \partial_z P(z_0) \tag{1.49}$$

where $G_{10}(z_m, z_0)$ is the z-derivative of $G_{00}(z_m, z_0)$ at z_m $G_{11}(z_m, z_0)$ is the z-derivative of $G_{01}(z_m, z_0)$ at z_m Combining gives

$$\begin{bmatrix} P \\ \partial_z P \end{bmatrix}_{z_m} = \frac{1}{2} \begin{bmatrix} G_{01} & -G_{00} \\ G_{11} & -G_{10} \end{bmatrix}_{z_m \leftarrow z_0} \begin{bmatrix} P \\ \partial_z P \end{bmatrix}_{z_0} \tag{1.50}$$

This matrix differs from the matrices which developed out of the transaural approach in that account is taken of the partial derivative with respect to the surface boundary of the region in question. This is needed not to create the correct pressure field on the surface, which can be done by just using propagation matrix G_{01}, but to insure that the interior of the region has the correct signal everywhere. The effect of this discretization of the continuous surface integral will be examined in the next chapter.

It is interesting to note that the approach of Romano was to create the correct directional vector as well as magnitude of pressure at certain points. A similar approach could be applied to get the correct directional derivative in the normal direction to the surface. This would require just two speakers per point, which makes sense since only one additional parameter is required and because the directional derivative part represents a dipole on the surface rather than a monopole. This in turn brings to mind the stereo formulation used in [85] where a common signal between the ears was found, and a separate filter was used to produce the correct ratio of pressures between the ears. Here one speaker can provide a correct pressure at a point, and an additional speaker can provide the correct directional derivative in one direction.

Another interesting similarity to note is the apparent relationship between time-reversal and least-squares solutions in the overdetermined case. Consider the least-squares formula

$$x = (A^* A)^{-1} A^* y \tag{1.51}$$

which can be rearranged as

$$A^* A x = A^* y \tag{1.52}$$

Here A represents the transfer matrix from the speakers to the sample points in the listening environment. A^* is the conjugate transpose (Hermitian transpose) which can be interpreted as the propagation from the sample points to the speakers, reversed in time due to the conjugation which inverts the phase, but still having the normal $\frac{1}{distance}$ attenuation through

space. With this interpretation the above formula says that the time-reversed propagation from the points to the speakers (A^*y) can be made the same for some x as the propagation of that x from the speakers to the points forward in time and back again backward in time (A^*Ax). The actual time-reversal approach propagated the target signal y back to the speakers and used that as the solution, i.e. $\hat{x} = A^*y$, with the hope that $A^*A\hat{x} \approx A^*y$. This will be the case if $A^*A \approx 1$ or $A^* \approx A^{-1}$. Iterating this process will work if

$$\lim_{n\to\infty} (A^*A)^n \to 1 \tag{1.53}$$

Chapter 2

A General Framework

2.1 Two Necessary Procedures

The recreation of a soundfield or a set of audio perceptions in a new listening environment requires at least two steps, or procedures.

Procedure 1: First, there must be a way to know about the desired soundfield. Ideally, the soundfield will be known at least as well as is needed.

Procedure 2: Second, a procedure should exist for approximating the appropriate aspects of the known factors of the desired soundfield in a new environment.

The dichotomy between these two can be seen even in the most simple mathematical model. Even in the case where

$$\begin{bmatrix} F_{1,1} & F_{2,1} \\ F_{1,2} & F_{2,2} \end{bmatrix} \begin{bmatrix} SP_1 \\ SP_2 \end{bmatrix} = \begin{bmatrix} P_1 \\ P_2 \end{bmatrix} \tag{2.1}$$

where $x = \begin{bmatrix} SP_1 \\ SP_2 \end{bmatrix}$ represents two speaker outputs,

$y = \begin{bmatrix} P_1 \\ P_2 \end{bmatrix}$ represents two target points and

$A = \begin{bmatrix} F_{1,1} & F_{2,1} \\ F_{1,2} & F_{2,2} \end{bmatrix}$ represents the transfer matrix from the speakers to the points,

the soundfield y must be known in procedure 1, and A^{-1} must be known in procedure 2 to calculate x from y. This general process is diagramed in Figure 2.1 where the procedure 1 receives signals and information in the form of measurements or actual recordings from an original environment or from a simulation, and calculates the desired soundfield. In this model, procedure 1 knows nothing about the listening environment where this information will be used. This decision is based on the desire to be able to use the target sound information in a variety of listening environments and therefore no information is included about any particular listening environment at this stage.

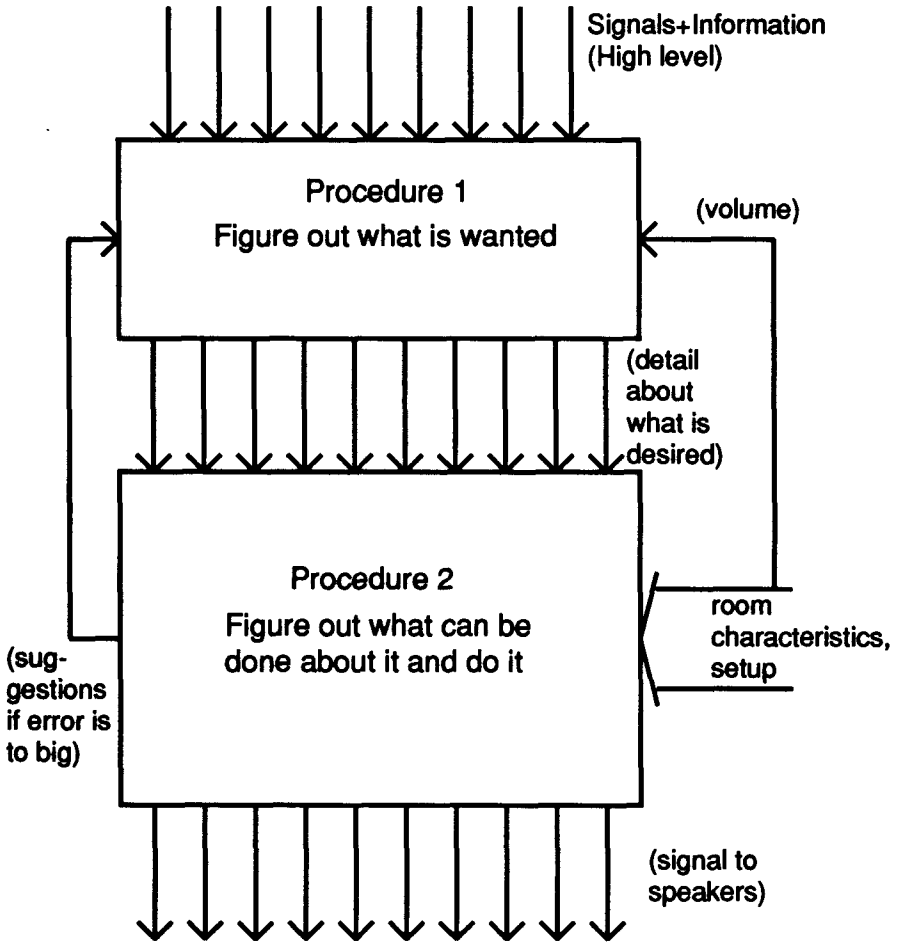

Figure 2.1: Showing how information might flow between two procedures

2.2 Knowledge of the Sound Field (Procedure 1)

Some information about the region of interest at frequencies for which reconstruction will be attempted must be known. Procedure 1 may be required to figure out what the nature of the desired field is from partial information and must put this information into a form that is useful. The partial information given to procedure 1 may be of many forms. Recorded samples or a set of measurements taken from the original environment, a set of options used in creating a computer simulation of a nonexistent environment, or a combination are among the possibilities. Ideally, procedure 1 should provide the most complete description of the desired soundfield qualities to allow procedure 2 the most flexibility in choosing a psychoacoustically relevant subset to emulate.

2.2.1 Target Generation

The output of procedure 1 must be in a form recognizable by procedure 2 which creates the best possible approximation in the listening environment. Procedure 1 does not necessarily have to calculate what the actual desired soundfield is. It just has to describe what is desired unambiguously. For instance, if a certain perception for an audience is desired, then the description of the target soundfield must include what is necessary for that perception, although some irrelevant information of the soundfield's structure may be ommitted.[1]

Procedure 1 can be thought of as forming some subset ITI of ideal target information about the original soundfield (whether real or not) from the set $I(SF)$ of all possible information about the original soundfield. In actuality, there will be error due to measurement inaccuracies and quantization. Although this might be more accurately described as noise added to ITI to create the actual target information ATI, it can also be viewed as a noise field N which is added to the original soundfield before the target information is taken. In this case the error is thought of as part of the target soundfield.

$$I(SF) \supseteq ITI \tag{2.2}$$

$$ITI + noise = ATI \tag{2.3}$$

$$I(SF + N) \supseteq ITI \tag{2.4}$$

Target information, once obtained, can be manipulated independently of the soundfields. Targets can be combined in various ways to create new targets without any need for corresponding soundfields in reality. Care must be taken that the target information is not self contradictory or problematic.

2.2.2 Information Content of a Soundfield

If a certain soundfield is to be reproduced exactly, the question arises as to how much information is required to describe it unambiguously up to a given frequency. A partial answer is given by the sampling theorem

Theorem: "A function $f(\vec{x})$ whose Fourier transform $F(\vec{\omega})$ vanishes over all but a finite portion of wave-number space can be everywhere reproduced exactly from its sample values taken over a lattice of points $l_1\vec{v}_1 + l_2\vec{v}_2 + \cdots + l_N\vec{v}_N, l_1, l_2, \cdots, l_N = 0, \pm 1, \pm 2, \cdots,$

[1] see Chapter 3 for examples where Procedure 2 creates new targets by selectively changing part of the original target's structure

provided that the vectors $\vec{v}_j$ are small enough to ensure nonoverlapping of the spectrum $F(\vec{\omega})$ with its images on a periodic lattice defined by the vectors $\vec{u}_k$, with $\vec{v}_j \cdot \vec{u}_k = 2\pi\delta_{jk}$." [79]

Strictly speaking, audio signals are not spatially bandlimited due to the 1/distance attenuation term, but any field not too close to a source is approximately bandlimited to the spatial frequency corresponding to the highest audio frequency present. If the desired signal consists of far away sources and the listening environment is quite some distance from the speaker, then this bandlimited requirement can be approximated since plane waves are bandlimited (for bandlimited audio signals) and the spherical waves of sources tend toward plane waves at larger distances. This spaital aliasing will be investigated in Appendix B. For now the effect is minor, and consists of slight variations in amplitude and phase between sample points which would not normally be noticable even if the error of the reconstruction was arbitrarily small.

Putting that aside for the moment, simple computations can be done which are most revealing. Since a 44.1k sampling rate translates into samples 0.77cm apart in space (based on the speed of sound = 340 m/s), a listening volume $1 \cdot 3 \cdot 3m^3$ would require 2.19 million spatial samples (if sampled rectangularly) to reproduce a bandlimited soundfield up to just over 20kHz, not counting a guard band. Even with the optimal equilateral triangle based sampling pattern 1.95 million samples would be necessary. However, thanks to the Kirchhoff-Helmholtz integral theorem described toward the end of the previous chapter, only the boundary of an area or volume must be considered, if the interior region contains no sources.

Assuming a constant speed of sound, the number and placement of the samples and the size of the listening environment together determine the highest spatial frequency. The highest spatial frequency in turn determines the highest audio frequency for which exact reconstruction would be possible if everything were bandlimited. A rule of thumb is that by rectangularly packing y samples in x N-dimensional units, the approximate distance between neighboring samples will be $(\frac{y}{x})^{\frac{1}{N}}$.

This rule of thumb works better with larger numbers of samples. For very few samples further analysis is of course required, but the rule remains somehow intuitive and may be of some use as an extrapolation. The Sampling Theorem shows that the smallest wavelength which can be unambiguously determined is just over two times the distance between samples. Considering a listening volume $1 \cdot 3 \cdot 3m^3$ we can use the rule of thumb to predict approximately what frequencies are detectable with varying approaches and numbers of samples. If we use rectangular sampling (which is suboptimal for higher dimensions but easier to start with) and fill the volume with samples, the result obtained is shown in Fig. 2.2. The label 'no. of speakers' on the graphs refers to the fact that for each additional speaker, one additional point in space can be controlled in theory. Thus, the number of samples required to determine the soundfield also has ramifications for how many speakers would be needed to reproduce the identical soundfield up to that frequency.

Unfortunately this is not very promising, although low frequencies may be achievable even with 2 or 4 speakers.[2] Even if much larger numbers of samples can be used, Fig. 2.3 shows how the benifit of this is minor.

Another approach to the same volume is to just place the samples on the boundary

[2]Based on the mentioned rule of thumb, and requiring further analysis since it is unclear how 2 or 4 speakers can be arranged in a cube

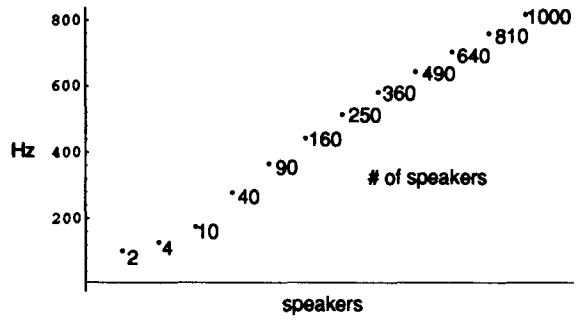

Figure 2.2: Approximate frequency attainable by N samples filling a $1 \cdot 3 \cdot 3m^3$ volume

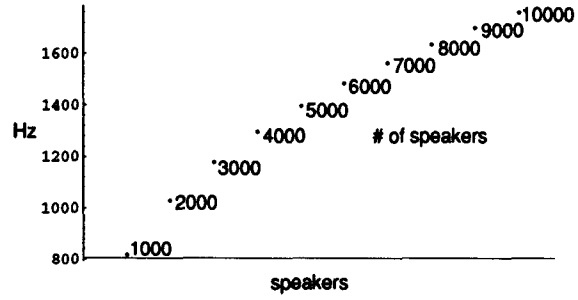

Figure 2.3: Approximate frequency attainable by N samples filling a $1 \cdot 3 \cdot 3m^3$ volume

surface. By the Kirchhoff-Helmholtz integral theorem, if there are no sources in the region, knowledge of the pressures and normal derivative at the boundary uniquely determine the soundfield inside. For consistency, the number of samples is in terms of pressure only. To sample the normal derivatives usually requires just as many speakers, and represents just as much data.[3] Also, 2-dimensional regions can be used as areas only (not bounding any region). In a $1 \cdot 3 \cdot 3m^3$ volume the surface area is $30m^2$ and if just the pressure is of concern the graph becomes that of Fig. 2.4. Twice as many speakers are needed if the normal derivatives must be sampled as well.

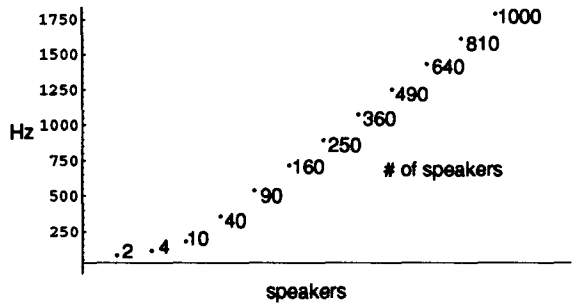

Figure 2.4: Approximate bandwidth attainable by N samples filling a $30m^2$ area

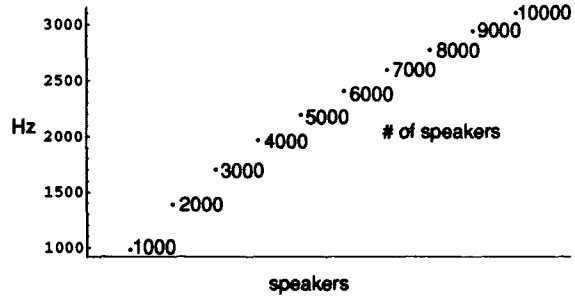

Figure 2.5: Approximate bandwidth attainable by N samples filling a $30m^2$ area

With many more speakers, Fig. 2.5 is possible when concerned about just the pressures. Considering just one cross-section of the original volume where everyone's ears are located, this $3 \cdot 3m^2$ surface can be sampled at various densities obtaining Fig. 2.6 and with large numbers of samples we have Fig. 2.7

Putting all the samples on a line produces the results in Fig. 2.8. Again for consistency, only pressure samples have been counted. To use the line as a boundary for a $9m^2$ area, the normal derivative must be approximated as well taking twice as much information from each sample location, and requiring twice as many speakers to reconstruct.

[3] in some cases the normal derivative can be made to be zero everywhere on the surface requiring no additional speakers or data

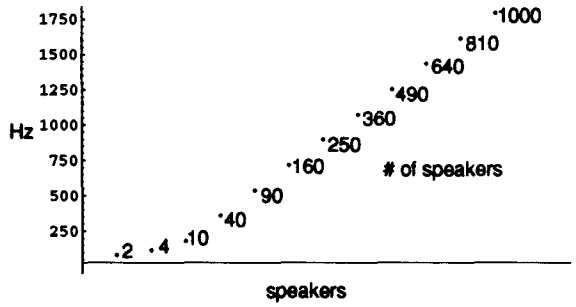

Figure 2.6: Approximate frequency attainable by N samples filling a $9m^2$ area

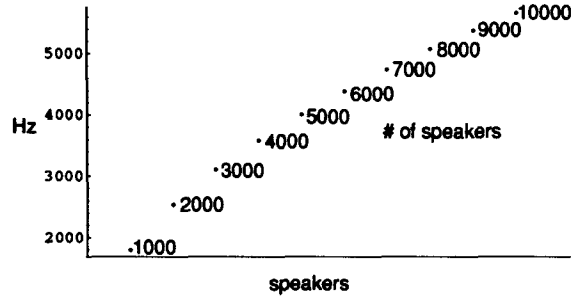

Figure 2.7: Approximate frequency attainable by N samples filling a $9m^2$ area

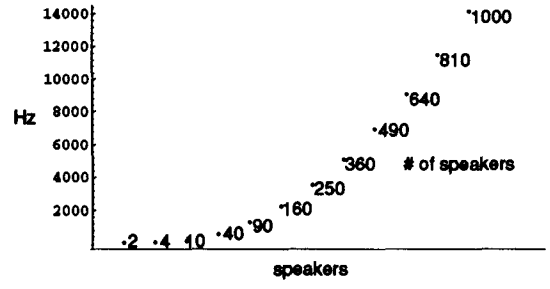

Figure 2.8: Approximate frequency attainable by N samples filling a $9m^2$ area

While the other graphs had a somewhat linear increase in frequency for a quadratic increase in the number of speakers used, since the rule of thumb does not take any root in the one dimensional case, the increase in frequency is quadratic with the quadratic increase in samples and linear with a linear increase in samples. The final graph (Fig. 2.9) shows that not too many thousands of samples are needed to achieve resolution beyond the limit of human hearing along a line.

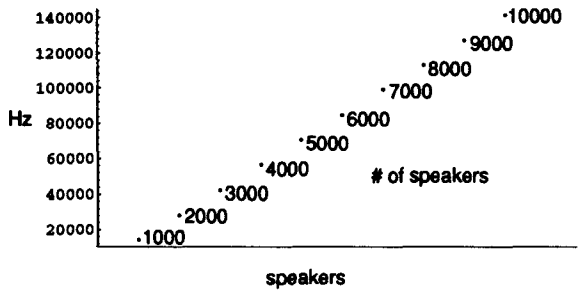

Figure 2.9: Approximate frequency attainable by N samples filling a $12m$ line

This seems promising but there are some tradeoffs. Twice as many speakers are required to use the line as a boundary when applying the Kirchhoff-Helmholtz integral theorem. All the speakers must be aligned on the same plane, or else the signal will enter the $9m^2$ area without crossing the boundary. This problem also arises when trying to simulate a source that is not co-planar, since to reproduce every sound, only conditions at the boundary can be used while such a sound "sneaks" in from above or below. If there is any reverberation in the actual room, it must be anti-symmetric above and below the listening region to cancel completely there, or else some signal will enter without being able to be compensated for at the boundary. Needless to say, this is not often the case. Finally, if some peoples' ears do not lie exactly in the boundary, then there will be some distortion due to the extra distance the sound has to travel resulting in phase misalignment. As a practical matter, if adjustable seats could be used to put everyone's ears in the same plane in an anechoic room, then this approach might be feasible, since the entire relevant soundfield can be related to samples along the boundary line.

2.2.3 Data Representation

The goal of procedure 1 is to describe the target information unambiguously to procedure 2. In order to figure out and describe the target, several considerations must be addressed. First, there is a dichotomy between real and simulated fields, even though there are many similarities. If the target field is based entirely on a computer simulation, then the measurement of the field is straightforward. The data can be measured anywhere without affecting the original field which is already quantized. If however, the target field is a real soundfield, then there is the additional problem of converting measurements of this field into an unambiguous target representation (hopefully without adding too much noise). Ideally, the data representation should be the same in either case. Procedure 2 should not need to care whether the soundfield it is trying to reproduce is real or simulated. In both cases, the information given to procedure 2 can be represented in at least four ways.

Complete Description (1)

This would entail a very large set of signals taken at enough sample points to give a description of the original soundfield in great detail. Thanks to the sampling theorem, a complete description of a bandlimited field is possible through careful sampling of the field. However, the amount of data required to give such a complete sampled description is unwieldy for even a moderately sized region, as discussed in the previous section on sampling (sect. 2.2.2). Also, because most soundfields are not spatially bandlimited, even though at every point the audio signal is bandlimited, there are some minor problems discussed in Appendix B. However, soundfields can be described this way when model-based coding (described below) is ineffective. Also, use of sampling is an intermediate description needed for reconstruction calculations by procedure 2. Therefore, this approach will be used in the next chapter.

Perceptual Description (2)

Here, enough data is transmitted so that what is perceived by any listener within the field is described. The perceptual description has many advantages. It allows for a reduction of data (compression) based on a perceptual model and is closely connected with the real goal of perceptual identity in the secondary environment. An example of this approach is the method of control used by Jot et al. [49] [50] in that the output, possibly transaural, is determined by entirely perceptually related parameters. Nonetheless, use of a perceptual description is difficult at this time, especially for larger soundfields. One invalid argument against such a description is that it does not meet the criteria of being unambiguous, since many soundfields may be perceived identically. However, such a description is unique with regard to perception, and that is the real goal, not to recreate a soundfield. The problem for a perceptual description is that there needs to be very accurate perceptual models which allow calculation from and to real soundfields. Procedure 2, which recreates the desired perceptions, must do so by creating a soundfield. There is no other way, short of plugging directly into the auditory mechanism, and this is not currently available. If a perceptual model is defined in such a way that a set of soundfields can be derived from a perceptual description, then a perceptual data set could be used as an output of the first procedure.

Model-Based Coding (3)

Here, the relevant features of a soundfield are described in terms of a model which can be used to recreate similar soundfields. Model-based coding is a description of the soundfield in terms of sources and 'affectors' which, in the real world, can absorb and reflect sound. In a computer simulation affectors can do anything to the sound which can be programmed. Although only point sources (monopoles) are needed to describe any soundfield in theory, it is better to model each source as a whole, or possibly a few parts, with its own radiation pattern dependent on both frequency and time. Some possible coding techniques for directivity patterns are discussed in Appendix C.

It is also better to base the model on real situations. Sources and affectors can be described to represent real objects so that the soundfield created by them is unambiguous. This model-based soundfield coding is similar to model-based image coding in its philosophy. The idea is to use certain known assumptions about the nature of objects in the world to reduce the amount of data [8]. Thus, one benefit from this approach is a vast reduction in the amount of data necessary as compared to a complete description. Another small

benefit is the elimination of spatial aliasing at this stage. However, getting the value for pressure fluctuations at any point in the original environment involves simulation of the soundfield at a given point rather than interpolation between samples. An example of a model-based coding system is the simulation environment described in Appendix A. Only a source's location, phase, amplitude and frequency are required data (directivity is not yet implimented), and the soundfield is described unambiguously.

If the model used for sound sources in a soundfield is applied to some object which does not fit the model, the result will be undesirable. It is difficult to describe very unusual soundfields which might nonetheless be desired. In such a situation a combination of model-based coding and a sampled description may be required.

Another problem with model-based coding is that when the desired soundfield is real, the model may not be able to recreate the results exactly. Even so, if the model is good, the resulting virtual soundfield will be substantially similar to the real soundfield even if at certain points the signal is substantially different between the original and the model. If the model is good, the resulting "virtual soundfield" will be good, and can be considered to be the desired field. To the extent that a real soundfield is unknown, it cannot serve as a goal. If a model is made which fits all the data collected from a real field, then one cannot prove that the model doesn't reproduce the real field exactly unless further measurements are taken, at which point the model could be improved (and the cycle starts over). It makes more sense to consider the simulated field produced by the model as the desired field because it is completely known, although it may not duplicate a real field on which it is based. Of course, in the case of pure computer simulation, the model-based approach is still valid for many virtual soundfields.

Just Collected Data (4)

Dealing with a real soundfield entails additional considerations. The main problem for a real field is measuring it in such a way that the data can be interpreted to be a unique soundfield. This approach is similar to sampling but broader since no restrictions are applied to the type or amount of collected data having some relevance, for example the level and rate of fluxuation of noise in an environment. If a real field is being observed, the measurements themselves can be given to a later procedure along with their context. Note that this approach does not necessarily describe a soundfield unambiguously, but just gives data necessary for another procedure to figure things out, and depending on the method, that procedure may achieve different results. The lack of interpretation at this stage limits the ability of procedure 1 to make the desired field known, which is its whole purpose. Without digressing too much on the nature of knowledge, the solution to a mystery is more than just the necessary clues, unless the clues make the solution so obvious that it is hard not to see. A similar assumption about the nature of knowledge of a desired field is at work here. Thus, just transmitting measurements which can be used to create an unambiguous solution may not be real knowledge of the field and hence not adequate output of procedure 1. Model-based coding, or even sampling do not provide complete knowledge of a real soundfield but are much closer to representing real knowledge since to get the pressure fluctuations or perception at any point is relatively straightforward.

2.3 Reconstruction of the Sound Field (Procedure 2)

2.3.1 Description

Assuming procedure 1 has done its job, there exists a description of at least some aspects of a desired soundfield. The purpose of procedure 2 is to figure out for a given environment and set of criteria, what signals should be sent to the speakers to come as close as possible to the desired result. Figures 2.2 to 2.9 show that the number of loudspeakers required to reproduce the identical soundfield is unwieldy for many frequencies, so exact reconstruction of the soundfield is unlikely. Since this approximation is to be evaluated by a real audience, the best criteria to use for minimizing the error are perceptual in nature. In order to measure the closeness of the approximation, some measure of distance is required so some norm(s) must be used to measure the resulting errors in the approximation for comparison. However just stating the goal of identical perception is not enough, since a perception has many aspects which may not all improve together. In addition, there may be some tradeoff between the perception of an individual or small group and the perception of larger groups in that the overall perception among many people may improve by making the perceptions of a few worse. Finally, some method must exist to figure out a good solution from all of this.

2.3.2 Role of Psychoacoustics

It seems reasonable that complete reconstruction of a desired soundfield is an ultimate goal of audio. However if a different soundfield is reconstructed which is perceptually identical to the desired soundfield, then for the audience there is no difference. Furthermore, using perceptual identity as a criteria does not make things any more difficult, since an exact soundfield reconstruction will have perceptual identity with the original soundfield. In fact, using perceptual identity as a criteria is likely to provide more flexibility to the system, since it is known that a wide variety of waveforms are perceptually identical and thus it is likely that a wide variety of soundfields will be perceptually identical. This point will be used and elaborated on in a specific context in the next chapter.

In order to use the concept of obtaining maximum perceptual closeness of soundfields, some psychoacoustic model must be available to evaluate the results. Without actual people, (which have the disadvantage of being hard to insert into a long iterative algorithm), some reasonable computation must be made to evaluate the likely effect on perception. This is an area where a great deal can be done, yet even simple perceptual models are of use. In addition, exactly what perceptual goals are most important must also be determined. Often the attainment of one perceptual goal gets in the way of another aspect of perception. For instance, it may be the case that getting the correct amplitude to the entire audience makes the localization bad for certain individuals.

2.3.3 Knowledge of the Listening Environment

A large part of reconstructing a soundfield is the ability to know what effect a signal sent to the speakers will have. This requires knowledge of the speakers response and directivity patterns (both amplitude and phase at all necessary frequencies) as well as knowledge of the room response. These items affect the impulse response between each speaker and each sample point and are transformed and incorporated into the transfer matrix which describes the effect of the speakers on the sample points in a straightforward fashion, as was shown

in the previous chapter. Unfortunately, the process of solving for the speaker output signal given a desired target often requires inversion of this matrix of transfer functions requires that the room response be invertible. Much work has been done on this (see [80] for references), but often inverting a room response is impossible to do exactly (see [78]), so some approximation must be made. A big problem is the change in room response when an audience is present. If the desired soundfield is closely reconstructed for the case in which an audience is absent, the presence of people will have a similar affect on the reconstructed and original soundfield, so not a great deal is lost. However, if psychoacoustic optimizations are done, then the effect of an audience may be quite different, since the soundfield is quite different. It remains to be seen whether or not this potential difference is a real problem in terms of perception.

2.3.4 Norms and Error

The goal of having the least error, regardless of the criteria, is closely linked to the notion of closeness to the desired result. However, there are many definitions of "closeness" depending on the definition of distance. To begin the problem of trying to reproduce a given soundfield, a goal is to come up with useful measures of distance in order to find a reasonable closest match.

It is important to be aware of the big picture. The error between the desired soundfield and the reconstructed soundfield has great theoretical value but is never perceived directly. Instead, the audience hears only the reconstructed soundfield and may not be aware of discrepancies between the reconstruction and the target. It is more important from a practical point of view to avoid unwanted artifacts which undermine the total experience, even if the moment to moment perception of the desired differs from the actual. Since near exact soundfield reconstruction is unlikely in the near future, choosing between various trade-offs can be informed by keeping in mind the final result. Therefore, the norm is only a tool to be used in calculations, but not the final nor even most important arbitrator of the result.

The best way to find perceptual trade-offs is to pursue the goal of minimal perceptual "distance" between the desired and reconstructed soundfields. Given some notation,

$s(t; x, y, z)$ = the original signal at point (x, y, z) at time t.
$\hat{s}(t; x, y, z)$ = the new signal at point (x, y, z) at time t.

the magnitude of the error, or difference between the two, at a given location and time is $|s(t_i; x_j, y_j, z_j) - \hat{s}(t_i; x_j, y_j, z_j)|$. Although in the real world s and $\hat{s}$ are real signals, it is convenient to consider them as complex signals (phasors) since relative phase and amplitude are not dependent on time. Note that for convenience everything is done in the same coordinate system, although the reconstruction will take place at a later time and different location. All the temporal and spatial relationships are preserved. It can be imagined that the original soundfield through time was simply shifted, rotated and delayed to match up with the reconstructed soundfield for comparison purposes.

A norm can be defined to create a measure of distance between the desired field and reconstructed field based on the error at each point and time. The general definition of a norm [53] is a real number associated with a vector such that

$$||x|| = 0 \iff x = \bar{0} \tag{2.5}$$

$$||\lambda x|| = |\lambda|\,||x|| \tag{2.6}$$

$$||x + y|| \leq ||x|| + ||y|| \tag{2.7}$$

Thus the norm has some important properties normally associated with length or distance, and can be used to measure error as the 'distance' between the target and achieved values. (Further restrictions must be applied to create a metric space [53].) The norms of interest include the family of L_p norms. The general L_p norm can be defined [53] as

$$||f(x) - \hat{f}(x)||_p = (\int |f(x) - \hat{f}(x)|^p)^{\frac{1}{p}} \tag{2.8}$$

which can be generalized to the time/space coordinates used here

$$error_p(t_0) = \left(\iiint |s(t_0; x, y, z) - \hat{s}(t_0; x, y, z)|^p dx dy dz\right)^{\frac{1}{p}} \tag{2.9}$$

The most studied L_p norms are the L_1, L_2 and L_∞ norms, where L_∞ is defined by the limit as p goes to ∞ in the above definition. The L_2 norm, or least squares norm is the most common choice for many engineering problems because of its simplicity of computation and general usefulness. Soundfield reconstruction is no exception, and use will be made of the L_2 norm where appropriate. The L_1 norm has certain desirable properties, and will occasionally be used to measure error. It will not be part of an optimization procedure since its less straightforward calculation outweighs its benifits. The L_∞ norm (often called the Chebyshev or minimax norm) is perhaps the best choice for sound, since that minimizes the largest local error and each ear perceives only a local stimulus.

It may be useful to consider how the error between a desired and reconstructed soundfield can be calculated in a simple example. To do this, additional notation is required.

$sp_j(t)$ = the signal (at unit distance from) speaker j.
d_{ij} = the distance between the location of speaker j and point i.
λ = the wavelength.
$os(t)$ = the signal (at unit distance from) original source.
dos_i = the distance between the original source and point i.

Almost all naturally occurring direct signals can be modelled as a set of point sources, and due to the linear properties of the medium, the effect of each point source can be considered separately and the results added. However, for now, assume that each speaker is a point source with unit directivity at all frequencies (a monopole). Now, if the phase change from $os(t)$ to p_i is

$$e^{-j(\frac{2\pi dos_i}{\lambda})} \tag{2.10}$$

and the change in amplitude is $\frac{1}{dos_i}$ (note that when this distance is 0 the amplitude is infinite, so the signal is normalized by taking it at unit distance) then

$$p_i = os(t) \cdot \frac{1}{dos_i} e^{-j(\frac{2\pi dos_i}{\lambda})} \tag{2.11}$$

Likewise the value of $\hat{p}_j$ is the sum of contributions from all the speakers with their respective phase and amplitude changes.

$$\hat{p}_i = \sum_{j=1}^{N} sp_j(t) \cdot \frac{1}{d_{ij}} e^{-j(\frac{2\pi d_{ij}}{\lambda})} \tag{2.12}$$

Although an exact solution to 2.9 is desired, in practice, integrals of such complexity can be approximated as summations i.e.

$$error_n = (\sum_{i=1}^{M} |p_i - \hat{p}_i|^n)^{\frac{1}{n}} \tag{2.13}$$

and $error_n$ is to be minimized for some n-norm with some chosen set of points.

Norms and Perceptual Criteria

Ideally, a perceptually based norm which measures the distance between what is perceived and the desired perception could be found. This perceptual norm would be very difficult to define completely, since how things are perceived varies from individual to individual and even varies within individuals depending on such things as mood, alertness, previous perceptions and so on. Even so, there are a number of psychoacoustic principles[4] which are sufficiently general to help in creating a "perceptually based" distance measure, even if it does not measure exact perceptual distance.

Unfortunately, the creation and use of such a norm would be difficult to deal with mathematically. Therefore, a somewhat different approach is taken here of applying a simple norm to perceptually relevant criteria. This approach has the benefit of giving an easy to calculate measure of error while still being perceptually based.

2.3.5 Individuals and Groups

One important consideration when deciding what to optimize for is the relative weight given to the individual listener in the context of an audience. Often, it is possible to improve things for the whole at the expense of certain parts. The least squares approximation proposed by Yanagida et al. in [100] does this implicitly, minimizing the total squared error at every person's ears, but not limiting the error at any individual ear. Such an approach can be thought of as a modified utilitarian approach, supplying the greatest total good for all ears. However one ear is half of what most individuals have, and it does not follow that the greatest good for all ears is the greatest good for all individuals, especially if the evaluation is based on perceptual results. Ears are not individuals. The smallest psychoacoustically useful measure is the individual. To provide the greatest total good for all individuals requires some reformulation of the problem in terms of individuals. An L_1 or L_2 norm applied to this reformulation over the whole audience will be justifiable under this philosophy. Conversely, emphasis can be given to individuals in the sense that each individual is not allowed to fare too much worse than everyone else. This can be seen to be akin to the L_∞ norm, or minimax where the maximum deviation is minimized among individuals. This approach may be good if it ensures that everyone gets a good result, bad if it ensures that everyone gets a bad result. Even beyond norms, these considerations about what the ultimate goal is should be firmly rooted since the philosophy effects the approach at every stage. It is often the case where some great benefit can be found by 'sacrificing' part of the target goals. The effect on individuals, the whole audience and effected subsets

[4]These psychoacoustic principles were discussed in Chapter 1, Section 2.3.

of the audience must be weighed against each other. Unfortunately, even if the relative importance of individuals, groups and the whole are known, they are not equally easy to implement. Instead, certain choices come up which must rely on a rough tradeoff between the few and the many, and it is better if such trade-offs are made in an informed manner.

2.3.6 Tweaking Phase

The perceptual considerations discussed can be used in formulating the perceptual goal of the soundfield reconstruction subject to constraints, such as trying to find the best solution by adjusting only the phase of separate frequencies. Some perceptually important potential goals for this 'tweaking phase' technique are to improve amplitudes and/or the relative phases of a soundfield without noticeably affecting the steady state output from the speakers. Such a tweaking of phase could be to optimize the relative phases at each frequency for a given speaker setup and listening environment by adjusting the phases of frequencies for each channel relative to the others, but not effecting the amplitudes. This approach could be used on many qualitatively designed pieces of music without effecting the perceived qualities of the output channels, and yet improve the perceived quality in the listening area.

It is apparent that adjusting the phase separately at individual speakers for a given frequency can effect the amplitude distribution of a sound field, since that adjustment effects where constructive and destructive interference occurs. However, it is much less apparent that adjusting the phase of individual speakers can also effect relative phase between points. Indeed the effect is a subtle one, and relies on the fact that since the amplitude decreases with distance from the speaker, not all sample points are affected the same amount by a change in phase at one speaker if there is an 'offset' due to the unchanging other speakers. Fig. 2.10 illustrates this effect. Even if a plane wave is adjusted there can still be a slight effect as long as the vector sum of the unchanged sources has different amplitudes at the sample points which are compared.

It is interesting to note the relationship between the affect on relative phase and on amplitude when adjusting the phase of one source relative to the others. The maximum theoretical change in amplitude due to such a phase adjustment corresponds to when the adjusted speakers phase is either the same as or 180 degrees against the phase of the other components. Thus, the theoretical maximum change in amplitude occurs when the relative phase is the same as that of the other components, as shown in Figure 2.11 (assuming non-zero, non-negative resulting amplitudes).[5] Of course, it is unlikely that the natural relative phase of the adjusted source would allow the theoretical maximum change in amplitude at both sample points simultaneously.[6]

Tweaking the phase is often permissible and achievable in certain situations. In many computer music pieces, phase is a free parameter. Although the frequencies and amplitudes have been determined by the composer, the actual relative phases of a given frequency on the channels are usually not specified by the composer. Instead, the actual phases on the channels are more an accident of the initial conditions and effect of whatever sound generation technique was used. Although some forms of sound generation, such as with oscillators, allow the user to determine the initial phase, rarely has any use been made of this

[5] Also, achieving the maximum theoretical relative phase difference will tend to increase the amplitudes due to some constructive interference, while achieving the theoretical minimum relative phase difference will tend to decrease the amplitudes.

[6] assuming the sample points are chosen independently of this property

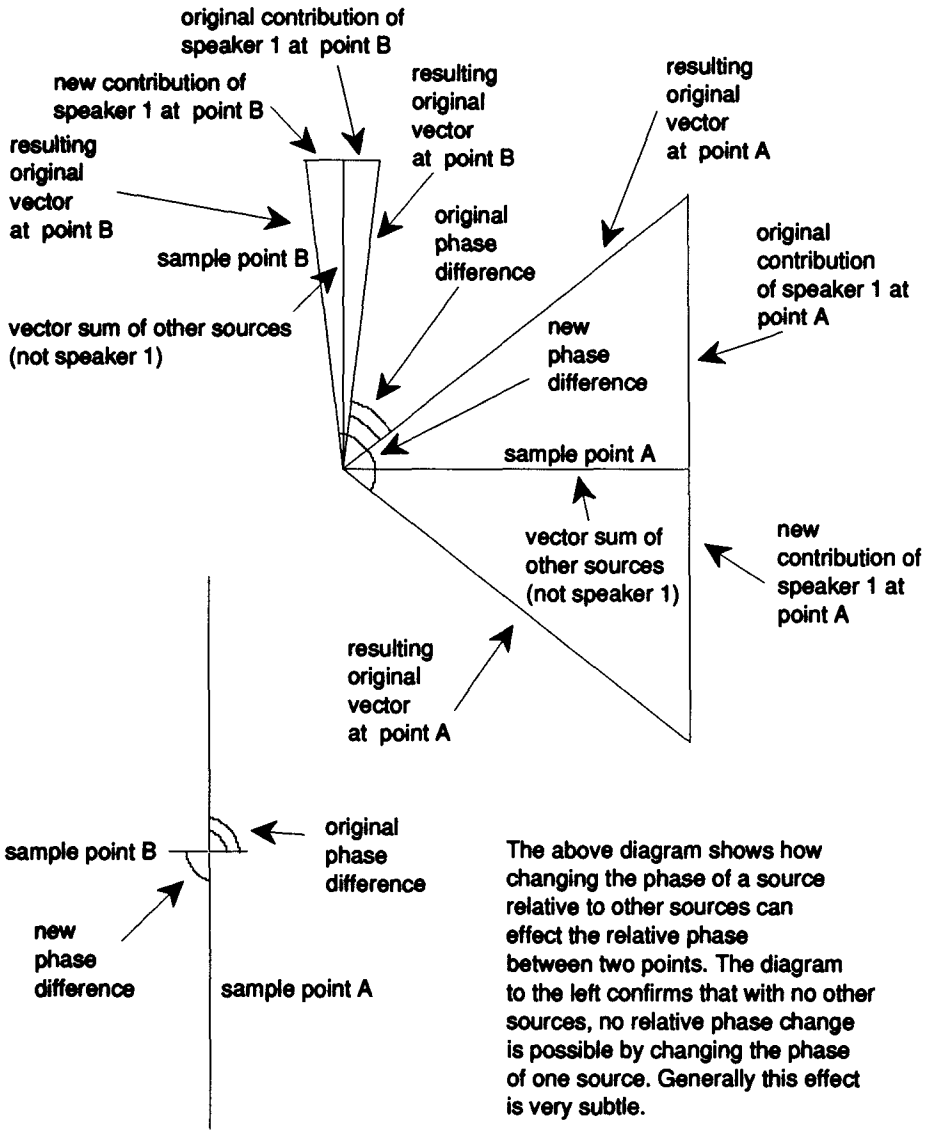

Figure 2.10: Demonstration of relative phase change at two points due to independent phase change at a speaker

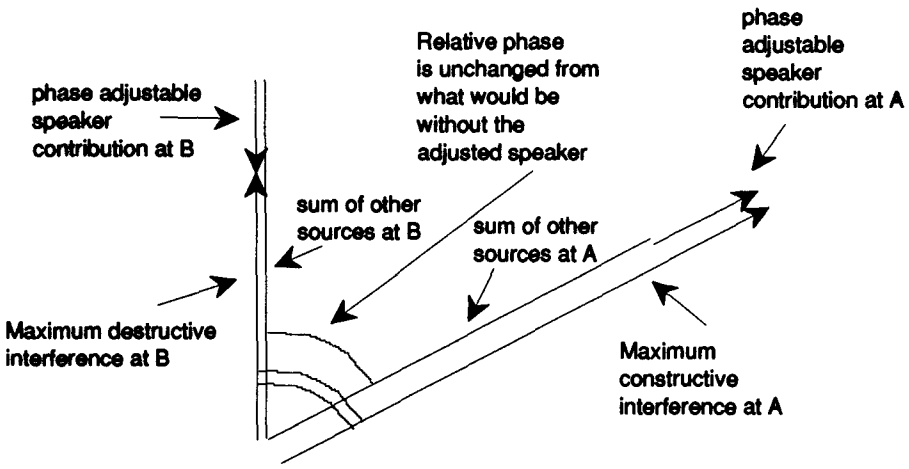

Figure 2.11: Maximum and minimum theoretical amplitude correspond to the same relative phase as if the adjusted speaker had no contribution

possibility, since the initial phase at a frequency does not usually affect perception except at attacks and other transients. However the relative phases between channels can effect perception, and as with good recording techniques, care should be taken to avoid unwanted cancelations in computer music. This avoidance of cancelations is usually not carefully optimized, although phase inversion is avoided between channels which are assumed to be equally distant from the listener, since most recordings are designed to be played in many different rooms with different speaker placements. If the listening environment is well known, then phase adjustments can be calculated which do not alter the perception of each channel independently, but improve the quality of the sound in the specified environment. This is especially straightforward for computer generated pieces using additive synthesis, where the piece can be re-computed with the only modification being changes in initial phase for the oscillators. For recorded music and other synthesis techniques, even if the optimal phases of the outputs are known, they may be hard to implement, though some sort of all-pass filter might be designed for each channel to improve the result.

For example consider the setup in Figure 2.12, where the speakers are at (25,25), (25,75), (75,25) and (75,75) with 9 sample points arranged in three rows in a square from (45,45) to (55,55). The original source is at (60,10). The wavelength used in this example is 12 units.

After a least squares approximation is found, phase adjustment is made to the first speaker which is at (25,25). The total error in the magnitude of the amplitude difference at the nine points is taken as a function of phase adjustment to the first speaker is shown for a changes between $-\pi$ and π in Fig. 2.13. This clearly shows that minor improvement through tweaking the phase is possible. In fact at least a local minimum could be found by finding the phase adjustment giving the nearest minimum for each of the speakers in turn and iterating. This would guarantee at least some improvement for whatever criteria are used.

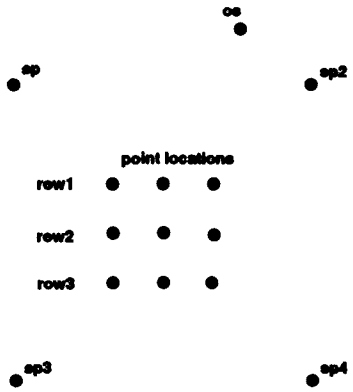

Figure 2.12: Initial setup (not drawn to scale)

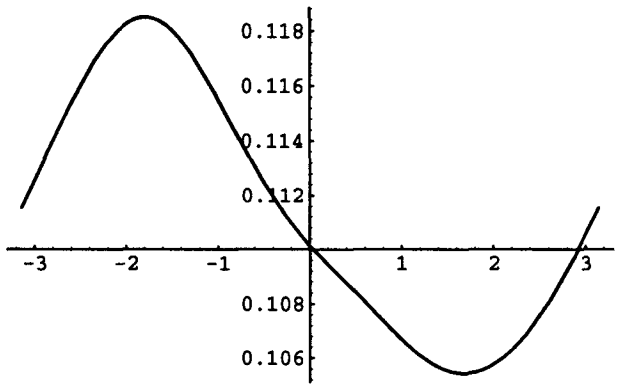

Figure 2.13: Effect of adjusting the phase of the first speaker on the absolute amplitude error

2.3.7 Time Optimization

In addition to the steady state techniques discussed in the next chapter, it is useful to consider events which are discrete in time separately. Instead of the frequency based approach to improving localization, it is better to have a model of localization based on the precedence effect and the effect of reverb direction. The time of arrival and amplitude of a sound at a point from each of the speakers is just a function of the distance to each speaker. Thus for any particular ordering and amplitudes of an event put through each speaker, the arrival times and amplitudes are determined. If the order of arrival and amplitude at each point is put in a model of perception, a cost function can be created to reflecting the error in terms of perceived localization or desirability. Finding a local minimum from an initial guess will produce a better result as long as the perceptually modelled cost function is good.

2.3.8 Standard Methods of Solution

Once a perceptually based cost function is determined, a method of finding a minimum should be implemented. One way is to do a grid search, where samples are taken of the cost function in some ordered fashion, and each time a minimum is found the resolution around that minimum increases.[7] Another useful approach is to use the Newton-Raphson algorithm, which uses a two-term Taylor series expansion of the cost function about a point to find what would be the minimum if the cost function were quadratic. This approach will quickly converge to a minimum in many instances. When the cost function is constrained in some way, use is often made of Lagrange multipliers, which adjoin the constraints to make a new cost function. A good reference on some of these standard techniques is [92]. Instead of applying these well known techniques, some new techniques are proposed in the next chapter.

[7] Such an approach is used in the next chapter to find phase adjustments to minimized a convex cost function.

Chapter 3

An Optimizing Approach using a Class of Iterative Methods

3.1 The Second Step

Procedure 2, as defined in the previous chapter, attempts to achieve a target soundfield as closely as possible in a given listening environment using information supplied by procedure 1 (which determined what the target soundfield is) in the context of audience perception and a hierarchy of perceptual tradeoffs. The main perceptual criteria considered here are

1. Loudness (Amplitude)

2. Perceived angle of arrival based on a simple model using relative phase

In addition, some non-perceptual criteria are considered as being useful mathematically in the methods used to reach the perceptual criteria. These are

1. Total absolute error (possibly squared)

2. Relative phase

3. Absolute Phase

3.1.1 Iteration/New Target Generation

Many methods of approximation which could be used by procedure 2 require iteration, such as the Newton-Raphson algorithm mentioned at the end of the previous chapter. A special class of iterative approaches is considered here, namely those consisting of new target generation. This refers to an iterative approach where some method is used to get an approximation to the target soundfield, but since only certain aspects of the target soundfield are important perceptually, a new target soundfield is created which allows the approximation method to get more accurate perceptual results. This type of iterative method has two main components which are repeated in order:

1. Soundfield approximation method

2. Soundfield adjustment method (new target generation based on perceptual factors)[1]

(3. Goto 1.)

A key for this to work is that after soundfield adjustment, the next soundfield approximation should improve whatever factor(s) are sought. This is not always easy to guarantee, so in some cases an attempt is made to create the best opportunity for improvement. If in the next iteration improvement does not occur, then either the old result should be kept and a different target tried or the iteration should continue to see where it leads, always retaining the best result. Another important point is that 2 maintains the desired perceptual aspects of the original target constant, and changes only the peripheral aspects.

This class of iterative approaches is very general and can be used with any useful approximation and adjustment methods. However the primary soundfield approximation method used here is least squares based on its ease of use. The set of soundfield adjustment methods are closely related to each other. The use of the simplest of this class in an iterative procedure will be referred to as Algorithm 1.

3.2 Algorithm 1: Replacing Phase to Improve Amplitude

People are not sensitive to phase except under unusual circumstances.[2] Since perceptually it is often better to get accurate amplitudes at the expense of phase accuracy, it would be nice to trade one for the other. Given a set of sampled locations at a given frequency, the amplitudes associated with those points are subject to the same sampling conditions as the entire signal (amplitude and phase), since the variations in amplitude are due only to the distance from the sources and interference between waves which is determined by the wavelength (along with the relative phase). Thus the sampling theory discussed in the last chapter applies to the amplitude field of a soundfield, and having the error in amplitude bounded at sample points bounds the error between sample points.

A very simple, yet very powerful algorithm exists to create a new soundfield which at discrete frequencies and sample points guarantees that the amplitude will improve or at least not get worse for whatever (good) approximation method is used. The following four points describe Algorithm 1 and its consequences.

1. Apply whatever approximation method to get a reconstruction of the original signal at the sample points.

2. Next compute a new soundfield using the phase information of the reconstructed field, but maintaining the amplitude from the original field. This produces a new target field which has the same amplitudes at each sample location as the original target, but the phases are those of the reconstructed field.

[1] Note that either 1 or 2 may be iterative processes in themselves.

[2] However the relative phases between the ears can be important for localization at frequencies below 1500 Hz.

3. By making this hybrid the target, and applying the same method to get as close as possible a reconstruction, any improvement in total error must be with respect to amplitude since the phases were made identical at all the sample locations.

4. The process can be iterated as many times as desired, each time improving the amplitude response. However the total error with respect to both amplitude and phase increases each time.

It is easy to see that any change produced by this iterative approach will be beneficial or at least not harmful. Since error is minimized at each iteration, the result can be no worse with regard to amplitude and the new phase. Since the original target's amplitude information can be reused at every step, there is no build up of error in the new target's set of amplitudes. Assuming the method applied to get the minimum total error (in some norm) actually works at every step, this algorithm can be used to apply the same error criterion to amplitude only, converging eventually to at least a local minimum. Figure 3.1 illustrates this process. Mathematically, this can be described as follows:

Let the desired amplitudes (indexed over k) be $a1_k$, where the initial target is $a1_k e^{-j\theta 1_k}$

The acheived minimum solution set for some norm is $a2_k e^{-j\theta 2_k}$

Then the new target is $a1_k e^{-j\theta 2_k}$, replacing the phase with that of the acheived solution.

Now by defining the difference in amplitude to be

$$\sum \|a1_k - a2_k\| = \sum \|a1_k e^{-j\theta 2_k} - a2_k e^{-j\theta 2_k}\| \tag{3.1}$$

and with the next solution set $a3_k e^{-j\theta 3_k}$ having the property

$$\sum \|a1_k e^{-j\theta 2_k} - a2_k e^{-j\theta 2_k}\| \geq \sum \|a1_k e^{-j\theta 2_k} - a3_k e^{-j\theta 3_k}\| \tag{3.2}$$

(because the minimum of $\sum \|a1_k e^{-j\theta 2_k} - a3_k e^{-j\theta 3_k}\|$ can't be worse than what we already have)

then it must be that

$$\sum \|a1_k - a2_k\| \geq \sum \|a1_k e^{-j\theta 2_k} - a3_k e^{-j\theta 3_k}\| \tag{3.3}$$

Now since

$$\sum \|a1_k e^{-j\theta 2_k} - a3_k e^{-j\theta 3_k}\| \geq \sum \|a1_k - a3_k\| \tag{3.4}$$

(with equality only when $\theta 2_k = \theta 3_k$, as in Equation 3.1)
then

$$\sum \|a1_k - a2_k\| \geq \sum \|a1_k - a3_k\| \tag{3.5}$$

3.2.1 Application with Least Squares

To illustrate the application to a least squares solution, consider the setup of Figure 3.2 where 121 sample points from (45,45) to (55,55) in a rectilinear grid are in the center of four speakers at (25,25), (25,75), (75,25) and (75,75). The original source is at (50,-10).[3] If only the difference in amplitude in a certain region is of concern, then looking at the amplitude of the difference between complex quantities is not the best way to go. Two complex values may have the same magnitude, but near opposite phase producing a difference with a large amplitude. However, the difference between the two amplitudes is zero. Thus, a new way of comparing signals over a region is to find the difference of their amplitudes. This new comparison for the situation shown in Figure 3.2 is shown in Figure 3.3.

In figures like Figure 3.3, grey represents zero since the difference can go either positive or negative (unlike before). White represents -1 and black +1 with resolution 1, meaning that the amplitude is not scaled (or is scaled by 1). Note that the area of optimization does not seem that much closer to zero than other regions, since this comparison shows only amplitude differences whereas the optimization took into account both amplitude and phase. By iterating the algorithm described in the previous section 100 times (taking just a few seconds in Matlab), the result shown in Figure 3.4 is obtained.

This result might seem unusual in that it emphasizes one speaker, but as will be seen later when this procedure is implemented using the short-time Fourier transform, certain frequencies get concentrated on particular speakers. Since phase is not an issue, if one speaker can approximate the amplitude best across the region, other speakers will only tweak it slightly. This may actually be bad for localization, requiring other remedies be used in conjunction with this algorithm. Also, if the amplitudes of the complex differences is plotted, the result is Figure 3.5. Comparison of Figure 3.5 with Figure 3.2 shows that by improving just the amplitude comparison, the total error (summing absolute magnitudes of the complex differences) has gotten much worse.

Actually, the placement of the original source in this example is unusual in the symmetry it creates with the region of optimization and speakers. This produces an unstable saddle point which shows that the algorithm need not always converge to the absolute minimum. Before further discussion of this special case, a more typical example is one without such symmetry, such as obtained by moving the original source location to (49,-10).

Using as an example a wavelength of 10.625 units, Figure 3.6 shows a plot of the sum of the absolute values of the difference between the amplitude of the original source and the approximation by the loudspeakers for every iteration through 100. Notice how quickly it converges. However, Figure 3.7 shows that the convergence is to a value higher than an obtained minimum. This undershoot is caused by the fact that the algorithm uses the least squares approximation at every iteration whereas the plot is of total amplitude difference. A plot of the sum of the squared amplitude differences does not have this problem, of course, but may be a less intuitive or less desired measure of accuracy (Figure 3.8). It is used just to show the steady convergence of the amplitude error, and because a least squares approach

[3] In the enviroment use to generate these pictures, the y-axis is flipped

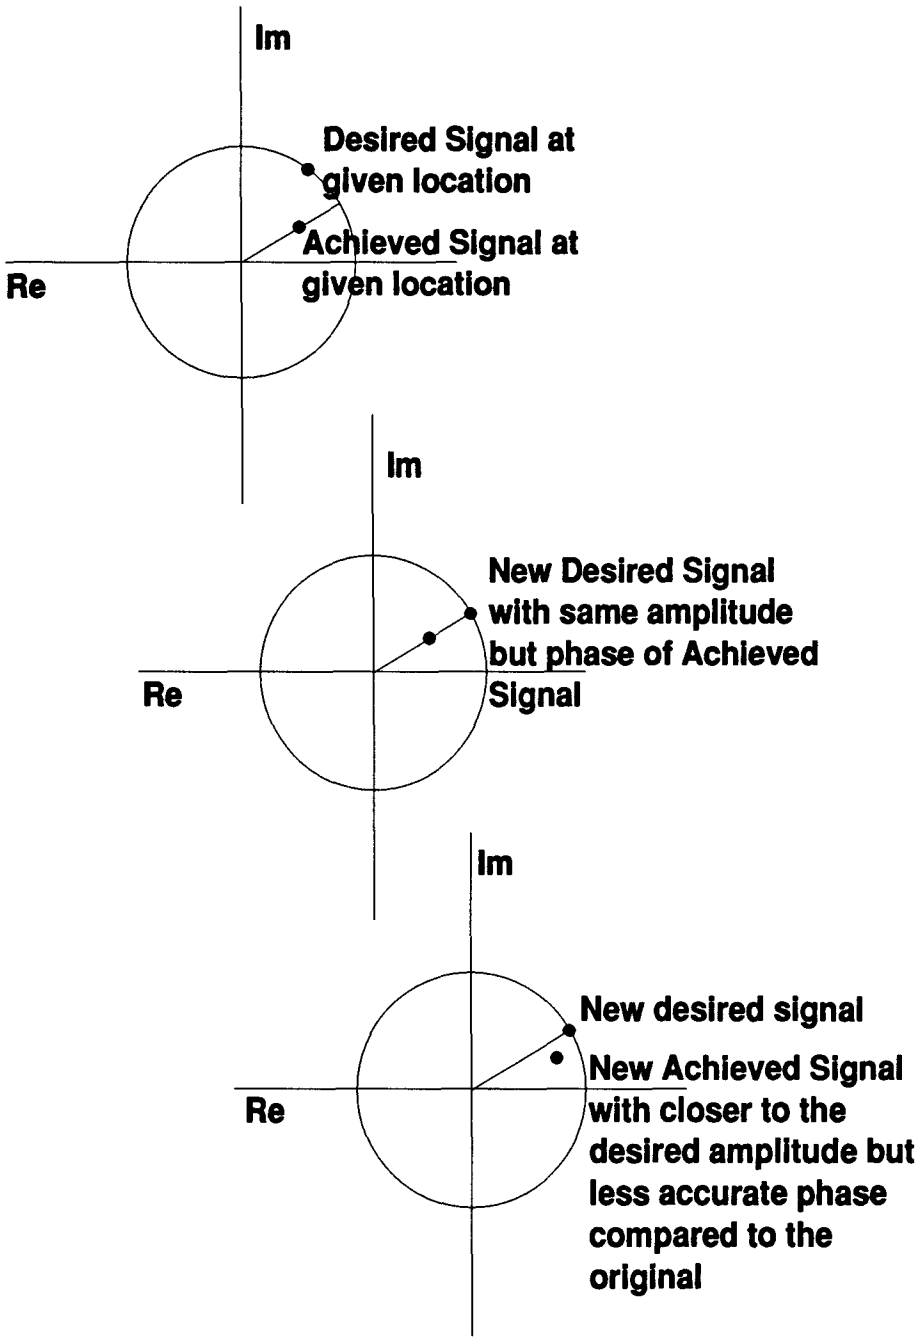

Figure 3.1: Illustration of Algorithm 1 in three steps, (the individual sample locations need not improve in amplitude but the total error will be less).

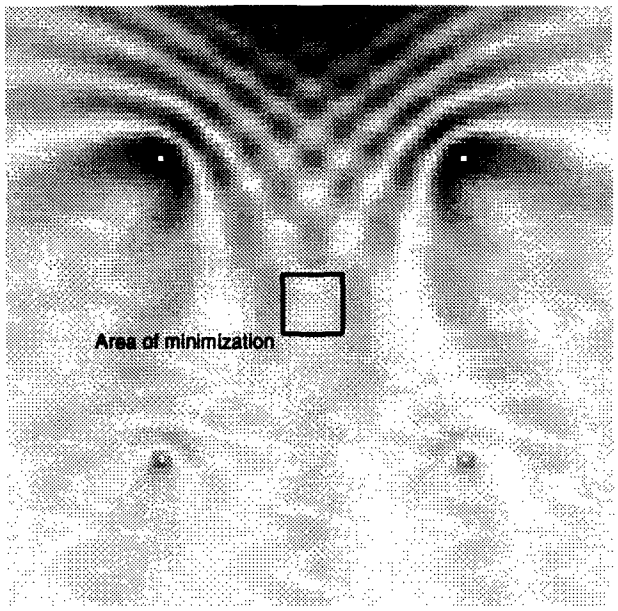

Figure 3.2: Amplitude of differences for least squares solution (resolution 10, i.e. black$\geq$0.1, white= 0).

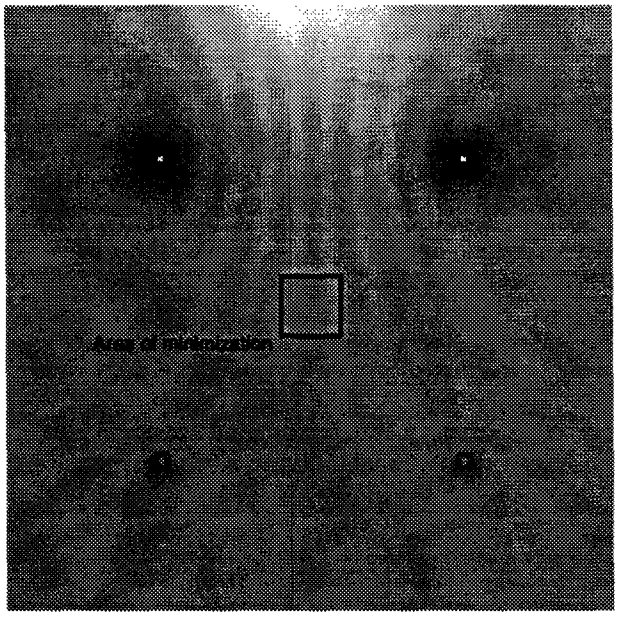

Figure 3.3: Difference of amplitudes after few iterations (resolution 10, i.e. black$\geq$0.1, white$\leq$-0.1).

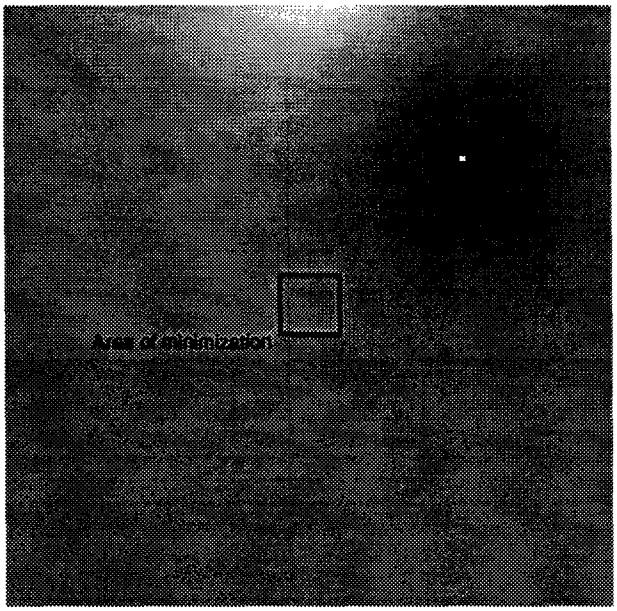

Figure 3.4: Difference of amplitudes after 100 iterations (resolution 10).

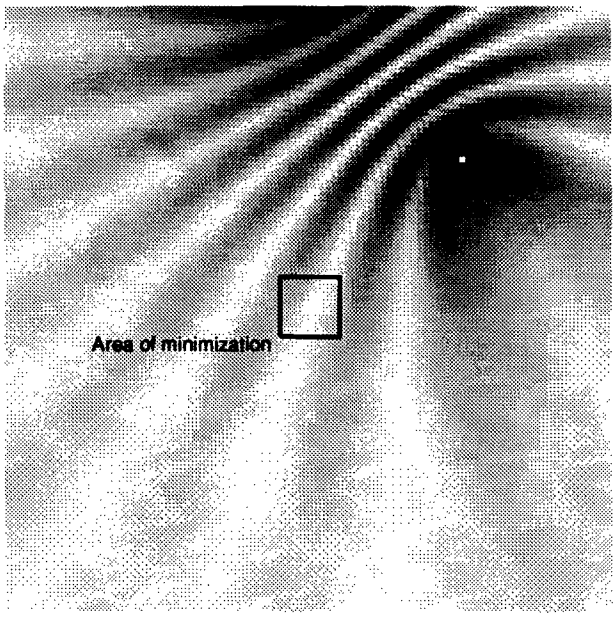

Figure 3.5: Amplitude of differences after 100 iterations (resolution 10).

is used for minimization in each iteration. A better measure of perceptual accuracy would be on a log scale, but that is not needed to demonstrate the points made here. However, the conversion to a log scale is just one additional step and is much more useful in conventional audio practice. If the least unweighted sum is desired, then the point of largest undershoot can be retained. It should be noted that undershoots of this sort can occur before a saddle point, so care should be taken to take the result before the first increase in total amplitude error just prior to the final convergence of the algorithm.

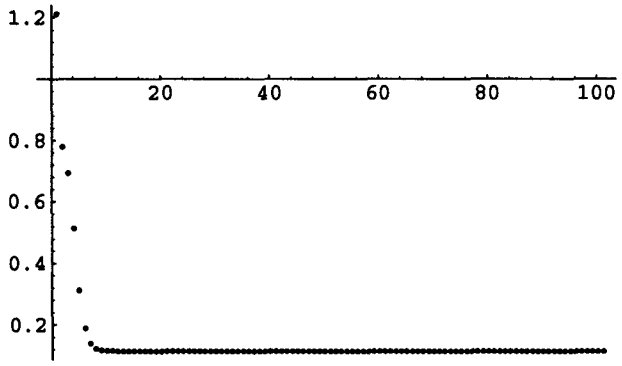

Figure 3.6: Sum of absolute value of amplitude differences, iterations 0-100.

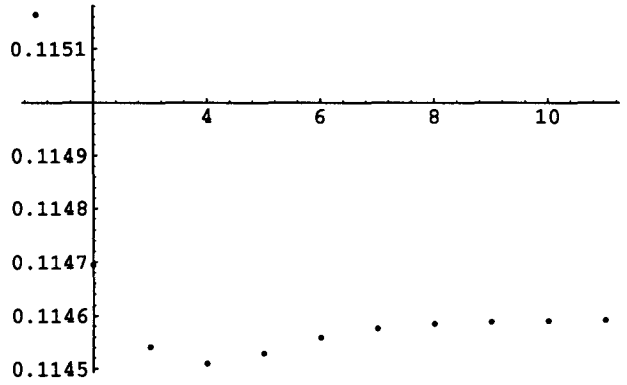

Figure 3.7: Sum of absolute value of amplitude differences, iterations 10-20 (undershoot).

3.2.2 Saddle Points

Returning to the symmetrical example of Figure 3.2, the original source is at (50,-10), splitting the loudspeakers. Here, the algorithm seems to converge first on a solution balancing the two closest speakers, before eventually falling to a better solution emphasizing one of the closest speakers more than the other.[4] Because of the symmetry of the setup, it is

[4]The undershoot is almost immediate, from an initial value of 1.214243 to 0.801318 to the lowest local value of 0.799987 at just the second iteration, before creeping up to near the saddle point at 0.800694 at iteration 17.

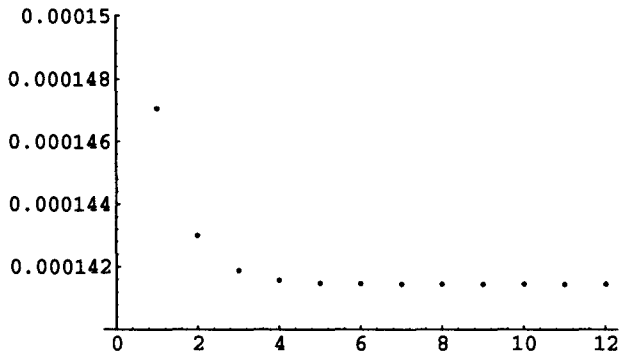

Figure 3.8: Sum of squared amplitude differences, iterations 9-20 (no undershoot).

easy to see how the initial symmetrical solution is arrived at, since no preference can be given to either the left or right forward speaker. Only when the process has sufficiently converged to the saddle point does the low level 'noise' due to rounding error perturb the balance, allowing a better solution to be found. The symmetry also demonstrates that two minimums are possible, since the better solution emphasizing one of the forward speakers could just as well emphasize the other with the same resulting error.

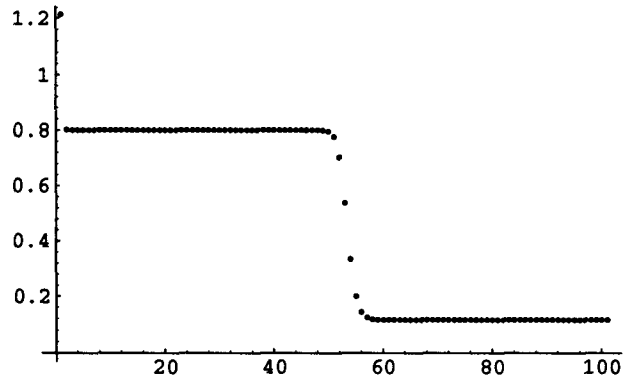

Figure 3.9: Initial convergence to a saddle point, followed by convergence at a (local) minimum.

In order to see the final convergence as a loss of symmetry, a comparison of the amplitudes of the speaker signals at each iteration shows what happens. In Figure 3.10, the two front speakers have the same amplitude as each other, and likewise two rear speakers mirror each other, until shortly after iteration 40, where differences become noticeable.[5]

To see how the symmetry is perturbed, it is necessary to "zoom in." One way to do this is to take successive differences in error amount between iterations. Thus, when converging, the differences in error amount become small. Taking the $\log_{10}$ of this allows the rate of convergence to be seen clearly even when the successive differences are minute. Looking

[5]The phases also mirror each other between pairs until after iteration 40.

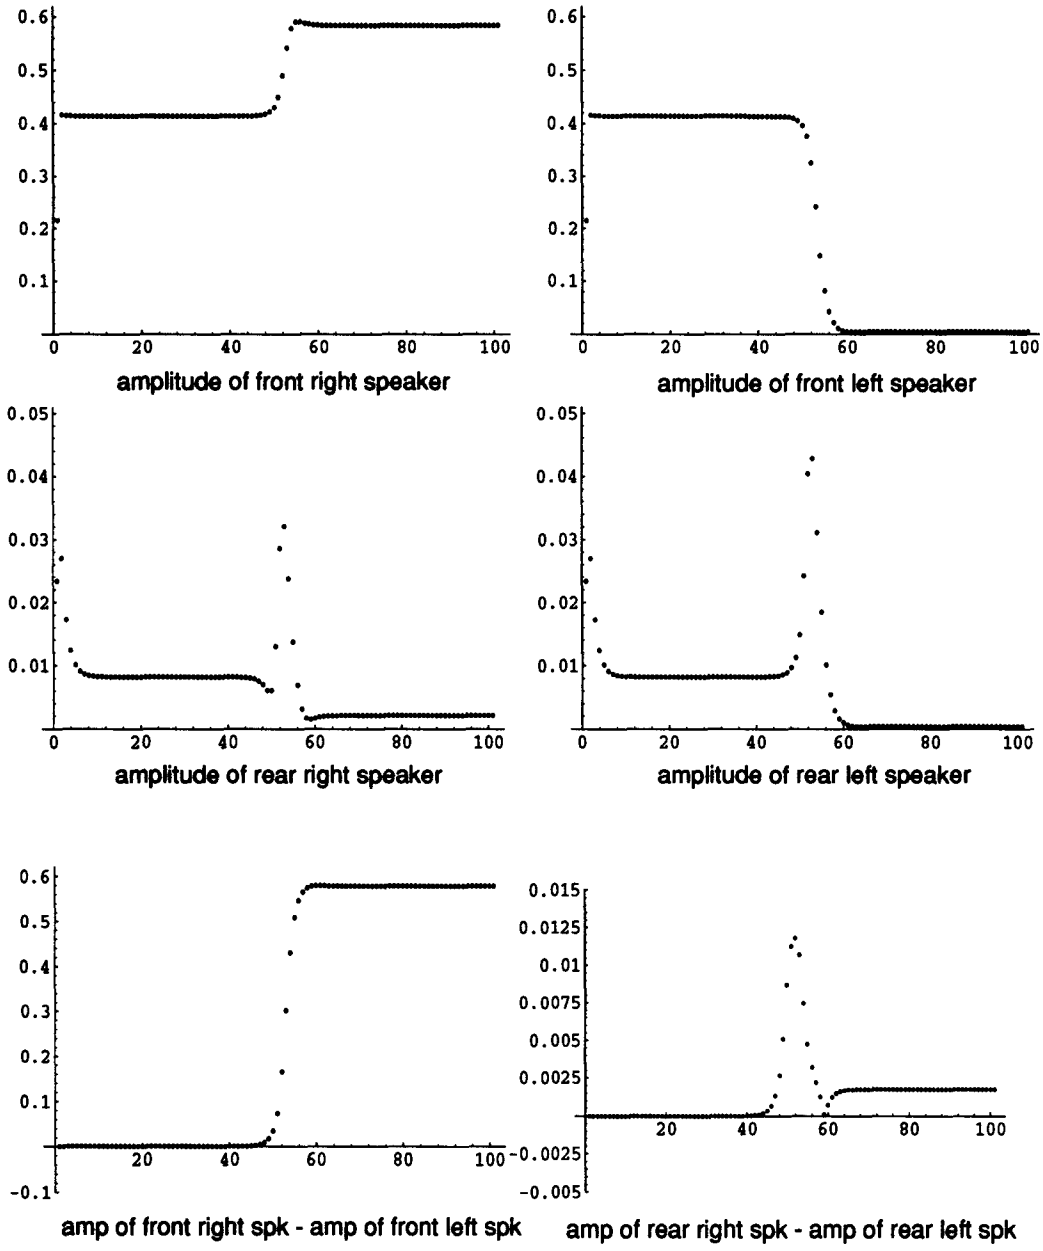

Figure 3.10: Amplitude of speaker signals at each iteration, and the differences between 'symmetric' pairs.

at the error at all 121 sample points shows some interesting patterns, but first Figure 3.11 gives the average of the $\log_{10}$ of the error difference between each iteration.

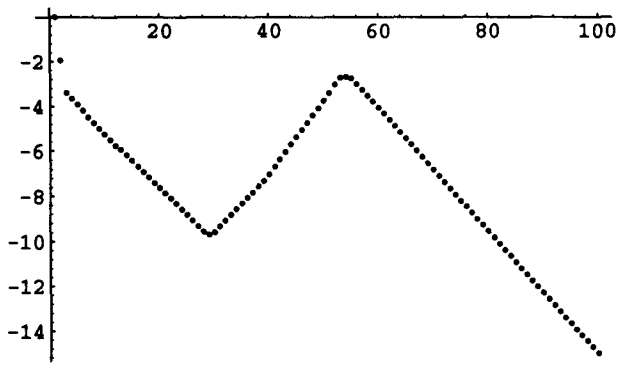

Figure 3.11: Average $\log_{10}$ error difference between iterations.

Comparison with Figure 3.9 shows (as would be expected) that when the successive average error difference is below 10^{-4}, the change in total error is not very noticeable on a linear scale. What is striking is how linear the average error difference is on a log scale. Figure 3.11 also provides clues to the loss of symmetry in its change in directions, especially by locating the first increase in average error difference at data point 29.[6]

Following the $\log_{10}$ of the individual error differences shows an interesting pattern. As the error differences decrease from iteration 1 to 28, there occur certain patterns in the sample points, where at certain points the change in error will decrease faster or slower. While these patterns are noted to give a general idea how the error changes at the sample points, the cause of these particular patterns is buried in the phase changing algorithm and least squares calculations. As the error reductions increase between iteration 29 and 54, the lag of the middle can be attributed to the remaining symmetry. Figure 3.12 shows the $\log_{10}$ of the individual error differences, arranged on the original grid (viewed from near the bottom left speaker), and as 121 data points.

The most important noticeable feature in these graphs is the loss of symmetry when the error differences become small. Although visible in some graphs, this feature can be brought out by reversing the grid over the line of supposed symmetry and subtracting. This produces a purely antisymmetrical graph which is 0 when the prior graph was symmetrical, and which has corresponding dips and hills where the prior graph was asymmetric. This procedure shows that the change in error amount between iterations is highly symmetric when the change is relatively large from iterations 1 to around 24 and from about 34 to about 47. After iteration 47, the speaker signals lose their symmetry which effects the rate of error change, so that it is no longer symmetric even though the change in error is relatively high. However, the near return to the symmetric situation between iterations 24 and 34 may be due to the fact that the changes in amplitude error between iterations are large enough not to produce numerical error, while the numerical errors produced between iterations 24 and 34 which gave rise to a slight asymmetry, were not large enough to cause the significant changes in the prevailing speaker symmetry, as seen in Figure 3.10 and Figure 3.9. Thus,

[6]This is the 28th iteration since the first data point on each figure is the least squares solution prior to the first iteration of the algorithm.

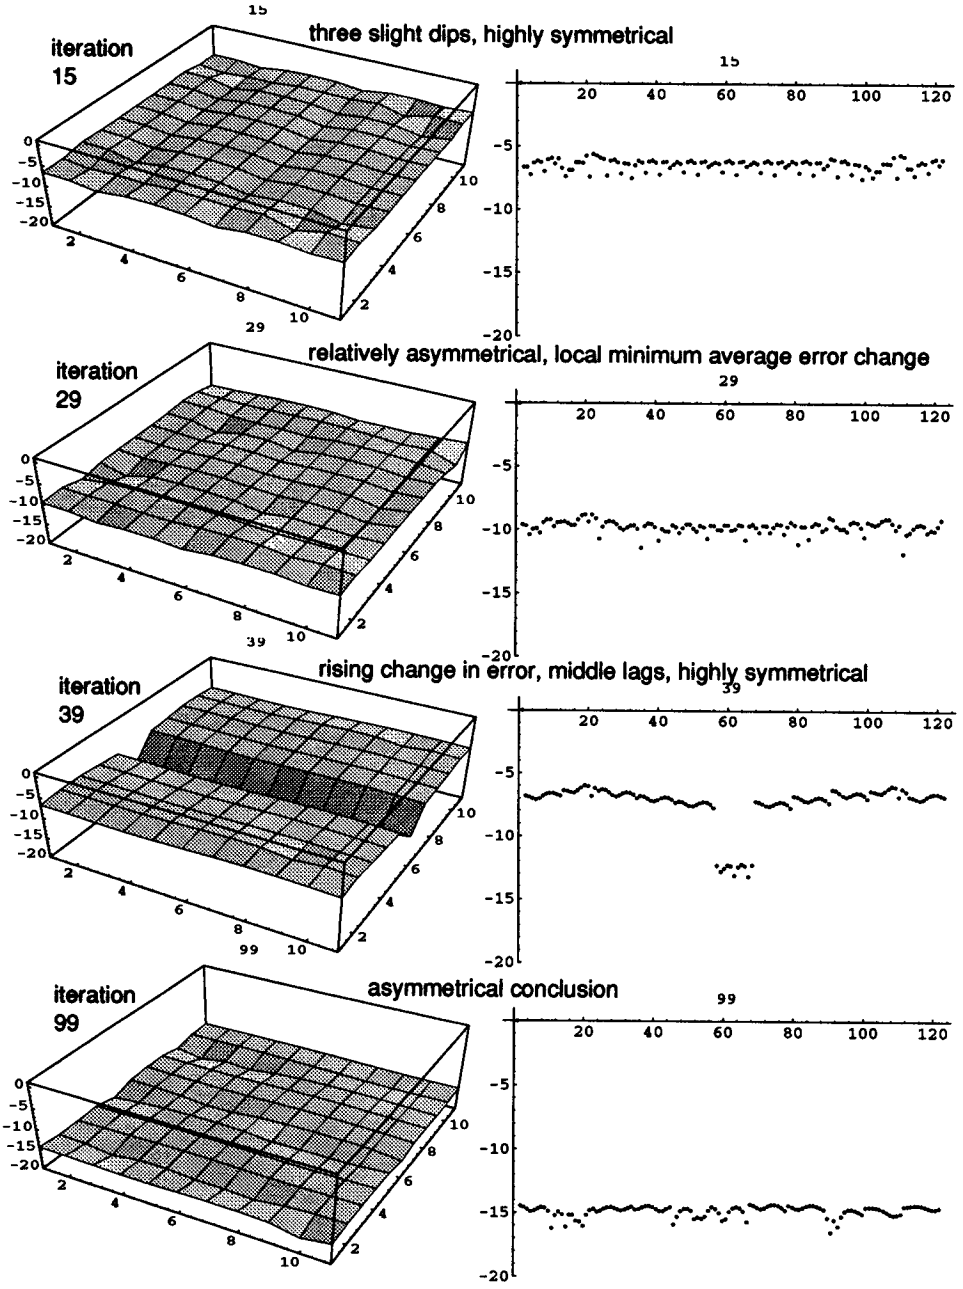

Figure 3.12: $\log_{10}$ error difference between iterations on 11 by 11 grid and as 121 data points for selected iterations.

there is the curious situation that by converging to a saddle point, the change in the amount of error becomes small, causing small numerical errors to create small asymmetries, causing the algorithm to begin to slide off the saddle point, increasing change in error reduction, which in turn makes the numerical errors small, so that the remaining near symmetry is maintained for many iterations with large symmetry in the rate of error decrease. The whole process of getting knocked off the saddle point is oddly almost symmetrical.

As an additional aid to finding points where this algorithm of replacing the phase of the target field with the achieved phase reaches a saddle point or a minimum, it is useful to notice that this will only be the case when replacement of the phase introduces no changes. In that case $\angle y = \angle Ax$, which implies

$$\tan^{-1} \frac{Im(y)}{Re(y)} = \tan^{-1} \frac{Im(Ax)}{Re(Ax)} \tag{3.6}$$

$$Im(y)Re(Ax) = Im(Ax)Re(y) \tag{3.7}$$

$$Im(Axy^*) = 0 \tag{3.8}$$

$$Re(Axy^*) = Axy^* \tag{3.9}$$

3.3 Variation 1: Replacing Amplitude to Improve Phase

Despite the possibility of finding a saddle point or local minimum, Algorithm 1 is very good in the sense that, in conjunction with any method to minimize distance from the achieved values to the target, it is guaranteed not to produce worse solutions with respect to the norm used, as was shown in the last section. However, replacing the amplitude field in order to improve the phase field is not the same in this respect. Luckily, the amplitude field is more important perceptually, but in order to approximate relative phase, the approximation of phase at the expense of amplitude is worthy of brief consideration as a potentially useful technique.

Here, the goal is to put the new target in such a relationship with the obtained result such that a new least squares solution is likely to reduce individual phase error. The phase of the original target is kept while the amplitude is altered at each sample to create the new target. Three approaches are apparent if the next iteration is assumed to minimize total distance.

1. Minimize the distance between the target's phase line and the achieved result.

2. Make the achieved result orthogonal to the difference between the achieved result and the new target by adjusting only the target's amplitude.

3. Make the target the same amplitude as the achieved result.

The effect of trying to minimize distance on any of these approaches is not as straightforward. To help see what effect change in distance has on phase error, the effect of traveling in a straight line at constant speed from (1,5) to (1,-5), shown in Figure 3.18 on phase and amplitude error relative to a target at (1,-5) is shown in Figure 3.19. This can be interpreted to represent what effect an equally weighted attempt at minimizing distance error has on amplitude and phase error starting away from the origin. If the property of

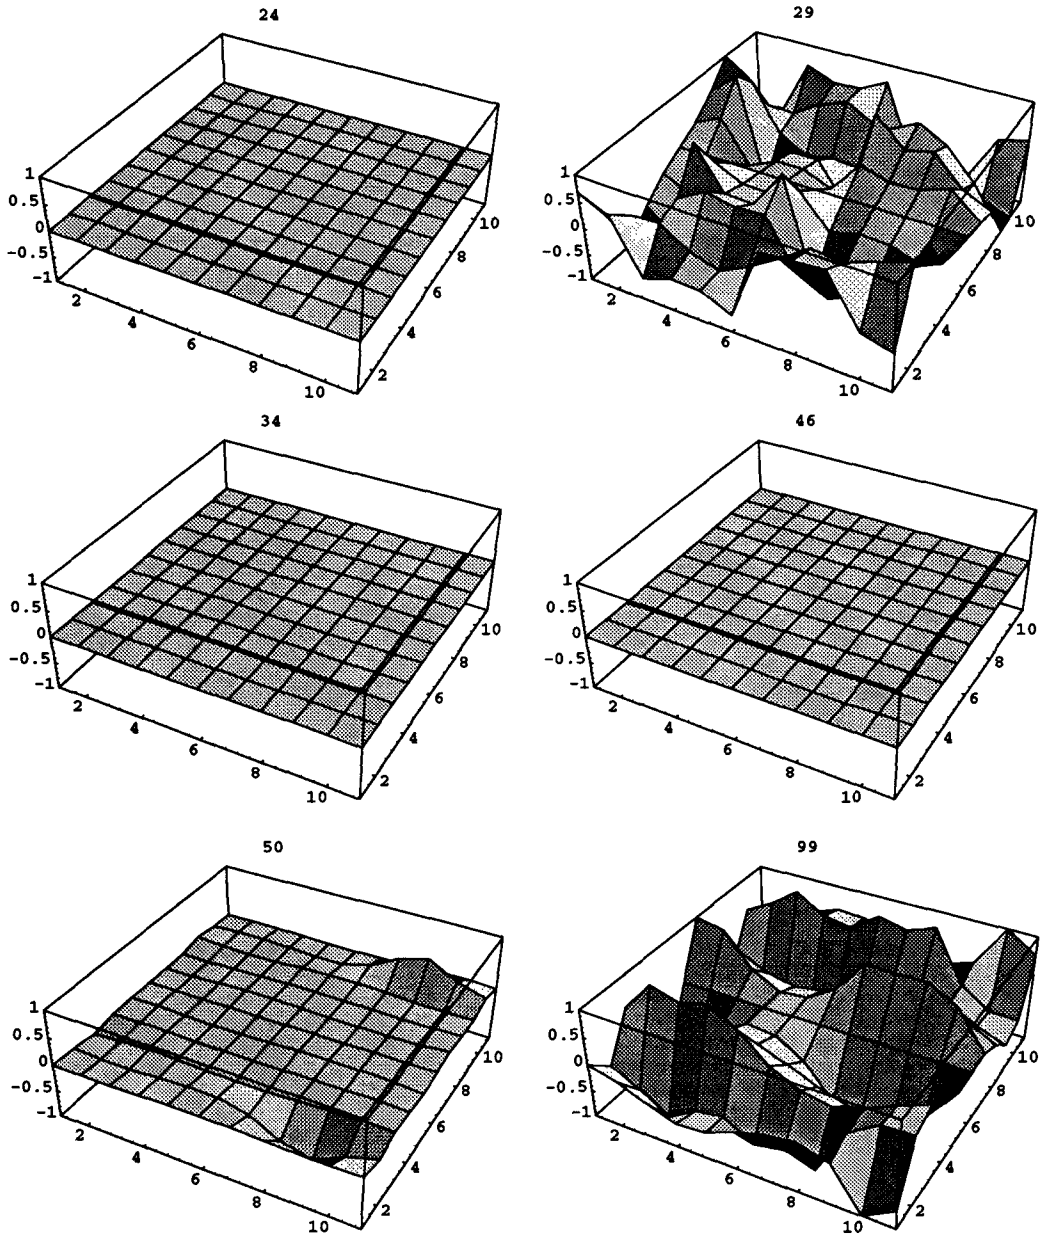

Figure 3.13: Graphs of relative amount and location of asymmetries of rate of error reduction on the 11 by 11 grid for selected iterations (axis across middle line from top-left to bottom-right).

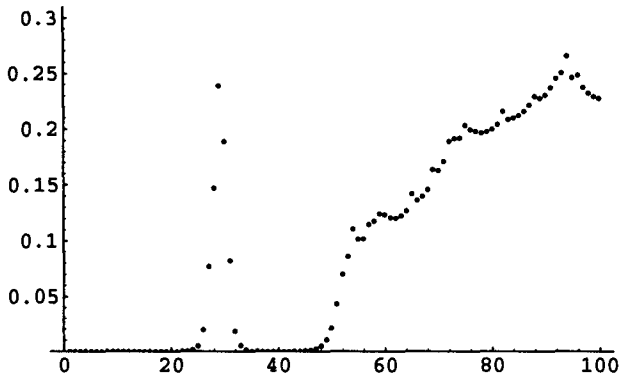

Figure 3.14: Average measure of left/right asymmetry over the 11 by 11 grid for $\log_{10}$ error difference between iterations. (0 implies perfectly symmetric).

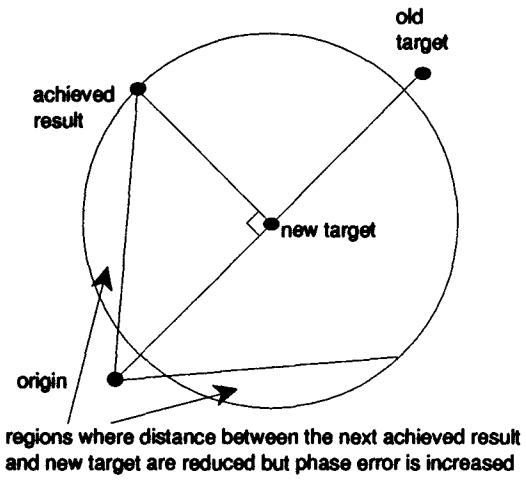

Figure 3.15: Adjusting the target to maintain phase and minimize distance from target phase line and acheived result.

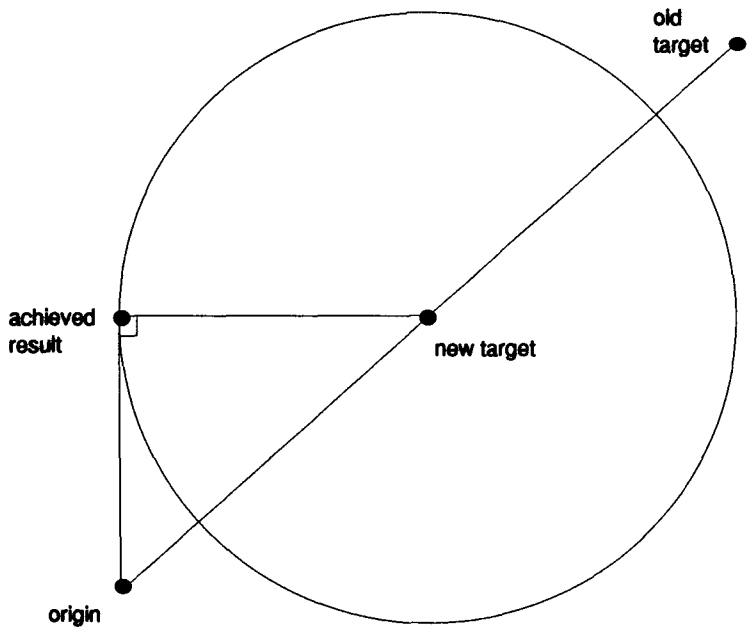

Figure 3.16: Adjusting the target to maintain phase and have the difference with the achieved result be orthogonal to the achieved result.

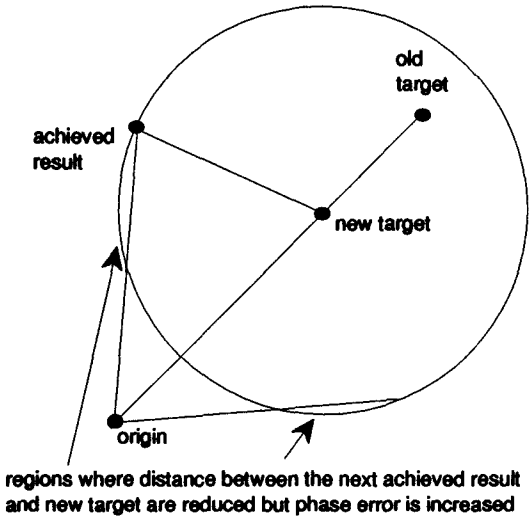

Figure 3.17: Adjusting the target to maintain phase and have the same amplitude as the achieved result

the squared weighting is represented by adjusting the time progression by the reparameterization $t \rightarrow (1-(1-t)^2)$ the situation changes to Figure 3.20. This new figure can be interpreted as showing how 'quickly' a parameter changes assuming the emphasis placed on larger distance errors by the usual square 'weighting'. Unfortunately, these figures give little more than an intuitive feel for what the optimization is 'trying to do', but in most situations the reduction in error will not be on such a linear path. The typical least squares situation is faced with many competing constraints. Still, the ultimate goal is to minimize the squared distances, which even if done in a linear route, can be see to have a nonlinear relationship with both amplitude and phase, so the actual solution will not be linearly related to these parameters.

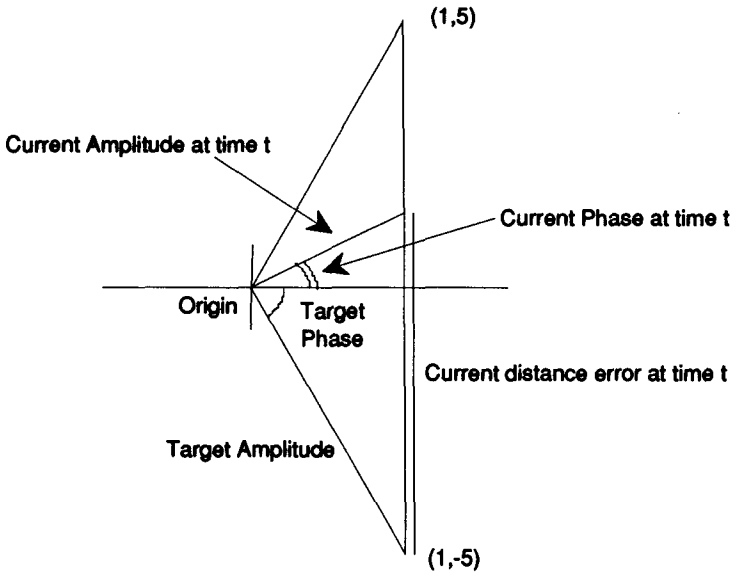

Figure 3.18: Setup for evaluation of the effects of minimizing distance on amplitude error and phase error.

Comparing these approaches with Figures 3.18 and 3.19 show that in approach 2 where the difference between the target and achieved value is orthogonal with the achieved value, then an attempt to minimize total distance is likely to improve the next achieved values phase at the maximum rate. Figures 3.16 to 3.17 show that only this approach guaranties that any improvement in distance will also improve phase. In approaches 1 and 3, there are regions where the distance to the new target can decrease but the phase error actually increase. Approach 1 has the property of increasing the rate of phase change the closer the achieved values get to the target. However, to maximize the potential phase improvement with distance, approach 3 is the best since it corresponds to placing the 'starting point' and target equally spaced about the middle of the phase curve where change is the greatest. Note that the actual improvement may not be largest in this case, nor does improving the phase approximation at each point necessarily translate into improving the relative phase

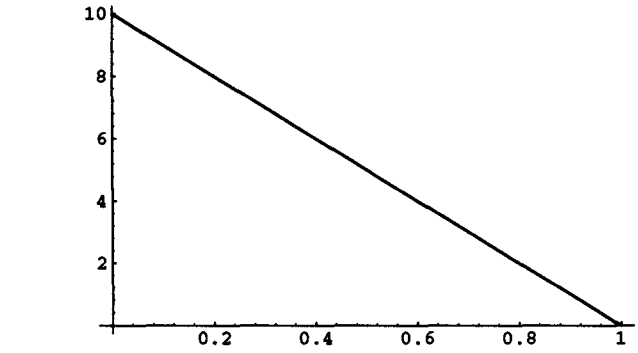

Change in distance from target parameterized as a function of time

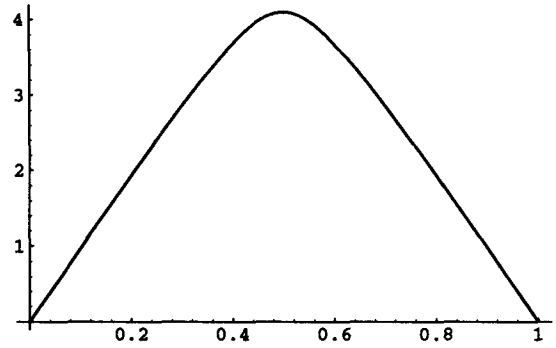

Change in amplitude error

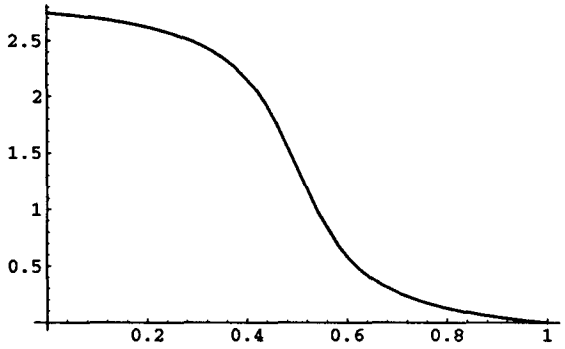

Change in phase error

Figure 3.19: Effect of change in distance along a straight path to a target on amplitude error and phase error.

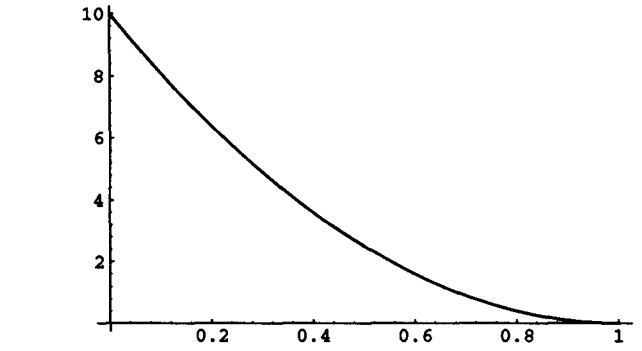

Change in distance from target parameterized as a function of time^2 to reflect square weighting

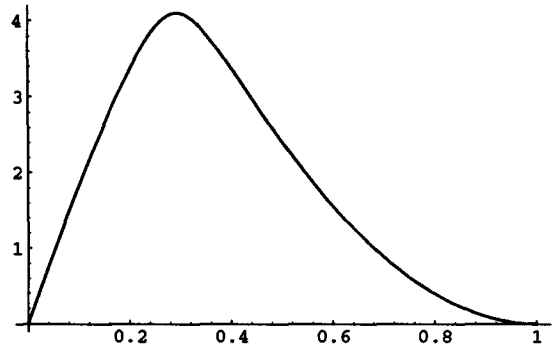

Change in amplitude error

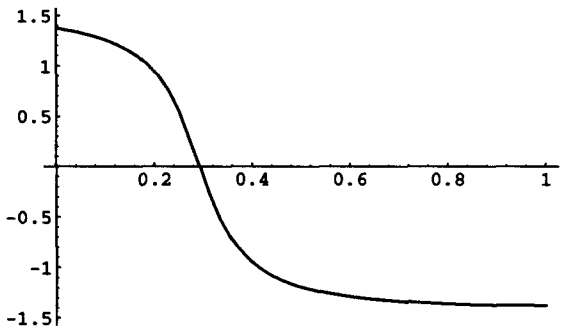

Change in phase error

Figure 3.20: Effect of reparametrized path to a target on amplitude error and phase error.

at each point, (and certainly not to necessarily improving perceived angle).[7] All that is guaranteed is that the least squares error relative to the new target will not get worse, since the target is brought closer in this sense to the achieved result.

3.4 Relative Phase

Since relative phase is important in localization at frequencies below 1500Hz [95], it would be nice to preserve relative phase between the ears. A different modification of Algorithm 1, which traded phase for amplitude, is to minimize some criteria over a set of points by changing the phase of the new target, but maintaining the original amplitude field and the relative phase between those samples of the original target. For example, in Figure 3.21 the distance between two targets and their respective attained values is reduced while mantaining the amplitude of the targets and the relative phase between the targets. Thus preserving selectively the relative phases in the target field allows some adjustment of the target.[8] In Algorithm 1, the phase was essentially thrown out at each point in each iteration. Another view is that the soundfield was replaced by a new target soundfield where the phase of the source was adjusted independently at each point such that the error between the achieved value and the desired value was minimized for some criteria. Any improvement in the next iteration would benefit the approximation with respect to amplitude. This idea of finding the minimum distance by adjusting the phase of the source independently for each sample point can be extended to a set of points. If this is done, then the relative phase of the target is maintained between the points in the set although the actual phases have been adjusted. The next iteration can again do no worse, and is likely to improve things, at least with respect to amplitude.

3.4.1 Creating Distinct Sets

Since the original target's relative phase is maintained on sets of sample points, the choice of these sets will have a big effect on the outcome of the entire process. In a situation where there is knowledge of where the ears will be, this information can be used to advantage. Such a case commonly occurs when an audience is seated to face the same direction. Most ears will be lined up in rows, so relative phase along the rows will have a big effect on everyone's localization of sound, while the relative phase between rows will not effect localizations. Rows are the most common example of audience arrangement, but any pattern of audience placement which allows knowledge of the likely ear positions of audience members can be exploited the same way. Relative phase between the ears is maintained as much as possible, at the expense of relative phase elsewhere in the field.

The only parameter being used in this case to generate the new target field can be viewed as changing the relative phase between distinct sets (relative phase being maintained within each separate set). The more distinct sets which can be made the better, since there is greater freedom in generating the new target. There are at least three reasonable ways to separate the sample points into sets in order to maximize the number of distinct sets while maintaining the likelihood of benefiting the audience.

[7] In fact when simulations of approaches 2 and 3 were attempted for the setup used in section 3.4.2, all the relevant error critera got worse. What was surprising was that the phase error itself also got worse, which may be due to the least squares weighting of the solutions.

[8] If relative phase is preserved globally, then no improvement can be made by adjusting it.

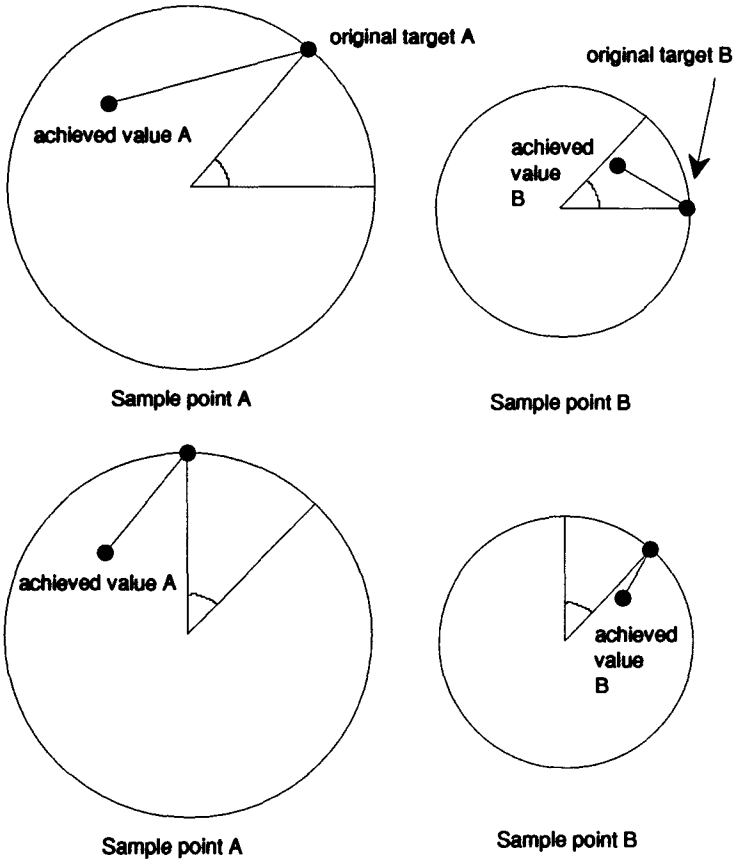

Figure 3.21: Maintaining the amplitude of two targets and their relative phase while reducing the total error (difference) with the achieved values.

1. In the situation where the audience is seated in lines, but the exact location of their ears is not known, the relative phase between these lines is certainly not necessary to keep.

2. Since most pairs of ears are about the same distance apart, the relative phase must be preserved at a certain period along the line rather than everywhere in the line. Figure 3.22 demonstrates this principle. Thus, each line can be broken into further sets provided that the sampling interval along that line is finer that the average distance between a pair of ears. The number of sets per line is simply the number of samples per the average ear separation. While this increases the number of sets, the benefit from these sets may not be very large, since there is very little room to maneuver between samples.

3. Another useful way to increase the number of distinct sets is to allow breaks along the line. Since the relative phase across such a break is 'discarded' to benefit relative phase elsewhere, chances are that the relative phase across any given break will worsen. Thus, the person whose ears cross a break will get worse relative phase than otherwise. However, if these breaks are chosen such that the error was very bad there anyway, then the worsening relative phase will not be a problem, since there was no way that that person could have good enough relative phase to matter. Furthermore the points between which relative phase is the worst are the ones which are holding up the relative phase error at the other points. Each iteration the target field for a set is replaced with one having the best fitting phase and if the worst relative phase points are removed from a set, the maximum benefit is achieved on the rest of the set.

3.4.2 Preserving Relative Phase in the Target

A simple example of this procedure was done for the setup in Figure 3.23, where the speakers are at (25,25), (25,75), (75,25) and (75,75) with 9 sample points arranged in three rows in a square from (45,45) to (55,55). The original source is at (100,50). The wavelength used in this example is 12 units.

A least squares solution is taken, and the squared error per row between this solution and the target as a function of the phase of the target is shown in Figure 3.24. In each row reduction is possible so the phase adjustments on the target are made to minimize total squared error per row (while maintaining relative phase per row). After adjustment, the effects of further phase adjustment on the target maintaining relative phase per row is shown in Figure 3.25. Finally, the least squares solution over the entire set is found using the new target, which is shown in Figure 3.26 where again the plot shows the effect of phase adjustments per row maintaining relative phase. The amplitude of the squared error per row is much smaller in rows 1 and 3, and about the same in the middle row. The resulting minimum points are very close to 0 phase adjustment. This indicates that although further iterations are possible, the gain on the next iteration is likely to be small, since very little adjustment of phase would be used to minimized the error. However, as was seen in Algorithm 1, a saddle point where very little improvement takes place each iteration can eventually give way to a much improved solution, so it may be worth iterating this expanded algorithm just in case.

Figure 3.27 shows the process another way. Here, information from the three sample points per row are superimposed by row. The circles indicate the target amplitudes, the

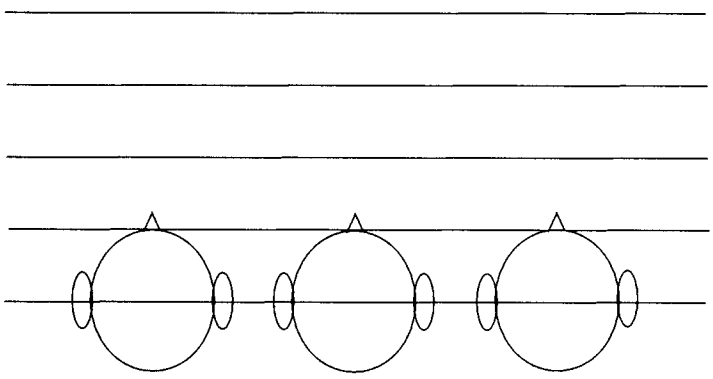

Equal phase contour of a plane wave

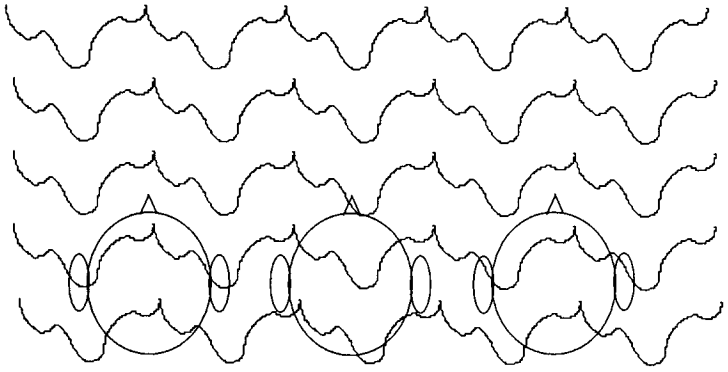

Equal phase contours periodic at ear separation intervals,
the relative phase between ears is the same as a plane wave

Figure 3.22: Maintaining relative phase between ears using periodic sets.

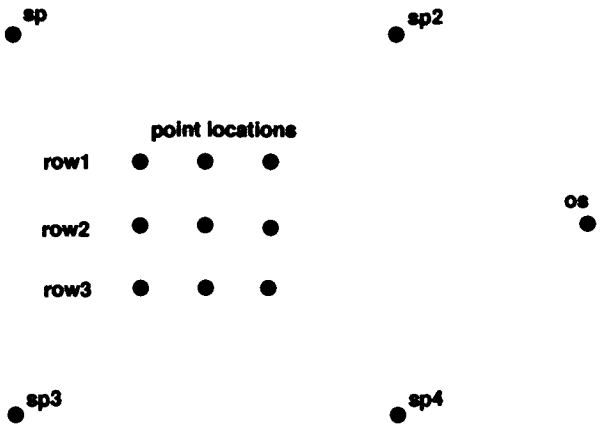

Figure 3.23: Initial setup.

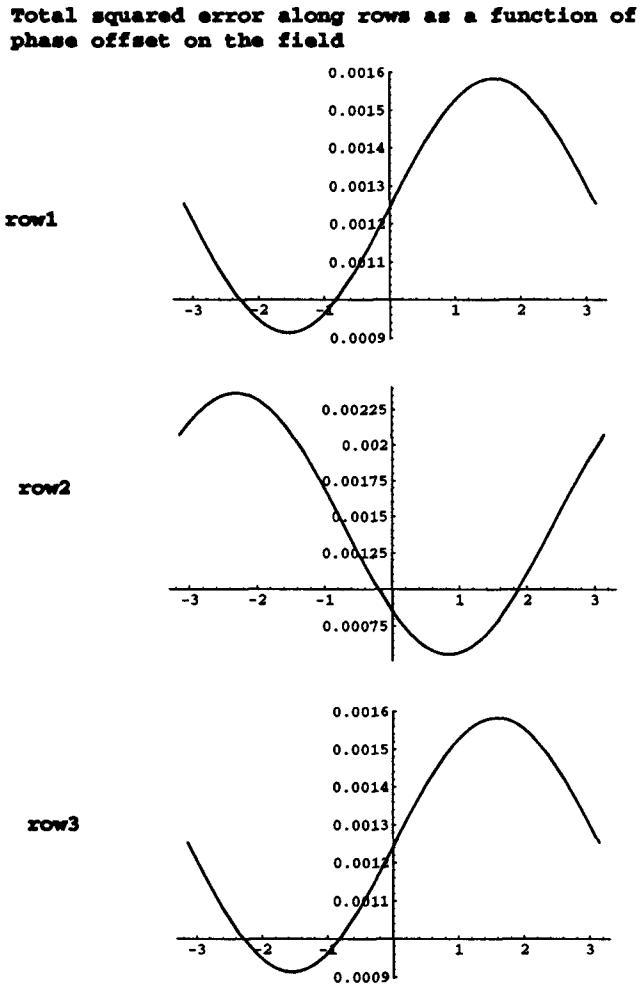

Figure 3.24: Squared error per row as a function of the phase of the original source.

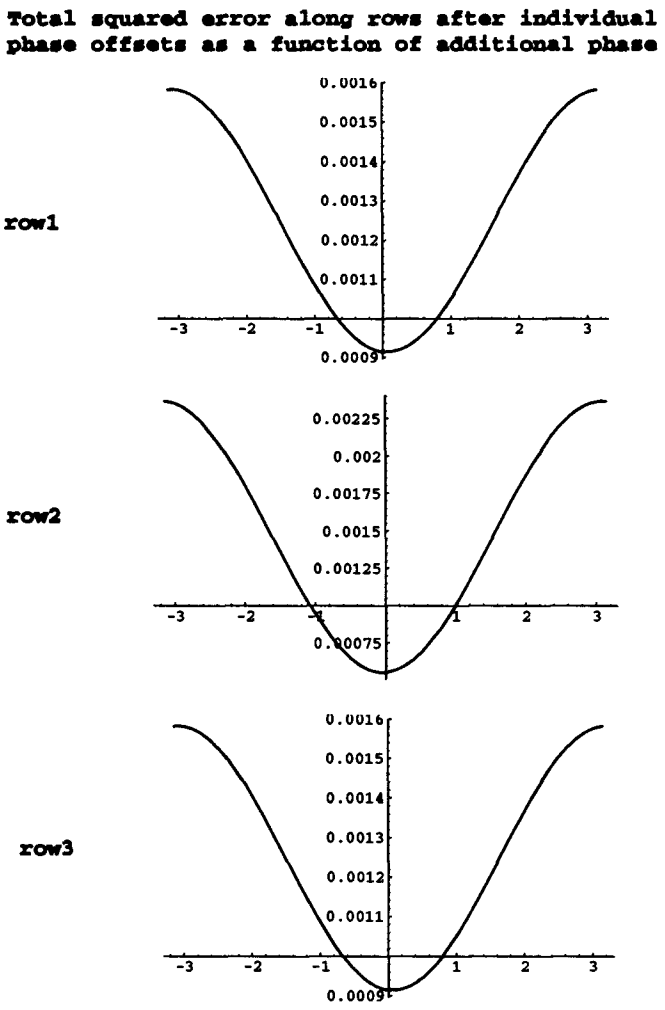

Figure 3.25: Effect of further phase adjustment after the target has been replaced.

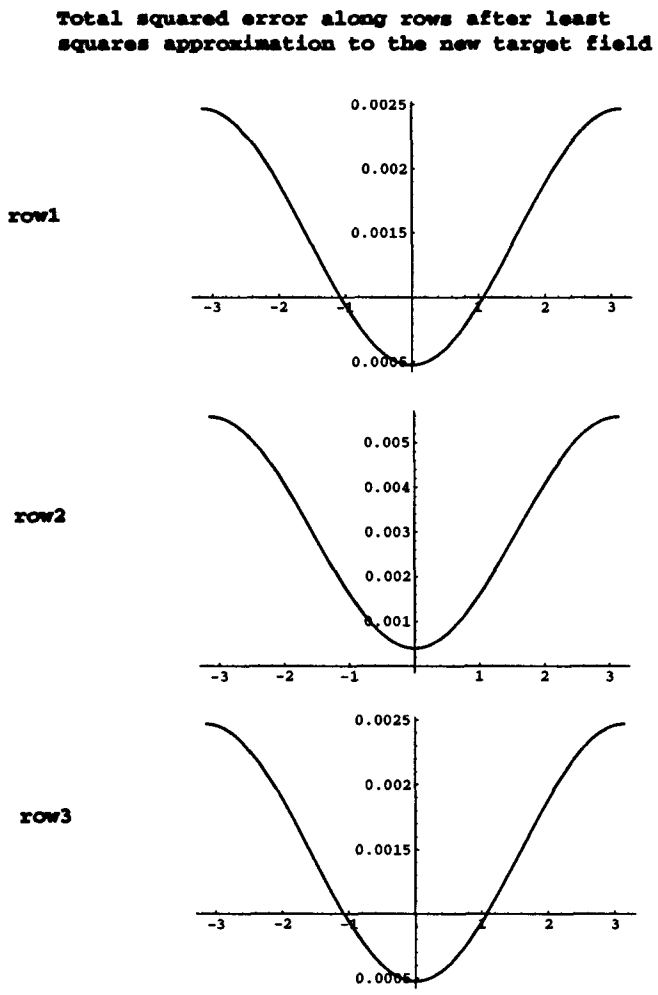

Figure 3.26: Squared error per row on the new target as a function of changing its phase (after a least squares approximation).

lines connect the achieved values with the target values. The first set shows the initial least squares solution. The second set shows the adjustment of the targets such that relative phase is maintained in each row, but the total squared error (length of the lines) is minimized. The final set shows the new least squares solution.

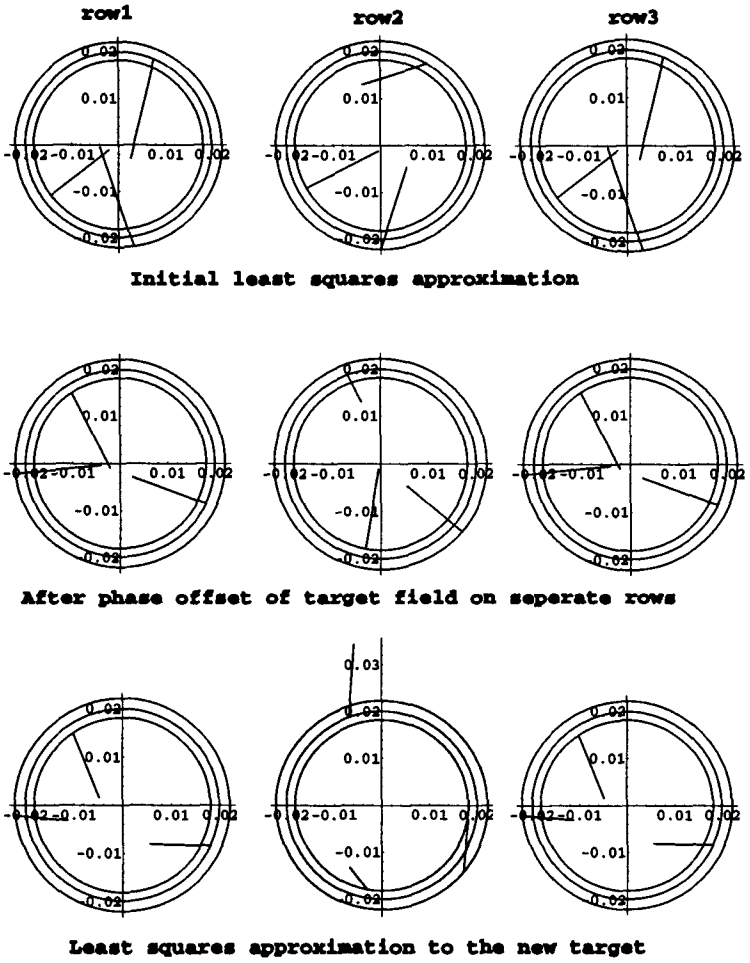

Figure 3.27: All the stages involved in one iteration of the expanded algorithm.

3.4.3 Perceived Source Direction.

The reason for maintaining relative phase among a set of points is to maintain or improve localization for each iteration. However, if just a straight least squares approach is taken, then the amplitude is the biggest beneficiary in most cases. The achieved relative phase among the sets may actually worsen, but only when the amplitude improves. Although what actually happens depends on the situation, a few examples seem to indicate that maintaining the relative phase among a set tended to maintain the overall phase error at about the same level, while allowing for improvement with respect to amplitude error. Compared with the situation where each individual sample point had its target phase replaced with its own

achieved phase, then the improvement in amplitude correctness was much better, but the degradation of relative phase was much worse.

In this particular case, the use of the above technique resulted in a very slightly improved average relative phase error, but the phase error got worse between four of the six pairs. This highlighted the need to attempt to put the error in terms of perception. A simple model of perception based on the relative phase of the paired sample locations was devised. Given two sample points distance d apart, the difference in phase between the two caused by a monopole at some distance location is due to the different path lengths to the source. If the source is far away, the distance in path length is approximately $d\cos(\theta)$, where θ is the angle to the source from one of the sample points and the wavelength is wl. The difference in phase between the two samples is then

$$\frac{d\cos(\theta)}{wl} \tag{3.10}$$

Thus, to find θ from the phases at the sample points, ϕ_1, ϕ_2 respectively, if

$$(\phi_1 - \phi_2) = \frac{d\cos(\theta)}{wl} \tag{3.11}$$

then

$$\theta = \arccos\left(\frac{wl(\phi_1 - \phi_2)}{d}\right) \tag{3.12}$$

In calculating the error $|\theta_{old} - \theta_{new}|$, the θ_{old} used was not the actual angle, but the perceived angle using this method of the original source. It was found that even though the average relative phase error improved slightly with one iteration, the average perceived angle error got worse. However, the main improvement was with regard to amplitude. It is also interesting to note that not all relative phases produce 'real' perceived angles. For all possible angles, the relative phase has a certain limit, which if exceeded, causes $\left|\left(\frac{wl(\phi_1-\phi_2)}{d}\right)\right| > 1$ giving a complex number when the inverse cosine is taken. This was dealt with by just taking the absolute value of the error (as was done for all the errors), since the result is a number greater than π which still gives some indication of how far off the result is in terms of this model of perception, since for perceivable results the error cannot exceed π.

3.4.4 Least Squares Error of Differences

Other approaches were also taken. An interesting comparison can be made between the expanded version of Algorithm 1 just described, and the algorithm it was derived from, where all the sample points have their phase replaced by the achieved phase to create a new target field. When done iteratively this process was called Algorithm 1, and like the expanded algorithm, was only iterated once, i.e., one new target field was created and approximated. Another approach tried was to do a least squares approximation to the difference between the points where relative phase is to be maintained. It seemed intuitive that such an approach might naturally preserve the relative phases between selected samples

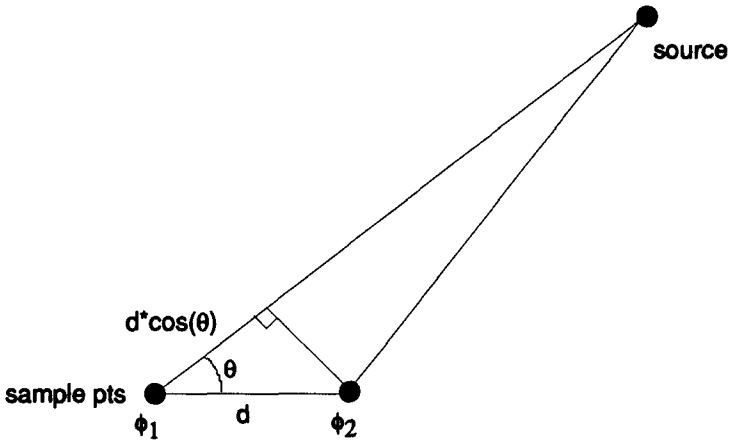

Figure 3.28: Diagram showing the setup for a simple model of perceived angle

in trying to minimize the difference error. In this particular case, the matrix

$$D = \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix} \tag{3.13}$$

was used where D was applied like a weighting to both sides of the overdetermined system $D \cdot Ax = Dy$, where A describes the transfer from speaker values x to more sample point values y. The resulting least squares solution is $x = ((DA)^*DA)^{-1}(DA)^*Dy$ or $x = (\hat{A}^*\hat{A})^{-1}\hat{A}^*\hat{y}$ where $\hat{y} = Dy$ and $\hat{A} = DA$. Notice that the number of simultaneous equations in this new formulation has been reduced from 9 to 6. In principle, this approach could turn an overdetermined system into a determined one or even an underdetermined system or vice versa depending on how many differences are taken. What is lost in this least squares of selected differences approach is any attempt to determine what the actual values are at individual samples as compared with the original target. Instead, only selected differences between samples are approximated to the corresponding differences in the target. This leaves no reason to assume that the individual values at the reproduced samples are anywhere near the target values. A solution to this is to combine the normal least squares formulation with this differencing version. This combination can be easily achieved by including an identity matrix I with the differencing matrix D to get

$$\hat{D} = \begin{bmatrix} I \\ D \end{bmatrix} \tag{3.14}$$

Using $\hat{D}$ the same way as D was used results in a least squares solution for the actual target values and selected differences simultaneously.

Table 3.1: Comparison of several least squares based approaches.

Method	$\|Ax-y\|^2$	$(\|Ax\|-\|y\|)^2$	$\frac{\|\Phi_{old}-\Phi_{new}\|}{k}$	$\frac{\|\theta_{old}-\theta_{new}\|}{k}$
straight l.s.	0.00333691	0.00221786	0.769307	0.993029
Algorithm 1	0.00114433	0.00109584	0.793131	1.13395
expanded Algorithm 1	0.00135572	0.00100008	0.76473	1.30951
difference matrix	0.00334448	0.00209985	0.757798	1.0122
difference and identity matrix	0.00334085	0.00213335	0.751794	0.989128

3.4.5 Single Iteration Comparisons

The results of these five methods of approximation are compared in the following table with respect to total squared error, total squared amplitude error, average absolute relative phase error between $k = 6$ pairs of selected points, and average absolute perceived angle error from the $k = 6$ pairs. In the case of the two algorithms, the evaluation of $|Ax - y|^2$ is based on the latest target field where the phases have been adjusted. If this wasn't done, the values of $|Ax - y|^2$ would get worse for the algorithms, but that isn't a fair or useful evaluation, since the first target field was not the target field they were going after. Comparing how close they come to obtaining the latest target is a useful comparison, since the latest target has all the 'important' aspects still in it, though what is deemed important can be different such as here with amplitude only in Algorithm 1 or amplitude plus relative phase along rows in the extended algorithm. The results for the previously described example are given in table 3.1.

These results are quite interesting. Comparing everything with straight least squares, only the iterative algorithms improved the total error of column 1. This was done by cheating, in that their comparison was made with a different target which already had less total error than the initial least squares solution. In the sense of least squared error relative to the original target, the least squares solution is by definition the best. Thus, the differencing matrix and hybrid differencing matrix had worse total error, since comparison was made with the original target. However, all four new approaches were better than least squares with respect to $(|Ax| - |y|)^2$, but the iterative approaches were significantly better.[9] Regarding relative phase between the selected samples, Algorithm 1 which just threw out the phase information got worse, which is no surprise. The iterative algorithm which maintained the relative phase between these points in its new target got an improved result, but only slightly. The best results in this category came from the differencing matrix and hybrid approach, with the hybrid(!) doing the best.[10] This indicates that maintaining relative phase in the target is not enough to improve relative phase, but tends to hold it at around the same total error while improving the amplitude error significantly. Finally, the perceived angle actually got worse by the greatest amount using the expanded algorithm! The other approaches did not affect the perceived angle much, but only in one case was the situation improved. This highlights the large distinction between improving a measure of average absolute relative phase error, and average absolute perceived angle error. The situation becomes even more involved when relative amplitude is included in the model of perceived angle of arrival.

[9] Here, it does not matter that the target has been replaced in the iterative approaches, since the amplitude of y is maintained.

[10] To be fair, the differencing matrix did better with respect to $(|Ax| - |y|)^2$ which more than compensates for the relative phase error in minimizing just the differences.

Table 3.2: Comparison of hybrid least squares based approaches.

Method	$\|Ax-y\|^2$	$(\|Ax\|-\|y\|)^2$	$\frac{\|\Phi_{old}-\Phi_{new}\|}{k}$	$\frac{\|\theta_{old}-\theta_{new}\|}{k}$
straight l.s.	0.00333691	0.00221786	0.769307	0.993029
exp. alg. 1 and diff. matrix	0.00142603	0.00102727	0.769499	1.23675
exp.alg.1 and diff./id. matrix	0.00139216	0.00101167	0.769189	1.25314

3.4.6 Combining Approaches

The previous table suggests that perhaps a combination of the extended Algorithm 1 with the alternative matrix approaches might be fruitful. Ideally, the algorithm would allow the reduction of amplitude error by adjusting the target, while the matrices involving the differencing matrix seem to work better at improving the average absolute relative phase error. It makes no sense to combine Algorithm 1 with the differencing matrices, since in creating the new target Algorithm 1 throws out the relative phase information that is the goal here. However, the expanded Algorithm 1 can be combined with the alternative matrix approaches, since it maintain the vital relative phase information. Table 3.2 shows what happens when the expanded Algorithm 1 is used in conjunction with the two matrix approaches for one iteration, where the new target has the same relative phase between relevant points and the amplitude of the original target is maintained. The solution for which the target adjustment routine minimizes error is the modified least squares solution and the final solution is also the modified least squares solution with the same modification.

This shows that the average relative phase error is maintained at about the same level as the straight least squares approach and worse than just using the matrix adjustments. However, the amplitude error was reduced to the low levels typical of the iterative algorithms. Again, the perceived angle has a complicated relationship with the relative phase error, and minimizing the average absolute phase error is not the right approach to improving perceptual location error, even with the simple perceptual localization model used here. Rather, the average absolute phase error has been included for comparison to show that none of the methods so far discussed improve relative phase errors significantly.

Part of the reason relative phase does not improve much in any of these examples is that the sole goal of the straight least squares approach is to minimize the squared error in terms of distance. While it is clear that this approach has some relevance for optimizing with respect to amplitude, the effect on phase is often not large. The effect of traveling in a straight line at constant speed from (1,5) to (1,-5), shown in Figure 3.18 on phase and amplitude error relative to a target at (1,-5) can be see in the Figure 3.19. This can be interpreted to represent what effect an equally weighted attempt at minimizing distance error has on amplitude and phase error. If the property of the squared weighting is represented by adjusting the time progression by the reparameterization $t \rightarrow (1-(1-t)^2)$, the situation changes to Figure 3.20. This new figure can be interpreted as showing how 'quickly' a parameter changes assuming the emphasis placed on larger distance errors by the usual square 'weighting'. Although the typical least squares situation is faced with many competing constraints, the ultimate goal is to minimize the squared distances, which can be see to have a nonlinear relationship with both amplitude and phase, so the actual solution will not be linearly related to these parameters. The situation becomes even more tricky when looking at relative phase.

Table 3.3: Other optimizations with original source at (100,50).

Method	$\|Ax-y\|^2$	$(\|Ax\|-\|y\|)^2$	χ	$\frac{\|\Phi_{old}-\Phi_{new}\|}{k}$	$\frac{\|\theta_{old}-\theta_{new}\|}{k}$
l.s.	0.00333691	0.00221786	1.11886	0.769307	0.993029
phase	0.00139859	0.00102727	0.365748	0.766569	1.28572
distance	0.00335237	0.0021614	1.2319	0.557017	0.584526
2ce distance	0.00135794	0.00100211	1.54849	0.765198	1.30393

Table 3.4: Other optimizations (50,0).

Method	$\|Ax-y\|^2$	$(\|Ax\|-\|y\|)^2$	χ	$\frac{\|\Phi_{old}-\Phi_{new}\|}{k}$	$\frac{\|\theta_{old}-\theta_{new}\|}{k}$
l.s.	0.00333691	0.00221786	1.11886	2.07717	1.0529
phase	0.00330047	0.0021621	1.3247	2.39911	1.18401
distance	0.00327129	0.00209975	1.22145	2.42247	1.19622

3.4.7 Minimizing Other Criteria

Algorithm 1 has many possible interpretations besides replacing the phase of the target field at each sample point. Another possible interpretation of Algorithm 1 is that of minimizing the squared difference between the achieved results and the target through phase adjustment. This interpretation was used for the expanded algorithm, since it guarantees that the next iteration can do no worse. It is also possible to view Algorithm 1 as minimizing the total (non-squared) error.[11] This approach has less basis in intuition, but does have the nice effect of minimizing total error. Also, Algorithm 1 can be seen as minimizing the phase difference at each sample point, in that case trivially, since it is obvious that replacing the target phase with the achieved phase produces zero phase error. Also, minimizing the total phase error can be seen as intuitively based if maximum improvement is sought regarding the difference in amplitude error. When each point is treated individually, all these approaches are equivalent. However, if relative phase and amplitude are maintained on a set of points then these approaches have subtle differences. Table 3.3 shows how creating the new target with these different criteria affects the results of one iteration. A new column is added, χ, which is the average absolute phase error at all the sample points.[12] Surprisingly the non-squared distance minimization produced the best relative phase and perceived result so far. Because of this a second iteration of that approach was taken which destroyed the good phase relationships but continued to improve the amplitude approximation.

To see if a similar first iteration effect for the distance optimization approach occurred with a different configuration, the original source was moved to (50,0). Because of the symmetry, the initial least squares solution is identical to when the source was at (100,50). The results are shown in table 3.4.

Here, the distance optimization did no better than the other approaches. The low relative phase error using this approach when the original source was at (100,50) seems to be a matter of that particular setup. Part of the reason relative phase does not improve much in most of these examples is that the only 'goal' of the straight least squares approach is to minimized the squared error in terms of distance. While it is clear that this approach has some relevance for optimizing with respect to amplitude the effect on phase is often not large.

[11] or error raised to the power of any other non-zero real number

[12] as opposed to relative phase between select samples

3.4.8 Weighted Least Squares

One approach to improving the least squares method is to create a weighting based on the perceptual error. By letting $y_2 = Dy_1$ where D is some weighting matrix, in these cases diagonal, a solution is sought where $DAx_1 = Dy_1 + De$ with $De = D(y - Ax)$ so

$$\hat{x}_1 = ((DA)^*DA)^{-1}(DA)^*Dy_1 = (A^*D^2A)^{-1}A^*D^2y_1 \tag{3.15}$$

when D is diagonal and real. This is similar to the application of the differencing matrix earlier.

If a weighted least squares process is to be used to reduce the perceived angle error, the question naturally arises as to how the weighting should be calculated. The difficulty is that perceived angle depends on relative phase while the weighting is done to individual squared errors. One simple solution is to apply the magnitude of the perceived angle error as a weight equally to both the samples which determine it. Thus the weight of a particular sample is the sum of the perceived angle error on both sides of it. If the sample is on an edge, then the only perceived angle error which it is related to is applied. If the arrived at weighting is also applied to the target replacement maintaining relative phase method to find the minimum point prior to using the weighting in the least squares solution, then the total weighted squared error is guaranteed to improve. However the perceived angle is not guaranteed and as can be seen in Figure 3.29, actually doesn't improve very well. These figures again relate only to the example with the original source at (100,50) and the same sample points as has been used throughout this chapter. An interesting pattern is created where the average magnitude of the perceived angle error travels almost predictably across three levels of value. The lowest value is achieved after the first iteration, at about 0.65. It seems as if this particular method can get stuck in a chaotic cycle.

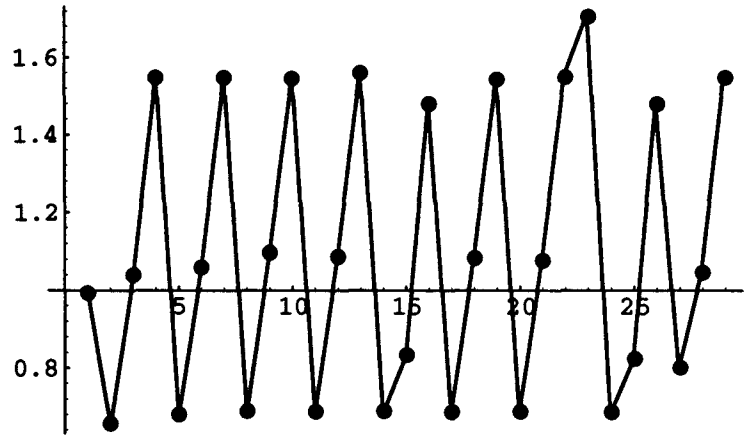

Figure 3.29: Thirty iterations of weighted least squares with weighted least squares minimization

If the weighting is not applied to the new target generation method, but instead the unweighted least squares error is minimized while maintaining relative phase among the rows, then a different pattern emerges. After some initial bad results, the algorithm seems

to become localized in an improved neighborhood, at least in this case over the first 60 iterations as seen in Figure 3.30.

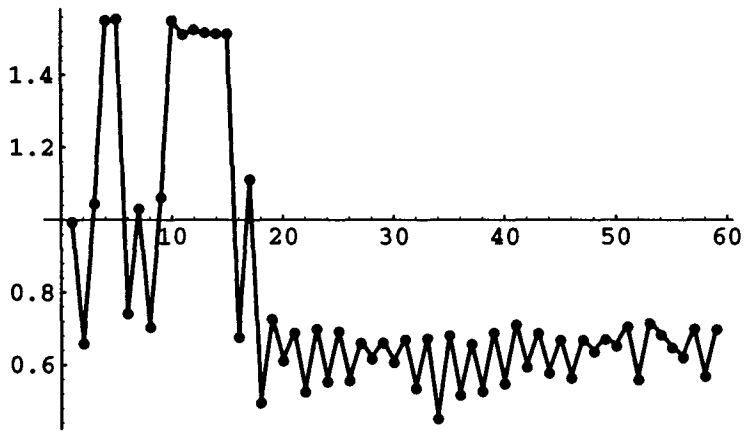

Figure 3.30: Sixty iterations of weighted least squares with un-weighted least squares minimization.

3.5 Conclusion

The methods mentioned here are not meant to be exhaustive. On the contrary, these simple ideas are meant to serve as a starting point for more sophisticated methods. More complex models of perception and a more complex environment are steps toward some useful implementation. Even so, these simple experiments show that it is possible to use new target generation based algorithms to improve perceptual factors. This seems especially promising when there is some control over the optimization procedure, such as weighting a least squares approach. Such methods are very good at improving the total amplitude error over the sample points. Relative phase, being related to perceived angle of origin for a sound, is useful to preserve because it tends to keep the relative phase from becoming too much worse, (as happens when relative phase is not preserved). However, since the relationship between relative phase and perceived angle is not linear, changes brought about for perceived angle can be unpredictable, even if the amount of relative phase error improves. By understanding the relationships involved between amplitude, relative phase and perceived angle, one can try to maximize the benefit at each iteration of the algorithm used.

Chapter 4

Implementations

4.1 Introduction

Once decisions have been made about the overall philosophy regarding perceptual tradeoffs, and appropriate methods are found to pursue these goals, certain issues of implementation still arise, some of which are examined briefly here. Numerical problems arise when trying to do a least squares solution on large ill-conditioned matrices, which will often result when the path differences from various speakers to various sample points is not large. An simple iterative technique is introduced in a non-rigorous fashion which may be useful in closely approximating a least squares solution. In addition, some extension from single frequency to continuous frequency is needed, (unless the audience is content to listen to a single tone). A straight forward implementation using the short time Fourier transform is proposed, and aspects such as the convergence of Algorithm 1 at the frequency bins and error cause by the different arrival time of windowed packets at the sample points is discussed. Finally some sound examples of these techniques are discussed.

4.2 Numerical Problems

Up until now, only systems with five or fewer speakers have been looked at in detail. This is fine for the current situation, where stereo is the norm, but to pursue of soundfield reconstruction requires as many speakers as possible, coordinated very carefully. Unfortunately, two related problem arise when trying to expand the theory to larger numbers of speakers.

(1) As the number of speakers increases, they necessarily get closer and closer together. Since they operate on the same set of sample points, this closeness causes the matrices involved to have increasingly similar columns, which makes the matrices less well-conditioned. This causes great magnification of any error when the matrices are inverted, (or the equivalent).

(2) Since the number of speakers is increasing the size of the matrices involved, there are many more computations involved, causing greater numerical error.

Thus, the situation becomes one of giving more and more numerical error to a system that is becoming less and less tolerant of error. So, while the theory is still valid for arbitrarily many speakers, the implementation becomes unwieldy.

To illustrate this problem, consider the following setup (also for comparison with the techniques of the next section). Twenty-five speakers are placed along the line $y = 10$. If given an initial amplitude of 1 and phase of 0, they look like Figure 4.1

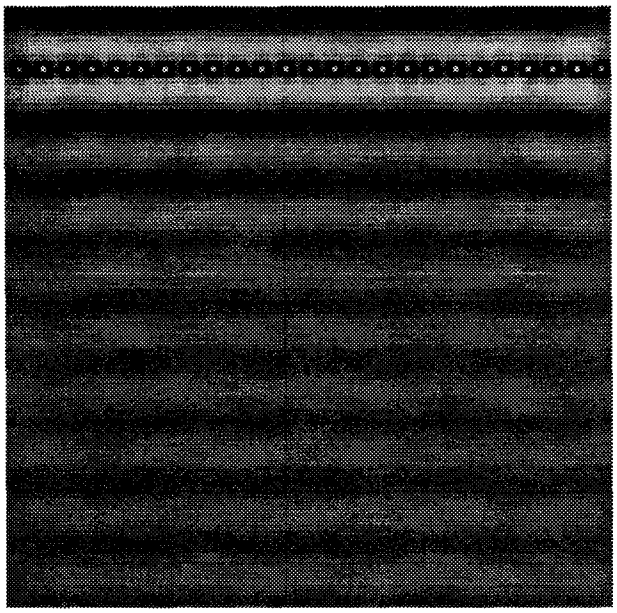

Figure 4.1: Diagram of setup.

Now, doing a least squares approximation of one wavelength (length 10) from a original source at (50,-10) on 121 sample points between (35,35), (35,65), (65,35), (65,65) in Matlab gives the result shown in Figure 4.2 (with resolution set to 1).

The amplitude plot at this resolution is almost completely black. Reducing the resolution to 0.01 shows that indeed the process seems to be trying to do the right thing within the sample area. Yet with a maximum error of around 104 times the original signal and an average error of 21 time the original within the region of interest (corresponding to about 40dB and 26dB respectively), the audience would be experiencing a much better approximation in a silent room.[1] The possible use of this technique in rock concerts has not be examined. The fact that this problem is so ill-conditioned for larger arrays can be related to the work of Putnam et al. [78] which shows that, even in the underdetermined case, the ability to invert a room response is related to the ratio of the largest singular value to the smallest singular value across all frequencies.

4.2.1 Projection Methods

Since the numerical problems involved in solving large matrix systems of the type needed are so bad, an alternate procedure is needed. One alternate approach is to project the target soundfield vector onto a vector produced by a single speaker. That speaker is made to match the target soundfield as closely as possible, then its contribution is subtracted to

[1] In fact, a least squares approximation should never do worse than turning all the speakers completely off!

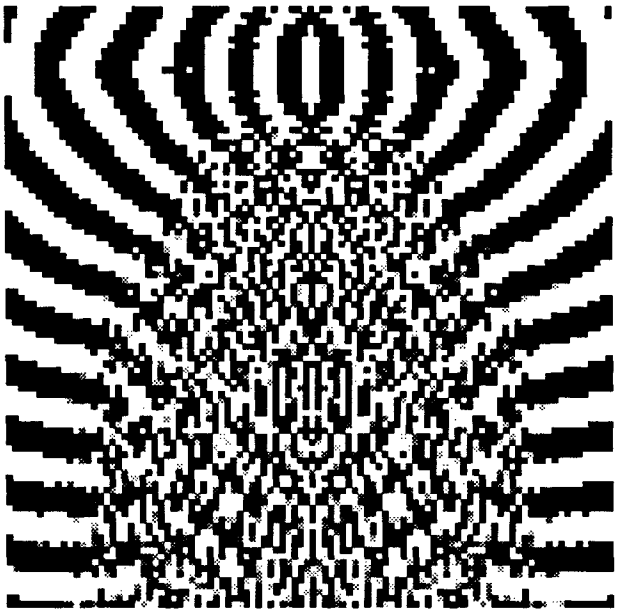

Figure 4.2: Real part of bad numerical error.

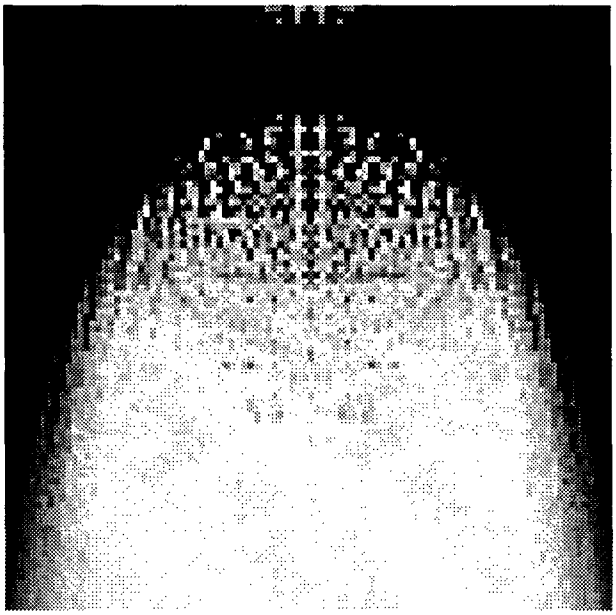

Figure 4.3: Amplitude of the difference with bad numerical error (0.01 resolution).

produce a new target soundfield. This process is repeated until no further improvement is obtained.

If the speaker vectors were orthogonal, then this process could remove the contributions of each speaker without affecting the contributions of the other speakers. Unfortunately this is not the case, but at least the speaker vectors are (almost always) linearly independent. Only contrived cases where two speakers have identical distances to all the sample points, or some scalar multiple of the identical distances, destroy this linear independence. Thus, in the case with an equal number of sample points as speakers, the speakers span the space, demonstrating what was shown with matrices in a previous chapter. Since the speaker vectors are linearly independent and span the space, they can be combined to create an orthonormal basis, for example by the well known Gram-Schmit process [42]. Unfortunately, finding an orthonormal basis is prone to the same numerical error described in the previous section.

Since the vectors are merely linearly independent, projection of the target vector onto a speaker vector allows removal of all components in that particular 'direction', but since other speaker vectors also have components in that direction, projection of the new vector onto other speaker vectors will add some of the directional component just removed (unless the other speaker vector is orthogonal with the first, which also only occurs in contrived situations, and very rarely in practice).

By way of a simple illustration, consider the vectors $(1,1)$ and $(1,0)$ trying to produce the target vector $(0,1)$. The solution is readily seen as $(1,-1)$ since $1\cdot(1,1)+-1\cdot(1,0)=(0,1)$, but to solve it by projections leads to the following. Projecting $(0,1)$ onto $(1,1)$ gives a 0.5 multiplier to $(1,1)$ since the vector from $(0.5,0.5)$ to $(0,1)$ is orthogonal to the vector $(1,1)$, and $(0.5,0.5)$ lies on the vector $(1,1)$, as shown in Figure 4.4. This component of $(0.5,0.5)$ is subtracted from the original vector $(0,1)$ to give $(-0.5,0.5)$. The projection of this new target vector onto $(1,1)$ is zero, i.e., all the $(1,1)$ direction has been removed. However since $(1,0)$ is not orthogonal to $(1,1)$, this subtraction has moved the target vector relative to its direction. Whereas the previous target of $(0,1)$ had a projection of 0 onto $(1,0)$, the new target $(-0.5,0.5)$ projects to $-0.5\cdot(1,0)$. Subtracting the projected component from the target gives a new target at $(0,0.5)$. This target again projects onto the vector $(1,1)$ but not $(0,1)$ since that directional component has just been removed. In fact the new problem is just a scaled version of the original problem! Since each time the target has been reduced by $\frac{1}{2}$, the limit is $(0,0)$ and by remembering how much of each component vector has been used to reduce the target vector each time, an infinite series is formed $(\frac{1}{2}+\frac{1}{4}+\frac{1}{8}+\frac{1}{16}...)(1,1)+(-\frac{1}{2}-\frac{1}{4}-\frac{1}{8}-\frac{1}{16}...)(1,0)=(0,1)+(0,0)$. The sum of these series gives the right answer $1\cdot(1,1)+-1\cdot(1,0)=(0,1)$.

4.2.2 Using Iterative Projection to Approximate Least Squares

Although graphically it is easy to see what is meant by this projection of one vector onto another, a clearer mathematical definition would be useful. The speaker vector being projected onto just determines a line, and that vector will be scaled to match the point of projection. Since the goal is to remove entirely from the target vector the directional component of the speaker vector, the difference between the two vectors should be orthogonal to the speaker vector. This orthogonal line representing the mismatch, or error, which cannot be removed by the speaker vector, is the least distance from the target vector (point in N-space) to the line. Since the line representing the error is orthogonal to the line on which the speaker vector lies, and connects this line with the target vector, a right triangle

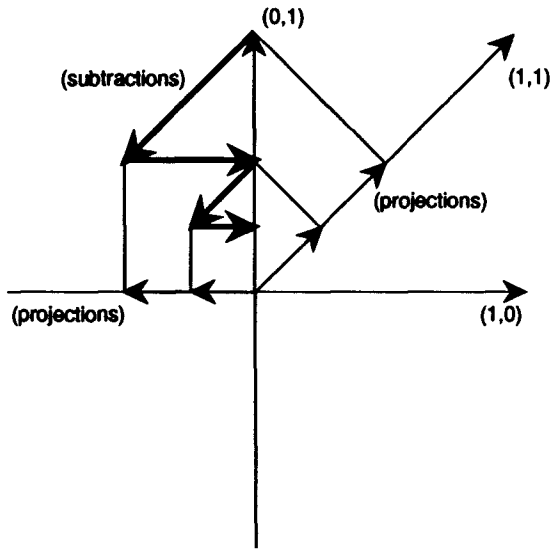

Figure 4.4: iterated projections on linearly independent nonorthogonal vectors.

is formed. By the Pythagorean theorem, the square of the target vector minus the square of the error vector gives the square of the correctly scaled speaker vector. Since the error vector has been determined to be the shortest from the target vector point to the line in which the speaker vector lies, the square of this error vector will also be the smallest. Thus, finding the projection of the target sampled soundfield vector onto a vector of a speaker's sampled soundfield amounts to finding the least squares solution. (see Figure 4.5)

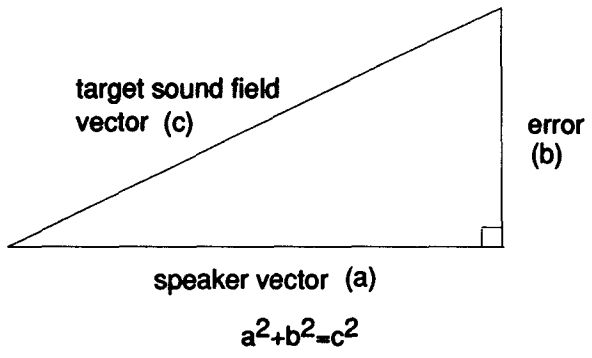

Figure 4.5: Projection of a soundfield vector onto a speaker direction when viewed in the same plane.

Subtracting the contribution of a speaker from the soundfield iteratively always reduces the error, (unless the contribution is zero). This is because each time the error, or remaining soundfield after subtracting a speaker's contribution, becomes the new soundfield. Each time the new error is smaller than the old error, which can be thought of as the hypotenuse

in Figure 4.5. Thus, the sequence will converge to the correct solution, though each speaker component may need to be projected onto infinitely many times. As an example, consider a setup with speaker at (25,25), (25,75), (75,25) and (75,75) and sample locations at (45,45), (45,55), (55,45) and (55,55) with wavelength 10. The original source location is at (60,1). Just taking the matrix inverse to solve the problem (numerical error is not a problem for this small setup) gives initial speaker signals of 0.6541 - 0.1581j, 0.4866 - 0.1899j, 0.1074 + 0.1900j and 0.1536 + 0.2673i respectively (see Figure 4.6). The same four digit precision values can be achieved by this projection method after 16 total projections. However, an infinite number of projections would be required to get exactly the right answer since each projection shifts things in the directions of the other speaker vectors.

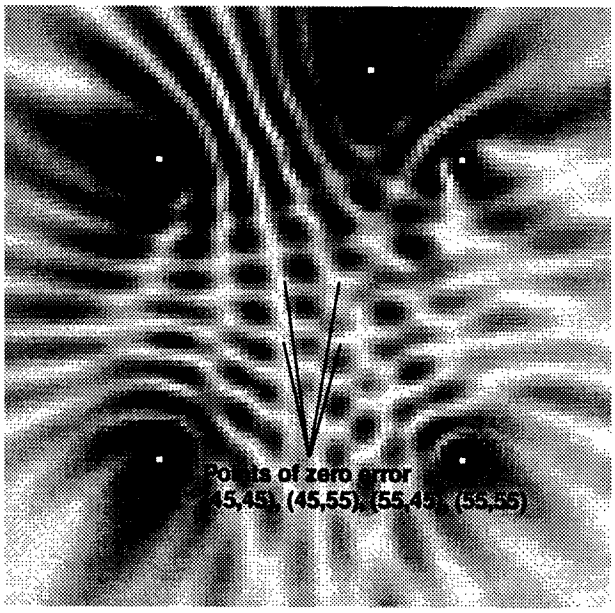

Figure 4.6: Amplitude of difference between speaker and source fields (resolution 10).

This convergence happens not only in the case where the number of speakers equals the number of sample points, but it also converges to the least squares solution when there are more sample points than speakers for the same reasons. In the limit, all directional components represented by the speaker fields are removed and the remaining error will be orthogonal to the space spanned by the speaker fields. Hence, the error vector will have the least distance to the hyperplane containing the speaker vectors which is equivalent to least squares as shown before.

The advantage to the projection method is avoiding numerical problems. Consider an example with twenty-five speakers, such as that in Figure 4.1. If the projection method is used to simulate a source at (50,-10), and iterated until the decrease in total error between projections is less than 0.0001, then the solution obtained is shown in Figure 4.7

Comparison of this solution with Figure 4.3 clearly shows that whatever numerical error occurs in iterating projection, it is much more tolerable than the error involved in matrix manipulations.

As a final example of the projection method, a comparison is made between the method used in chapter 1 to get Figure 1.13 (a rough treatment of a method by Berkhout [15]) and

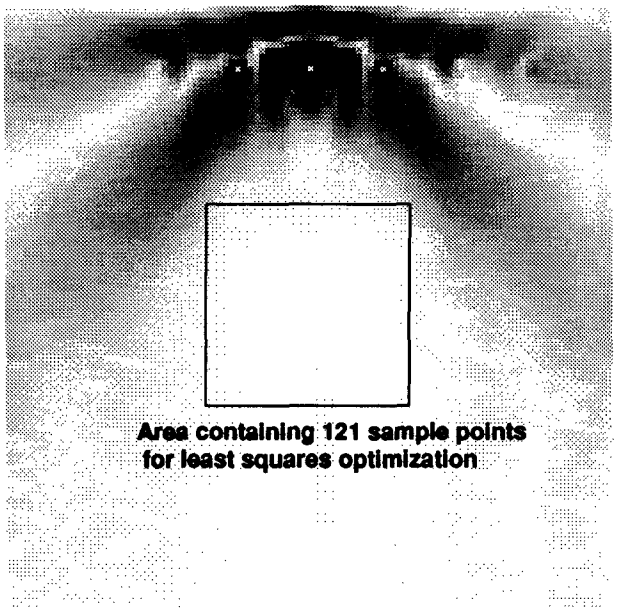

Figure 4.7: Amplitude of difference between speaker and source fields obtained by projection method (resolution 10).

iterative projection with the same setup, designed to approximate a least squares solution on the 121 points shown. A contour plot of the result from the projection method is given in Figure 4.8. The error of the amplitude is given in Figure 4.9

4.3 Applying the FFT

It is straight forward to get the least squares solution frequency by frequency. By using a descrete frequency decomposition of the signal, the appropriate adjustments can be made at each frequency, and the results recombined to provide the signals necessary at each speaker. One standard way to get such a frequency decomposition is by the short time Fourier Transform (STFT).

Up until now this dissertation has been focused on manipulations at a single frequency. In order to implement a system to do wavefront reconstruction, it is necessary to deal with non-monochromatic sounds. One way to use these single frequency manipulations is to dissect the signal into discrete frequencies using the STFT. The use and theory behind this powerful method has been well documented, so the aspects focused on here are the applications to wavefront reconstruction. The basic idea is to take a short duration discrete time signal, transform it into discrete frequencies, apply some of the operations discussed in earlier chapters to individual frequencies, inverse-transform the results for each speaker to get the signal at each speaker.

There are several difficulties with this approach. First, there is the problem that not all windowed packets arrive at the sample points at the same time. The following analysis of arrival times applies not only to windowed sounds, but to any change in the filters applied to each speaker. For instance, if a set of filter banks are applied to the speakers then it is

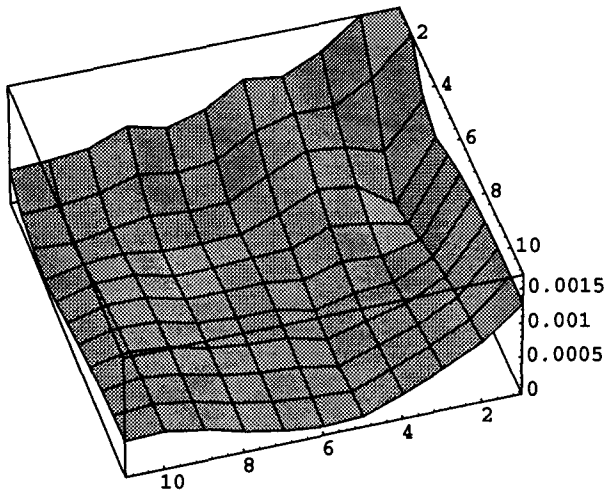

Figure 4.8: Amplitude of difference between speaker and source fields obtained by projection method.

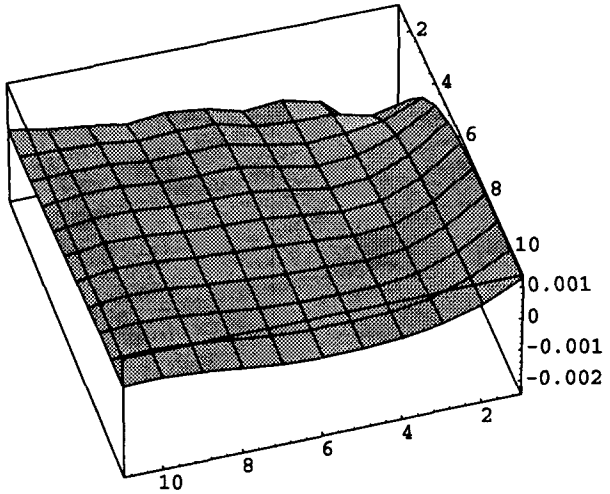

Figure 4.9: Difference in amplitude between speaker and source fields obtained by projection method.

useful to know what the effect of changing the filters at a certain time has in the soundfield.

Finding the DFT of a windowed discrete signal actually assumes the windowed information is periodic for all time since the sinusoids in the decomposition of the DFT are periodic. However, only one window length is used for each DFT. Thus, if the windowed information arrives at a point out of sync, the effectiveness of whatever approximation was used will be further reduced. Therefore, it is necessary to evaluate the relative timing of the windowed information within the region of interest. When the signal changes dramatically with each window, this analysis will indicate the amount of potential error due to window mismatch at a given point. Conversely, if the signal doesn't change much with successive windows, then the problem of window mismatch at points may not be very severe at all.

The analysis of this for rectangular windowed packets is straightforward. Consider a point in space with distances d_1, d_2, d_3, and d_4 from the speakers, a quadraphonic example. If the beginning of a window packet is released from the speakers simultaneously at t_0, then the arrivals of the beginning part of that windowed information at the point will be at $t_0 + d_1/c$, $t_0 + d_2/c$, $t_0 + d_3/c$ and $t_0 + d_4/c$ respectively where c is the speed of sound. Thus, the difference in arrival time is due just to the different amounts of time needed for the sound to travel from each speaker to the point in space. Likewise, if the end of the windowed information occurs at the speakers simultaneously at t_1 then the arrival of the end of the windowed information at the point occurs at $t_1 + d_1/c$, $t_1 + d_2/c$, $t_1 + d_3/c$ and $t_1 + d_4/c$. The duration in which the correct window (of information) is at the point in space is the earliest ending minus the latest beginning (see Figure 4.10). The duration of the window itself, i.e. $t_1 - t_0$ can be denoted as w to simplify the general result, that the duration of having the correct windows at a given point is

$$w + \frac{d_{closest}}{c} + \frac{d_{farthest}}{c} \tag{4.1}$$

Although this analysis is good for any number of speakers, its use is unfortunately limited to rectangular windowing. The shape of the window used is not taken into account if this type of analysis is done for non-rectangular windows. An alternative interpretation of the preceding rectangular window overlap analysis is to think not of only the difference between the latest beginning and earliest ending of a particular window, which gives the time boundaries, but rather the multiplication of all the corresponding windows together. The result for the rectangular window is the same, but for other windows shapes the inclusion of all corresponding windows in the multiplication gives a more useful result. This result reflects the relative amplitudes of the windows at a given time. The result, just as in the rectangular case, can be looked at in time to show when and how much the corresponding windows still correspond at a given point. In Figure 4.11 an example is given with a Hann window.[2] These amplitudes can be further weighted by the relative amplitude (due to distance) each has at a given point if desired. If the variation in amplitude is not large, this refinement makes computation needlessly cumbersome, since the normalization at each point must be relative to the best scenario for that point, and cannot be done globally. Also, the amplitude attenuation with distance has already been factored into whatever optimization was done in the frequency domain, so in some sense such a consideration may be redundant.

[2]Frequently this window is called 'Hanning,' although the inventor's name was von Hann. The use of this term may have something to do with emphisizing the similarity to the Hamming window. It is interesting to note that Hamming uses the term 'von Hann' [46]

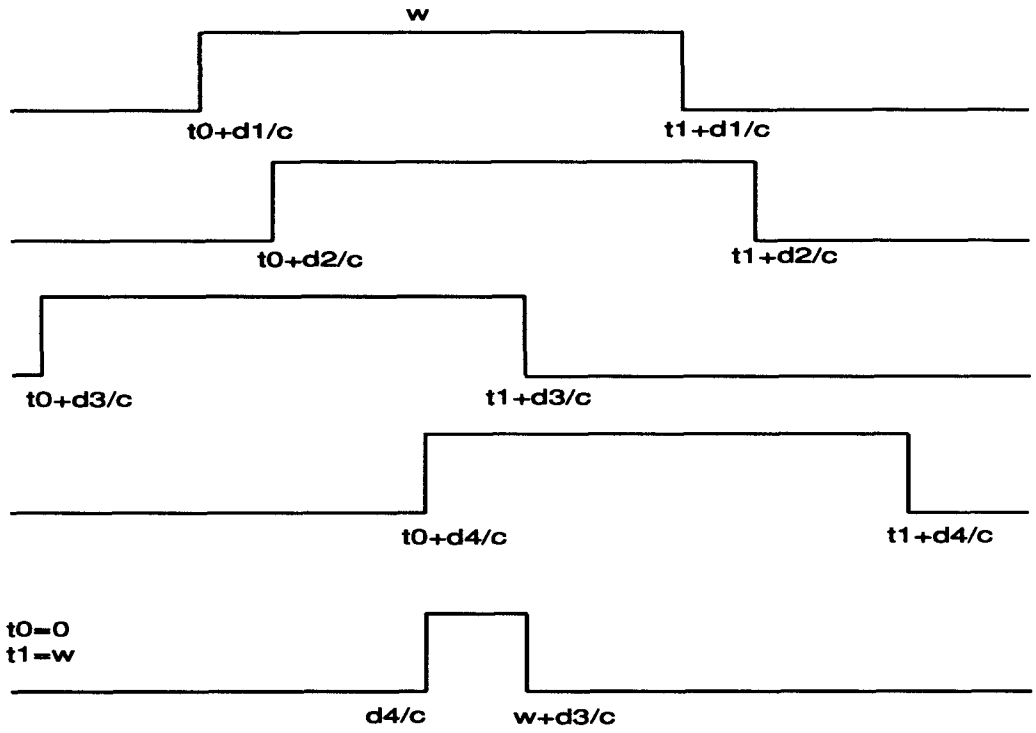

Figure 4.10: Window packet arrival times.

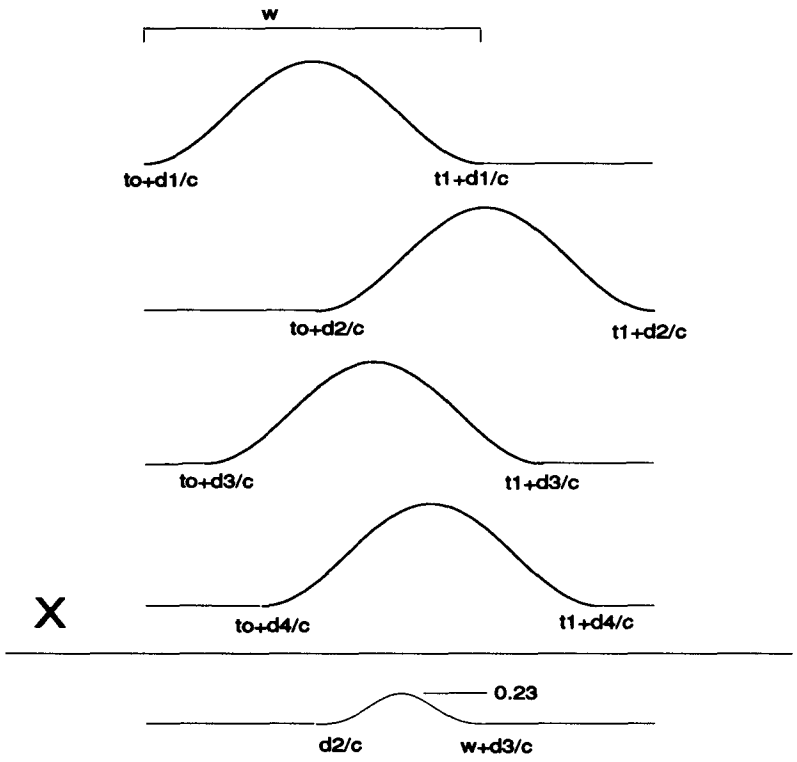

Figure 4.11: Window packets multiplied.

By integrating over time, a ratio with the best case of exactly corresponding arrival can be made (even if the best case does not actually occur anywhere in the field!). This ratio reflects the error caused by the desynchronization of windows of any shape, and can be interpreted as a percentage of synchronization accuracy. A plot of these percentages for the 121 points in the quadraphonic setup is given in Figure 4.12

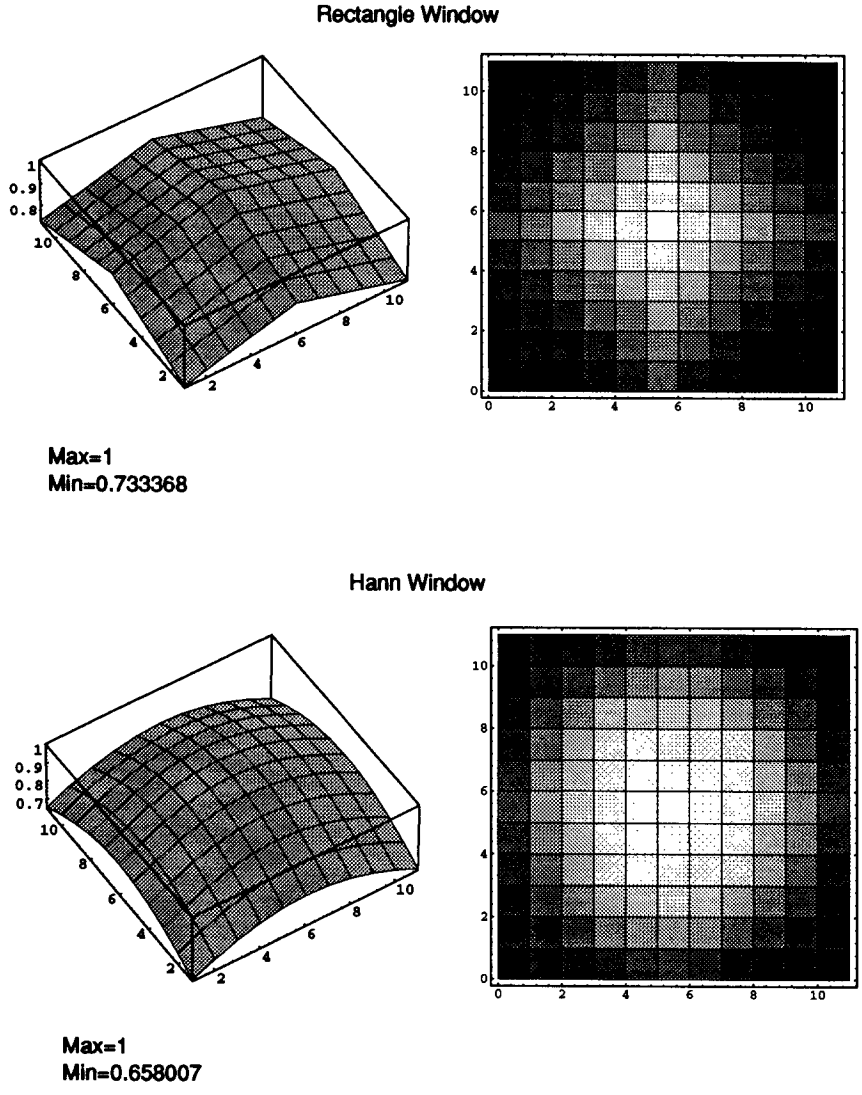

Figure 4.12: Percentage of correct window correspondence.

Since the window correspondence factor is very relevant when the signal changes dramatically from window to window, it would be nice to optimize the amount of correct window overlap with in a given region, and this can be done using standard techniques. For the examples which follow, the speakers were left synchronized.

4.3.1 Amplitude Error at Different Frequencies using Algorithm 1

Unfortunately convergence of Algorithm 1 at different frequencies can differ widely, even though the convergence of very close frequencies is usually similar. As was seen in a previous chapter, saddle points exist which not only cause some solutions to stick at a non-optimum result, but also effect the rate of improvement for at nearby frequencies. A wavelength of 53.125 units as well as the wavelengths corresponding to its first 63 harmonics is shown in Figure 4.13.[3] The fifth wavelength corresponds exactly to the single wavelength example used in a previous chapter to demonstrate the possibility of saddle points, 10.625 units.

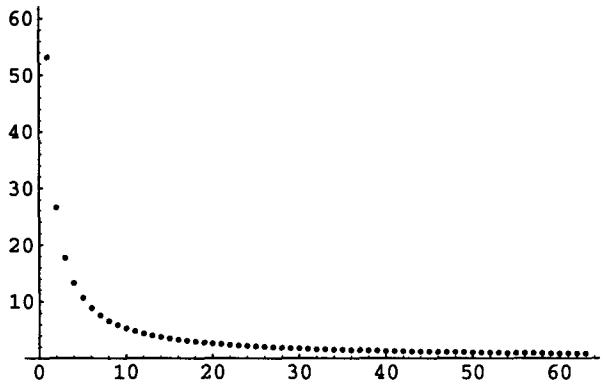

Figure 4.13: Wavelengths of 63 harmonics.

Algorithm 1 iterated 100 times produces varied results at different wavelengths, although similar wavelengths often have similar results, as shown in Figure 4.14 and Figure 4.15 Since the samples in the grid are 1 unit apart, wavelengths smaller than 2 units may find solutions on the grid which do not correspond well to the solution over the entire sampled region due to aliasing.[4] thus only the first 26 wavelengths avoid this. Even so, there is variety in the effects of the algorithm, as seen in Figure 4.16 and Figure 4.17

4.4 Examples

To see how well using the short time Fourier transform would work, some simple simulations were done. The process used is diagramed in Figure 4.18. Since the 121 point sampling grid in the center of the room was not dense enough to prevent aliasing at higher frequencies, a technique which was used of scaling the sampling grid so that it was critical at the Nyquist frequency is shown in Figure 4.19. Fig. 4.20 show a display of a system beep called Frog.snd frozen in time during propagation, and some efforts at reconstruction over four speakers.

Although these figures do not convey the resulting sound, several interesting phenomenon were noticed. Many different short sounds were tried with different microphone placement and approaches based only on applying the least squares solution, often in conjunction with Algorithm 1, to the bins of a short time Fourier transform taken with overlaping windows. Although pick-up points for sound could be put anywhere in the simulated

[3]Later the units are taken to be 0.1 meters and $c = 340m/s$ so that the wavelengths given above correspond to frequencies at the bins of a STFT

[4]The algorithm itself does not care about aliasing

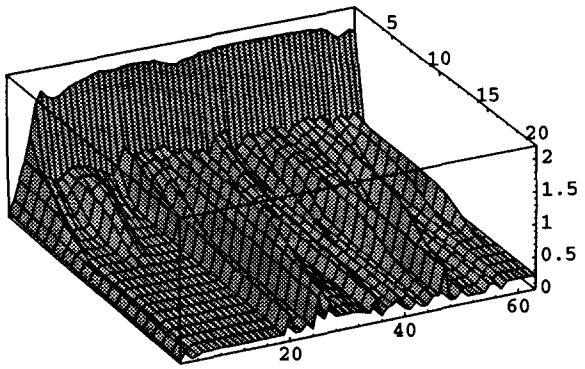

The first 20 iterations

Figure 4.14: Sum of the absolute value of the difference in absolute values of the reconstructed soundfield and target at 121 sample points over 63 frequencies.

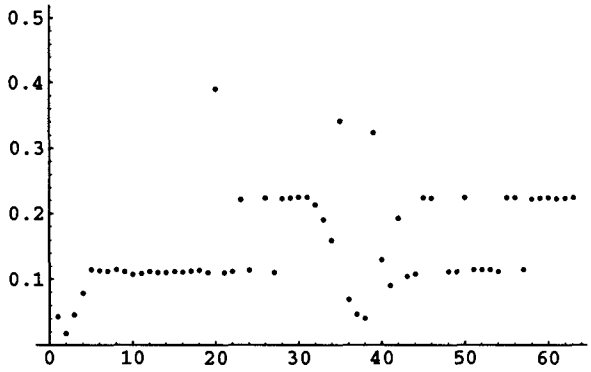

After 100 iterations

Figure 4.15: Sum of the absolute value of the difference in absolute values of the reconstructed soundfield and target at 121 sample points over 63 frequencies.

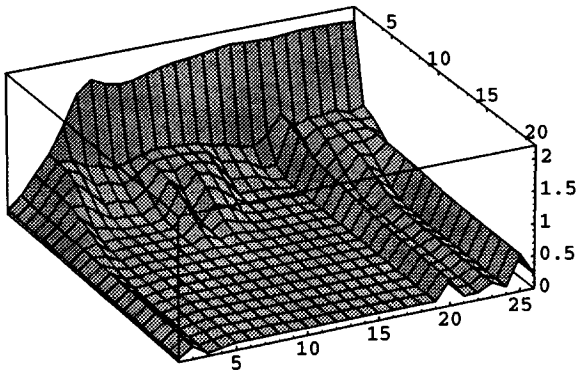

Figure 4.16: Sum of the absolute value of the difference in absolute values of the reconstructed soundfield and target at 121 sample points over 26 frequencies.

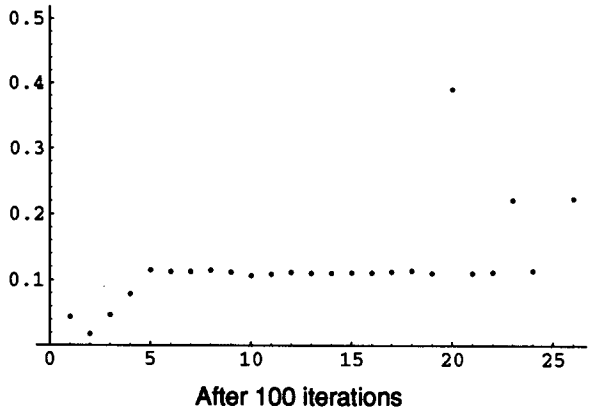

Figure 4.17: Sum of the absolute value of the difference in absolute values of the reconstructed soundfield and target at 121 sample points over 26 frequencies.

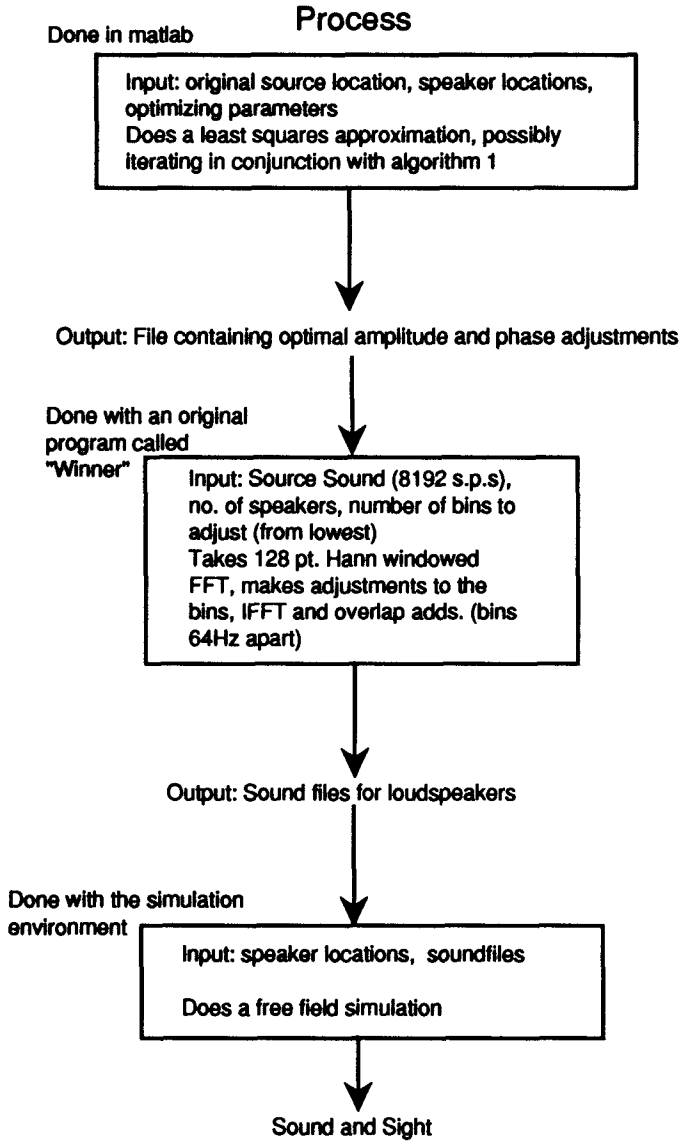

Figure 4.18: Process of sound manipulation.

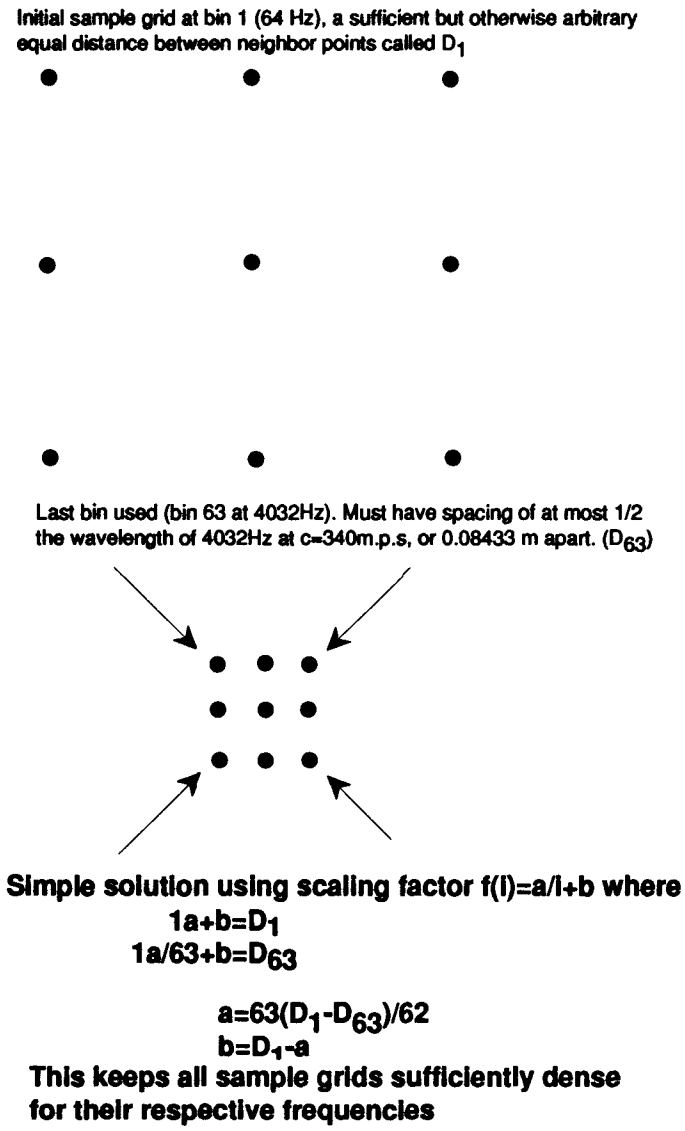

Figure 4.19: Process of scaling sampling pattern to avoid aliasing.

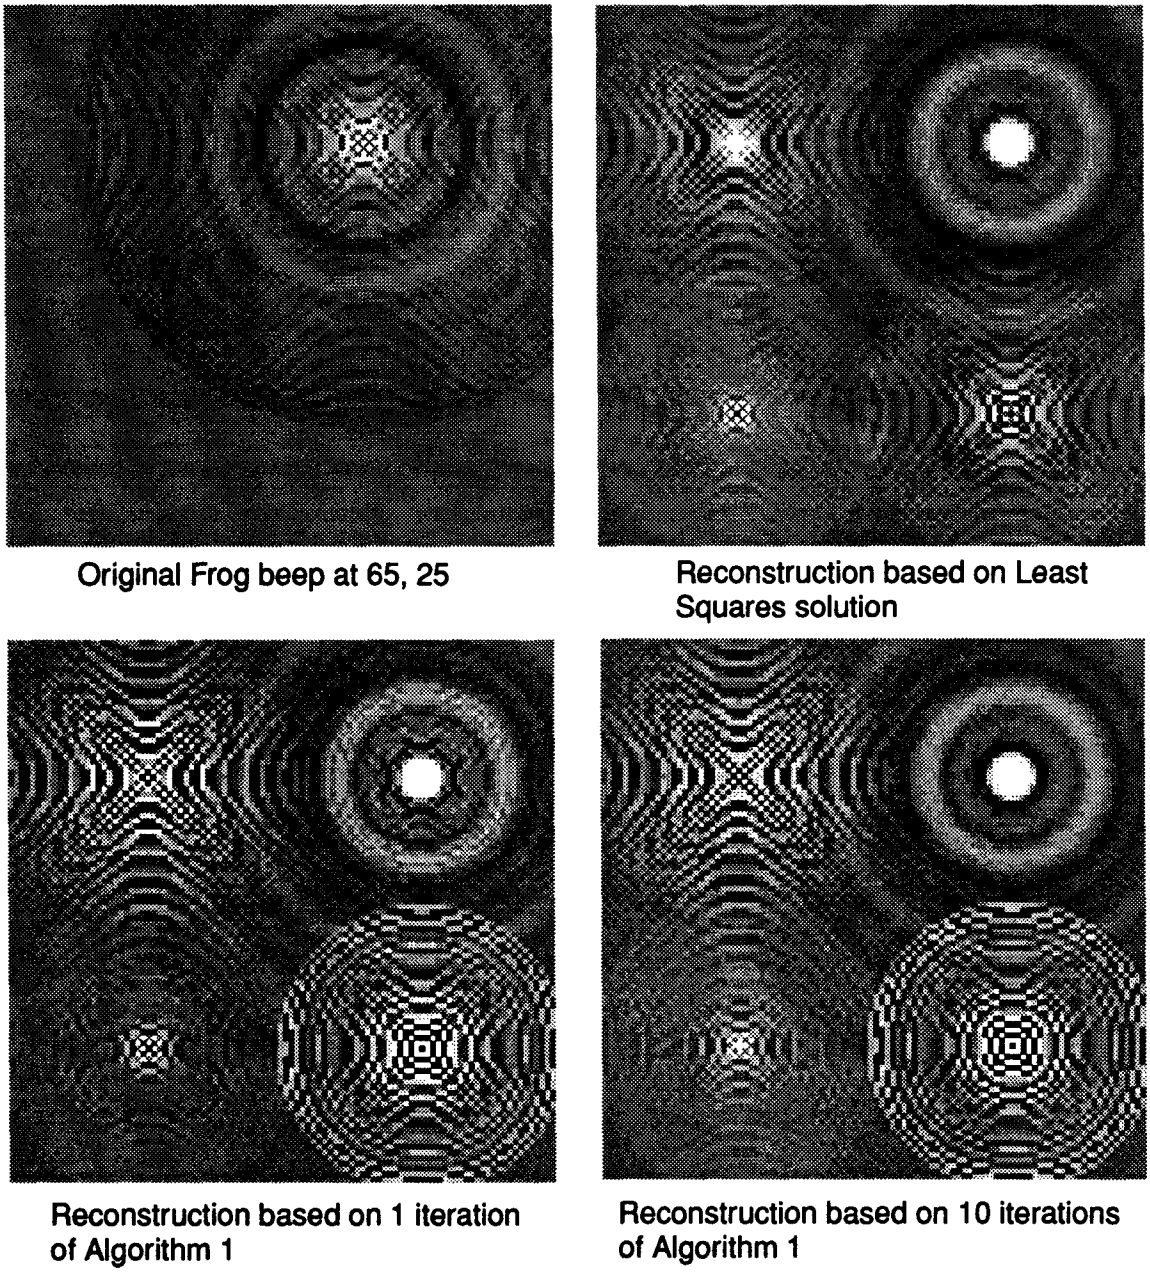

Figure 4.20: View of original Frog beep and reconstructions based on least squares and Algorithm 1 applied to short time Fourier transform bins.

2-D world, the sounds described here were taken from (50,50), right in the center of the speakers and sampling grid, not to mention the only point where the windowed packets are completely synchronized. In otherwords, the following discriptions represent the very best this method can achieve as opposed to an average result. The first thing noticed was that applying the least squares solutions to the frequency bins often resulted in very bad sounds. These sounds were similar to what might be achieved through a tight lowpass filter. This might have had something to do with the uneven scaling of the sampling grid, but probably more to do with increasingly poor correspondence between the original location's soundfield and the speaker's field at higher frequencies. Another observation was that in general, scaling the sampling grid to prevent additional aliasing did not have a large effect over not doing this, possibly because aliased solutions might not be that much worse than non-aliased ones in many cases. Finally, application of Algorithm 1 produced remarkably better sounding results than just using a least squares approximation, which shows that trading phase for amplitude is useful perceptually. Although mildly careful listening could detect the differences between the original and reconstruction, the general quality of the sound seemed relatively close.

4.5 Other Possible Approaches

In addition to the simple STFT implementation, several other possibilities exist. Ideally, a time domain filter using just delays, sums and multiplies could be constructed with the correct response at all frequencies, eliminating the need for taking the STFT and its inverse. However, the nature of the problem, when approached from a least squares perspective, is not analytic. Finding $x = (A^*A)^{-1}A^*y$ (see Equation 1.21) where A^* represents the conjugate transpose, requires combining conjugates with non-conjugates producing non-analytic functions. One possible remedy, touched on briefly in Section 1.5 is separating out the conjugated filters, and feeding the signal through backwards to implement the conjugation on a normal filter. Although in principle the regular and conjugated parts of the transfer matrix $(A^*A)^{-1}A^*$ could be separated out and dealt with in this way, it would probably be impractical since separating these parts of $(A^*A)^{-1}$ gets quite messy, and may be prone to numerical difficulties due to the large number of separate filtering operations. Of course, since time is 'reversed' for many filters, this could not be done in real time (except as a course approximation obtained by chopping up the signal and feeding small sections rapidly through the 'conjugate' filters backwards). Perhaps a better approach would be to find the desired filter response using a dense set of discrete frequency responses solved for by the methods discussed. Filter design techniques such as the Carathéodory-Fejér method [45] which approximate both desired magnitude and frequency response can be used in this regard. Of course, the room response needs to be taken into account [78] [80], as well as the directivity pattern of the loudspeakers, when any real implementation is attempted.

Chapter 5

Conclusion

In the past, most 3-dimensional sound techniques over loudspeakers depended on creating an illusion of localization. The goal of soundfield reconstruction has been too easily dismissed as requiring too many loudspeakers. With some care and attention paid to psychoacoustic factors, modest improvements can be made in almost any listening environment, even with very few speakers. In environments where many speakers can be used, approaching soundfield reconstruction from a perspective informed by perceptual goals may lead to even better results. As all current soundfield reproduction techniques are based on transfer matrices which directly or indirectly describe the effect of the loudspeakers on the environment, various different approaches can be combined for optimization of specific soundfield aspects. Choosing carefully the perceptual goals and tradeoffs is the first step to a successful result. Iterative algorithms can be used to find minimums with respect to certain norms and perceptual parameters where a closed form solution may not be available. Even if the minimums are not global, the resulting soundfield will be improved to some degree.

The leap from the simple anechoic monopole monochromatic examples examined to useful implementations is large but feasible. Before any soundfield is recreated, some knowledge of a target soundfield must be obtained. The simple examples demonstrate some techniques for approximating certain aspects of the soundfield in a new environment at individual frequencies. Filters can be made to have the right frequency response at many individual frequencies, which will come close to an ideal filter for a given situation. Often the ideal filter response is unattainable for a least squares based solution, since the conjugation involved makes ideal filter response non-analytic in the z-plane. The knowledge of the room response and directivity of the loudspeakers is necessary for an adequate implementation. Room invertablity is a big issue [80], especially in light of [78], where the numerical problems involved are shown to be worse than origially supposed. In addition, the effect of the audience itself may be unpredictable, but it is also unpredictable for the original soundfield. Although psychoacoustic optimizations may create a different soundfield which is affected differently by the audience, it is not clear how noticeable this effect might get.

While this discussion has hopefully built some context for soundfield reconstruction and perception, much more work needs to be done. The extension to reverberant environments with non-monopole speakers in odd arrangements is relatively straightforward in theory only. Knowledge and representation of a soundfield created from many complex, non-steady state and possibly moving sources is another difficulty. Unfortunately, soundfield reconstruction is best attempted out of real time. Actual implementation issues such as bandwidth, amount of processing required out of real time, numerical problems and

efficiencies must also be dealt with. Nonetheless, as processing gets faster and cheaper, bandwidth increases, and demand for virtual reality in audio continues, there are likely to be systems based on similar ideas sometime in the not-so-distant future.

Appendix A

The Simulation Environment

In order to facilitate the testing of ideas a simulation environment was created, running under NeXTStep[1]. This environment was designed around the goal of achieving results both visually and audibly, as well as to do some simple analysis mathematically of how well a given approach is working. In order to make calculations less cumbersome and the display more straightforward, a 2 dimensional simulation was chosen. The simulation environment allows speakers to be placed anywhere, and will render a representation of a two-dimensional slice of the environment. An animation feature allow the creation of a movie showing how the sound propagates. There are also features to record data for further analysis. Currently all sound sources are monopoles, although plans are made to allow arbitrary directionality pattens.

There are two modes in which the environment currently works. There is a single frequency mode which propagates complex sinusoids, and a sound mode which allows propagation of only real values. The sound mode works with NeXT soundfiles to create pictures of the propagation at a given time and allows a omnidirectional pickup to be placed anywhere within the environment, creating a new sound file which can be compared with a target.

In both modes, the pressure at any given point and time is just the sum of all the pressures at that point and time from all the speakers. There is no non-linear interaction, and no reflections, thus the model is of an anechoic room, or free-field condition. The design of the program is flexible enough that these situations could be added without inordinate effort, but there are no current plans to add these features.[2]

The view of the soundfile is actually a 100 by 100 TIFF file[3], generated from converting a 100 by 100 grid of floating point numbers to 256 gray levels. The mapping of the floating point values onto the 256 greys can be adjusted by a resolution factor, which is simply multiplied with the sound field floats before the conversion.

There are four basic views that can be shown in the single frequency mode, real imaginary, amplitude and phase, which simply show the corresponding aspect of the complex values at each point. Amplitude can be shown in three ways, either total amplitude, the log of the total amplitude or a comparison (difference) of the total amplitude of two different sound fields. Phase is shown between π and $-\pi$.

[1] NeXTStep is a computer operating system

[2] The current code is currently freely available to anybody who may wish to add such features.

[3] Tagged Image File Format

In the sound mode, only the real display and the amplitude display are used. The primarily purpose is to generate sounds by simulating sound propagation for further study and comparison outside the simulation environment. Fig. A.1 shows the entire environment with a typical display.

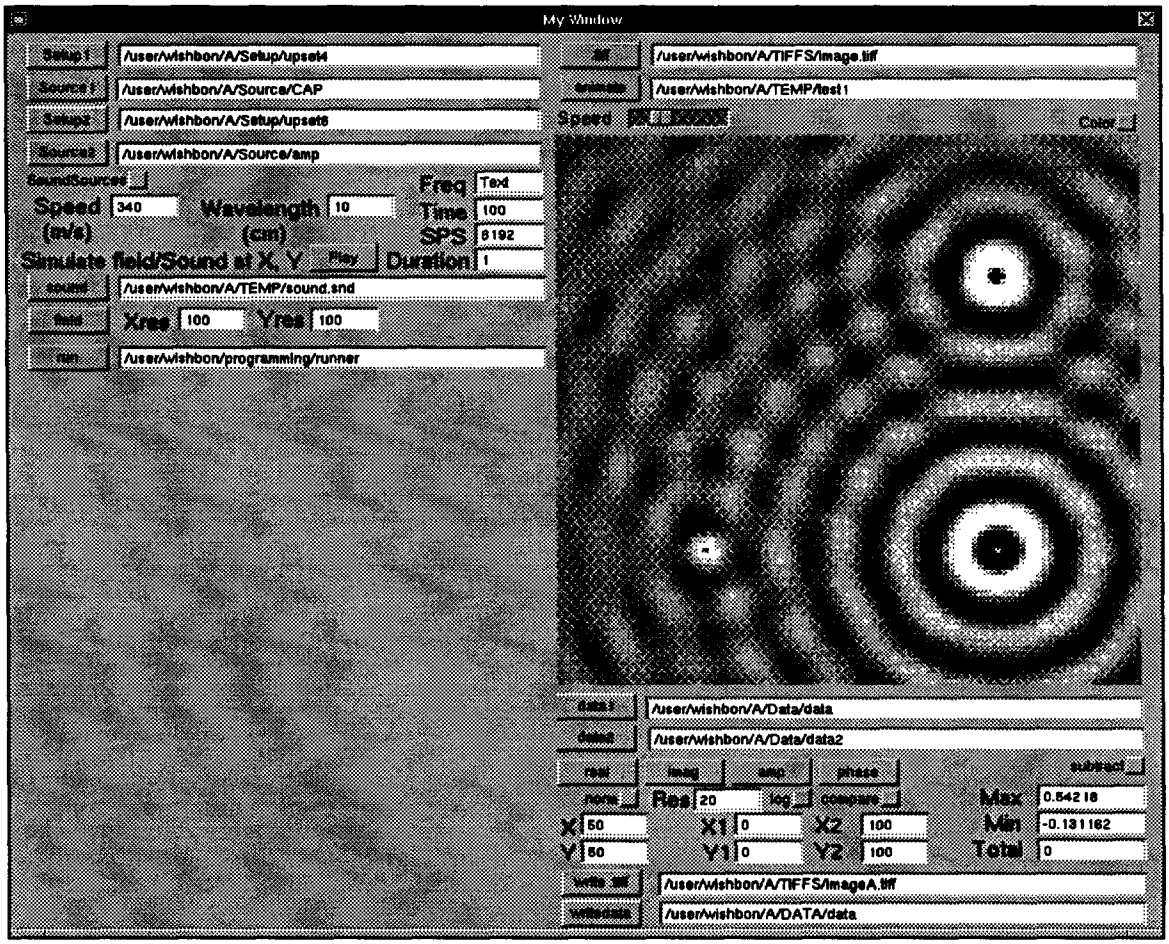

Figure A.1: Difference of amplitudes (resolution 10).

One very useful option gives the ability to compare two soundfield by taking the difference. The amplitude of the difference and the difference in the amplitude are both options[4] Similarly the phase of the difference can be shown with black indicating π and white indicating $-\pi$. Whereas the amplitude difference will not change in the steady state, the phase difference will cycle between π and $-\pi$ as the waves propagate. However, the relative phases between two sources or two points remains constant (modulo π). The time chosen for time dependent plots will always be when the original source is at zero phase unless otherwise indicated.

Finally, the projection of these complex sinusoids onto the real plane or imaginary plane can be graphed. This corresponds to the actual pressure fluctuations through space at a

[4]there is no squaring function to find the squared difference currently

Figure A.2: Phase of the difference.

given instant, as opposed to the amplitude or phase. Differences between soundfields can also be taken in the real mode. [5]

[5]Projection can also be done onto the imaginary plane

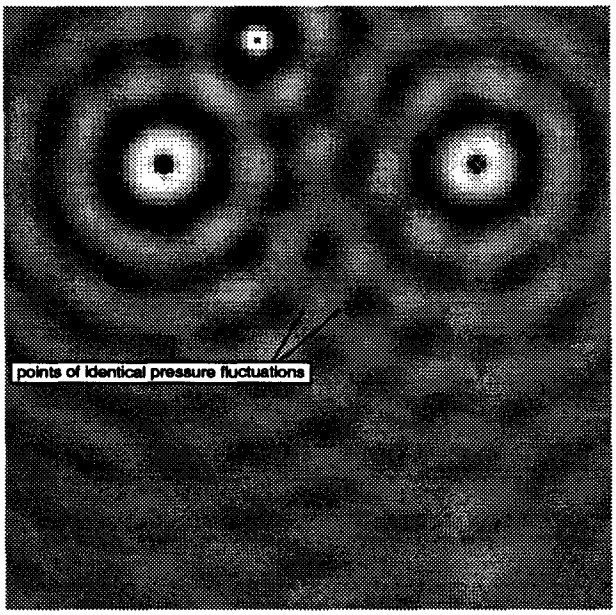

Figure A.3: Real part of the difference.

Appendix B

Spatial Aliasing

Even though the original audio signal is bandlimited with respect to time, the spatial signal is not bandlimited due to the change in amplitude with distance from the source. The 1/distance attenuation is not a bandlimited function, and when multiplied by the bandlimited source output in the spatial domain, a spatial signal results which is not bandlimited. both for the desired signal and for the speaker in the listening environment where this signal is approximated.

This can be treated mathematically as follows: Consider a bandlimited audio signal $x(t)$ going through a speaker and becoming a function of distance so that at time t_0 the signal at distance d is $\frac{1}{d}x(t_0 - d/c)$ where $c =$ the speed of sound. Let $\hat{x}(d) = x(t_0 - d/c)$. To get the spectrum of this, take the Fourier transform:

$$F\left(\frac{1}{d}\hat{x}(d)\right) = F\left(\frac{1}{d}\right) * F(\hat{x}(d)) = F\left(\frac{1}{d}\right) * \hat{X}(\omega) \tag{B.1}$$

Since $F\left(\frac{1}{d}\right)$ is not a bandlimited signal, its convolution with $\hat{X}(\omega)$ cannot be bandlimited, even though $\hat{X}(\omega)$ is bandlimited.

The fact that the spatial signal is not bandlimited may or may not be a problem. Any time limited signal cannot also be bandlimited, but this does not bring the world to a halt. The key issue is how the final result is perceived. The effect of spatial aliasing can be thought of as an inaccuracy in the spatial amplitude field of the sound field. The phase field can also be affected slightly. Thus, the perceived effect would only occur if the new amplitude field or phase field was different enough between a person's ears to affect perception. This issue is rarely considered when far field assumptions are appropriate, since sources approximate plane waves at such distances. Although far field assumptions are useful in optics and at higher audio frequencies, care must be taken at lower frequencies where only a few wavelengths separate the detectors from the source.

Since the 1/distance term tends toward zero as distance increases, the only signals which will not have spatial aliasing problems will be plane waves which can be thought of as originating from sources at infinity. Thus, if the sampled sound field consists of plane waves and the audio signal is bandlimited then exact reconstruction is possible from adequate sampled data since the spatial signal is bandlimited.

It seems intuitive that as distance from a source increases, the amount of spatial aliasing due to that source decreases. This can be shown to be the case by looking at the Fourier Series of the function $\frac{1}{d}$ for a given interval between d_1 and d_2. Expanding the function $\frac{1}{d}$ into its Fourier series will give an initial idea of its spectrum, but only roughly since the

process interprets the function as periodic, and thus discontinuous which may exaggerate higher frequencies. The complex form of the Fourier series is defined as

$$f(t) = \sum_{n=-\infty}^{\infty} \hat{F}_n e^{jn\omega_0 t} \tag{B.2}$$

where $\omega_0 = \frac{2\pi}{T}$ and $\hat{F}_n = \frac{1}{2}(a_n + jb_n)$ with

$$a_n = \frac{2}{T}\int_T f(t)\cos(n\omega_0 t)\,dt \tag{B.3}$$

and

$$b_n = \frac{2}{T}\int_T f(t)\sin(n\omega_0 t)\,dt \tag{B.4}$$

Since $\frac{1}{ct}$ is a real signal, $\hat{F}_{-n} = \hat{F}^*_{+n}$. Fig. B.1 gives the amplitude of the coefficients $F(\omega_n)$ of the complex Fourier series expansion of $\frac{1}{d}$ over the interval where $d_1 = 1$ and $d_2 = 2$.

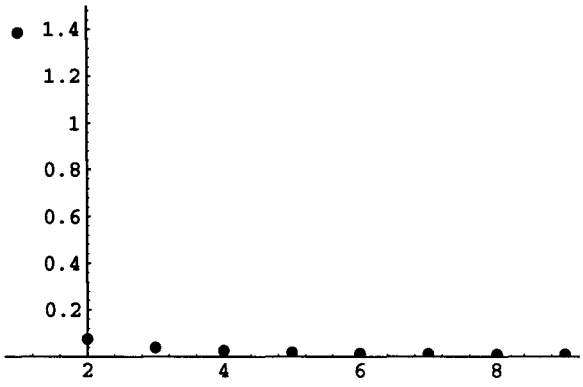

Figure B.1: Amplitude of Fourier Series coefficients of $\frac{1}{d}$ near origin on interval [1,2]

Fig. B.2 likewise gives the amplitude of the coefficients of the Fourier Series expansion of $\frac{1}{d}$ over the interval where $d_1 = 2$ and $d_2 = 3$.

Although this analysis gives some insight, additional windowing is desirable to avoid the discontinuity. This windowing can reduce the sidelobes of the implicit rectangular window used, with the cost of increasing the main lobe width. This is not a problem for functions with strong isolated components such as many sounds, but for a function without separation of main frequencies this may be a problem. A partial help in this case is to remove the D.C. component prior to windowing. This process is shown in Figure B.3. The amplitudes of the resulting Fourier series coefficients are shown in Figure B.4.

Another technique used in signal processing to aid in discrimination of peaks in the spectrum is interpolation via zero padding. Analogously in the continuous case, the zero function can be appended to the windowed attenuation function, and the Fourier series based on this longer function. Since the integral is over the zero function for the region outside the original range, an equivalent formulation is to leave the range the same, but change the Fourier kernel to reflect the larger period with the 'zero padding'. If m is the

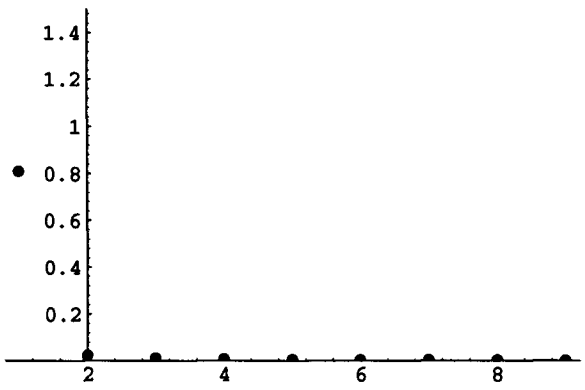

Figure B.2: Amplitude of Fourier Series coefficients of $\frac{1}{d}$ farther from origin on interval [2,3]

period length and k is an integer

$$\hat{f}(t) = \begin{cases} f(t) & t_1 + mk \leq t \leq t_2 + mk \\ 0 & t_2 + mk \leq t \leq t_1 + m(k+1) \end{cases} \tag{B.5}$$

$$a_n = \frac{2}{mT} \int_{mT} \hat{f}\,(t) \cos\left(n\omega_0 t\right) dt = \frac{2}{T} \int_{T} f\,(t) \cos\left(n\omega_0 t/m\right) dt \tag{B.6}$$

and

$$b_n = \frac{2}{mT} \int_{mT} \hat{f}\,(t) \sin\left(n\omega_0 t\right) dt = \frac{2}{T} \int_{T} f\,(t) \sin\left(n\omega_0 t/m\right) dt \tag{B.7}$$

After using this to simulate the effect of zero-padding by a factor of four, the result is shown in Figure B.5. Here the first peak corresponds to the first Fourier coefficient without the zero-padding. The first three peaks are discernible. Use of these methods shows that there are significant higher frequencies which can cause spatial aliasing, and oversampling spatially is desirable to reduce this effect.

Since the boundary of the listening environment is closest to the sources, spatial aliasing will be greatest near it. While adequate sampling is necessary throughout the environment to ensure accurate approximation of audio frequencies (using that approach), it might be worthwhile to add more sampling in a way that deflects this increase in spatial aliasing. The exact distribution of the extra sampling would be dependent on the particular setup, but the idea would be to increase the number of samples in a given distance in proportion to the rate of attenuation due to distance from the virtual source or the loudspeakers. This approach of putting more samples near the boundaries can be seen as related to the Kirchhoff integral approach where all the sampling is done on the boundary.

The amplitude of the frequencies in the Fourier Series expansion drops by half with each doubling of distance in the sense that twice as long an interval is needed to preserve the same ratios with an interval half the distance from the source. To look at it another way, doubling the distance from the source removes half of all spatial frequencies, so each doubling of distance causes a spherical wave to move half the distance toward the goal of becoming a plane wave in some sense.

While this analysis shows how spatial aliasing comes about, a more practical approach to understanding the significance of the pitfalls associated with spherical waves is to compare them with plane waves directly. Since non-aliased sampling only occurs with plane waves

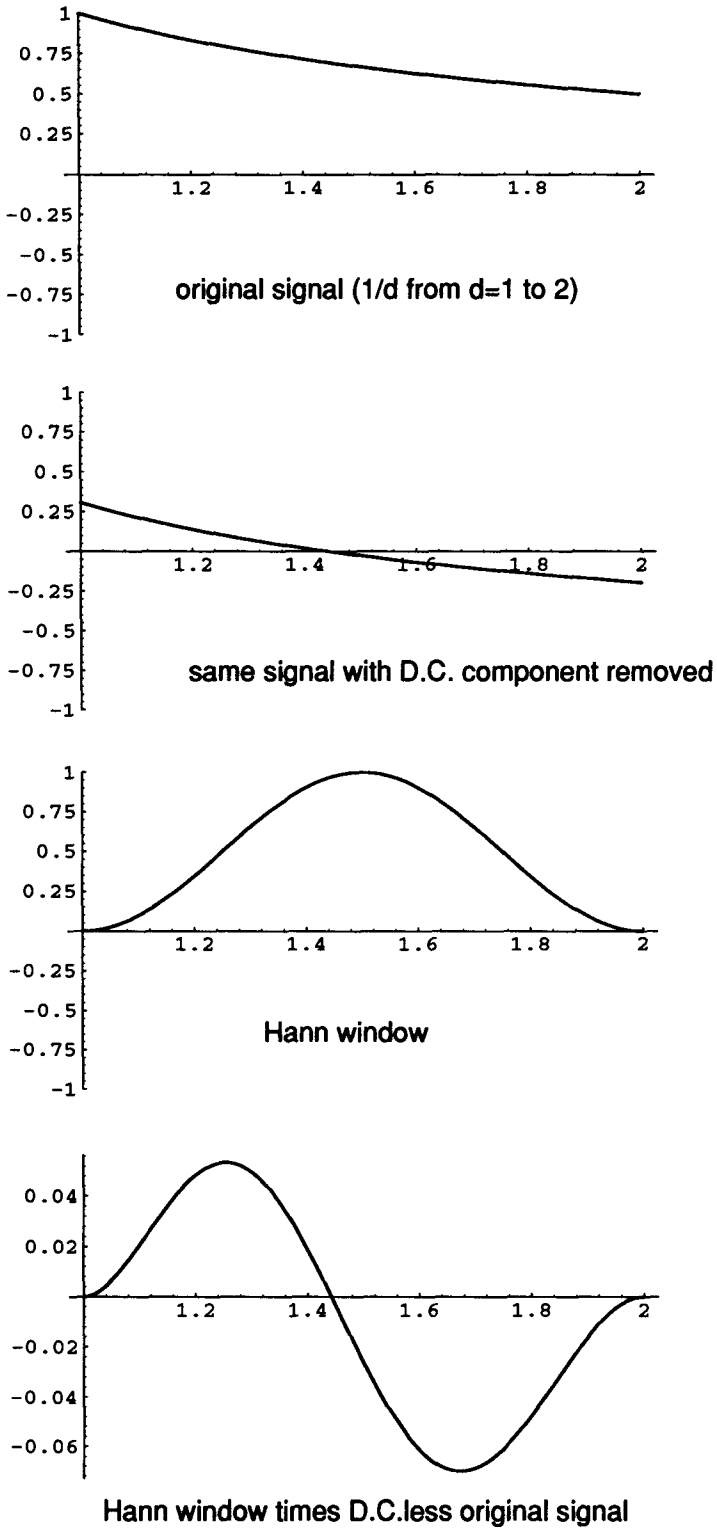

Figure B.3: Removal of the D.C. component from the attenuation before applying a window function.

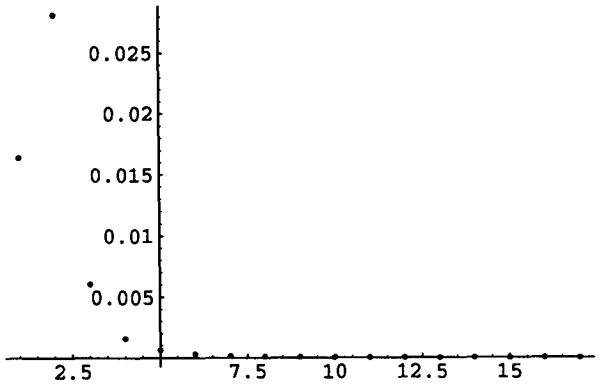

Figure B.4: Amplitude of Fourier Series coefficients of $\frac{1}{d}$ near origin on interval [1,2].

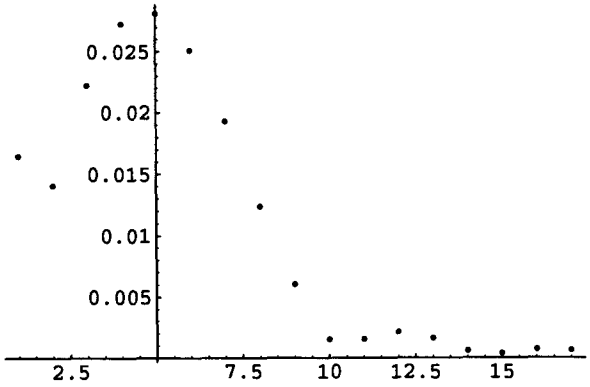

Figure B.5: Amplitude of Fourier Series coefficients of $\frac{1}{d}$ near origin on interval [1,2] with zero padding.

or combinations of plane waves below the Nyquist frequency, then comparing a given wave with a plane wave of the same frequency reveals the nature of these errors. Based on the formula for fluctuations due to spherical waves at a given frequency

$$p = \frac{1}{dist} e^{-j(\frac{2\pi dist}{w})} \tag{B.8}$$

it can be seen that the error related to distance terms can manifests itself as an error in amplitude and/or phase (time of arrival). This says nothing new since the error signal at a given point is just another signal with a certain phase and amplitude. Consider two sample points, both distance R from a source and distance l from each other. If the comparison is made between a wave from the source and a plane wave of the same frequency such that the signal is the same at the two points, then looking along the line determined by the points, it can be seen that between the points the signal from the near source is ahead in phase, (arrives early) and larger in amplitude than the plane wave. Outside the points, the signal from the near source arrives late and is smaller in amplitude. Fig.B.7 in a latter section diagrams this situation in more detail.

At any given point in space, the audio signal due to the pressure fluctuations is simply the sum of pressure fluctuations at that point from all sources, (loudspeakers, reflections, unwanted insects etc.). However those individual pressure fluctuations are dependent on the distance from their source. Since the desired wavefront reconstruction will, in general, have different virtual sources from the actual N sources in the reconstruction, the signal between the chosen N points in space (where there is an exact match) will have error in the form of incorrect amplitude and time of arrival. Since this is true for each of the sources, the total at such a point may be the sum of many signals, each slightly wrong in phase and amplitude, causing an unwanted (linear) filtering effect. This is true no matter how large N is or how the sample points are situated. Since the spatial frequencies are not band-limited, the sampling theorem does not hold and error will occur at locations other than the sample points.

B.0.1 The One Source Case

Looking at the geometry involved in a single point source trying to mimic another point source can help illustrate the nature of this spatial aliasing. Comparison of spherical waves to plane waves can be treated as a special case of this. A set of sources can be visualized by picturing all sources as radiating spheres. All the points on such a sphere have equal amplitude and can be thought of as representing a single point in time of an audio signal. If the audio signal is a delta function, a single sphere travels from the real source and the virtual source. Note that these delta functions do not have to be released at the same time. As these spheres expand at the speed of sound, they will run into each other. The points where each arrives at the same time are the points of time equality in the sense that any audio signal put through these two sources with the same time difference will arrive "in sync" at these points.

Mathematically this requirement can be described as the points where

$$c(t - t_1) = \sqrt{(x - x_0)^2 + (y - y_0)^2} \tag{B.9}$$

and

$$c(t - t_1) = \sqrt{(x - x_1)^2 + (y - y_1)^2} \tag{B.10}$$

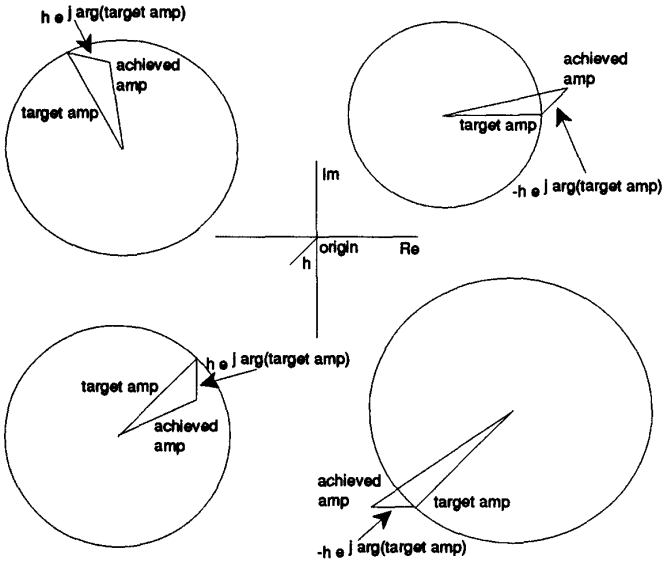

Figure C.7: possible sample values with reconstruction error vectors h, -h.

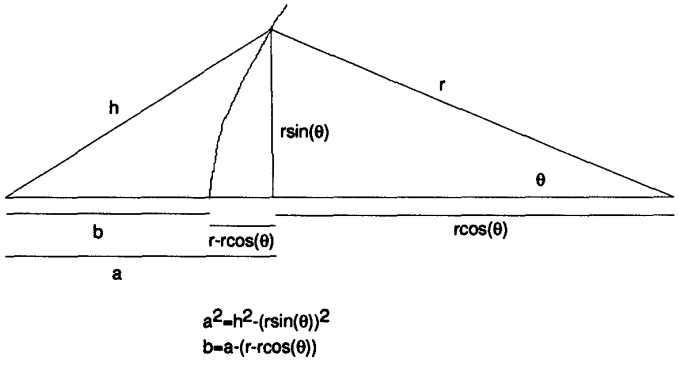

Figure C.8: Outward error vector h.

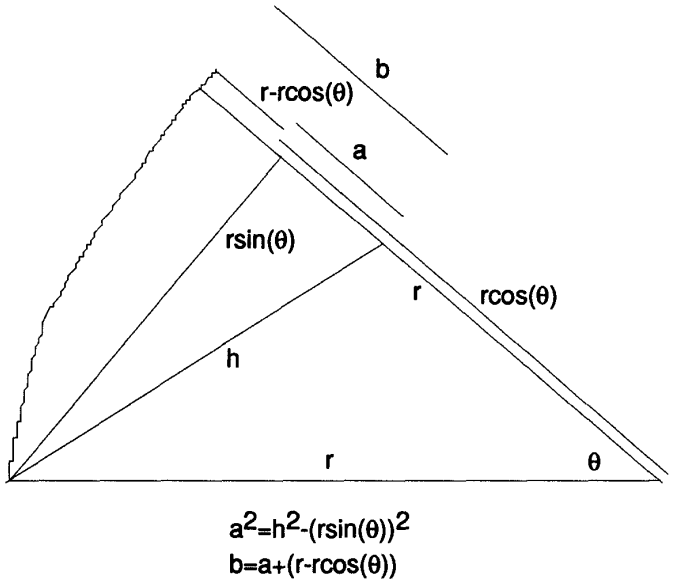

Figure C.9: Inward error vector h.

Unfortunately, this reduction in error is undesirable here since the goal is to find an extremal basis, and error reduction may make the chosen point not the right one to include even though the amplitude difference was maximum at that point. The can be seen in Figure C.11 where the absolute real error on the basis decreases after iteration 3. The final decreases before the steady state is reached may be attributable to the course grid of 0.5 used to find the maximum amplitude difference to substitute into the new basis, but such a large hiccup in the middle of the approach is not easily dismissed.

The whole process can be viewed as an attempt to minimize the maximum difference of the achieved result with a surface of revolution which denotes the target amplitude. Clearly this is only desirable in cases where phase is not important or can be overlooked without great cost. The use of an initial target with a source of the same frequency provides only initial conditions which are assumed to be closer to ideal than some other initial conditions representing a foreign frequency with the same target amplitude.

This proposed process was simulated for an original source located at (0,-0.5) where 5 speakers at (-2,0), (-1,0), (0,0), (1,0) and (2,0) were to try to minimize the maximum error with respect to amplitude along the straight line from (-2,1) to (2,1). The de la Vallée Poussin algorithm was used (single substitution into the basis), and the algorithm for eliminating the imaginary part was applied three times to try to eliminate the imaginary part of the error. Not surprisingly, for frequencies just above D.C., the adjustments to the basis increased the error as desired until convergence. Beyond a certain frequency of around 3.7, the progress of the basis error was not steady, and often backtracked. Nonetheless, eventually the basis converged on what seems a reasonable solution. Although no guarantee is made the it is the optimal solution, the basis points are the extremal points for the solution obtained despite the backtracking in the de la Vallée Poussin algorithm iterations.

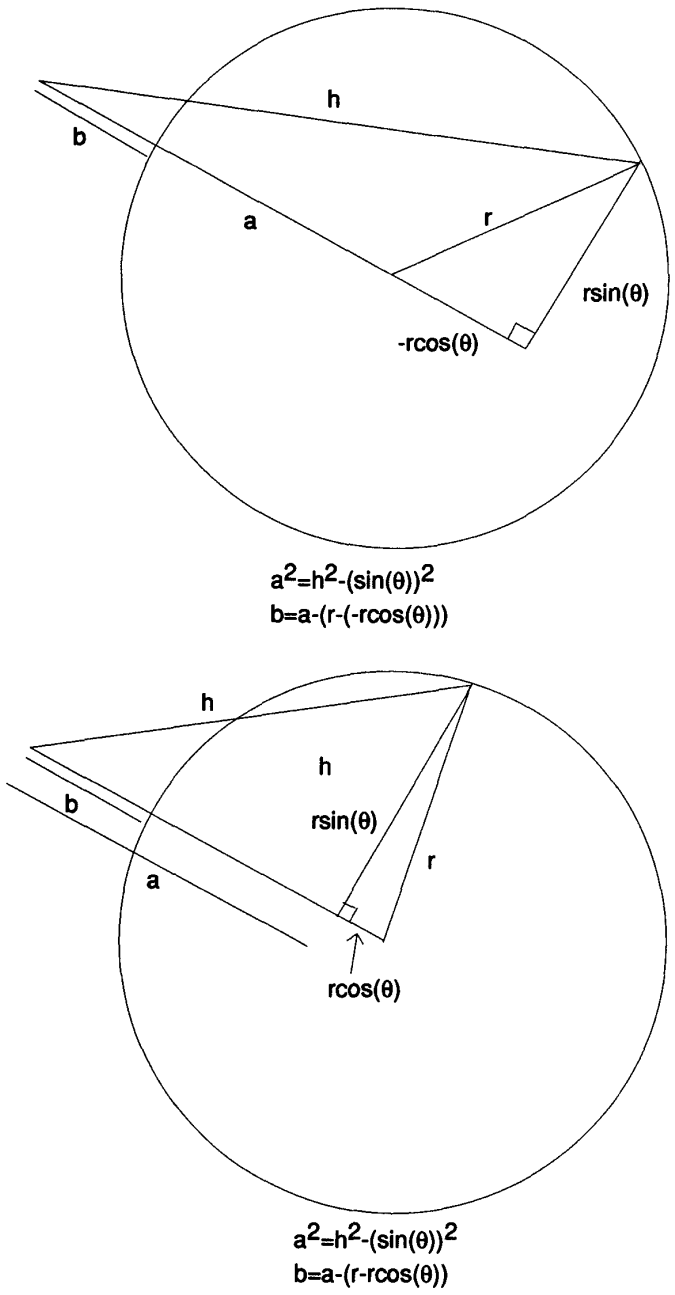

Figure C.10: Inward error vector h.

The progress of the error function is shown in Figure C.12 through Figure C.15 for various frequencies.

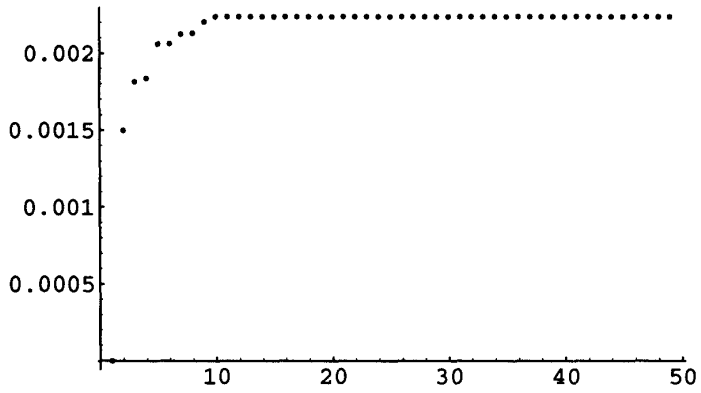

Figure C.11: Progression of basis error at D.C.

freq=3.7

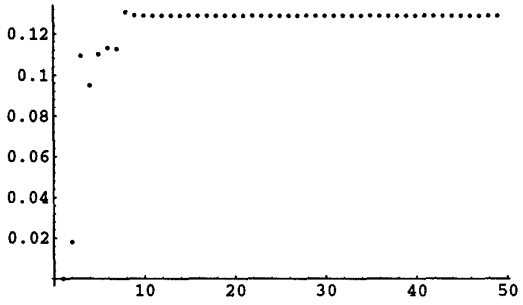

Absolute Value of the real part of the basis error per iteration

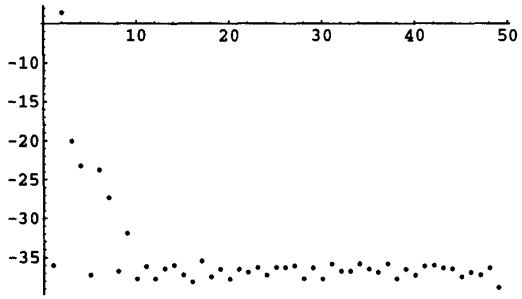

Log of the absolute value of the imaginary part of the basis error per iteration after using Algorithm 1 three times to remove the imaginary part

Figure C.12: Progression of basis error at frequency 3.7

freq=3.8

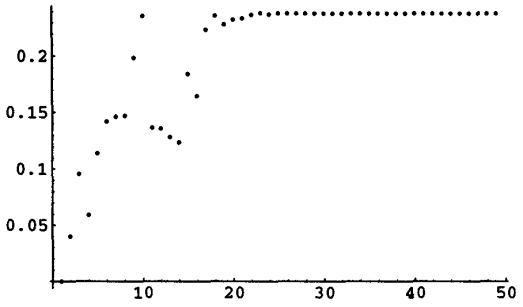

Absolute Value of the real part of the basis error per iteration

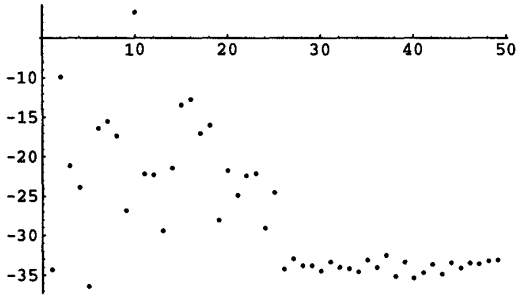

Log of the absolute value of the imaginary part of the basis error per iteration after using Algorithm 1 three times to remove the imaginary part

Figure C.13: Progression of basis error at frequency 3.8

freq=10

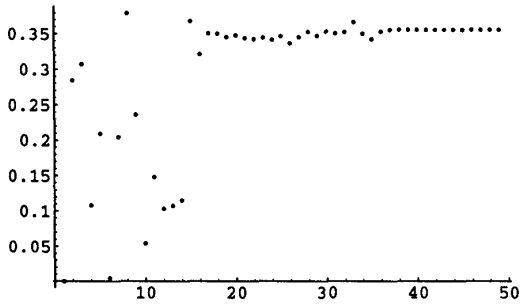

Absolute Value of the real part of the basis error per iteration

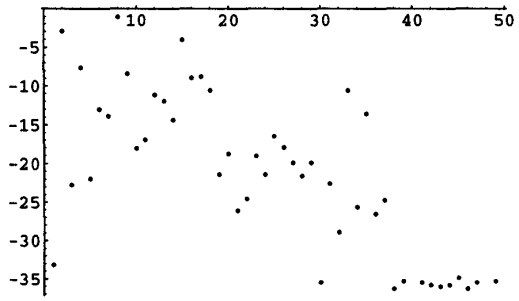

Log of the absolute value of the imaginary part of the basis error per iteration after using Algorithm 1 three times to remove the imaginary part

Figure C.14: Progression of basis error at frequency 10

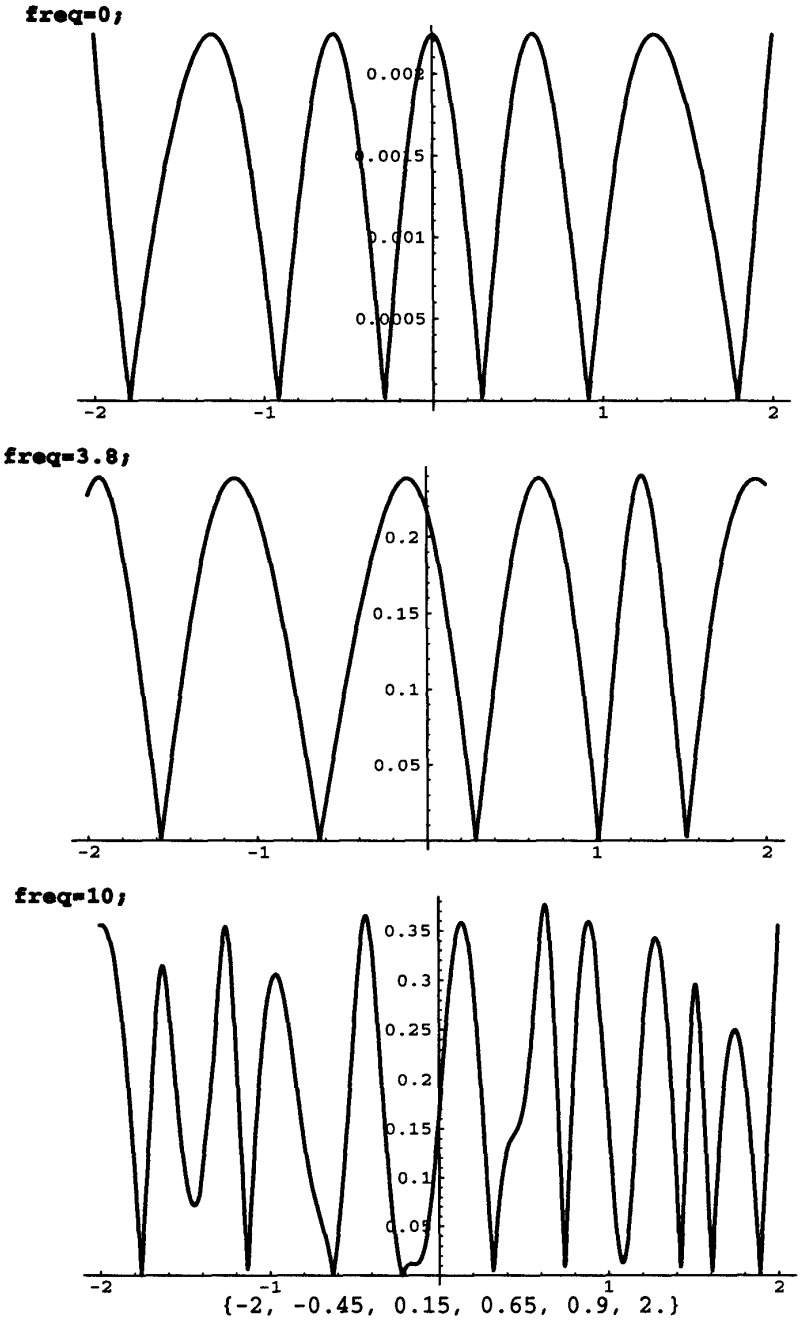

Figure C.15: Final results at three frequencies, for frequency 10 the final basis is noted since the grid search was not fine enough (0.05 spacing was used) to avoid slim peaks slipping through

Appendix D

Minimax Directivity Plots

In order to use loudspeakers effectively in the context of sound field reconstruction, detailed knowledge of their directivity pattern(s) for both phase and amplitude over all used frequencies must be known. While this precise knowledge does not often come with the loudspeaker, testing can be done to find a great amount of data about a particular loudspeaker and interpolation can be done to find approximate intermediary values. Ideally, such a directivity pattern could be extrapolated to be correct in the near-field as well as the far-field of the loudspeaker, since at low frequencies the near field can extend quite a distance. In this respect, probably the best approach to recording a directivity pattern is nearfield holography. This technique allows interpolation to both the near and far field from a single set of data on a boundary [65] However, the amount of data used to denote a directivity pattern and the ease of which the amplitude and phase adjustments for a given direction at a given frequency are also issues.

To code directivity patterns in a useful way, it would be nice to have a representation which allows control over the error and is efficient. Furthermore, it would be good if some of the data could be disregarded in such a way as to minimize the effect on the reconstruction with the remaining data. Such an approach for a one dimensional directivity pattern is to represent the data in Fourier series. Fig. D.1 shows that in polar coordinates, the Fourier basis functions represent a suitable orthogonal basis for directivity patterns. In addition to the obvious least squares approach of truncating a function's Fourier series, a truncated F.S. can also be found such that it minimizes the weighted maximum error. The approach used is similar to the one used by Parks and McClellan in [77]

$$\begin{bmatrix} 1 & \cos 2\pi F_0 & \cdots & \cos 2\pi nF_0 & \frac{1}{W(F_0)} \\ 1 & \cos 2\pi F_1 & \cdots & \cos 2\pi nF_1 & \frac{-1}{W(F_1)} \\ \vdots & & \ddots & & \vdots \\ 1 & \cos 2\pi F_{n+1} & \cdots & \cos 2\pi nF_{n+1} & \frac{(-1)^{n+1}}{W(F_{n+1})} \end{bmatrix} \begin{bmatrix} d_0 \\ d_1 \\ \vdots \\ d_n \\ \rho \end{bmatrix} = \begin{bmatrix} D(F_0) \\ D(F_1) \\ \vdots \\ D(F_n) \\ D(F_{n+1}) \end{bmatrix} \tag{D.1}$$

Here, the d_n represent filter parameters solved for, ρ is the unweighted error, and $D(F_n)$ is the current basis, used as described in the previous appendix. This equation is evaluated iteratively with new basises arrived at through the Remez exchange algorithm. Instead of using this formulation for filter design, it can be used on odd functions, in particular odd directivity patterns. The condition for oddness of a directivity pattern is just that it be symmetrical around the zero angle point. This idea can be extended to a weighted Fourier

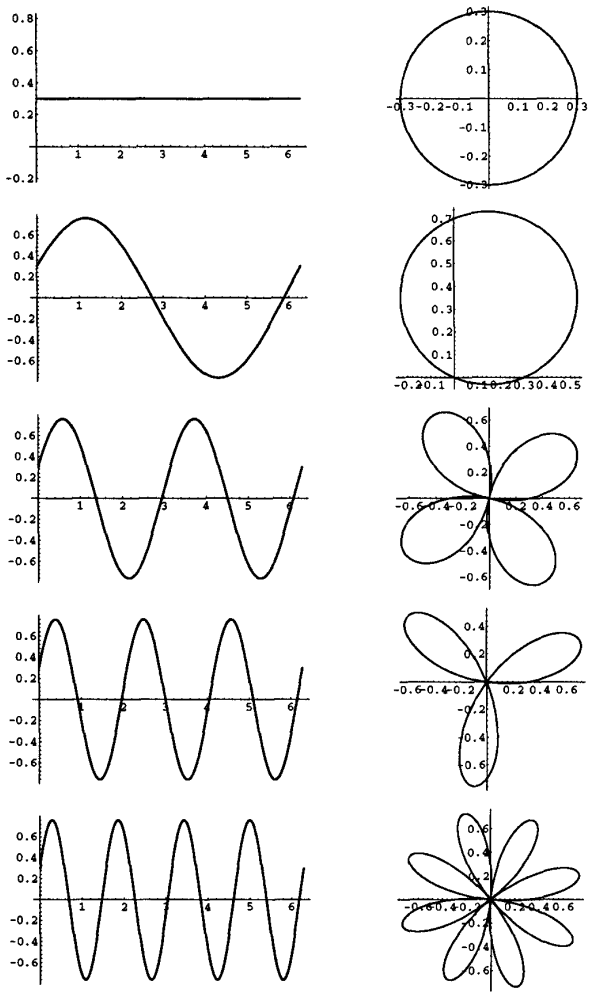

Figure D.1: Fourier Series basis functions in Cartesian and polar coordinates.

Series variant, which can handle nonsymmetrical directivity patterns.

$$\begin{bmatrix} \frac{1}{W(\Theta_0)} & 1 & \cos 2\pi\Theta_0 & \cdots & \cos 2\pi n\Theta_0 & \sin 2\pi\Theta_0 & \cdots & \sin 2\pi n\Theta_0 \\ \frac{-1}{W(\Theta_1)} & 1 & \cos 2\pi\Theta_1 & \cdots & \cos 2\pi n\Theta_1 & \sin 2\pi\Theta_1 & \cdots & \sin 2\pi n\Theta_1 \\ \vdots & & & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{(-1)^{n+1}}{W(\Theta_{n+1})} & 1 & \cos 2\pi\Theta_n & \cdots & \cos 2\pi n\Theta_n & \sin 2\pi\Theta_n & \cdots & \sin 2\pi n\Theta_n \end{bmatrix} \begin{bmatrix} \rho \\ a_0 \\ a_1 \\ \vdots \\ a_n \\ b_0 \\ b_1 \\ \vdots \\ b_n \end{bmatrix}$$

$$= \begin{bmatrix} X(\Theta_0) \\ X(\Theta_1) \\ \vdots \\ X(\Theta_n) \\ X(\Theta_{n+1}) \end{bmatrix} \tag{D.2}$$

This version was used on a simple directivity pattern shown in Figure D.2 using the de la Vallee Poisson algorithm described in the previous appendix. Here $X(\Theta_i)$ is the current basis and ρ is the unweighted error. Other than a_0, which represents the D.C. term, the a_i and b_i are the current coefficients of the Fourier series' cosines and sines respectively. Using a D.C. term plus 15 cosine and 15 sine terms, the directivity pattern was approximated by Figure D.3, with the error shown in Figure D.4 using an unweighted approximation.

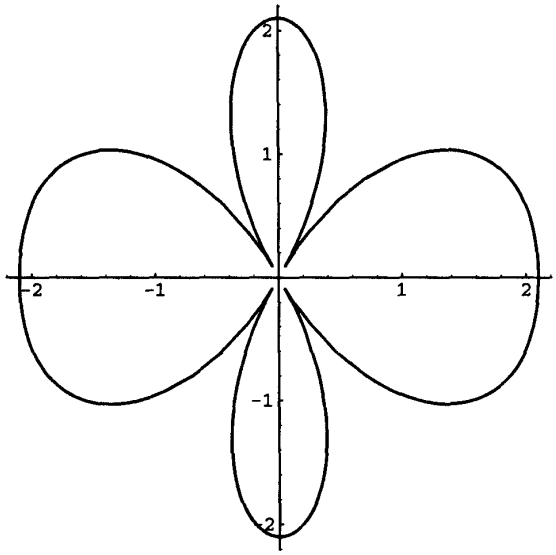

Figure D.2: Target directivity pattern (amplitude only).

As an example of a weighted solution, the directivity pattern was approximated by minimax with a weighting proportional to amplitude. In most cases this is not a good

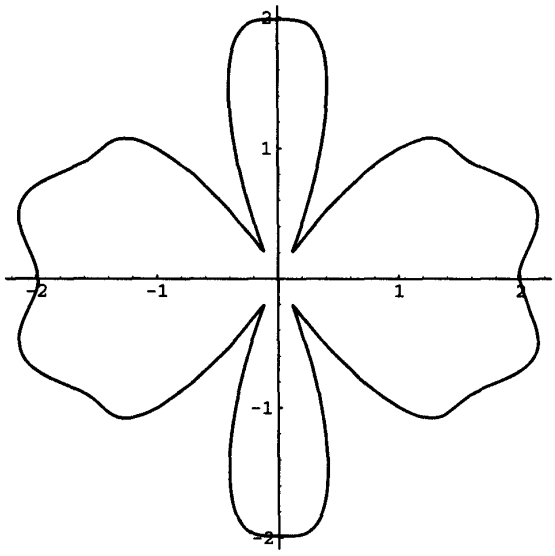

Figure D.3: Reconstructed directivity pattern using 31 terms.

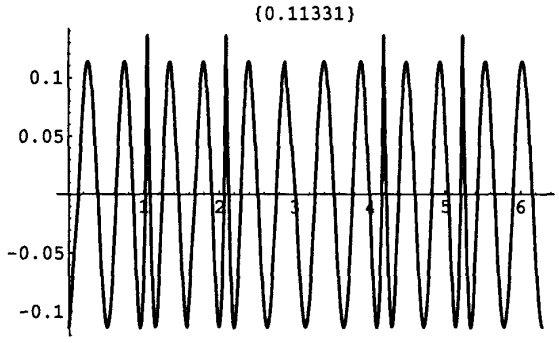

Figure D.4: Error in the reconstructed directivity pattern.

weighting but here it allows the effect of the weighting to be easily seen. This resulted in Figure D.3.

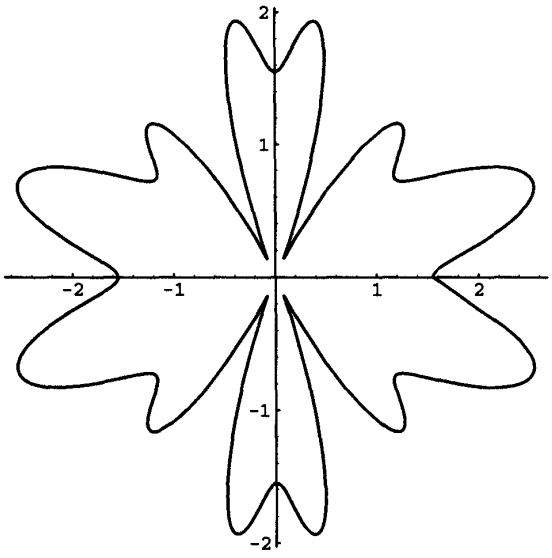

Figure D.5: Reconstructed directivity pattern using 31 terms and a weighting function.

Although these examples are obviously not adequate for audio, the approach of using minimax Fourier series to approximated directivity patterns may still have some benefit. The flexibility involved with weighting allows careful control over the resulting representation, and since the resulting coding is based on sinusoids, they are easy to decode from the coefficients. Unfortunately, it may be necessary to code the variation of phase with directivity separately. Also, this planar representation is only useful for radiation patterns which are surfaces of revolution if all three dimensions are to be used. Some extension to a irregular three dimensional directivity pattern using two dimensional Fourier series is likely to be possible.

Appendix E

Minimax Sound Field Projection

Since the best fitting speakers are found one at a time, instead of applying a least squares fit of the best speaker to the source, a minimax fit can be used instead. However this should be done with a note of caution, since at each stage the best fit is selected with a criterion more closely related to a least squares fit.

Upon finding the speaker with the best fit under whatever criterion, that particular speaker's value can be made to be its best fit in the Chebychev sense. To illustrate, consider the problem $\alpha a = b$ where α is a complex value representing the speaker, and a, b are vectors of complex values representing the transfer relationships from the speaker to the samples and the desired value at the sample points, respectively. Finding the best value α in the Chebychev sense is a convex problem, both in the case where error is measured as the difference in amplitudes or the amplitude of the difference.

Since only one speaker is being made to optimally match data, the relative phases between speakers is a nonexistent issue, so finding the minimax difference in amplitudes does not involve phase at all. However, since other speakers will be added in an attempt to reduce the error, it is important to get the phase to match as closely as possible, but this problem is completely separate from finding the amplitude at the speaker which produces the minimax difference with amplitude at the sample points. Mathematically, this is because $max||\alpha a| - |b|| = max||\alpha||a| - |b||$, since the absolute values are real quantities, the problem is one-dimensional. Further, if $x = |\alpha|$, only positive values of x can be explored. The answer will be positive since if $d, e > 0$ and are constant, and x is real, then xd describes a line of positive slope through the origin, $xd - c$ shifts that line below the origin and $|xd - c|$ creates a line with negative slope where $xd - c < 0$. Since the minimal point of these convex sets is always positive, the minimal point of their intersection will also be positive. This can be seen if Figure E.1 and Figure E.2.

Consider an example where $a = [0.5 + 0.5j, 1 + j, 2 + 2j, 3 + 3j]$ and
$b = [3+2j, 3+j, 3, 3-j]$. The plot of $max(|xa| - |b|)$ from $-3 < x < 3$ is given in Figure E.3

The minimum is found using a grid method, where a set of 20 equally spaced points are taken and the point producing the minimal value is found. Next the minimal value is sought on 20 points equally spaced between the points of the previous set neighboring the minimal point. Thus for each iteration, the accuracy improves by one decimal place, and the process can be carried on until the limit of the precision of the sample data is reached. The example given, this algorithm finds the solution to be at $|\alpha| = 1.36731$ which gives $||\alpha||a| - |c|| = [2.63872, 1.22861, 0.867331, 2.63872]$.

Next, since the phase is arbitrary, it would be good to get this quantity as close as

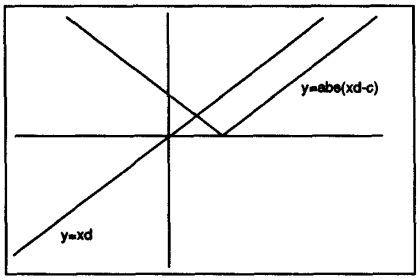

Figure E.1: Demonstration that the solution point must be positive given constraints.

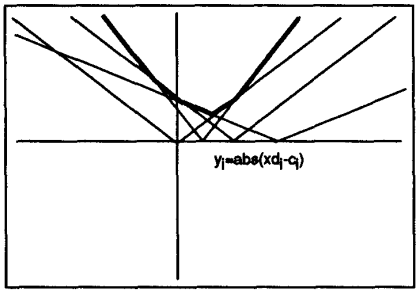

Figure E.2: Positive solution point is the minimum of the maximum of the lines.

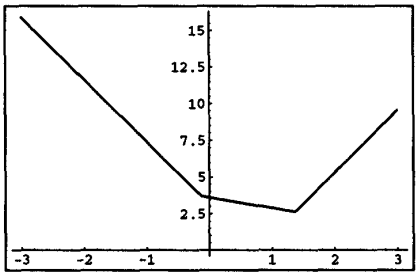

Figure E.3: Plot of $max(|xa| - |b|)$ from $-3 < x < 3$.

possible to the original over the region. Unfortunately, this is not always a convex problem, as shown by the following counter example. If $q = [1, 1, 1, 1]$ and $p = [I, 1, 1, I]$ then $\alpha = 1$ gives $max||\alpha||q| - |p|| = 0$ but the plot of $max|\alpha e^{(j\omega)} \cdot q - p|$ over $0 \leq \omega < 2\pi$ clearly shows two local minimums(Figure E.4).

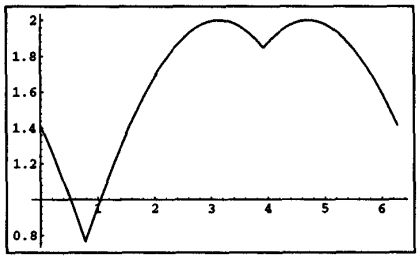

Figure E.4: Counter-example for convex phase hypothesis.

Nonetheless, the algorithm given above can be used with good chance of finding the global minimum or a near optimal local one. Continuing the prior example, by good fortune the problem is convex as seen in Figure E.5.

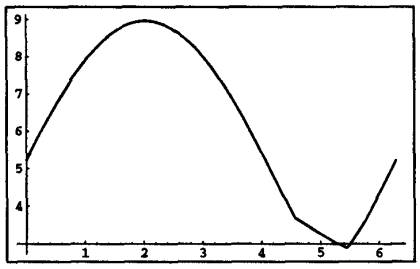

Figure E.5: Example of convex phase.

Applying a grid search to find the value of ω which minimizes the maximum phase difference gives $\omega = -0.265346 \cdot \pi$ where $-\pi < \omega \leq \pi$. Thus $\alpha = 1.36731 \cdot e^{j(-0.265346 \cdot \pi)}$ which minimizes $max(||\alpha a| - |b||$ and minimizes $max|\alpha a - b|$ given $|\alpha|$. In this case, $|\alpha a - b| = [2.88563, 1.5287, 0.882737, 2.88563]$.

Another approach is just to look at the fit of $max|\beta a - b|$, where the independent variable has been renamed for later comparisons. Since $max|\beta a - b|$ will be continuous (each value of $|\beta a|$ is continuous and the maximum value will change when these cross the current maximum value), then if each element is convex, the maximum value will also be convex. Each element has the form $|\beta a_i - b_i|$. By varying the real and imaginary components of β separately, it is apparent this is quadratic in the real and imaginary directions, since the absolute value is defined in terms of summing the squares of the real and imaginary parts of a number. Taking the square root of this to get the absolute value does not affect the convex nature, since the square root of a larger number will always be bigger than the square root of a smaller number. The surface defined by β is a convex bowl shape, shown in Figure E.6

for the vectors $a = [0.5 + 0.5j, 1 + j, 2 + 2j, 3 + 3j]$ and $b = [3 + 2j, 3 + j, 3, 3 - j]$ (same as before).

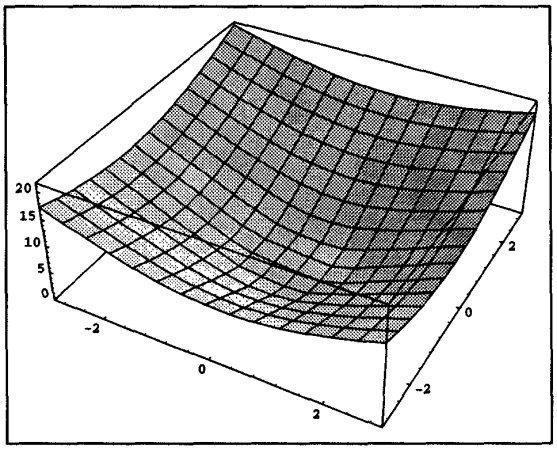

Figure E.6: Convex minimization of maximum surface.

To find the minimum point, a grid method, analogous to the previous real-only method, was used. It finds the minimum value along the real and imaginary directions alternately, and each iteration zooming in by a factor of ten, searching from the minimum point to the points on either side. Figure E.7 illustrates that procedure. On the left and bottom sides of the figure are successively shorter lines indicating the search region at each iteration. Each line contain 20 equally spaced samples.

Using this method, the value minimizing $max|\beta a - b)|$ is $\beta = 0.92506 - j$, which produces $|\beta a - b| = [2.88142, 1.52019, 0.863231, 2.88141]$ which can be compared with the earlier method version where $\alpha = 0.919115 - 1.0123j$. In polar terms $\alpha = 1.36731 \cdot e^{j(-0.265346 \cdot \pi)}$ while $\beta = 1.36225 \cdot e^{j(-0.262385 \cdot \pi)}$. Thus the two methods produce very close results in this case. Still $max|\beta a - b| = 2.88142 < max|\alpha a - b| = 2.88563$, so if the primary criterion is minimizing the absolute value of the differences, then β is slightly better. Conversely, if the goal is to minimize the maximum difference of the absolute values, then α is slightly better since $max||\alpha a| - b|| = 2.63872 < max||\beta a| - |b|| = 2.64229$.

The choice between these two methods must be considered taking into account that other speakers need to be coordinated as well.

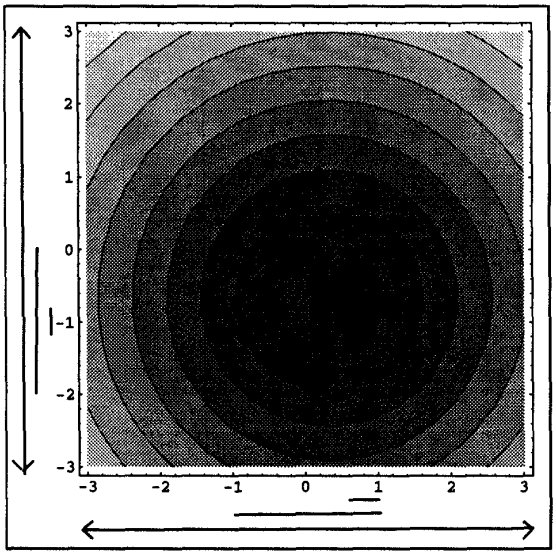

Figure E.7: Convex minimization of maximum surface.

Bibliography

[1] "Basic Dolby Surround vs. Pro Logic," accessed from *Dolby Web Page, http://www.dolby.com/index.html*, Nov. 1995.

[2] "Surround Sound Past, Present, and Future," accessed from *Dolby Web Page, http://www.dolby.com/index.html*, Nov. 1995.

[3] "Dolby Surround AC-3: Questions and Answers," accessed from *Dolby Web Page, http://www.dolby.com/index.html*, Nov. 1995.

[4] "The (Near) Future of Multichannel Sound," accessed from *Dolby Web Page, http://www.dolby.com/index.html*, Nov. 1995.

[5] "Ambisonic Surround Sound FAQ Version 2.5," accessed from *Ambisonic Surround Sound Web Page, http://www.omg.unb.ca/ mleese/ambison.html*, Nov. 1995.

[6] Wolfgang Ahnert, "Complex Simulation of Acoustic Sound Fields by the Delta Stereophony System (DDS)," *Journal of the Acoustical Society of America*, vol. 35, no. 9, pp. 643-652, Sept. 1987.

[7] Hiroshi Asayama, Sho Kimura, and Katsuaki Sekiguchi, "Revised finite sound ray integration method based on Kirchhoff's integral equation," *Journal of the Acoustical Society of Japan, (E)* vol. 10, no. 2, pp. 93-100, 1989.

[8] Aizawa, Kiyoharu, Huang, and S. Thomas, "Model-Based Image Coding: Advanced Video Coding Techniques for Very Low Bit-Rate Applications," *Proceedings of the IEEE*, vol. 83, no. 2, pp. 259-271, February 1995.

[9] Frederick J. Ampel and Ted Uzzle, "Multichannel auditory perspectives, a historical view of Harvey Fletcher's forgotten contributions and their ramifications," presented at *129th meeting of the Acoutical Society of America, Washington D. C.*, June 1995.

[10] Aoki et al., "Stereo Reproduction with Good Localization over a Wide Listening Area," *Journal of the Audio Engineering Society,* vol. 38, no. 6, pp. 433-439, June 1990.

[11] D. W. Batteau, "The role of the pinna in human localization," *Proceedings of the Royal Society of London*, vol. 168, series B, pp. 158-180, 1967.

[12] D. R. Begault, *3-D Sound for Virtual Reality and Multimedia*, New York: Academic Press, 1994.

[13] Durand Begault, "Challenges to the Successful Implementation of 3-D Sound," *Journal of the Audio Engineering Socitey*, vol. 39, no. 11, pp. 864-870, Nov. 1991

[14] J. C. Bennett, K. Barker, and F. O. Edeko, "A New Approach to the Assessment of Stereophonic Sound System Performance," *Journal of the Audio Engineering Socitey*, vol. 33, no. 5, pp. 314-321, May 1985.

[15] A. J. Berkhout, "A Holographic Approach to Acoustic Control," *Journal of the Audio Engineering Socitey*, vol. 36, no. 12, pp. 977-995, Dec. 1988.

[16] A. J. Berkhout, *Applied Seismic Wave Theory*, Amsterdam: Elsevier, 1987.

[17] Jens Blauert, "Sound localization in the median plane," *Acustica*, vol.22, pp. 957-962, 1969.

[18] Jens Blauert, "Lateralization of jittered tones," J*Journal of the Acoustical Society of America*, vol. 70, pp. 694-698, 1981.

[19] Jens Blauert, *Spatial Hearing*, trans. John S. Allen, Cambridge Mass.: MIT Press, 1982.

[20] Jeffrey Borish, *Electronic Simulation of Auditorium Acoustics* PhD dissertation, Stanford University, CCRMA, Dept. of Music, 1984.

[21] Marina Bosi, "A Real-Time System for Spatial Distribution of Sound," Dept. of Music, Report No., STAN-M-66, 1990.

[22] Henry D. Brant, *Orbits*, (Sound recording for 80 trombones and 30 basses), L.P., Composers Recordings CRI SD 422, prod. 1980.

[23] Albert S. Bregman, *Auditory scene analysis: the perceptual organization of sound*, Cambridge, Mass.: MIT Press, 1990.

[24] A. M. Bruneau, "Amplitude and Phase Measurements of Vibrations of Radiating Surfaces in Order to Determine the Emitted Sound Field," *Journal of the Audio Engineering Socitey*, vol. 31, no. 12, pp. 907-913, Dec. 1983.

[25] R. A. Butler and K. Belendiuk, "Spectral cues utilized in the localization of sound in the median sagittal plane," *Journal of the Acoustical Society of America*, vol.61, pp. 1264-1269, 1977.

[26] Edward C. Carterette and Morton P. Friedman, eds. *Handbook of Perception, Volume IV, Hearing*, New York: Academic Press, 1978.

[27] Didier Cassereau, Najet Chakroun, Francois Wu, Mathias Fink and Frederic Datchi, "Synthesis of a Specific Wavefront using 2D Full and Sparse Arrays," *IEEE Ultrasonics Symposium Proceedings*, pp. 563-568, 1992.

[28] Didier Cassereau and Mathias Fink, "Time-Reversal of Ultrasonic Fields-Part III:Theory of the Closed Time-Reversal Cavity," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 5, pp. 579-592, 1992.

[29] Knud Bank Christensen, "The Application of Digital Signal Processin to Large-Scale Simulation of Room Acoustics: Frequency Response Modeling and Optimization Software for a Multichannel DSP Engine," *Journal of the Audio Engineering Socitey*, vol. 40, no. 4, pp. 260-274, April 1992.

[30] John Chowning, "The Simulation of Moving Sound Sources," *Computer Music Journal*, vol. 1, no. 3, pp. 48-52, 1977.

[31] John M. Chowning, *Selections: Turenas, Stria, Phone, Sabelithe* C.D., Mainz, W. Germany: Wergo, prod. 1988.

[32] P. M. Clarkson, J. Mourjopoulos, and J. K. Hammond, "Spectral, Phase, and Transient Equalization for Audio Systems," *Journal of the Audio Engineering Socitey*, vol. 33, no. 3, pp.127-131, March 1985.

[33] Duane H. Cooper, "Problems with Shadowless Stereo Theory: Asymptotic Spectral Status," *Journal of the Acoustical Society of America,* vol. 35, no. 9, pp. 629-642, Sept. 1987.

[34] Duane H. Cooper and Jerald L. Bauck, "Prospects for Transaural Recording," *Journal of the Audio Engineering Society*, vol. 37, no. 1/2, pp. 3-19, Jan./Feb. 1989.

[35] Duane H. Cooper and Jerald L. Bauck, "Generalized Transaural Stereo," Presented at *93th Convention of Audio Engineering Society, San Francisco*, Preprint 3401 (J2), Oct. 1992.

[36] Don Davis, and Carolyn Davis, *Sound System Engineering*, 2nd ed., Indianapolis: Howard W. Sams & Co., 1987.

[37] V. F. Dem'yanov and V. N. Malozemov, *Introduction to Minimax*, trans. D. Louvish, Israel: Keter Publishing House Jerusalem Ltd., 1974.

[38] W. E. Feddersen, T. T. Sandel, D. C. Teas, and L. A. Jeffress, Localization of high-frequency tones. *Journal of the Acoustical Society of America*, vol.29, pp. 988-991, 1957.

[39] Mathias Fink, "Time-Reversal of Ultrasonic Fields-Part I:Basic Principles," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 5, pp. 555-566, 1992.

[40] Giovanni Gabrieli, *Sacrae symphoniae, Selections, The antiphonal music of Gabrieli*, L.P., Columbia Records, prod. 1969.

[41] Michael A. Gerzon, "Ambisonics in Multichannel Broadcasting and Video," *Journal of the Audio Engineering Socitey*, vol. 33, no. 11, pp. 859-871, Nov. 1985.

[42] Gene H. Golub, *Matrix computations*, Baltimore: Johns Hopkins University Press, 1989.

[43] Joseph W. Goodman, *Introduction to Fourier Optics*, San Francisco: McGraw-Hill, 1968.

[44] Toshiyuki Gotoh, Yoichi Kimura and Akitoshi Yamada, "A new sound image localization control system for stereophonic recording," *Journal of the Acoustical Society of Japan, (E)*, vol. 5, no. 2, pp. 85-94, 1984.

[45] Martin H. Gutknecht, Julius O. Smith and Lloyd N. Trefethen, "The Carathéodory-Fejér Method for Recursive Digital Filter Design," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-31, no. 6, pp. 1417-1426, Dec 1983.

[46] R. W. Hamming *Digital Filters*, 2nd ed., New Jersey: Prentice-Hall, Inc., 1983.

[47] Heringa et al., "The Acoustics and Sound System for Hemispherical Film Projection," *Journal of the Audio Engineering Society*, vol. 35, no. 3, pp. 119-129, March 1987.

[48] Takayuki Hidaka, "Sound-Field Simulator for Room Acoustic Design and Assessment-Introduction of Wave-Theoretical Treatment to Synthesized Sound," *Journal of the Audio Engineering Socitey*, vol. 41, no. 11, pp. 914-919, Nov. 1993.

[49] Jean-Marc Jot, Veronique Larcher, and Olivier Warusfel, "Digital Signal Processing Issues in the Context of Binaural and Transaural Stereophony," presented at *98th Convention of Audio Engineering Society, Paris*, Preprint 3980 (I6). year????

[50] Jean-Marc Jot, and Olivier Warusfel, "Spat : A Spatial Processor for Musicians and Sound Engineers," Presented at *CIARM'95, Ferrara, Italy*, May 1995.

[51] Stephen Julstrom, "An Intuitive View of Coincident Stereo Microphones," *Journal of the Audio Engineering Socitey*, vol. 39, no. 9, pp. 632-649, Sept. 1991.

[52] Stephen Julstrom, "A High-Performance Surround Sound Process for Home Video," *Journal of the Audio Engineering Society*, vol. 35, no. 7/8, pp. 536-549, July/August 1987.

[53] L. V. Kantorovich and G. P. Akilov, *Functional Analysis in Normed Spaces*, D. E. Brown trans., A. P. Robertson ed., New York: Pergamon Press Ltd., 1964.

[54] Gary S. Kendall, William L. Martens, and Shawn L. Decker, "Spatial Reverberation: Discussion and Demonstration," in *Current Directions in Computer Music Research*, Max Mathews and John Pierce ed., Cambridge: MIT Press, pp. 65-87, 1989.

[55] Mendel Kleiner, Bengt-Inge Dalenback, and Peter Svensson, "Auralization-An Overview," *Journal of the Audio Engineering Socitey*, vol. 41, no. 11, pp. 861-874, Nov. 1993.

[56] Setsu Komiyama, "Subjective Evaluation of Angular Displacement between Picture and Sound Directions for HDTV Sound Systems," *Journal of the Audio Engineering Society*, vol. 37, no. 4, pp. 210-214, April 1989.

[57] Setsu Komiyama, Akira Morita, Kohichi Kurozumi and Katsumi Nakabayashi, "Distance control of sound images by a two-dimensional loudspeaker array," *Journal of the Acoustical Society of Japan, (E)* vol. 13, no. 3, pp. 171-180, 1992.

[58] Y. Kubota and E. H. Dowell, "Asymptotic modal analysis for sound fields of a reverberant chamber," *Journal of the Acoustical Society of America*, vol. 92, no. 2, pt. 1, pp. 1106-1112, Aug. 1992.

[59] Y. Kubota and E. H. Dowell, "Asymptotic modal analysis for sound fields of a reverberant chamber," *Journal of the Acoustical Society of America*, vol.92, no.2, pt.1, pp. 1106-12, Aug. 1992.

[60] Fernando Lopez-Lezcano, "Quad sound playback hardware and software," unpublished document, 1991.

[61] Fernando Lopez-Lezcano, "A four channel dynamic sound location system," unpublished document, 1991.

[62] C. J. MacCabe, and D. J. Furlong, "Virtual Imaging Capabilities of Surround Sound Systems," *Journal of the Audio Engineering Society*, vol. 42, no. 1-2, pp. 38-49, Jan./Feb. 1994.

[63] A. H. Marshall, C. W. Day and L. J. Elliott, "The Acoustical Design of a 4000-Seat Church," *Journal of the Acoustical Society of America*, vol. 35, no. 11, pp. 897-906, Nov. 1987.

[64] Max Mathews, and John Pierce, eds. *Current Directions in Computer Music Research* Cambridge, MA: The MIT Press, 1989.

[65] J. D. Maynard, E. G. Williams and Y. Lee, "Nearfield acoustic holography: I. Theory of generalized holography and the developement of NAH," *Journal of the Acoustical Society of America*, vol. 78, no. 4, pp. 1395-1413, Oct. 1985.

[66] David G. Meyer, "Computer Simulation of Loudspeaker Directivity," *Journal of the Audio Engineering Society*, vol. 32 pp. 294-315, 1984.

[67] Franklin Miller, Jr., *College Physics*, 2nd ed., New York: Harcourt, Brace and World, Inc., 1967.

[68] Akira Mochimaru, "A Study of the Practicality and Accuracy of Impulse Response Calculations for the Auralization of Sound System Design," *Journal of the Audio Engineering Socitey*, vol. 41, no. 11, pp. 882-893, Nov. 1993.

[69] Moore, F. Richard, "A general model for spatial processing of sounds," *Computer Music Journal*, 7(3), 6-15, 1983.

[70] Masayuki Morimoto, and Yoichi Ando, "On the simulation of sound localization," *Journal of the Acoustical Society of Japan, (E)*, vol. 1, no. 3, pp. 167-174, 1980.

[71] A. D. Musicant and R. A. Butler, "Monaural localization: an analysis of practice effects," *Perception and Psychophysics*, vol. 28, pp. 236-240, 1980.

[72] P. A. Nelson and S. J. Elliott, *Active Control of Sound.* New York: Academic Press, 1992.

[73] P. A. Nelson, "Active control of acoustic fields and the reproduction of sound," *Journal of Sound and Vibration*, vol. 177, no. 4, pp. 447-477, 1994.

[74] Ole Kirkeby and P. A. Nelson, "Reproduction of plane wave sound fields," *Journal of the Acoustical Society of America*, vol. 94, no. 5, pp. 2992-3000, Nov. 1993.

[75] Alec Nisbett, *The Technique of the Sound Studio: for Radio, Recording Studio, Television, and Film*, 4th ed., New York: Focal/Hastings House, 1979.

[76] J. M. Nuetzel and E. R. Hafter, "Lateralization of complex waveforms: effects of fine structure, amplitude, and duration," *Journal of the Acoustical Society of America*, vol.60, pp. 1339-1346, 1976.

[77] Thomas W. Parks and James H. McClellan "Chebyshev Approximation for Nonrecursive Digital Filters with Linear Phase," *IEEE Transactions on Circuit Theory*, vol. CT-19, no. 2, March 1972.

[78] William Putnam, Davide Rocchesso and Julius Smith, "A Numerical Investigation of the Invertibility of Room Transfer Functions," presented at *IEEE ASSP Workshop on Applications of Signal Processing to Audio and Acoustics, New Paltz, New York*, Oct. 1995.

[79] Daniel P. Petersen, and David Middleton, "Sampling and Reconstruction of Wave-Number-Limited Functions in N-Dimensional Euclidean Spaces," *Information and Control*, vol. 5, pp. 279-323, 1962.

[80] J. N. Mourjopoulos, "Digital equalization of room acoustics," *Journal of the Audio Engineering Society*, vol. 42, no. 11, pp. 884-900, Nov. 1994.

[81] Anthony Romano, "Three-Dimensional Image Reconstruction in Audio," *Journal of the Audio Engineering Society*, vol. 35, no. 10, pp. 749-759, Oct. 1987.

[82] Mario Rossi, *Acoustics and Electroacoustics*, Trans. Patrick Rupert Windsor Roe, Norwood MA: Artech House Inc., 1988.

[83] Thomas D. Rossing, *The Science of Sound*, 2nd ed., New York: Addison-Wesley Publishing Company Inc., 1990.

[84] N. Sakamoto, T. Gotoh, and Y. Kimura, "On 'out-of-head localization' in headphone listening," *Journal of the Audio Engineering Society*, vol.24, pp. 710-715, 1976.

[85] N. Sakamoto, T. Gotoh, T. Kogure, M. Shimbo and A. Clegg, "Controlling Sound-Image Localization in Stereophonic Reproduction," *Journal of the Audio Engineering Society*, vol. 29, no. 11, Nov. 1981.

[86] E. A. G. Shaw, and R. Teranishi, "Sound pressure generated in an external-ear replica and real human ears by a nearby point-source," *Journal of the Acoustical Society of America*, vol. 44 pp. 240-249, 1968.

[87] C. W. Sheeline, *An investigation of the effects of direct and reverberant signal interaction on auditory distance perception*, PhD dissertation, Stanford University, Department of Hearing and Speech Sciences, 1982.

[88] M. R. Schroeder, B. S. Atal, "Computer Simulation of Sound Transmission in Rooms," *IEEE International Convention Record,* pt. 7, pp. 150-155, 1963.

[89] M. R. Schroeder, *Number Theory in Science and Communication.* New York: Springer-Verlag, 1984.

[90] M. R. Schroeder, "Digital simulation of sound transmission in reverberated spaces," *Journal of the Acoustical Society of America*, vol. 47, pp. 424-431, 1970.

[91] M. R. Schroeder, "Natural Sounding Artificial Reverberation," *Journal of the Audio Engineering Society*, vol. 10, no. 3, pp. 219-223, July 1962.

[92] Robert F. Stengel, *Optimal Control and Estimation*, New York: Dover Publications, Inc., 1994.

[93] Tsuyoshi Usagawa, Seizi Nishimura, Masanao Ebata and Josuke Okda, . "Analysis of moving sound sources -Separation of multiple sources-," *Journal of the Acoustical Society of Japan, (E)* vol. 10, no. 4, pp. 189-195, 1989.

[94] Peter Vogel, *Application of Wave Field Synthesis in Room Acoustics*, DRIR dissertation, Technische Universiteit te Delft, Netherlands:Groot Hertoginnelaan, 1993.

[95] William A. Yost, and Donald W. Nielsen, *Fundamentals of Hearing*, 2nd ed., New York: Holt, Rinehart and Winston, 1985.

[96] Rhonda Wilson, "Equalization of Loudspeaker Drive Units Considering Both On- and Off-Axis Responses," *Journal of the Audio Engineering Socitey*, vol. 39, no. 3, pp. 127-139, March 1991.

[97] D. Wright, J. H. Hebrank, and B. Wilson, . Pinna reflections as cues for localization. *Journal of the Acoustical Society of America* vol. 56 pp. 957-962, 1974.

[98] John M. Woram, *The Recording Studio Handbook*, 1st ed., Plainview, N.Y.: ELAR Pub. Co., 1981.

[99] Yoshio Yamasaki and Takeshi Itow, "Measurement of spatial information in sound fields by closely located four point microphone method," *Journal of the Acoustical Society of Japan, (E)* vol. 10, no. 2, pp. 101-110, 1989.

[100] M. Yanagida, O. Kakusho, and T. Gotoh, "Application of the least-squares method to sound-image localization in multi-loudspeaker multi-listener case," *Journal of the Acoustical Society of Japan, (E)*, vol. 4, no. 2, pp. 107-109, 1983.