

The Pennsylvania State University
The Graduate School
Graduate Program in Acoustics

MEASUREMENT AND CALCULATION OF INDIVIDUAL
HEAD-RELATED TRANSFER FUNCTIONS
USING A BOUNDARY ELEMENT MODEL
INCLUDING THE MEASUREMENT AND
EFFECT OF SKIN AND HAIR IMPEDANCE

A Thesis in
Acoustics
by

Brian F.G. Katz

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Doctor of Philosophy

May 1998

©1998 Brian F.G. Katz

Abstract

This research investigates various aspects of the Head-Related Transfer Function (abbreviated HRTF), which is a description of the acoustic frequency filtering performed by the geometry of the head as a function of incident angle. The effects of this filtering are used in the brain to determine the location of sound sources in space. Initially, various methods for measuring the HRTF are examined, as well as several means of normalizing or equalizing the data. One method is chosen which best represents the informational content of the measured data for comparisons between experimental methods. The question as to whether the acoustic properties of skin and hair contribute to the HRTF is also examined. Measurements are made of the acoustic absorption and impedance of various skin and hair samples using a plane wave tube and two microphones. The limitations of this technique and published standards are also included. Finally, an individual HRTF is calculated using an optically generated surface mesh and a numerical boundary element (BEM) solution. The results of the impedance measurements are included in the calculations. Final analysis consists of comparing the various calculated HRTFs and measured HRTFs. Geometric variations in the head mesh such as removal of the pinna are also included. Good agreement is found given the assumptions made in the generation of the computational model (*i.e.* lack of torso) throughout the frequency range of the model, which extends from 1–6 kHz. Computational speed and size of the numerical problem limit the work to this region.

Contents

List of Figures	vii
List of Tables	x
Acknowledgements	xi
Chapter 1. Introduction	1
1.1. Research Overview	1
1.2. Thesis Organization	2
Chapter 2. Literature Review	3
2.1. Localization	3
2.1.1. Justification	3
2.1.2. Historical Review – Up to 1960	3
2.1.3. Historical Review – 1960 to 1973	12
2.1.4. Current Literature	13
2.1.5. Anatomical and Neurological Processes in Localization	16
2.1.6. Anatomical and Neurological Processes of Localization for a Monaural Listener	19
2.2. Absorption and Impedance Measurements	20
2.3. HRTF Calculations	20
Chapter 3. Acoustical Parameter Measurements	22
3.1. Measurement Technique	22
3.2. Apparatus	23
3.2.1. Microphone Calibration	24
3.2.2. Microphone Distance to Termination	26
3.2.3. Spacing Between Microphones	30
3.2.4. Planar Surface Assumption	33
3.3. Measurements	35
3.3.1. Reference Terminations	35
3.3.2. Effect of distance errors	38
3.3.3. Skin	40
3.3.4. Hair	41
Chapter 4. HRTF Measurement	47
4.1. HRTF Equalization and Labelling	47
4.2. Session WPAFB	50

4.2.1. Apparatus	54
4.2.2. Individual HRTF	54
4.2.3. Pinna Replica HRTF	57
4.2.3.1. BAMAR Head	57
4.2.3.2. Pinna Replica Construction	57
4.2.3.3. Measurement	58
4.2.4. WPAFB Results	58
4.3. Session CRE	73
4.3.1. Apparatus	73
4.3.2. Individual HRTF	75
4.4. Measured HRTF Data Comparison	80
Chapter 5. HRTF Calculation	86
5.1. Boundary Element Method Approach	86
5.2. BEM Model Generation	88
5.2.1. Head Geometry Acquisition	88
5.2.2. Mesh Modification	89
5.2.2.1. Data Format	91
5.2.2.2. Closure	91
5.2.2.3. Coarsen Mesh	92
5.2.2.4. Refine Mesh	93
5.2.3. Field Mesh	93
5.3. Calculation	96
5.3.1. Calculation Speed	96
5.3.2. Various Meshes	98
5.4. Results	113
5.4.1. Spherical model	113
5.4.2. Rigid models	116
5.4.3. Impedance Condition	121
Chapter 6. HRTF Comparisons	127
Chapter 7. Conclusion	132
References	134
Appendix A. MatLab Scripts	138
A.1. <code>anspace.m</code>	138
A.2. <code>polynul2.m</code>	139

A.3. <code>whatc.m</code>	142
Appendix B. C Code	144
B.1. <code>decimate.c</code>	144
B.2. <code>edgealign.c</code>	158
Appendix C. Sysnoise Related Command Files	162
C.1. Example setup command file for head mesh	162
C.2. Hair definition command file	162
C.3. Example setup command file for sphere mesh	163
C.4. Example computation command file	163
C.5. Unix <code>sysmon</code> script	164
C.6. Example <code>sysmon</code> log file	165
Appendix D. Data	168
D.1. WPAFB source locations	168
D.2. CRE source locations	173

List of Figures

Figure	page
2.1. Investigation of perceived frontal elevation of band limited noise	13
2.2. Apparatus and results for investigation of perceived elevation	14
3.1. Impedance tube setup	23
3.2. Impedance tube measurement configurations	25
3.3. Distance determination: pressure data	27
3.4. Distance determination: transfer functions	28
3.5. Distance determination: interpolation results	29
3.6. Effect of asymmetric data on null interpolation routine	31
3.7. Separation determination: phase calculation results	33
3.8. Reflection coefficient reference measurements for idealized terminations	36
3.9. Anechoic termination impedance measurement	36
3.10. Reflection coefficient reference measurements for open termination	37
3.11. Example of the smoothing of measured impedance data	38
3.12. Effect of error in s and l on Z	39
3.13. Reflection coefficient measurements for skin	40
3.14. Reflection coefficient measurements for hair	43
3.15. Absorption coefficient measurements for hair and rockwool.	44
3.16. Impedance calculations for hair	46
4.1. Diffuse field equalization function for various measurements	49
4.2. HTRF Equalization methods: Free field	51
4.3. HTRF Equalization methods: Diffuse field	52
4.4. HTRF Equalization methods: Front-source	53
4.5. WPAFB microphone design and placement	55
4.6. WPAFB spherical source array	56
4.7. WPAFB BAMAR mannequin	58
4.8. WPAFB pinna mold experimental setup	59
4.9. HRTF Subject 1 with symmetry sources front	61
4.10. HRTF Subject 1 with symmetry sources rear	62
4.11. HRTF Subject 1 with symmetry sources opposite ear side	63

4.12. HRTF Subject 1 with symmetry sources listening ear side	64
4.13. HRTF Subject 2 with symmetry sources front	65
4.14. HRTF Subject 2 with symmetry sources rear	66
4.15. HRTF Subject 2 with symmetry sources opposite ear side	67
4.16. HRTF Subject 2 with symmetry sources listening ear side	68
4.17. HRTF Subject 1 with symmetry sources the front median plane	69
4.18. HRTF Subject 2 with symmetry sources the front median plane	70
4.19. Front source equalized HRTF Subject 1 with symmetry sources the front median plane	71
4.20. Front source equalized HRTF Subject 2 with symmetry sources the front median plane	72
4.21. Variation in pinna mold HRTF between subjects	74
4.22. CRE Source on listening ear side	76
4.23. CRE Source opposite listening ear side	77
4.24. CRE Source rear	78
4.25. CRE Sources in frontal median plane	79
4.26. Equalization functions for three experimental measurements	82
4.27. Experimentally measured HRTF for 1L with source: left	83
4.28. Experimentally measured HRTF for 1L with source: right	84
4.29. Experimentally measured HRTF for 1L with source: rear	85
5.1. Cylindrical scanner omission effect	90
5.2. Mesh coarsening methods	92
5.3. Mesh refining methods	94
5.4. BEM WPAFB field mesh	95
5.5. BEM calculation space and time requirements	97
5.6. BEM head mesh: Isometric view	100
5.7. BEM head mesh: Front view	101
5.8. BEM head mesh: Side view	102
5.9. BEM head mesh: Right ear close-up	103
5.10. BEM head mesh: Underside view	104
5.11. BEM head mesh with Hair subset defined	105
5.12. BEM head mesh with pinna removed: Isometric view	106

5.13. BEM head mesh with pinna removed: Front view	107
5.14. BEM head mesh with pinna removed: Side view	108
5.15. BEM head mesh with pinna removed: Top view	109
5.16. BEM small spherical mesh: diameter = inter-auditory spacing	110
5.17. BEM large spherical mesh: volume equal to head mesh	111
5.18. BEM spherical mesh with hair element set	112
5.19. Point source on a sphere: theory <i>vs.</i> BEM	114
5.20. Effect of varying overdetermination element location and number	115
5.21. BEM rigid model solutions- sources 9 and 12	117
5.22. BEM rigid model solutions- sources 1 and 6	118
5.23. BEM rigid model solutions- sources 18 and 86	119
5.24. BEM rigid model solutions- sources 52 and 57	120
5.25. BEM hair model solutions- sources 9 and 12	122
5.26. BEM hair model solutions- sources 1 and 6	123
5.27. BEM hair model solutions- sources 18 and 86	124
5.28. BEM hair model solutions- sources 52 and 57	125
6.1. Front source equalization functions for measured and calculated HRTFs . . .	127
6.2. HRTF measurement and calculation comparisons for sources left and right . .	129
6.3. HRTF measurement and calculation comparisons for high and low frontal median plane sources	130
6.4. HRTF measurement and calculation comparisons for rear source position . . .	131

List of Tables

Table	page
3.1. Microphone distance calculations.	32
3.2. Measured skin areas	40
3.3. Hair sample details	41

Acknowledgements

There are many numbers of people who have helped make this work possible. Let me first thank those who donated their time and/or facilities so that it was possible to perform the work as desired. The facilities necessary to carry out this research were beyond those available at Penn State, or any other single institution. The following people and laboratories agreed to provide equipment and assistance at no cost other than the recognition and thanks which they deserve.

At the Wright-Patterson Air Force base I would like to thank Richard McKinley for the use of the BioAcoustics Lab's facilities; Dennis Allen, David Ovenshire, and Robert Bolia for their help with the acoustic data collection; Angie McCavitt and Linda Morris for making the pinna molds; Kathleen Robinette for the use of the optical scanner; and finally Sherri Blackwell and Tina Brill for making the head scans. At Aureal Semiconductor Inc. I would like to thank Scott Foster for the informative discussions and use of the Crystal River Engineering facilities, Martin Walsh for running the equipment, and Dale Sparling for assisting with the data decoding. At LMS Numerical Technologies (LMS-NIT) I would like to thank Sysnoise and Tim Coats for providing me with additional copies of the Sysnoise acoustics modelling package so that I could complete my research in a timely fashion. In addition, at the Applied Research Lab I would like to thank Richard Stern for supporting me through the final stages of this work and also Timothy Kulbago for his guidance in the initial stages of the `decimate` routine.

In the acoustics program I would like to thank Karen Brooks for her patience in listening to my many rantings and ravings, Carolyn Smith for supporting me through some of said rantings, Maya Freeman for her exhaustive grammatical review of this work (though she didn't get to see all of it so those errors are my fault), and Tim Leishman for his continued encouragement in pursuing this new field and for being a sounding board for many of the novel approaches used in this work.

My supervisor, Jiri Tichy, deserves a great deal of thanks, and congratulations, for allowing me to pursue this topic and for not throwing me out of his office during the many heated discussions we had. I would also like to thank the other members of my committee who provided a great deal of guidance and support throughout the many stumbling blocks encountered throughout this research.

And finally, I would like to thank my wife Marjut, for giving up more than I will ever be able to repay in order that I finish this research. She has been my inspiration and my muse throughout this work. I would surely have fallen apart if it were not for her.

Chapter 1.

Introduction

Localization of sound is a very complex phenomenon in humans. It serves in a variety of ways, from pinpointing a potential attacker to perceiving fullness and envelopment in a concert hall. Determining a mental picture of the surrounding environment without the need to look around, and move around, is an ability which may be unique to humans. One of the key factors in this process which has come to light over the past few decades is the Head-Related Transfer Function, or HRTF. The head, and especially the pinna (or outer ear), through its complicated shape changes the frequency content of sounds reaching the ear drum. These changes are dependent on the sound's incident direction. The HRTF is the set of filters representing these changes which are described as function of position around an individual. When used correctly, this function is capable of reproducing spatial acoustic scenes over headphones or loudspeakers. It is the first step in understanding how people localize, and what further processes are involved. As the function is based upon the geometrical characteristics of the head and ears, to fully understand the HRTF it is necessary to relate the structural characteristics with the acoustic results. This research hopes to provide another tool for analyzing this relationship. In order to obtain real knowledge of the geometric effects, the ability to change the geometry and monitor the results is almost essential. As it is rather difficult to alter an individual's head and pinna shape, it is proposed that a computer model of the surface of an individual be modelled, and the HRTF be calculated. The computer model has the property that changes can be made to a person's head shape with ease. Such changes could be as drastic as removing the ears entirely. If this model can be verified, using real measurements of the same individual, it could provide a new method for understanding the Head-Related Transfer Function, and in so doing results in a better understanding of human localization. This is the goal of this work.

1.1. Research Overview

This research entailed the calculation of an individual HRTF from geometrical data. The calculation was based on a boundary element model (BEM) simulation, disregarding any internal structure of the head. The geometrical surface model mesh was obtained by using a laser scanning device on an individual subject. In order to evaluate the calculation, HRTF measurements of individuals were also made. Several methods for obtaining this measurement were used.

One component of the system investigated was the effect of the acoustical impedance of the skin and hair of the subject. A measurement system was built, capable of measuring the acoustical impedance of human skin (without surgical removal) and also of hair samples (removed) for frequencies over which the simulation was utilized. Various conditions of skin

were measured, including variations in thickness, presence of bone, and variability between individuals. Hair was also studied under various conditions, including packing density and variations between individuals. The results of these measurements were included in the calculation and subsequent effects were observed.

An additional physical HRTF measurement was made in which the impedance of the individual was made more uniform and rigid. To achieve this, pinnae molds of the subjects were made which were placed in an adjustable rigid head. The HRTFs measured in this system were an approximation of a purely rigid head and were also compared to the various other measurements and calculations.

The comparison of the various HRTF measurement results was informative on its own. Using three different techniques and two totally separate apparatus, all with the same subjects, it was possible to examine the variations not only between people (which is well known though not understood) but the variations due to experimental methods, conditions, and inherent assumptions. This was of interest as the HRTF is the key component in many virtual reality auralization simulations, and is assumed to reflect individual specific acoustical characteristics of the subject. Therefore it should not be measurement specific. These variations indicate a possible additional cause of error in auralization systems.

1.2. Thesis Organization

The work contained in this thesis consisted of several separately definable, though very interconnected, research projects, separated into chapters. Though each project can stand on its own, it is important to keep the concept of the entire research in mind throughout the work.

The next chapter addresses the historical literature related to human aural localization research, human surface acoustical impedance measurements, and attempts to generate computational models of the effects of geometric structure on human listening. The extensive literature review concerning research into human localization begins before the turn of the century. The reason for this was to present a clear view of the evolution of this research and how this research relates to the progress in the field.

The third chapter details measurements made of the acoustical impedance of skin and hair. This entailed the design of a high frequency impedance plane wave tube. One result of this work, in addition to the acoustical data, was an investigation and realization of the errors and limitations of such an experiment.

Chapter four investigates several methods for measuring head-related transfer functions. Analysis includes the effects of different measurement techniques, and various means of examining and reviewing HRTF data. The fifth chapter details the computational calculation of the head-related transfer function. This includes the process used to generate a computationally ready model of an individual subject. The results of the acoustical impedance measurements are included in the computations. Finally, chapter six compares the results of the measured and calculated head-related transfer functions under the various conditions used throughout the work. A number of appendices are included which contain source code written during the course of this research. In addition are some details of experimental geometries.

Chapter 2.

Literature Review

2.1. Localization

2.1.1. Justification

In order to understand fully the reasoning behind the current research, it is important to examine previous and current knowledge of human aural localization ability and methods. It is the background of research efforts in this area that have prompted, and been the justification of, the pursuit of this research endeavor. For over 100 years scientists from many disciplines –physics, psychology, physiology, and audiology– have attempted to determine what the mechanisms are that enable humans to localize sound sources in three-dimensional space. It is the hope of this researcher that this work will be a real contribution to this ongoing area of study.[†]

2.1.2. Historical Review – Up to 1960

The earliest paper published concerning human’s ability to localize sounds was entitled “Our Perception of the Direction of a Source of Sound.” This paper was published in 1876 by Lord Rayleigh. [36] In this simple experiment the test subject was placed on an open lawn, blind-folded and surrounded by a number of people standing in a circle. Using sound sources of either human voice or tuning forks (up to 256 Hz), the ability to locate the sound was examined. It was found that people are very good at locating voices, needing only a word or vowel, to accurately place the source within errors of only a few degrees. Pure tones were much more difficult to locate; only general left/right determination was possible. There was no accuracy, and many instances of sounds in the front being localized to the rear, and vice versa. In addition, it was determined through calculations that the difference in levels at each ear, based on the head being a fixed sphere, would not be sufficient to determine the position (left versus right) of the source. From this experimental evidence the strong statement was made that the ability to distinguish from front to back “would thus appear to depend on the compound character of the sound in a way that is not easy

[†] In reviewing the literature it was quickly found that the number of works is tremendous when considering such a large expanse of time and effort throughout history. The research which will be presented in this review is very much an overview and should not be considered as all-encompassing. In addition, there are many works that could not be included or even reviewed due to lack of availability and/or translations. This is most evident in the period of the 1930’s–1960’s when there was a great deal of work done in Austria and Germany that has not been translated.

to understand, and for which the second ear would be of no advantage.” [36][p.33] Further work was done using higher frequency sounds [37] in which a gas bag powered whistle was used. It was noticed that for these high frequencies there was a shadowing effect of the head.

In 1877, one of the real pioneers in localization research published his first paper, “On Binaural Audition,” by Silvanus Thompson. [51] The experiment consisted of isolating the subject in one room and feeding rubber tubes to each ear from other rooms where sources were located at the other end of the tubes. This would ensure that the signals the subject heard were from the tubes, that no sound energy would propagate by any other path from the source to the subject. The tubes were placed either very close to the ears or in the ears. The sources used were matched tuning forks, one slightly mistuned. The results were that the subjects heard beat frequencies (up to 2 beats/sec.) resulting from the mismatched forks, but not the difference tones, which would be heard if both forks were placed in front of the same tube.

The following year, after seeing Alexander G. Bell demonstrate his new invention—the telephone—Thompson decided it was a good way to present separate sounds to subjects. It was an ideal setup for 180° phase reversals which would remain steady, as opposed to the tuning fork beat frequency method. This experiment provided very drastic results;

I arranged a . . . microphone . . . with two Bell telephones . . . Placing the telephones to my ears, I requested my assistant to tap on the wooden support of the microphone. The result was deafening. I felt as if simultaneous blows had been given to the tympana of my ears. But on reversing the current through one telephone, I experienced a sensation only to be described as someone tapping with a hammer on the back of the skull from the inside. [52][p.386]

Variations of the sound level with the phase reversal in place resulted in the sound image moving along the back of the head towards the ear in which the source was louder, until the image was at that ear. These effects were independent of pitch.

After these results, additional experiments in phase were deemed necessary, especially the effects of varying phase. This again required the use of tuning forks and tubes. Instead of the mismatched forks, a single tuning fork was used. The signal from the tuning fork was led to the ears via tubes, and phase effects were created by varying the lengths of the tubes in reference to each other. This resulted in a mix of perceptions, with part of the sound image in the ears and part at the back of the head. An additional setup, using a curved wire with the ends placed in the ears and the tuning fork placed at different positions on the wire, provided the same results. Further experiments involved using organ pipes as sources, being better suited for sustained playing. It was again noted that difference tones are not perceived when the signals are fed separately, but do exist when mixed before presentation to the subject. [53]

The conclusion of this was that the brain combines the signals in its processing, resulting in the perception of beat frequencies and phase. This was in contrast to the views expressed by Helmholtz in *Sensations of Tone*. Helmholtz stated that the perception of complex sounds is only dependent on differences in amplitude, and in no way on

differences of phase. [52] Additional comments were made by Taylor in *Sound and Music*, “The ear being deaf to differences of phase in partial-tones, perceives no distinction between such modes of vibration...but merely resolves them into the same single pair of partial-tones”. [52][p.388] No attempts were made by the author to justify this claim or to understand whether the basis of these phenomena, were physical, physiological, or psychological.

A ground breaking experiment performed by Rayleigh in 1882 looked at the localization ability of someone deaf in one ear. [38] The responses, though containing many errors, were not random enough to attribute the correct answers to chance. Interestingly, the subject made mistakes that are not known to occur in binaural hearing subjects. Primarily there were left/right confusions, which are not present in subjects with normal hearing.

In 1894 an in-depth experimentation on localization of symmetric sources and effects of subject motion was performed by Münsterberg and Pierce using loudspeakers placed upon a 1 meter diameter spherical cage. [32] It was found that placing the sources at $\pm 90^\circ$ azimuth (with straight ahead being 0°) resulted in externalized sound images at 0° and 180° , but never images in the head. The image would sometimes move from 0° to 180° and back. Subsequent positioning of the sources showed consistent results for any symmetrical spacing. In addition, front/back symmetry spacing could be used with the same results, *i.e.* $+20^\circ$ and -110° .

Using two symmetrical sources, the radial distance was increased for one source. For this experiment the image would start in the center and move towards the ear which had higher sound levels. No front/back reversals occurred during radial source motion. Using two sources at 0° and 180° and radially increasing the distance of one source resulted in the image at either 0° or 180° with no front/back reversals. Using two front/back symmetrical sources on one side of the head, it was observed that the image would appear at either source position or the median of the two positions.

Subject rotation and its effect on localization was tested using a constant tone (non-moving) source, with the subject rapidly rotated in one direction. After an abrupt stop, the sound image still appeared to rotate around the subject. Another interesting experiment, in which the subject (in a non-rotating position) looks to one side without moving their head, shows an unexpected result. The subject here is blindfolded and there is only one source. The experimental results put the image at least 10° to the rear.

The different theories of the day (and their authors) were summarized as such. (Stumpf) The actual input to the left and right ears are different, from experience of association the position is determined. (Preyer) Spatial sensory nerves in the ear canal exist which respond to the direction of sound. (Kries and Block) Source location is derived from the level intensity difference between left and right ear. In reviewing the results of their experiments Münsterberg and Pierce proposed the following new theory regarding localization. There exists a union of sensations of movement and of sound. From observation, it was seen that the instinctual reaction of subjects when hearing a sound is to turn towards the source. Over time, associations are built which correlate the sounds heard to the expected position of the head location, removing the need for actual motion.

Observing the differences in localization ability of different people, and in an attempt to quantify localization in some way, Münsterberg and Pierce also make the observation

that people are built differently. “The differences in the shape of ear and head, as well as the differences in the hair, beard, etc., are so marked and are such important factors in determining any particular localization of sounds within a few feet of the head that no one individual’s special localization can be considered a standard for those of another.” [32][p.466]

In 1901 Rayleigh repeats Thompson’s experiment of 1878 (without credit) with the addition of taking care to ensure no propagation of sound from opposite telephone speakers. [39] By using very low level sounds and creating the same effects, Rayleigh concluded that the beat effects could not be due to sound passage through the bones or soft tissue in the head.

One of the most extensive investigations of single ear localization, using mono-deaf or “monaural” subjects was done in 1901 by Angell and Fite. [1] The primary subject was 30 years old, and had been deaf in one ear since the age of four. A 1.3 meter sphere cage was used to place the source. The subject had an accuracy of 75%, compared to a normal listener who typically has 60–75% with complex sounds such as speech. The majority of errors placed sources which were located in the center (median plane) as images closer to the good ear. The subject was allowed more time than the typical listener, and was also allowed to replay the sound at their convenience.

In an attempt to determine how the subject was able to localize, several attempts at masking were done. One theory (presented in a fashion previously under Preyer) suggested that there are sensory nerves in the head and neck which indicate the direction of a sound source. To mask these effects, the subject was covered with a very heavy cloth in which there was only a small opening for the subject’s good ear. This had no effect in reducing localization ability. Another theory was that the subject was using visual knowledge of the room and acoustic room effects, such as reverberation and reflections to help determine the location of the source. To mask this, the subject was rotated and moved to random points in the room while blindfolded. This also produced no reduction in localization ability.

An analysis of the errors showed that front/back determination was better than that of the average listener. Localization on the side of the functioning ear was good, with results on the opposite ear being more varied and uncertain. A comparison was made using pure tones. The results were a degradation of localization to 14%, but it had already been reported that pure tones are difficult to localize. It was noticed during the experiment that obstacles like clothing, when obstructing the direct source-ear path, induced errors. From this, it was determined that the quality of the sound is a factor in localization. “Sounds which are complex in nature undoubtedly undergo modification through the dampening [sic] and reinforcing of their partial tones by the pinna, the external meatus and the head, in a manner which must vary somewhat regularly with variations in the spatial position of the object from which the sound emanates.” [1][p.244]

Further tests were performed on monaural subjects who had been deaf in one ear for a period ranging from one to more than 26 years. It was noticed that the longer the subject had been deaf, the better their localization ability. Localization practice also increased localization ability. Without informing the subjects of their results during the experiments, the subject who had been deaf for just 1 year increased from 12% to 25% accuracy. This improvement was attributed to simple concentration on the act of localizing.

Range localization was also tested, having the source and subject in an open field. The results were comparable to a normal listener.

More and Fry, using different length tubes, as in the earlier works of different authors, reproduced the older phase experiments, but with one addition. [29] Instead of simply asking the subjects where the image was, they inscribed a circle on the floor on which the subject was to indicate the direction. For this experiment the subjects all reported that the sound image was in the space, and not in the head as previous works have reported. In analyzing their results, especially for sound sources directly in front of the subject, More and Fry make the following statement: “As neither a difference in intensity nor in phase exists, we were led to believe that a complex tone produces a different sensation according to the direction the sound waves entered the outer ear. Certainly the shape of the outer ear is well adapted to modify the short waves for the higher harmonics, and thus to change the character of complex tones.” [29][p. 454]

In 1907 Rayleigh, after performing additional tuning fork experiments, came to the decided conclusion that phase can not be ignored as a factor in localization and hearing. [40]

It seems no longer possible to hold that the vibratory character of sound terminates at the outer ends of the nerves along which the communication with the brain is established. On the contrary, the processes in the nerve must themselves be vibratory, not of course in the gross mechanical sense, but with preservation of the period and retaining the characteristic of phase- a view advocated by Rutherford, in opposition to Helmholtz, as long ago as 1886. [40][p. 223]

Following this, Rayleigh experimented with front/back confusions. [41] Using a long corridor, speech sources, and blindfolded subjects, results showed that front/back confusions existed, and were more evident the lower the tone of the voice. After moving the experiment to an open field the same results were observed. Another interesting observation, while standing in front of the listener, but facing away, was that the subject placed the image behind them.

A somewhat remarkable study was performed by More in 1909 using normal hearing and monaural subjects. [30] Using normal listeners, a tuning fork, and tube apparatus, it was observed that localization by phase (in the tuning fork setup) became difficult above 1 kHz. When localization tests were performed on monaural subjects with the same apparatus, no localization ability was found. Going to an open field, monaural subjects were tested using a tuning fork (256 Hz) with a resonator box. Localization ability was found to be good on the functioning ear side and greatly reduced on the opposite side. In addition, when the fork was struck and the resonator box was slowly covered the image moved from the good ear side towards the deaf ear. Several conclusions were made from these results. One conclusion was that bone conduction could not be a significant factor in the localization process. Bone conduction paths from a normal source would travel many structural acoustic paths to the inner ear. This created a collage of perceived sounds, making the localization of sounds a very complicated matter. Secondly, it was concluded that localization ability is based primarily on loudness evaluation between the ears. “The use of two ears is probably necessary for any but the most rudimentary determination of directions of sounds”. [30][p. 316] More’s conclusions are also based on a theory which

suggests it is easier to localize pure tones than complex tones. What is remarkable about this work was that it ignored almost all of the previous experiments, observations, and conclusions regarding localization.

Physiological attempts at describing localization ability were presented in a thorough manner in 1926. [4] In this theory which considered hair cell motion as a physical explanation for the method of phase detection, experiments were performed determining maximum phase difference perceptions between the ears. The time delay necessary for a subject to hear two distinct sources was found to be 1.7–2.6 ms. In this time, Banister states, sound will travel about 57 cm, which is the average distance from ear to ear around the head. In comparing to measurements made during this current work, a distance of 38 cm was typical, comparable to a time delay of 1.1 ms. In addition to this, Banister observed that for a continuous sound, the pitch varied with location.

In 1928 an in-depth experiment was performed where the ear input signals were switched, left \rightleftharpoons right. [61] This was accomplished using the “pseudophone”, which was comprised of two funnels (typically used by the hard of hearing) attached to tubes which fed the signal to the ear. The funnels were placed so that sound from the right was fed to the left ear, and vice versa. The actual experiment was performed by a single subject (the author) who wore the device for about 21 days in total, keeping a journal of observations. The results were interesting. The sound image would sometimes be transposed front/back or left/right, but not always. When the source was in view, the sound image would relocate to the correct real source position. In some instances, some sources would be transposed and some not in the same setting. If the subject closed his eyes, non-transposed images tended to flip left/right. Three days of the experiment consisted of constant wearing of the device and the subject walking around town. When the subject slept, the ear canals were blocked. As more confusions occurred, it was noticed that vision became the sole source of object localization, with no error or enforcement by sound. But when the object was not in view, localization was still reversed. With the eyes closed, head motion was able to be used to correctly localize a source while wearing the device. When the device was removed at the end of the three days, normal localization immediately returned. One of the major conclusions of this investigation was that localization involves neuromuscular adjustments with respect to the source. “The physiological conditions of localization are to be sought in pattern changes of muscular tonus”. [61][p.429] This agreed with the earlier works of Münsterberg and Pierce.

In an attempt to further explain how localization was possible, and the role of head motion, Young [62] performed another experiment. The funnels (or “remote pinna”) were placed in another room, with the signals fed (not transposed) to a subject in a separate room. The funnels were spaced apart at an appropriate distance in free space (*i.e.* no ‘head’). A source was attached to a rotating ring so that it could be directed from any angular location. The results showed no left/right confusion. All images, though, were placed outside the field of vision and in the horizontal plane (*i.e.* no elevation). This lack of elevation was not seen as a problem, but as a way to simplify the experiment to 2D. Variations in the direction the funnels faced was examined in an attempt to force the image to the front with no success. Different funnel sizes and shapes (more symmetric) were also used with no real change in results. Variations in the funnel separation (increased by

3 cm) restricted the image to a few degrees left and right of the median plane. Young's conclusion was that "when head movements are rendered ineffective to change the binaural stimulus pattern by the method of detaching the 'pinna' acoustically from the head, there is no physical basis for objectively accurate discriminations of up, down, front, back, and intermediate angular directions". [62][p. 109] Monaural localization was also tested by closing one tube and only using one funnel. No localization was possible. The conclusion was that head movements are here shown to be necessary for localization. Young refutes previous experiments showing monaural localization ability, specifically the work of Angell and Fite, by saying that head motion was not restricted enough. Angell and Fite used a head rest and instructed the subjects not to move their head. The research and results of this work are definitely useful in examining how adept the localization process is for people, but the validity of the conclusions drawn is rather poor in that many assumptions regarding the acoustic effect of the pinna, head, and body (already discussed by several researchers 20–30 years prior) was ignored.

Up to this point very little work had been done regarding the frequency limitations of localization. It was generally understood from experiments typically using tuning forks with frequencies lower than 1 kHz, that pure tones were difficult to localize. Complex sounds were tested using clicks or speech, with good localization results, but the frequency content was not analyzed. The extent of frequency analysis was that localization ability was generally understood to rapidly deteriorate around 1 kHz, and be very difficult by 3 kHz. It is at this frequency that phase discrimination becomes impossible for the subjects. [4] In 1936, Stevens and Neuman tested this theory. [49] Using primarily pure tones, but also some clicks and hisses, all presented over a loudspeaker, localization was tested beyond this limit. It was found that, as previously determined, localization deteriorated from 1–3 kHz, but then slowly returned at 4 kHz and was fully accurate at 10 kHz. Their only reasoning as to why this had never been found before was that no one ever tried to measure localization ability beyond the 3 kHz range. In measuring localization errors, they did make one interesting assumption. For any front/back reversal, they re-flipped the answer before computing the error. This means that, if the source was at $+20^\circ$, and the image was at $+160^\circ$, the error would be $+10^\circ$.[†] These reversals were found to be much more prevalent at frequencies below 3 kHz, at a rate that would be comparable to chance. Their conclusion was to divide localization ability into two regions, low frequency and high frequency. Low frequency localization was attributed to phase differences between the ear signals. High frequency localization was attributed to differences in signal level intensity.

The contribution of subject motion to localization was examined in great detail by Wallach. [56] Two very interesting experiments were performed. The premise of both experiments was the necessity of head movement for the adequate discrimination of elevation in localization, as postulated by Young. [61] The first experiment was rather complicated in design. An array of loudspeakers was arranged in an arc, at 3° intervals, in front of the subject. A rotary switch, also with 3° spacing, was attached to the head so that when the

[†] This practice is commonplace among many current localization researchers. Front/back confusions are considered as a separate type of error, apart from localization accuracy. This results in some difficulty when comparing various researcher's data, as not all authors report whether or not this assumption is made, and some authors do not report the amount of front/back confusions which existed in the data at all.

head moved (horizontal rotation) the contact of the switch moved. The contacts of the switch were wired to certain speakers on the arc. The subject, seated in a swivel chair, was rotated back and forth by an assistant. Several arc spacing and wiring configurations were calculated for certain image positions, thereby giving different audio cues than the rotation cue. The first trial used an arc spacing and wiring scheme such that for every 3° of subject rotation, the active source moved by 6° . This was intended to place the image behind the subject. For all five subjects the results were positive. The second trial placed the speakers at $4^\circ 30'$ separation for the 3° of actual head movement. This was designed to place the image at an elevation of 60° . These results for this configuration were also very good. Comparisons between blindfolded and not blindfolded subjects showed better results when the subjects could see. The third trial used a speaker separation of 3° . This should result in the image being overhead, as it is the only source position for which rotation has no effect on the perceived sound. Results for this trial were reasonable when blindfolded, and very good when vision was not obscured.

The second experiment focused on the subject's impression of their own motion. The subject was seated in a fixed position, with a chin rest, inside a cylinder which rotated about them at a rate of about 0.15 Hz. Regular markings on the cylinder and chair gave reference points for the vision of the subject. After some time in the cylinder, the subject would have the impression of self-rotation in the direction opposite to the actual rotation direction of the cylinder. A speaker was placed on a boom outside the cylinder. A variety of trials were performed to repeat the first experiment in principle. The first trial placed the physical source directly in front of the subject. Instead of the image being directly in front of the subject, the subjects placed the image directly overhead. The next trial placed the physical source at 90° (this source position can be completely determined by phase differences). The resulting image was correct. For the subject it seemed to rotate in unison. The final trial rotates the source about the subject in an attempt to reproduce the state where the image appears in the rear. Using the subject's impression of motion, the desired source position must be rotated about the cylinder in the opposite direction, but at the same speed. If the cylinder is rotating to the right, the subject has the impression of rotating to the left. With the source also rotating to the left, the desired image should appear directly behind the subject. The desired result occurred with none of the subjects. If the subject localized the source from its actual physical position it would seem to rotate about the subject in the same direction as he was, twice as fast as he was. The image actually appeared to switch from directly in front to directly behind depending on the source position. For some subjects the image was at rest and they were rotating about it, others were uncertain. For other subjects the image appeared to move slowly in the direction opposite to the actual direction of source rotation.

Clark and Graybiel performed a basic repeat of the Münsterberg and Pierce rotation experiment. [13] Subjects were placed on a rotating chair. The subject was spun, then abruptly stopped. The image was found to be displaced in the direction of rotation. The displacement had a typical value of 17° at the outset. The amount of deflection decreased with time after the subject was stopped, taking about 30 sec. to return to normal position.

In addition to reviewing research performed, it is also informative to review what was being presented in text books and taught to young acousticians. The role of the pinna in

localization theory is an important factor which has changed through time. The following is a passage from a textbook printed in 1940 concerning the outer ear and its role.

The pinna now plays little, if any, part in the process of hearing. In animals it is movable and can be used to collect sound and to locate the direction of its source. Man has now almost completely lost this capacity for moving the pinna. It may play some slight part in shielding the ear from sounds from behind, thus making easier the discrimination between sounds coming from behind and sounds coming from in front, and the more obtrusive the ears are the better will they serve this purpose, but in actual practice the discrimination is nearly always made by turning the head so as to present one ear to the source of sound and leave the other shielded by the head. [60][p. 457]

Contemporary to this text is the fundamental acoustics textbook of Kinsler & Frey. In the first edition of Kinsler & Frey (1950) the function of the pinna is described as follows:

The pinna of the outer ear serves as a horn to receive acoustic energy and lead it into the auditory canal. In a human being the pinna is a relatively ineffective device, and from an acoustical point of view is almost useless. In some animals, however, it supplies an appreciable gain, particularly over certain frequency ranges. [26][p. 362]

Finally, one of the earlier texts on audiology (1958) described the function on the pinna and outer ear as such:

Of all the parts of the ear the pinna is the most prominent and the least useful. It serves the purpose of directing sound waves into the external meatus in a more concentrated fashion than would otherwise be possible. The function of the pinna in relation to the external canal can be likened to that of cupping the hand behind the ear in a difficult listening situation... In modern society, the pinna is almost purely ornamental, inappropriate though the word may seem, applied to something as homely as the ear! [33][p. 18]

A review of the literature up to this point shows several interesting points. Although a good deal of work, both acoustically and psychologically based, had been done regarding human aural localization this research was either not accepted or not viewed by other researchers and scientists who were interested in the same field. From the beginning statements are made concerning the effect of the pinna (as acoustically important in shape and complexity), the need for complex sources in order to localize, the changing of tone with respect to source position, and even the possible effects of hair and clothing on one's ability to localize a sound source in space. But, as is evident, these earlier works were not considered by many, and certainly not by the academic instructional community. An explanation may lie in the fact that research in this area was (and still is) being done by researchers in a wide range of fields, from physical acoustics to psychology to biology. These fields do not typically share the same publications and therefore some degree of lack of communication can be expected.

2.1.3. Historical Review – 1960 to 1973

Up to this point in history the majority of knowledge and research had been based primarily on subjective testing. In 1967 “The Role of the Pinnae in Human Localization” was published by Batteau, which proposed that the sounds entering the ear canal were different for different incident directions was published. [5] This research differs greatly from previous works in that instead of subjective impressions, structural anatomy was the primary focus for determining the method of localization. This work ushered in a new direction in localization research whereby acoustics, and to a lesser extent psychology, was employed as the research tool. In addition, the pinna was becoming the accepted means by which localization was possible.

Batteau’s work described how the various surfaces of the pinna would provide the listener with organized reflections, delayed 20–300 μ s from the incident signal. The pattern of these reflections would depend upon the incident direction of the sound source. These ordered reflections would be decoded by the brain, resulting in the perception of the source location. At the time of this work, it was not possible to experimentally measure these predicted reflections.

Soon after this work the pinna was investigated as a truly acoustic system. In 1968 Shaw and Teranshi mapped the resonance patterns of a pinna replica (with ear canal). The results were well defined modal shapes at approximately 3, 5, 9, 11, and 13 kHz. [45] This showed for certain that the pinna was of a sufficient size and complexity that it could contribute to and affect sound waves.

Around the same time subjective testing was being used to investigate the spectral factors in localization. Two research groups used arrays of loudspeakers (vertical hidden array in front of the subject [42](see Fig. 2.1) and hemispherical array over the subject [7](see Fig. 2.2)) and presented various band limited signals. It was determined that certain frequency bands, when dominant, create a sound image at different locations. Using the vertical array the image moved from +20° elevation to –10° with the frequency (source was pulsed sinusoids) decreasing from 7.2 kHz to 0.25 kHz. With the hemispherical array (all sources presenting the same signal) a random noise signal filtered with peaks remaining at 250–500 Hz and 2–4 kHz appeared to be located towards the front. A primary peak in the region of 8 kHz resulted in an image overhead, and peaks at 1 and 12 kHz resulted in a sound source image to the rear. These results show definitive evidence that frequency content of the perceived signal affects the location of the sound source image. There was a great deal of work being done in this area at the time, too much to attempt to describe here. Without dealing with the actual results, the above presents an idea of the current state of knowledge.

In 1969 one of the first dummy head recordings was made. [14] Localization results using the mannequin recording were poor, especially in terms of elevation discrimination. As a final milestone in this era of pinna research was a study by Gardner and Gardner. [18] This study examined the effect of progressively filling cavities of the pinna on the ability to locate sources in the median plane. It was observed that as the pinna progressed more towards a flat surface (finally being covered with a block with a carved ear canal), the ability to judge elevation deteriorated rapidly. This work showed for certain that the




Figure 2.1: Perceived frontal elevation of band limited noise. Source [42]

contours and intricate shape of the pinna is essential to sophisticated source localization (*i.e.* more than left/right).

Academic literature in hearing science was still not current with the scientific knowledge which had been gained over the almost 90 years of investigation. One text described the function of the outer ear only in the following context;


In lower animals, the auricle [external ear] is able to be directed for better collection of sound. However, in man this ability is of little purpose. The hairs at the entrance of the ear canal [function to] trap dust, insects, and other small particles. [19][p.18]

The lack of information in the academic community regarding the documented subjective knowledge, and recent experimental knowledge, of the role of the pinna and frequency filtering in human localization is somewhat remarkable.

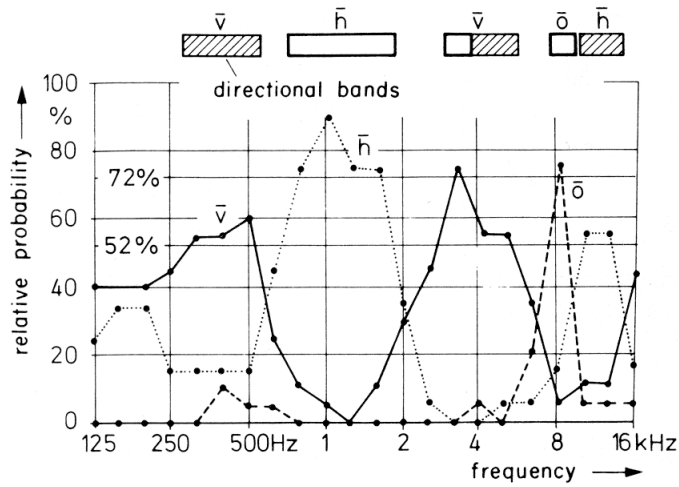
2.1.4. Current Literature

Over the past 20 years there has been an explosion of research in the area of spatial hearing and localization. This has been due to advances in signal processing and the miniaturization of sensors. Further work has been directed towards the creation of virtual reality environments, where it is desired to recreate the three-dimensional experience of a listening environment. The continuing increase in computational ability of computers and digital signal processing hardware makes this goal more feasible every day. Due to the vast amount of work done, only a few key works will be mentioned which describe the development of the knowledge. This is by no means a comprehensive or exhaustive summary. For a detailed review of the literature see references [6] and [7].

In 1974, Hebrank and Wright considered the experimental evidence of frequency dependence on source localization and combined this with the theory put forth by Batteau. [22] The result was a conversion of Batteau's ordered reflection from the pinna contours in the time domain to spectral filtering in the frequency domain. It was possible to coincide, to a rough extent, with predicted frequency nulls in a measured pinna response.



a. Apparatus to determine elevation of band limited noise. Indicated are the locations considered front, v , overhead, o , and rear, h in the subjective test results given below.



b. Results of hemispherical array for perceived elevation using location notation from (a).

Figure 2.2: Experiment and results showing the perceived elevation of band limited noise. Source [7]

Using miniature microphones, Searle *et al.* made individualized recordings of sources in the median plane within the ear canals of several subjects. [43] These recordings were played back to the original subjects, and it was found that individuals who listened to their own recordings were able to determine the position of the original sources. Following this Butler and Belendiuk repeated their experiment with the additional listening experiment of mismatching original recording subject and playback subject. [9] When listening to recordings made from other individuals the localization of the recorded sources was greatly deteriorated. This showed that the pinna effects are individual in nature. It was also reported that certain recordings were good for many subjects while other were poor for many subjects. “Hence, individual differences in...localization may, in large part, be due to differences among pinnae configurations of people.” [9]

By the late 1980’s technology made it possible to measure individual HRTFs and incorporate them into complex virtual acoustic environments. In 1989 Wightman and Kistler performed a similar experiment to Butler’s using individually measured HRTFs. Sources from various directions were simulated to individuals using their own HRTFs and others. Experiments involved the measurement of localization ability with individual HRTFs and the effects of using non-individual or other’s individual HRTFs in the processing. The results showed that when using an individual’s own HRTF the localization errors were small. When using another’s HRTF the errors increased, with the primary error being the reduction in elevation information. In addition, the HRTF of a subject with poor natural localization ability imparted the same poor ability onto those who used their HRTF. This showed that the information to localize is contained in the HRTF, and, more interestingly, that there are individuals who, given their own natural HRTF, are poor at localizing and that this can be attributed to their HRTF. This implies that there are certain head and pinna configurations which do not provide the cues necessary for good localization. When the subject with poor localization ability used HRTFs from other individuals their localization ability never improved. These studies show that the HRTF is based on physical geometry, describing the spectral changes to the incident sound as a function of incident angle which are then used by the brain to locate sound sources (see section 2.1.5 below).

As with the historical periods mentioned previously, the academic publications used to educate young acousticians is lacking in current knowledge. For the final reference on this topic the 1982 edition of Kinsler & Frey’s *Fundamentals of Acoustics, 3rd Ed.* is presented.

The *pinna* of the outer ear serves as a horn collecting sound into the *auditory canal*. In a human the pinna is a relatively ineffective device, but in some animals it supplies an appreciable gain over certain frequency ranges. [27][p. 257]

The only variation from the 1950 text is the omission of the phrase “from an acoustical point of view is almost useless” used in describing the human pinna. This is not the only text in use, but being one of the foundation texts for current acoustic education it serves to show that not all of the academic and research community understands the complexity of human localization and how it is a function of the anatomy of the outer, as well as the inner, ear. The possible relations of the outer to inner ear in function, and their effects on human localization is discussed in the next section.

2.1.5. Anatomical and Neurological Processes in Localization

The ability of the human to localize sound is a complex matter, requiring a great deal of signal processing and coding throughout the auditory pathway and nervous system.[†] Initially, sound signals are frequency coded by the filtering properties of the head and pinna. This filtering imparts specific frequency shaping, dependent on the location of the source with respect to the listener. In addition, the physical spacing between the ears provides different inputs to each ear, including time delays, which can give some information regarding the position of a source. Once the sound approaches the ear canal (within a few mm) the coding for spatial information is complete. [21] Propagation through the ear canal provides for additional filtering (a function of the length and diameter of the ear canal) but this does not provide any directional information.

Before discussing the physiology of localization, it is important to review a few points regarding localization ability. Much historical work has been done to determine the limitations of localization, external effects on localization, and some of the important auditory cues in localization. For sources which consist of low frequency components only, below about 1 kHz, localization is based upon the time-of-arrival or phase difference between the two ear input signals. This difference provides for general location cues (left/right) but it is limited to a spatial cone whose axis is the interaural axis. Sources whose components contain high frequency contributions, above 3 kHz, are localized much more precisely using the frequency information obtained through the filtering properties of the head and pinna. [7,49] This spatial filter is called the Head-Related Transfer Function (HRTF), and it contains the spatial filtering information for any source point (usually on a sphere) to the ear canal. Sources whose frequency components are between 1–3 kHz are difficult to localize. These results were all obtained without any knowledge or consideration of the internal workings of the auditory system.

The ear canal and ear drum, though not providing any directional signal processing, do effect the localization functionality of the auditory system. The ear canal, which behaves basically as a duct from the pinna to the ear drum, can have acoustic modes. The first resonance of the ear canal is around 3.5 kHz. The first resonance of the ear drum is approximately 2 kHz. Once the ear drum vibrates at frequencies above resonance, the exact input to the middle ear, especially with regards to phase, can not be straightforward or ideal. It is also interesting to note that the path length from ear to ear is on the order of the wavelength for a 1.2 kHz acoustic wave. All three of these frequencies are within the poor-localization region. With multiple wavelength propagation between the ears, the phase difference does not provide reliable information. Ear drum resonance definitely removes the ability to discern usable phase information, and ear canal resonances provides for poor transmission to the ear drum.

Transmission of the acoustic signal from the ear drum to the cochlea fundamentally consists of an impedance transformation from the outer to inner ear. This impedance transform is not linear, since it provides a higher impedance at lower frequencies, thereby

[†] The majority of information regarding this subject was derived or extrapolated from the reference by Durrant and Lovrinic [15] which discussed the anatomy and physiology of hearing, but does not contain any reference to localization. Additional specific information was obtained from the cited references.

reducing the perceived level of the lower frequencies. This is accomplished through the active components of the middle ear which respond to differing pressures with differing tension. With the important information for localization being in the higher frequency region, it can be postulated that, in addition to protecting the inner ear from damage due to high energy low pressure waves, the reduction of low frequency information allows for the acquisition of more high frequency information. This permits more detailed processing of the high frequencies, useful in localization and speech processing.

Once inside the cochlea, frequency decomposition of the signal occurs. This is accomplished through the cochlea, where different hair cells respond to different frequencies of excitation. The details of this process are not crucial to the current discussion, but there is an important constraint within the system which does relate to localization. The hair cells function by basically converting motion to electrical activity. Motion of the hairs stimulate the associated auditory neuron to transmit a signal. After the transmission of the signal, the neuron needs time to repolarize in order to have enough voltage to transmit another signal, defined as the refractory period. This action takes time, on the order of 0.001 sec. The result of this refractory period is a maximum spike rate for the hair cell neurons of 1 kHz. For signals with frequencies above 1 kHz, the neurons will not be able to fire in phase with the excitation signal. Given a signal of 5 kHz presented to both ears with a phase difference, the neural activity associated with, for example, the peaks of the signal, will not coincide with the actual peaks because the resolution is not enough. The result of this is that the neural signals from each ear can not be compared reliably to obtain a time difference. Again, there is a frequency limitation of phase discrimination in the region of 1 kHz.

Once the auditory nerve signal is transmitted, it is fed into the cochlear nuclei, which constitutes the beginning of the central auditory pathway. This is done along the auditory nerve, or VIIIth cranial nerve, transmitting all these signals into the Central Nervous System (CNS). These nuclei provide additional neural impulses based on properties of the auditory nerve signal. These new signals correspond to such factors as excitation duration, signal onset, and signal termination. Further processing of the auditory signal will utilize these complex signal components.

The typical structure of the brain is such that information from the left side of the body is processed on the right side, and vice versa. This is true for the majority of the cochlear nuclei, but not all. This leads to the concept that there are neural processes which make use of the signals from both ears. This is indeed the case. There are auditory nerves, at and above this crossover point, which show responses dependent on combined information. The auditory cortex (for which there is one for each side) holds these complicated processing nerves and is the terminating point of the auditory pathway. One group of neurons responds to time delays or phase differences between the signals. The neural signal changes in amplitude according to the time delay. There are also neurons which respond to intensity level differences between the auditory signals. This processing is done by a special type of neuron, excitory-inhibitory. These neurons basically react to the difference between input signals. The output of the neuron is a function of the excitation due to one ear's input and inhibited by the other ear's input. Therefore, a neutral output, where all neurons have a constant base level of output, would indicate equal signals (at least for

the types of frequency dependent properties delivered along the two specific neural inputs) at each ear. It is these cross-correlation type signals which provide a great deal of information for determining the location of sound sources. In addition, there are connections between the temporal lobes which run through the cerebellum which provide higher order interactions and processing of binaural signals. Research into the auditory cortex (in animals) has shown that it is necessary for certain auditory processing functions, but not all. Without the auditory cortex it is still possible to recognize the onset of a sound, changes of intensity and frequency, and changes in location. Use of the auditory cortex allows for more integration and interpretation in processing the signals. Removal of the auditory cortex though, inhibits that ability, leading to the inability to discriminate tonal patterns, sound duration, and localize sounds in space. Therefore, it is concluded that here, in the auditory cortex of the temporal lobe of the brain, a significant part of the localization of sound sources takes place. Other parts of the temporal lobe, such as secondary auditory cortices may also contribute to localization, especially due to the large amount of data being processed.

The localization of sounds is a process by which the frequency filtering cues of the head and pinna are mapped, in the brain, to relative coordinates in space. This is a complicated cortical process. Researchers have suggested that this mapping response is learned by associating auditory cues with visual perceptions of source locations. Over time, this mapping is learned and the location of sources from auditory cues alone is possible. If the listener's sense of their own spatial orientation is disrupted (*i.e.* through abrupt changes in rotation causing the sensation of spinning or dizziness), then the location of sound sources seems to incorporate the perceived orientation. This has been shown in great detail by a number of researchers. [56] This indicates that the localization of sources is not simply an auditory process, but includes higher order brain functions which combine learned responses, complex pattern matching, and cross referencing with other senses in the brain. The result is a unified (though not always correct) perception of the location of a sound source.

Comparisons of localization accuracy between humans and animals has shown that humans have the most acute localization ability. Considering animals that are capable of moving their pinna, if this ability is restricted the result is a reduction in their localization ability. Humans do not need to move their pinna or head to localize a sound source. The complex shape of the pinna provides detailed frequency filtering with respect to source direction, and the complexity of the human brain provides a very sophisticated acoustic analyzer with a large data base for pattern matching. This enables the almost immediate determination of the location of a sound source without the need to move. The entire auditory system appears to be dealing with the apparent physical limitations of the parts, from the size of the head to the response time of neurons. All of these physical, anatomical, and physiological parameters are represented in localization functionality and seem to be related. Whether the neural response time is a function of the diameter of the head, or vice versa, is another matter. What is interesting is how all the parts seem to be coadapted to ignore faulty signal information and provide accurate (on the order of a few degrees) resolution for spatial localization.

2.1.6. Anatomical and Neurological Processes of Localization for a Monaural Listener

The previous section discussed the aspects of binaural localization throughout the human anatomy. In the event of unilateral hearing loss, some of these functions are no longer present. Given the inner ear being the damaged part, it is assumed that the auditory nerve, auditory cortex and brain functions are intact.

With the neural connections still operating fully, at least on one side, there still exists all the frequency filtering cues and processing ability which exists for binaural listeners. This information provides the majority of localization information for high frequency sources. Missing information are phase differences (useful only in localizing low frequency sounds) and the second ear input (which provides additional cues to reinforce the pattern matching processing for determining the location of a source). Having two ears allows for a type of independent confirmation of source localization, as each ear has separate signals which correspond to a source location. If both cues indicate the same position, the confidence in that position would be greater. Therefore it should still be possible, as has been shown in numerous studies, that someone with a unilateral hearing loss can localize. [1,2,7]

If the case of median plane localization is examined, many of the cues typically absent from an individual with unilateral hearing loss signals are also missing for a binaural listener. There are no phase difference signals, as the source is equidistant from both ears. Additionally, the pinna filtering cues should also be the same, assuming mirror image pinna. In this way, no additional information is gained through cross-correlation processing, other than that the signals are the same. Even so, median plane localization is possible, and is also a standard test when analyzing localization ability. The non-symmetry of an individual's pinna and the effects on localization is a current topic of debate, and may provide clues as to how median plan localization is better in some subjects than others. But, for listeners with unilateral hearing loss, localization using only frequency filtering cues is possible. Variations in this type of listener's ability depends a great deal on when the hearing loss occurred. As localization of this kind is fundamentally a learned process, the greater the period of the hearing loss the more time exists over which the skill can be acquired, resulting in improved localization ability. This has also been shown in research studies. [1,2] The longer the listener has to develop the spatial map for auditory signals, the more adept at localization they will become. Many studies which refute the localization ability of subjects with such a hearing loss do not reference this fact. In addition, many localization tests have been performed with tuning forks, which generate pure tones. This type of signal produces no frequency cues, only phase differences and intensity differences between the two ears. For a subject with unilateral hearing loss this is an impossible task, which led many researchers to believe that binaural hearing was required for localization. Finally, the location of the damage may be crucial to localization ability. If, instead of in the inner ear, damage exists in the auditory nerve or auditory cortex, the processing ability necessary to form and use a frequency spatial map may be damaged. If this were the case, the subject would have very poor localization ability, regardless of any hearing loss. As the method of hearing loss is almost never mentioned in the studies, it is difficult to determine the reasoning behind some listener's inability to localize with monaural hearing loss.

2.2. Absorption and Impedance Measurements

There have been relatively few research efforts investigating and quantifying the acoustical properties of skin and hair over the audible frequency range. The majority of work regarding skin tissue measurements has been in the ultrasonic range. A number of mechanical studies have been done by which the mechanical impedances of various bone structures have been tested using impact hammers, but this work is not directly applicable to the research here. [17]

The primary work in the area of acoustical measurements of skin and hair impedances was done in the early 50's. Skin measurements were performed on the upper arm and thigh of human subjects using a plane wave impedance tube and plane wave excited rod. [20] The experimental design analyzed the change in the fundamental mode resonance frequency due to the change in surface impedance. The plane wave tube was used for low frequencies (100–1500 Hz) while the rod was used, using the same technique, for higher frequencies (1.5–20 kHz). Results showed absorption coefficients ranging from 0.04–0.004. These results were on the same order as the mechanical impedance results determined in [17].

Measurement of skin and hair experiment utilized haired and hairless mice. [16,55] Two techniques were used. This first was to place the test animal very near to a high powered source (a 162 dB siren was used). The amount of sound absorption was calculated by measuring the increase in body temperature of the animal. This technique was used over a frequency range of 6–22 kHz. This measurement lasted for about an hour, by which time the animal died (another part of the experiment). From these measurements, the hairless mice had an absorption coefficient, α , of 0.01–0.13 over the frequency range where the amount of absorption increased with frequency. In contrast, the haired mice have values ranging from $\alpha = 0.5 - 1.1$ over the same frequencies. A second method utilized a reverberant chamber in which the variations in the reverberation time were calculated due to the presence of the mice. Hairless mice has values of $\alpha = 0.04 - 0.1$ while haired mice had $\alpha = 0.4 - 0.8$.

2.3. HRTF Calculations

The amount of research into computational calculations of the sound field around the human head is very limited. A substantial paper was written by Weinrich in 1984. [57] The goal of this work was directed towards a better understanding of the near head sound field for the purpose of hearing aid design and evaluation. The work divided the calculation, and geometry, into several parts. As an introductory experiment in using computational methods for this purpose many assumptions were made. The ear canal was modelled as a collection of cylinders of varying diameter. The response of the canal was calculated using transmission line theory. A very rough two-dimensional version of the pinna, only the perimeter was modelled using four lines, was created and the response as a function of elevation was calculated using a finite difference time domain approach. The results replicated some of the very basic HRTF effects, primarily the shifting of the primary high frequency notch (around 10 kHz) with varying elevation. The head was modelled using a surface mesh and the field was calculated using a boundary element method solution. The response was calculated for the horizontal plane. In designing the mesh, the pinna were

not included, resulting in a very coarse mesh with smooth sides. The mesh contained 212 elements, with the maximum spacing between nodes being 5 cm. At four nodes/wavelength, the absolute minimum, this yields a mesh which is usable up to 1.7 kHz. The results, compared to measurements made on a physical replica of the BEM model mesh, were good below this frequency. As the model did not contain the pinna, and was very coarse, only very general comparisons were possible to real HRTF measurements. Even so, this work showed that numerical calculations of the human listening response were possible. With the increase in computing power since 1984, the size and complexity of the calculations can be greatly increased, and therefore result in more sophisticated models which may better represent a real individual HRTF.

An additional work in 1988 approximated the head using a sphere or oblate spheroid and calculated the variations between the two using an analytical solution. [25] The result showed little variation between the sphere and spheroids, on the order of 0.5 dB. Comparisons to real HRTFs were not made in this work.

Chapter 3.

Acoustical Parameter Measurements

The primary goal of this research is an increased understanding of the Head-Related Transfer Function. Determination of the contributing factors towards the HRTF, in addition to examining geometrical parameters, includes the investigation of the acoustical parameters of the structural system. These parameters include the behavior of skin and hair in the audible frequency range. As a sound wave travels from the source to each ear, the complex impedance of the skin and hair modify the spectral content of the sound before it enters the ear canal. In order to investigate these effects, acoustical reflection coefficient and impedance measurements of skin and hair were undertaken over a section of the audible frequency range.

3.1. Measurement Technique

The measurement method used to determine the reflection and impedance values of the various surfaces utilized a two microphone impedance tube technique. [8,10,11,12,44] This system provides for the measurement of the reflection coefficient and complex impedance of a material sample. Several assumptions and limitations exist regarding this method and shall be explored. These conditions include the assumption of a planar surface, normal incidence, and the accurate estimation of the distance from the microphones to the sample. The justification for normal incidence measurements instead of grazing incidence measurements is presented in section 5.3.2.

This method utilizes the measured transfer function (as a function of frequency), H_{12} , between two microphones to separate the incident (p_i) and reflected (p_r) pressure waves. From this, it is possible to calculate the complex reflection coefficient R at the measured surface. The complex impedance Z can be derived from this calculation. A plane wave tube was utilized in this research, configured as shown in Fig. 3.1.[†] The theory behind this measurement can be described in the following manner.

The measured transfer function is defined by

$$H_{12} = \frac{P_2}{P_1} = \frac{P_{2i} + P_{2r}}{P_{1i} + P_{1r}} \quad (3.1)$$

where P is the Fourier transform of $p(t)$, and P_{ni} and P_{nr} are the pressures due to the incident and reflected waves at sensor n . It is then possible to define the following

$$H_{12i} = \frac{P_{2i}}{P_{1i}} \quad H_{12r} = \frac{P_{2r}}{P_{1r}} \quad R_1 = \frac{P_{1r}}{P_{1i}} \quad R_2 = \frac{P_{2r}}{P_{2i}} \quad (3.2)$$

[†] Different researchers define these parameters in slightly different ways, along with the definition of $\pm x$ directions.




Figure 3.1: Impedance tube setup. Define s as the spacing between the sensors and l as the distance from the second sensor to the impedance surface.

After a little algebra the reflection coefficient can be shown to be

$$R_1 = \left[\frac{H_{12} - H_{12i}}{H_{12r} - H_{12}} \right] \quad (3.3)$$

Since the measurement is in a plane wave tube, below the first cross mode frequency the transfer functions H_{12i} and H_{12r} are simply e^{-jks} and e^{+jks} respectively. In addition, translation of the calculated reflection coefficient to the surface $x = 0$ is only a multiplier of $e^{j2k(l+s)}$. When combined the result is

$$R = e^{j2k(l+s)} R_1 = -e^{j2k(l+s)} \left[\frac{H_{12} - e^{-jks}}{H_{12} - e^{-jks}} \right] \quad (3.4)$$

The absorption coefficient is calculated simply as $\alpha = 1 - |R|^2$. The surface impedance of the material is finally determined using Eq. (3.5) where Z_s is the acoustic impedance of the surface and Z is the corresponding specific acoustic impedance.

$$Z = \frac{Z_s}{\rho_0 c} = \frac{1 + R}{1 - R} \quad (3.5)$$

Using these definitions, the experimental measurements of the acoustical properties of skin and hair were conducted.

3.2. Apparatus

The parameters of the system, tube diameter and sensor locations, were designed to obtain optimum results over the frequency range of interest. For use in the computational calculation of the HRTF, the frequency range of interest was determined to be 1–6 kHz. The center frequency of this range was 3.5 kHz. The justification for this range is presented in section 5.2.

The measurement method requires that only plane wave propagation occurs in the tube. For this reason, the upper frequency range is limited by the first cross mode of the tube. The ASTM standard [3] requires the diameter to be less than $0.586c/f_{max}$ which, for an upper limit of 6 kHz, gives a maximum tube diameter of 3.35 cm. A 1" PVC tube was chosen. The internal diameter of the tube was 2.60 cm, with a wall thickness of 0.354 cm. The tube length is specified to be long enough to ensure that only plane wave propagation is present and to provide enough signal level. A 1 m long tube was used. It is desired to have the microphones as close as possible to the termination to obtain the greatest signal to noise ratio. The distance to the termination is limited by the presence of evanescent waves coming from the impedance termination, which is not necessarily perfectly planar. The standard suggests placing the near sensor no closer than one tube diameter away. A distance of two internal tube diameters was chosen here for the distance between $Ch2$ and the termination, resulting in a desired distance of $l = 5.2$ cm. The sensor separation distance is defined in the standard such that $s \ll c/2f_{max}$, resulting in a maximum spacing of 2.86 cm. The optimal spacing, resulting in the minimal amount of error, is defined by choosing $s = c/4f_{center}$. [8] Using 3.5 kHz as the center frequency of interest results in a desired distance of $s = 2.45$ cm.

The system utilized three Sennheiser microphone capsules, a 2.5 cm (1 in) diameter speaker, electronic preamplifiers, and a multichannel Hewlett Packard spectrum analyzer. The sensor microphones used have an exposed diaphragm diameter of approximately 1 mm. The sensors are flush mounted in the tube, using a clay seal on the outside to ensure against leakage. Several terminations were constructed to test and calibrate the system, in addition to the experimental terminations. A diagram of the apparatus and the setup for the various experimental measurements is shown in Fig. 3.2.

3.2.1. Microphone Calibration

The two microphones used in the impedance tube were calibrated relative to each other using a switching technique. For this method, the transfer function was measured for the microphones in the measurement position, then the sensor positions were switched and the transfer function was measured. This was used to ensure that any magnitude and phase variations used in the measurement were a function of the sound field, and not due to differences between the sensors. The geometric mean of the transfer functions was calculated as per Eq. (3.6) and used in all subsequent calculations. This procedure was performed prior to each measurement session, resulting in a separate calibration function for each data set to account for any possible effects due to weather or time dependent drift of sensor performance.

$$H_{cal} = \sqrt{H_{cal_{12}} \times H_{cal_{21}}} \quad (3.6)$$




Figure 3.2: Impedance tube measurement apparatus configuration. (a) Tube setup with termination sensor *Ch3* for use in measurements involving sensor distances. (b) Tube setup during skin measurement, note how the tube is pulled slightly away from the normal skin position. (c) Tube setup with hair sample and wire mesh in place, note that l is increased by the mesh insertion distance d so that l is the distance from *Ch2* to the surface of the material. (d) Anechoic termination.

3.2.2. Microphone Distance to Termination

According to the ASTM standard for the two-microphone impedance tube measurement, the distance from the end microphone to the surface sample must be known to an accuracy of 0.1 mm. [3] In the standard this requirement is waived for highly absorptive, roughly textured surfaces. This could apply to hair, but definitely not to skin measurements. For calibration with a rigid termination, and for skin which is assumed to be reasonably rigid at this point, this exception does not hold. Therefore, it is necessary to determine the distance from this microphone to the rigid to semirigid termination. The requirement is precision to 0.1 mm. The diaphragm was aligned parallel to the walls of the tube, and had a diameter on the order of 1.0 mm. Therefore, even if it was possible to measure the physical distance with an optical measurement method, such as a caliper or laser scheme, the precision required was on a scale smaller than the size of the diaphragm. It was necessary to determine the distance from the *acoustic* center of the closer diaphragm to the termination, for which this method would not suffice.

A standard way to determine this position is to move the position of the microphone, and scan for pressure minima. This method contains possible errors in determining the exact acoustic position of the end of the tube, requiring end corrections to be determined and included. From this position, it is possible to calculate the distance quite easily. In designing the system, the position of the microphones was determined, and then fixed. In addition, due to the small tube diameter, the insertion of a probe microphone inside the tube would most likely disturb the sound field in the tube such that the plane wave assumption would be compromised. Therefore, a method which allows for the fixed position of the microphone must be used. Taking the traditional method of scanning the sound field for the pressure minima at a given frequency and transforming the problem to the frequency domain, it should be possible to perform the same measurement by sweeping frequency and determining the exact frequency at which the fixed microphone is at a pressure minima. The necessary frequency precision was determined as follows: the approximate distance to the termination was measured to be 5 cm. Using 343 m/s as an approximate sound speed, this would result in a first null (for a quarter wavelength) to be at 1715 Hz. A variation of 0.1 mm would yield a distance of 5.01 cm, and a quarter wavelength frequency of 1711.6 Hz. With the required minimal accuracy being 0.1 mm, precision of 0.01 mm was attempted. Introducing this variation to the approximate measurement of 5 cm yields 5.001 cm, which has a corresponding quarter wavelength frequency of 1714.7 Hz. From these two rough calculations, a frequency resolution of 0.3 [0.01 mm resolution] to 3.4 Hz [0.1 mm resolution] is necessary, depending on the confidence in the measurement and the desire to remain within the standard requirements.

Sweeping the frequency at this resolution is possible, but due to the fact that the exact null is hidden within the noise floor (which results in a null that spans approximately a 10–15 Hz range), precise determination of this frequency would not be possible by simply choosing the frequency at which the minima occurs. Instead, interpolation of the data surrounding the null was used to extrapolate the frequency null and avoid contamination of the prediction by the noise floor. In doing this, it was also possible to sweep frequencies at a resolution of only 4 Hz as the curve shape, not the exact frequency of any one specific data point, is important.

For confidence in the interpolation, the shape of the frequency spectrum data should be smooth in the region before and after the null. The raw pressure data, measured as a function of frequency for either of the two measurement microphones, contains resonances of the tube in addition to the nulls due to the distance from the termination as shown in Fig. 3.3. It is clear to see that exact determination of the null position would be difficult using the raw pressure data, due to the influence of the tube resonances. These resonances are a function of the termination impedance, tube length, and driver, but not a function of the sensor location. In order to remove this corruptive influence so that the measurement contains only distance dependant information useful for determining the distance to the termination, a microphone was placed flush with the termination. A special termination, with a mounting hole for the microphone, was used for this measurement. The microphone was sealed in position and, assuming no real effect on termination impedance due to the diaphragm replacing a small portion of the rigid termination, resulted in no appreciable change in the acoustic response of the system. This sensor, located at the termination, contained the resonance information of the system, but no distance effects, as can be seen in Fig. 3.3.




Figure 3.3: Sound pressure data for impedance sensor and termination sensor for use in calculating microphone distances. *Ch3* corresponds to the termination sensor as shown in Fig. 3.2(a).

A measured frequency response between one of the measurement microphones and the termination microphone yields the necessary data (see Fig. 3.4). As there are two microphones in place for the impedance measurement, it is possible to use data from both microphones to determine their acoustic distance from the termination. Several nulls exist over the selected frequency range, each one capable of being used to determine the distance. This allows for error checking and redundancy in the distance calculation. Use of the transfer function between the impedance sensors, containing distinct peaks and nulls, is not suitable for this measurement. As seen in Fig. 3.4, the peak and null positions for this measurement are shifted up in frequency from the correct positions for use in distance determination, as seen in both the raw pressure measurement and the transfer function with respect to the termination sensor.

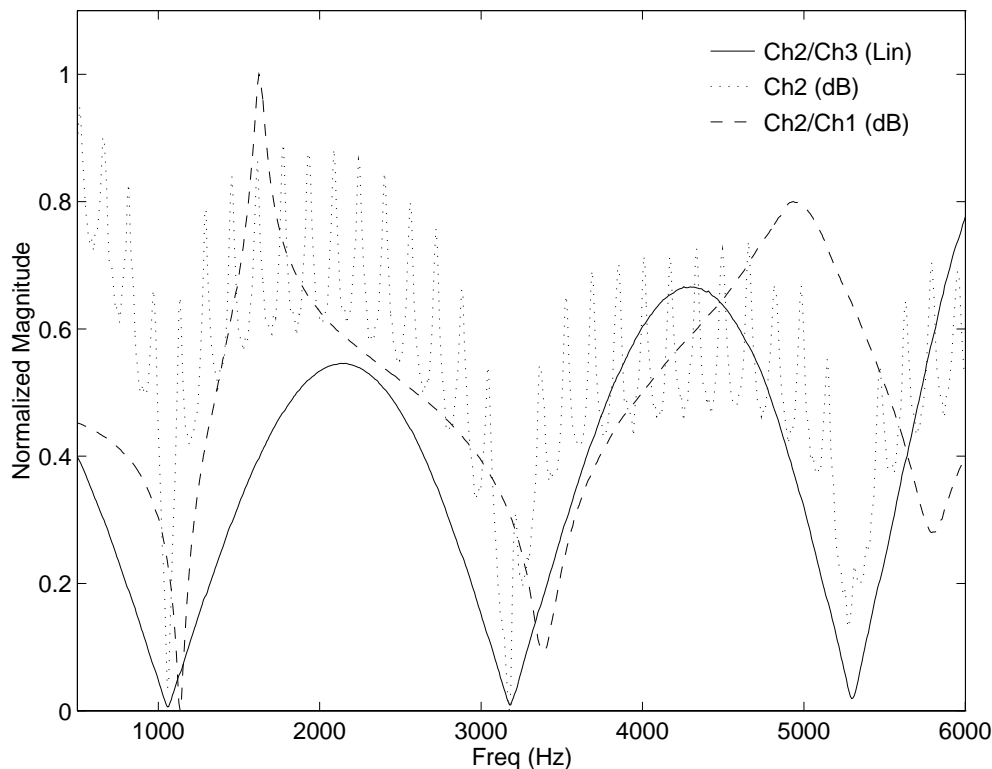


Figure 3.4: Distance calculation data: transfer function between the impedance microphone and the termination microphone, raw impedance sensor sound pressure, and transfer function between the two impedance microphones. The y-scale has been normalized for both the linear and log data so that all the data sets can be visually compared.

Interpolation of the frequency minima was determined (with only the distance dependant information) using `polynul2.m` (see section A.2). This script prompted the user for the null in question (from a graphical plot), from which it then used data points to either side of the initial guess to perform a 2^{nd} order polynomial fit for the minima. This result

was then used as the initial guess for a second polynomial fit to obtain the final minima position. An example of a `polynul2.m` result is shown in Fig. 3.5. The zoom plot shows a detailed view of the selected null area. The data points are shown, as well as the two interpolated null estimates. It can be seen that with the second estimation the null seems to be accurately determined. Repeated tests for various initial guesses showed that only two interpolation loops were necessary to provide accurate results. From this position (in frequency) the distance was calculated depending on which null is selected ($1/4$, $3/4$, or $5/4$ wavelength null) as a function of atmospheric conditions in the experimental chamber at the time of measurement. The speed of sound was calculated taking into account the temperature, barometric pressure, and relative humidity using `whatc.m` (see section A.3), which uses a polynomial approximation of published data and atmospheric effects on sound speed as described in Pierce. [35]




Figure 3.5: Results of interpolation to determine precise frequency null, and consequent microphone position. The dotted line(+) indicate the first pass interpolation and the solid line(o) indicate the final interpolation result.

Whether the above interpolation method accurately determines the null position of this type of curve can be investigated with a simple example. The question arises when the asymmetry of the curve about the null is taken into account. A simulated data set was generated using sinusoidal functions which had an obvious asymmetry (see Fig. 3.6(a)) in

which the first lobe of the data set is half the width of the second. Using the interpolation routine on this data resulted in the null prediction given in Fig. 3.6(b). It can be seen that the null position is slightly skewed in the direction of the curve with the steeper slope. For the degree of asymmetry used in this example, the deviation from the true null was approximately 10 Hz. This error is due to the fact that the simulated data, as well as the measured data, is sinusoidal in nature while the interpolation uses a polynomial function. As the accuracy required for the null position determination is 3.4 Hz, and the asymmetry of the real data is much less than this simulated case, it was assumed that any resulting errors are within the necessary range.


The interpolation method used to determine the frequency null minima was repeatable for a given frequency null within approximately 1.5 Hz. Several trials were performed and various termination conditions were used (results given in Tbl. 3.1). As the hair measurements required some means of being contained within a fixed area of the measurement system, a coarse wire mesh was used to keep the hair from moving, and was backed by the rigid termination (see Fig. 3.2(c)). To account for the difference in distance with this setup, the distance measurement was performed with the mesh in place, backed by the rigid sensor termination. As shown in section 3.3.1, the inclusion of the mesh has no appreciable effect on the measured quantities, and can therefore be assumed to be transparent over the frequency range of interest. For this reason, the distance used in the impedance calculation was measured from *Ch2* to the sample surface, not the screen surface. In addition, various sealing methods for the sample container were used to determine the presence of any leakage. A small leak or gap could result in an error in the acoustical distance to the termination. Effects of the various sealing methods can be seen in Tbl. 3.1.

3.2.3. Spacing Between Microphones

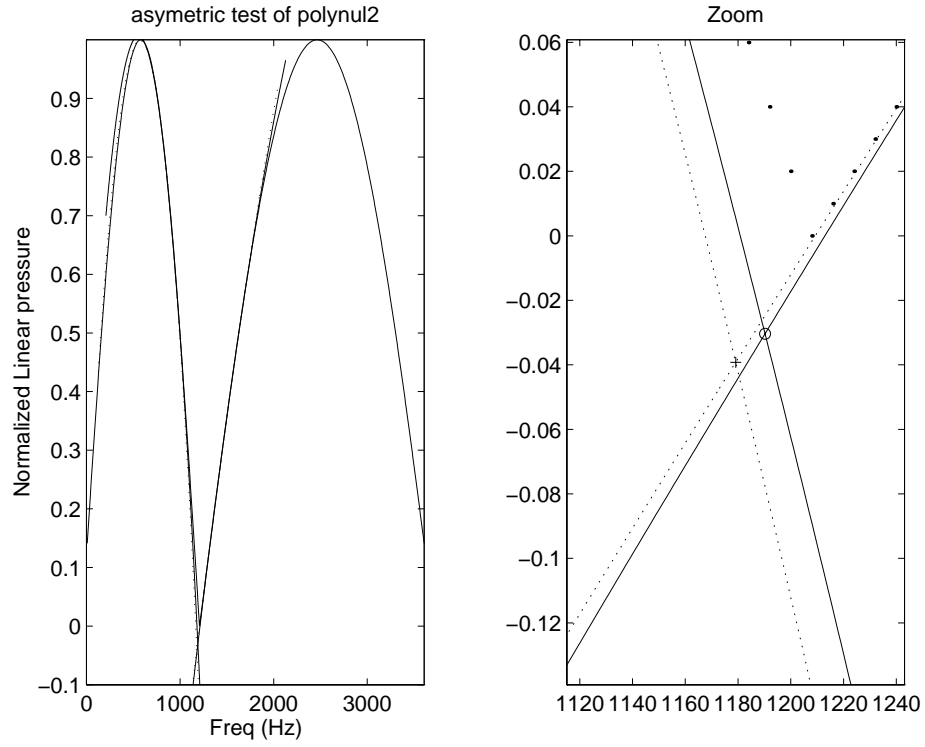
Using the distance to termination data from both microphones, it is also possible to determine the spacing between the microphones, s , a distance necessary for the acoustical calculations. Using the results from Tbl. 3.1, the separation distance was determined to be 2.365 ± 0.004 cm, taking an average for the distances determined from each null. This distance can also be calculated by using an anechoic termination and determining the phase delay between the two sensors, again using the transfer function. [8] This method of measurement is totally independent of the previous method and was therefore used as a reference check.

Using an anechoic termination, the transfer function was measured for the sensors in their normal position, and also with the positions switched. Using `anspace.m` (see section A.1) to calculate the spacing, the sensors were first calibrated as per Eq. (3.6), and the acoustic transfer function was determined using Eq. (3.7). The microphone spacing was then calculated using the phase delay Φ between the two sensors, normalized with respect to frequency, as defined in Eq. (3.8).

$$H = \frac{H_{anechoic}}{H_{cal}} = \frac{H_{12}}{\sqrt{H_{12} \times H_{21}}} \quad (3.7)$$



a. Simulated asymmetric test signal for interpolation routine.



b. Results of interpolation routine using asymmetric test signal.

Figure 3.6: Effect of asymmetric data on null interpolation routine.

Table 3.1: Microphone distance to termination calculations using `polynul2.m` for each null in measurement ($Ch1 - 3$ nulls, $Ch2 - 2$ nulls). `disne`: termination glued to end of tube. `discne`: termination glued and with clay around seal. `dissne`: termination glued with screen glued in place and clay around seal.

Datafile	$Ch1(\text{cm})$			$Ch2(\text{cm})$	
dis1e	7.669	7.667	7.675	5.294	5.324
dis2e	7.667	7.665	7.673	5.292	5.321
dis3e	7.667	7.664	7.670	5.291	5.320
disc1e	7.687	7.677	7.677	5.302	5.326
disc2e	7.687	7.677	7.677	5.301	5.362
disc3e	7.685	7.674	7.676	5.301	5.325
diss1e	7.685	7.806	7.822	5.359	5.446
diss2e	7.678	7.805	7.823	5.360	5.446
diss3e	7.688	7.807	7.823	5.361	5.446

$$s = -\frac{\Phi(H)}{2\pi f} \quad (3.8)$$

As the termination is not completely anechoic at low frequencies, there is some variation/oscillation in the result as a function of frequency. The results of the separation calculation are shown in Fig. 3.7. The thin curve shows the separation distance calculated as a function of frequency. The thick curve shows the same calculation with a running average over frequency (for details see section A.1). To determine a single spacing distance, an average of the separation values was determined, being calculated over a limited frequency range. The horizontal dotted line shows the average value; the frequencies over which the average was computed are indicated by the horizontal solid line. The vertical line shows the optimum frequency for the given microphone spacing which is equivalent to a quarter wavelength, $f_{opt} = c/4s = 3.7 \text{ kHz}$.

Results of this measurement are a separation distance of $2.304 \pm 0.006 \text{ cm}$, an accuracy in terms of repeatability well within the measurement standard. In comparison to the method devised above, though, there is a definite difference on the order of 0.06 cm , which is above the standard threshold for determining the distances. The exact nature of the discrepancy is unclear. For a typical measurement, accuracy of 0.06 cm would probably be sufficient. In this case, it is unclear how one would go about measuring an acoustic distance to any greater accuracy and be confident in the results. The results from the new method for distance measurements, using the frequency null interpolation, were chosen to be used in subsequent calculations as both s and l are required, and the anechoic spacing measurement only determines s . The resulting distances used are $s = 2.365 \text{ cm}$ and $l = 5.312 \text{ cm}$.

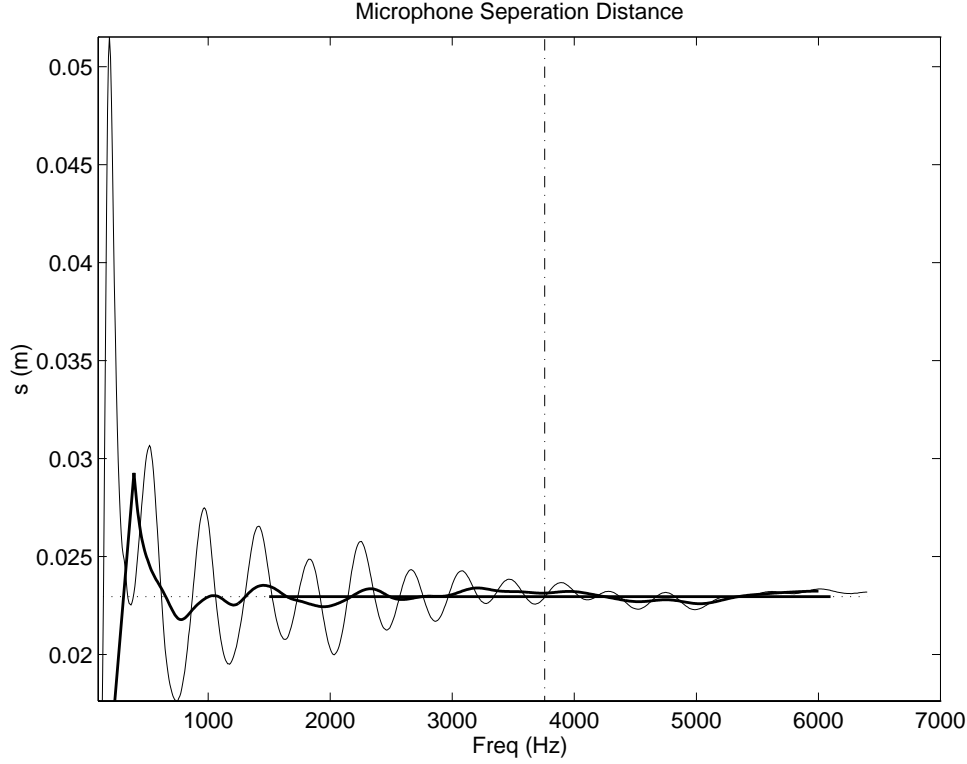


Figure 3.7: Microphone separation distance calculation measured using phase delay.

3.2.4. Planar Surface Assumption

For semirigid to rigid impedance terminations the distance to the termination is crucial, as described in the previous section. Therefore, it is important that the sample be flat to roughly the same degree of accuracy at the standard requires. In measuring the properties of skin, the measurement tube was placed against the area of skin to be measured. It was quickly observed that a curvature of the skin was induced which was proportional to the pressure applied with the tube. As discussed in the literature, and as would be expected, the impedance characteristics of the skin will vary with the applied pressure as the tension on the area will vary, creating a different mechanical system. It would be difficult to create a measurement system which would reliably use the same applied force for each trial, simulate normal skin condition, and provide a good acoustic seal. The following procedure was therefore conceived to minimize these problems. The impedance tube was attached to the skin area with a thin layer of nontoxic rubber cement, providing a good acoustic seal. Once cured, the tube was pulled slightly so that the skin sample was visibly raised from the skin surface area as represented in Fig. 3.2. The result of this procedure was that the skin sample was under little tension in the plane of the skin, keeping close to the normal skin position situation, without effects due to the applied pressure of the tube. Having the skin slightly pulled away from the normal surface removed the effect of bowing

or curvature due to the pressure of the tube against the skin, or any normal curvature present on the sample area. The result was a planar skin sample with reasonably typical skin tension in the planar direction. It was acknowledged that the tension in the direction normal to the surface would not be the same as the typical position, but this was assumed to have much less of an effect than the other conditions the procedure was designed to alleviate. To calibrate for the distance to the sample, the rigid sensor termination was attached to the tube in the same manner, by using rubber cement. To verify the adequacy of the seal, additional distance measurements were made by including an external clay ring with the reference termination. Results showed no appreciable effect and therefore the rubber cement seal was assumed to be sufficient for use in the skin measurement (see section 3.2.2).

3.3. Measurements

3.3.1. Reference Terminations

To verify the system design and calibration, several measurements can be made. As there are no materials which have an absolutely defined reflection coefficient or impedance, there are limited tests which can be performed. The case of an infinitely rigid termination is theoretically possible, but in all practical purposes can not be achieved or measured experimentally due to electrical noise, losses through the tube walls, or other factors. Measurements of a rigid termination, though, aid in showing the limitations of the specific experimental apparatus. The same holds true for an anechoic termination. It would be rather difficult to make a totally anechoic termination, but an experimental effort towards such can be measured and used to further quantify the behavior of the measurement system.

The impedance of an ideally rigid termination would be infinite. It is therefore not possible to compare the measured impedance to the theoretical value. On the other hand, the reflection coefficient is well defined and is suitable to use for these evaluations. The ideal rigid termination would have a reflection coefficient $|R| = 1$. The ideal anechoic termination would have a corresponding reflection coefficient $|R| = 0$. The experimental measurements of the physical implementation of these idealized terminations is shown in Fig. 3.8. (Explanation of the use of the wire screen is given in sections 3.3.4 and 3.2.2.) There are some losses in the system since the tube is not infinitely rigid, as seen in the low frequency limit of the rigid termination. Additional errors can be seen in the low frequency limit of the anechoic termination, which, when considered with the rigid and duct size limit, restrict the functional frequency range of the system to 1–6 kHz (see Fig. 3.8). Over this frequency range, the results of this measurement for the rigid termination produce a mean reflection coefficient of $\overline{|R|} = 0.97$, which for all practical purposes can be considered rigid. The anechoic termination measurement results in a mean reflection coefficient of $\overline{|R|} = 0.05$ over this range. For all practical purposes this anechoic termination is sufficient.

The impedance for an anechoic termination is well defined. The impedance of a totally absorptive material should be simply the impedance of the propagating medium, in this case air. Plotting the acoustic impedance, defined as the specific acoustic impedance normalized by the characteristic acoustic impedance of the medium ($\rho_0 c$), results in Fig. 3.9. The mean values for the anechoic impedance measurement over the selected frequency range are $\overline{|Z|} = 1.02$, $\overline{Re\{Z\}} = 1.01$, and $\overline{Im\{Z\}} = -0.02$. Ideally, the impedance should be purely real and equal to unity. These results are sufficiently close to the theoretical values.

The only termination which has a well-defined solution, other than the idealized rigid and anechoic terminations, is an open ended termination. The reflection coefficient for an unflanged thin walled circular duct has an analytical approximation given in Eq. (3.9), where a is the tube cross-sectional radius and $\ln(\gamma) = 0.5772$. [28] Comparison of the measured and theoretical value for the open ended tube termination can be seen in Fig. 3.10. The error in the theoretical approximation is less than 3% in the region of $ka = 1$, and




Figure 3.8: Reflection coefficient reference measurement showing the measured values for a rigid and anechoic termination. Measurements were made both with and without the wire screen in place.




Figure 3.9: Impedance measurement of the anechoic termination. The impedance Z is plotted as $Z_s/\rho_0 c$. An ideally anechoic termination would have an impedance of $|Z| = 1$, $Re\{Z\} = 1$, and $Im\{Z\} = 0$.

reduces as ka becomes greater or less than 1 with respect to the full Bessel function representation.

$$|R| = \begin{cases} e^{-\frac{1}{2}(ka)^2} \left(1 + \frac{1}{6}(ka)^4 \left[\ln \left(\frac{1}{\gamma ka} \right) + \frac{19}{12} \right] \right) & ; ka < 1 \\ e^{-ka} \sqrt{\pi ka} \left(1 + \frac{3}{32} \frac{1}{(ka)^2} \right) & ; ka > 1 \end{cases} \quad (3.9)$$




Figure 3.10: Reflection coefficient reference measurement showing the measured and theoretical values for an open ended tube. Measurements were made both with (OpenSc) and without (Open) the wire screen in place.

The measured reflection coefficient agrees well with the theory for a open ended un-baffled circular duct. The frequency at which there is the minimal amount of error in the measurement (as described in section 3.2.3) is 3.7 kHz. This is the frequency at which the sensors are at $\lambda/4$. At this frequency, and down to about 2.2 kHz, the measured response agrees very well with the theory. It is expected that over this range, the error in further measurements will be small, increasing as the this range is exceeded. From the results of the reference measurements of the rigid, anechoic, and open terminations there can be confidence in the measurement of the unknown acoustical properties of the test samples.

3.3.2. Effect of distance errors

Errors in determination of the microphone spacing and distance to termination greatly effect the measured acoustical parameters. Variations in the measured impedance due to errors in these distances is an important factor of consideration. Minimization of these errors has been the goal of the previous sections. In order to easily compare the impedance measurements, the resulting impedance data was smoothed. An example of the data smoothing result, as compared to the raw data result, is given in Fig. 3.11. Smoothing is performed as a running average as per the method used in section A.1.






Figure 3.11: Example of the smoothing of the measured impedance data. Results shown are for a rigid termination.

Variations due to fluctuations in the microphone spacing distance s are shown in Fig. 3.12(a). Three values for the spacing are shown. These values correspond to the result of the anechoic phase delay measurement, the frequency null interpolation measurement (for which the value was used in all subsequent calculations), and another, overestimated value for s which was the intended spacing in the initial design of the apparatus. As can be quite easily seen, the results for the $Re\{Z\}$ vary quite substantially with small variations in s , on the order of $\Delta s \leq 1.0$ mm. The value determined from the frequency null interpolation method provides for the least amount of fluctuation in the impedance values, resulting in a higher level of confidence.

Variations due to fluctuation in l are shown in Fig. 3.12(b). Values shown include the results of the frequency null interpolation method, and values varying by $\Delta l \leq 1.0$ mm. It is easy to see that small changes in l , even on the order of $\Delta l \leq 0.1$ mm, result in large changes in the impedance values. It is more difficult to determine which result to use as the ‘correct’ value, as there was no other method for determining the distance to the



a. $Re\{Z\}$ fluctuations due to variations in microphone spacing s , given in cm. Distance indicated with '*' is the value determined in section 3.2.3 and used in subsequent calculations. The value of $s = 2.304$ cm was the result of the phase delay measurement technique.



b. $Re\{Z\}$ fluctuations due to variations in microphone spacing l , given in cm. Distance indicated with '*' is the value determined in section 3.2.2 and used in subsequent calculations.

Figure 3.12: Effect of errors in s and l on the calculation of surface impedance Z .

termination. For this reason, the direct result from the experimental measurement was used, having established confidence in the procedure.

3.3.3. Skin

The main purpose of the acoustical measurement of the skin was to determine if skin behaves as a rigid material, or exhibits more complex behavior in the selected audible frequency range. To this aim, several measurements of skin were made for various skin circumstances. The measurements listed in Tbl. 3.2 were performed using the method described in section 3.2.4. The results for all these measurements are given in Fig. 3.13.

Table 3.2: The various acoustical measurement conditions for the skin tests.

Forearm
Upper arm
Center of forehead
Stomach (under last rib)
Cheek with teeth clamped (mouth closed)
Cheek with teeth loose (mouth closed)
Cheek with mouth open
Cheek puffed




Figure 3.13: Reflection coefficient reference measurement for a number of different skin locations and conditions.

The average values for the reflection coefficient for these measurements were $\overline{|R|} \simeq 0.97$. In comparisons to previous works (see section 2.2), these results agree well. For all practical purposes this can be assumed to be rigid, as the reference measurement for the rigid termination also had the same result. It can therefore be assumed that the effect of the skin impedance on the HRTF is the same as would be for any acoustically rigid material. In light of this, the model parameters for the skin area were not altered from the rigid boundary condition assumption (see section 5.3).

3.3.4. Hair

Measurements of the acoustical properties of hair were made on samples from four individuals. A fifth sample contained a second sample from one of the first four individuals. In addition, two random samples of two different foams were included for comparison. Details of the samples are given in Tbl. 3.3. Of initial interest was the variation in hair diameter from sample to sample, and especially the variation within a sample. Hairs $h1$, $h2$, and $h4$ have little variation in diameter within the sample, while $h3$ has a large variation (a factor of $\times 10$ more variation). One explanation for this could be the method of hair acquisition. Samples $h1$ and $h2$ were taken from a single location on the head, grouping the hair in a small bunch and cutting the bunch with scissors. This could ensure a more uniform distribution. Samples $h3$ and $h4$ were taken from clippings from a hair cut. It is possible that over the entire head (and neck) hair diameters vary and that these various hair types were included in $h3$. Why they were not present in $h4$ is probably a function of the individual or the type of hair cut.

Table 3.3: Hair physical properties including total mass and hair diameter of the indicated samples. Also included are data for foam samples (kf and wf) with approximate fiber diameters.

Sample	mass (g)	fiber diameter (mm)
$h1$	1.5811	0.047 ± 0.002
$h2$	1.0415	0.045 ± 0.001
$h3a$	1.7204	0.055 ± 0.01
$h3b$	0.6645	0.055 ± 0.01
$h4$	1.9630	0.057 ± 0.002
kf	0.2768	≈ 0.0508
wf	0.1590	≈ 0.0635

In order to accurately measure the acoustic properties of the hair, it was necessary to confine the samples to a definite space and to have something of a defined surface distance which was planar. A coarse wire mesh was glued to the end of the tube. The mesh wire had a diameter of 0.24mm with mesh hole size on the order of 1 mm^2 . To ensure that the screen caused minimal acoustic interference with the measurements, the

reference termination measurements were repeated with the screen in place. As shown in Fig. 3.8, the inclusion of the screen has no effect on the high impedance surface of the rigid termination or the anechoic termination. The measured impedance of the open termination also displayed no change, as can be seen in Fig. 3.10.

In addition, the distance to the rigid sample with the screen in place was repeated to accurately determine the new value for l . Using the final set of data in Tbl. 3.1, the modification of the distance to the termination with the inclusion of the screen can be determined. It is interesting to note that the inserted distance increases with frequency. The high frequency nodes measured show a screen distance insertion of 0.145 cm, the mid-frequency nodes show a distance of 0.121 cm, and the first and lowest frequency nodes show almost no effect on distance: 0.0–0.05 cm. It is difficult to determine which result to use, and how to interpret this discrepancy. It may throw possible doubt on the measurement scheme, or it may point out a low frequency limitation of the method. The latter is more plausible reasoning as it has been shown in a number of places that the low frequency response is not ideal (seen in Fig. 3.7 and Fig. 3.8). It is also plausible that the location of the acoustic center of the microphone is frequency dependent. Measurement of the distance to account for this would be very difficult. As can be seen in Fig. 3.7, there are oscillations in the measured distance due to other factors, such as losses in the tube and a non-ideally anechoic termination, which would need to be removed before such a detailed measurement could be made. As the region of interest lies between the second two nodes, the screen insertion distance will be calculated as an average between the two results, giving a new distance to termination value of $l = 5.446$ cm. Measured reflection






Figure 3.14: Reflection coefficient measurement for the hair and foam samples.

coefficients for the hair samples is given in Fig. 3.14. Data for two random foam material samples (*kf* and *wf*) are included for comparison.

It is quickly apparent that hair is not a rigid surface. There is some diversity in the hair results, while the results for the foam are almost identical. Examining the physical properties of the samples (see Tbl. 3.3) shows that the reflection coefficient generally increases as the mass of the hair sample decreases. Though, with similar fiber diameter and significantly less mass, the foam samples have comparable acoustical performance. This discrepancy is attributed to the complex nature of fibrous absorption and its relation to fiber length, packing geometry, fiber material, *etc.* There are theoretical prediction methods for determining the absorption coefficient for fibrous material. [34] These prediction equations require measurements of parameters such as flow resistivity. In addition to theoretical predictions, there are data available for common material like rockwool, which can be given as a function of bulk density (shown in Fig. 3.15(a)). These measurements were for rockwool of ‘practically infinite’ thickness, ranging from 0.5–1 m. A comparison exercise can be done between the mineral fiber data and the measured hair data. Given the mass of the hair samples in Tbl. 3.3, the bulk density can be calculated for a given sample. Using sample *h1* for this purpose, and knowing that the samples when measured are packed into a cylinder with a diameter of 2.60 cm and length of 1.5 cm the bulk density can be estimated as 200 kg/m³. Typical mineral wool fibers have a diameter of 0.005 mm, which is on the order of 10 times smaller than hair. Bulk densities for the rockwool data in Fig. 3.15(a) range from 50–200 kg/m³, for which the hair sample is at the upper limit. Measurement of the absorption coefficient α , where $\alpha = 1 - |R|^2$, is given in Fig. 3.15(b) for the hair and foam samples.



a. Absorption coefficient measurements for rockwool of thickness exceeding 0.5 m. (Source [34][p.211])



b. Absorption coefficient, α , measurements for hair and foam.

Figure 3.15: Absorption coefficient measurement comparisons between (a) hair measurements and (b) published data on commonly used absorbing material rockwool.

The rockwool samples were much thicker and had a smaller fiber diameter, both resulting in increased absorption. The rockwool with the greatest bulk density, assumed to have one of the largest fiber diameters, is seen to be quite low in absorption relative to the other samples. This can not be used as the only factor, as a few samples with much lower bulk density have lower absorption. There are other factors involved in the absorption of such materials, including packing density and bonding materials within the fibers. All can affect the acoustical properties in a variety of ways. What can be noted, though, is that the hair data follows a similar trend as the mineral wool, showing that there can be a level of confidence in the measured hair data. Taking into account the variation in thickness and fiber diameter, the hair is of comparable absorption, at least within the order of magnitude, as mineral wool.

The results of this measurement vary quite significantly from the work on furred mammals, which presented results in which α ranged from 0.2–0.4 over the frequency range of interest here (see section 2.2), though bulk densities of the hair were not cited. [16,55]

The final calculations regarding hair were the acoustic impedance measurements. Given the reference measurements made with the system, and the trend comparison with other measurements of absorbing material for the absorption coefficient, the results for the real and imaginary components of the impedance given in Fig. 3.16 can be used with confidence. With the variation between samples, and the obvious effect of hair density, it is not possible to conclude with a definitive value for the impedance of hair. For use in the HRTF calculation model, the inclusion of some realistic values for the impedance will suffice for this work, as it is in its initial stages. For this reason, an average value for the impedance is calculated from the data results as a function of frequency and implemented into the model (see section 5.3).

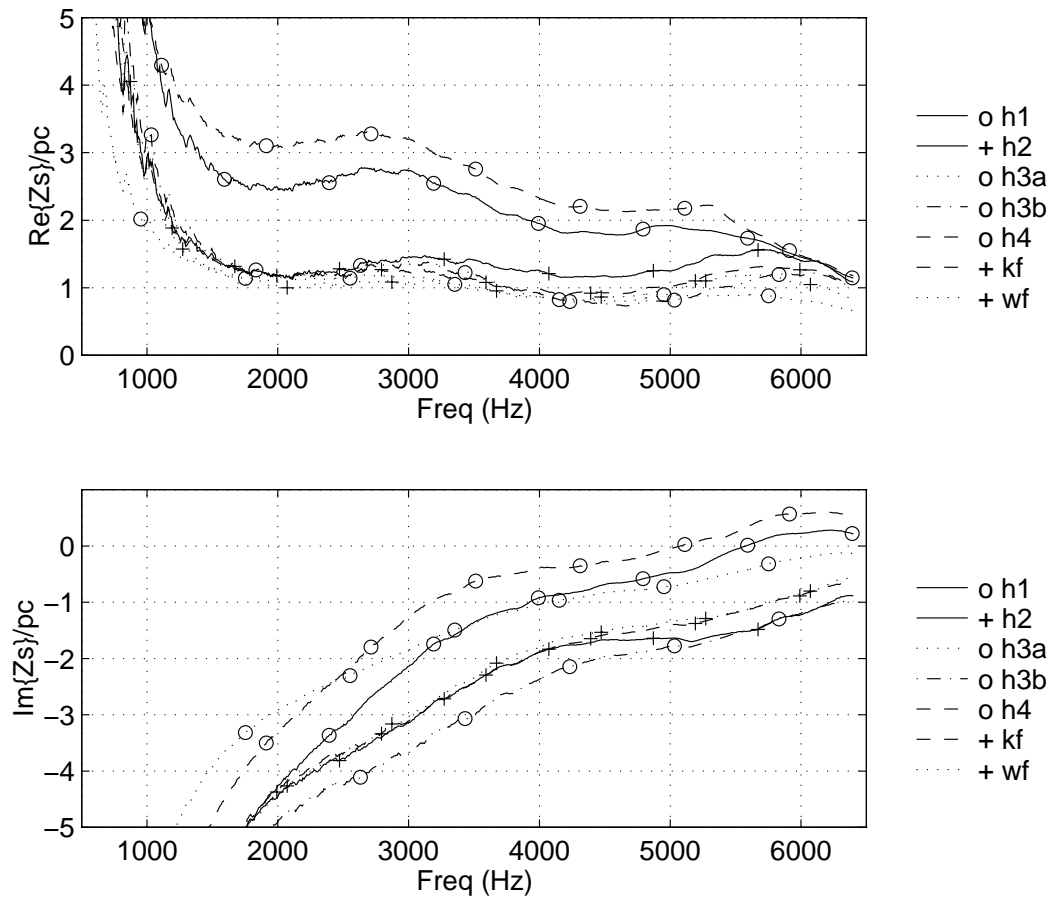


Figure 3.16: Measured $\text{Re}\{Z\}$ and $\text{Im}\{Z\}$ for hair and foam samples. $Z = Z_s/\rho c$.

Chapter 4.

HRTF Measurement

In order to evaluate the accuracy of the calculated HRTF in the next chapter, real individual HRTF measurements were made. These experimental measurements are later compared to the HRTFs calculated from the geometrical data. Two different measurement methods were used for obtaining real HRTF data. Due to the equipment and space necessary to make such measurements, these experiments were carried out two different research facilities. The first set up was at the Wright-Patterson Air Force Base (WPAFB). The second was at the headquarters of Aureal Semiconductor Inc. using the Crystal River Engineering system (CRE). Both locations agreed to cooperate with this research in making measurements and supplying hardware/software during the course of this research.

4.1. HRTF Equalization and Labelling

In order to facilitate the plotting of such a large number of experimental conditions, a labelling scheme for experimental data was used which utilized three parameters. Data were identified by subject number, measurement method, and ear. There were a total of six subjects (no. 1–6) used throughout the work, though not all measurements were made with all subjects. For example, only subject 1 was used for the computational work, so for the simulations no subject number was necessary. Measurement methods were as follows; WPAFB real subject = R , WPAFB pinna mold = P , and CRE real subject = C . Finally, ear designation was L or R for left and right ear measurement respectively. All simulated measurements used only the left ear, so again no designation was necessary. An example of the labelling scheme would be $1PL$, which would correspond to subject 1, WPAFB pinna mold measurement, left ear. This scheme may also be divided among different parts of a figure, so that the main plot could be of $1L$, where different data sets in the plot could be P , R , and C .

There are a variety of methods for looking at HRTF data which have been used by researchers. These methods include displaying the time series impulse response, diffuse field equalized, free field equalized, and front equalized frequency response. [7,45,46] These are different from presentation normalization or equalization methods, which concern themselves with the various headphone equalization methods (diffuse, free-field, front). The impulse response is typically used in signal processing for presenting virtual sources, but is difficult to interpret visually. Free-field equalization results in a comparison between the transfer function with the subject present and with the head removed. This is, ideally, what the HRTF measures. Further equalization methods are used to try to separate the directional from non-directional components and/or simplify signal processing when used for presenting virtual sources. The diffuse field equalization calculates direction dependent

components of the HRTF using Eq. (4.1), where D is the directional component of the HRTF, and H is the free-field HRTF component for direction θ_i . When used in a virtual source presentation system this implementation specific HRTF, which is less complex than the full HRTF, is then combined with a generic diffuse field transfer function, which is assumed to be the same for all subjects. [48] The reason for this assumption is to decrease the computational time of a real-time acoustic presentation system. Though this equalization method results in a purely directional component of the HRTF, it can only really be used to compare HRTFs measured in the same manner. The reason being that depending on the source number and distribution, the diffuse components can differ, making comparisons between multiple measurement methods difficult.

$$D_{\theta_i} = \frac{H_{\theta_i}(\omega)}{\left[\sum_k H_{\theta_k}^2(\omega)\right]^{\frac{1}{2}}} \quad (4.1)$$

As seen in Fig. 4.1 the diffuse field equalization function (the denominator of Eq. (4.1)) follows a general trend, but is not constant between individuals, measurement methods, or even between ears. Differences in measurement methods have a large effect on the calculated diffuse field. Other variations exist on the order of 5–10 dB, the largest being in the frequency region from 8–10 kHz. The fact that the two ears of a single subject do not have the same resulting diffuse field is interesting, implying that the diffuse field calculation does incorporate some characteristics which vary, from side to side, on the same person. These characteristics would be items other than body size and very general head shape. It may imply that the diffuse field calculation, though not containing ‘direction’ dependant information, does contain information which aids in the processing and decomposition of binaural signals through additional left/right discrimination. This variation suggests, assuming the error is not experimental, that each pinna has a general frequency filtering function which is imposed on any signal regardless of direction, and that these functions are not necessarily the same for both ears of a single subject. This could aid in resolving localization ambiguity and median plane localization. The variations in diffuse field calculations from the measured data were even greater for the CRE data sets, which do not have nearly the same number or locations as the WPAFB data sets (see sections 4.2.1, 4.2.2 and 4.3.1 for details). These missing locations, and corresponding spectral contributions, would then be missing from the diffuse field calculation, which would then result in different effects on the equalized HRTF. These differences would not be due to directional differences, but omission of data from the diffuse field calculation.

A final method, which can not really be used for virtual presentation purposes, is to normalize the HRTF by the response at 0° elevation and 0° azimuth, or directly straight ahead. The resulting equalized HRTF will then show how the directional component changes within the HRTF set. This is very useful for comparing different measurement methods, as is required here. As so much of the measured system is changing between the various methods, *i.e.* removal of the pinna, removal of the torso, removal of the arms and legs, it is clear that the diffuse field would not be the same. But, we can see from the front-equalized HRTF how these many changes affect the perceived field. Applying front-source equalization to diffuse field equalized data still has the desired effect and can

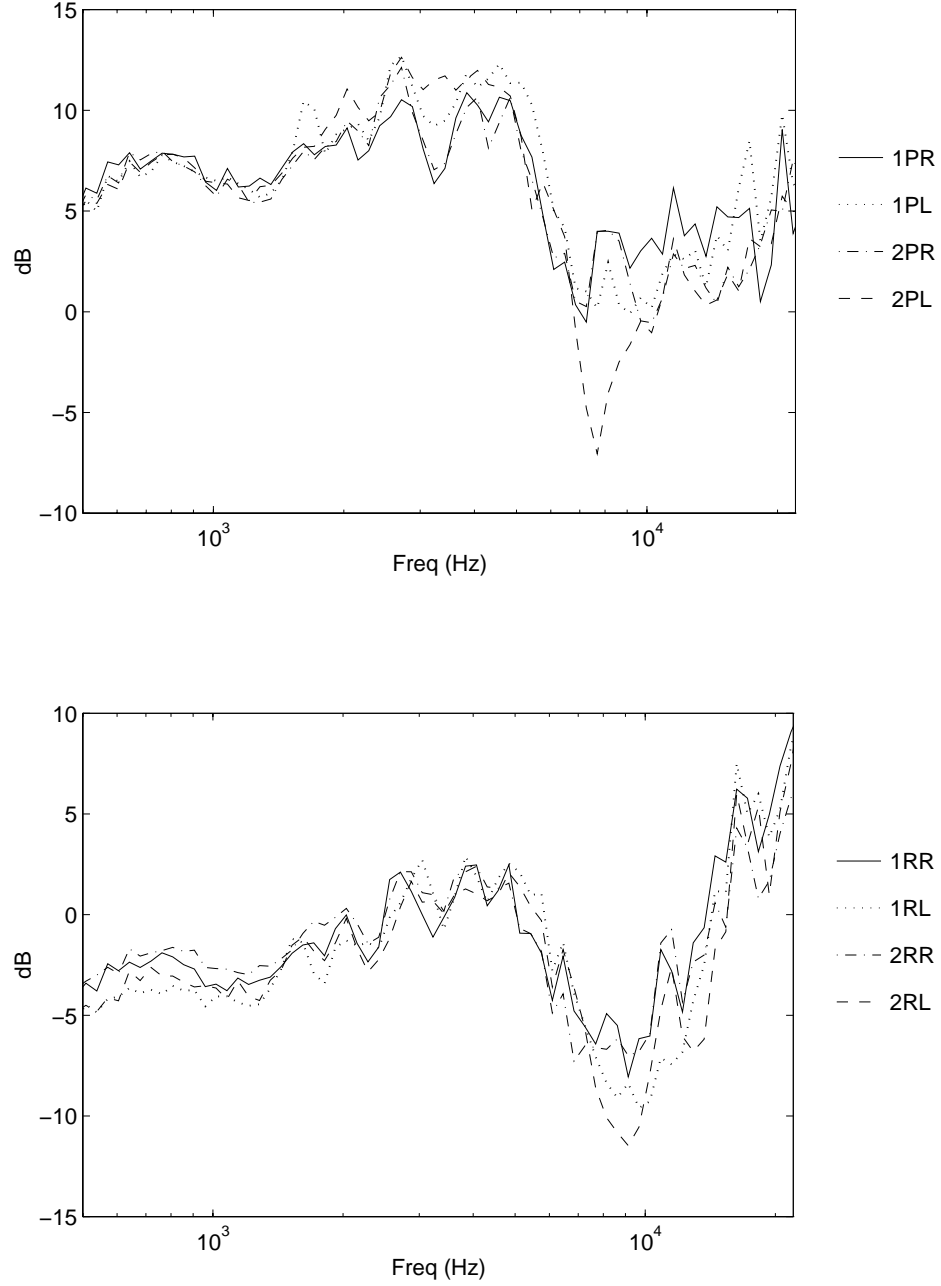


Figure 4.1: The diffuse field equalization function (the denominator of Eq. (4.1)) for two subjects using both real and pinna mold data for each ear.

be used. Measurements and calculations made using the WPAFB array setup do not contain a source position directly in front of the subject (see section D.1 for the detailed coordinates of all WPAFB points). Instead, there is a point 7.5° to the left and right. For front equalization, the mean of these two HRTF measurement results was used. Examples of diffuse field and front-source equalized data are shown in Fig. 4.2–Fig. 4.4 with the corresponding equalization functions.

It can easily be seen that the free-field equalized data show little agreement between the pinna mold and real data set in both level and shape, while for both the front-source and diffuse field equalization there is some definite low frequency agreement. This is expected, since for frequencies with large wavelengths (below a few hundred Hz) there should be minimal difference between the real subject and the mannequin head and torso.

Exact interpretation of the two equalization methods is important to understand when comparing data and choosing which method to use. The diffuse field method compares the response at any one source direction or location to the general diffuse/average response for all source positions. Therefore, inherent in the data is an assumption of the general sound field around the subject. When comparing different measurements (or calculations) it must be assumed that the diffuse fields for the different measurements are either identical or at least comparable. The diffuse field should be presented as well to have confidence in such an approach. As shown in Fig. 4.1, the calculated diffuse field varies from ear to ear and person to person on the order of 5–10 dB. Variations due to measurement method, real *vs.* pinna mold, are also evident. When the diffuse fields are different between measurements, interpretation of variations in the data can be difficult.

The front-source equalization method compares the response at any one direction to the ‘normal’ direction of straight ahead. The meaning of this comparison is to examine how, under the assumption that the front source HRTF is universal (at least in idea) do the different locations change with respect to spectral content as a function of frequency. Of course, it is important to compare the front source of the different data sets, just as it is important to compare the diffuse field of various data sets. But, this method makes no assumption regarding the measurement method, such as placement or number of points. The only assumption is that the front source HRTF is ‘properly’ measured. If the source left point (front-source equalized) for two data sets agree, this shows that whatever the general (one could almost say diffuse) field properties are, the directional variations between the two measurements are the same. This implies that the directional component of the HRTFs are comparable. Examination of the unequalized front sources for the two data sets would show possible variations due to other effects.

4.2. Session WPAFB

The equipment used for this measurement was located at the WPAFB Armstrong Laboratory in Dayton, Ohio. The measurement of an individual HRTF in this facility was a very time consuming experiment. For this reason, the laboratory was experimenting with using pinna replicas, made from real individuals, from which the HRTF measurement was made. In order to validate this method, HRTF measurements were made for both real subjects and pinna molds of those individuals.

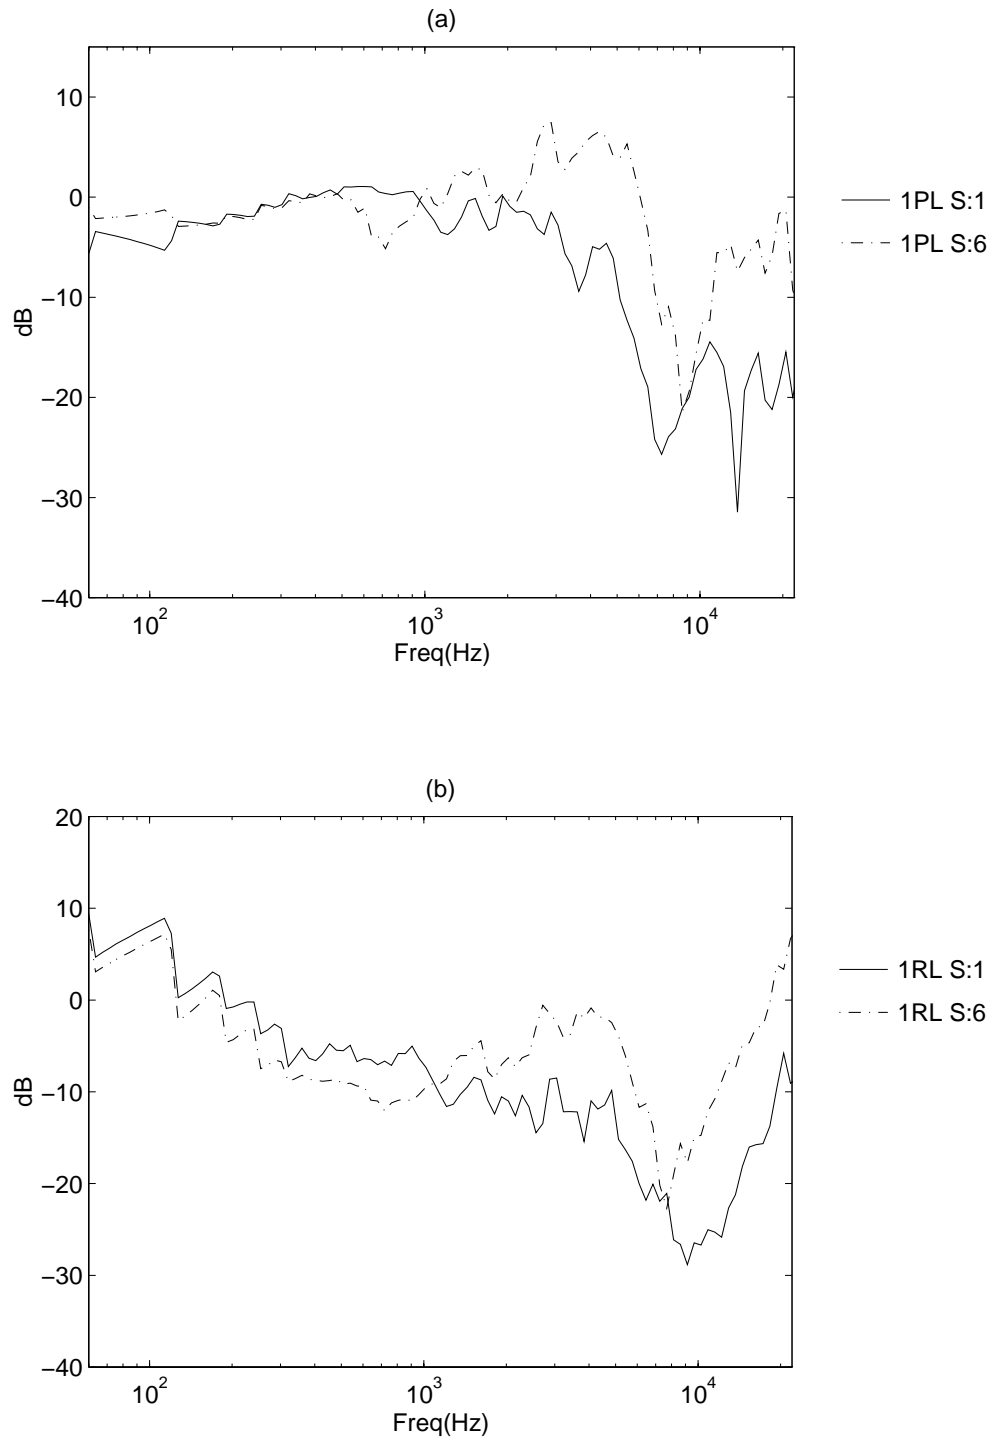


Figure 4.2: Free field equalized data showing sources front (WPAFB Source:6) and right(WPAFB Source:1) for (a) 1PL and (b) 1RL.

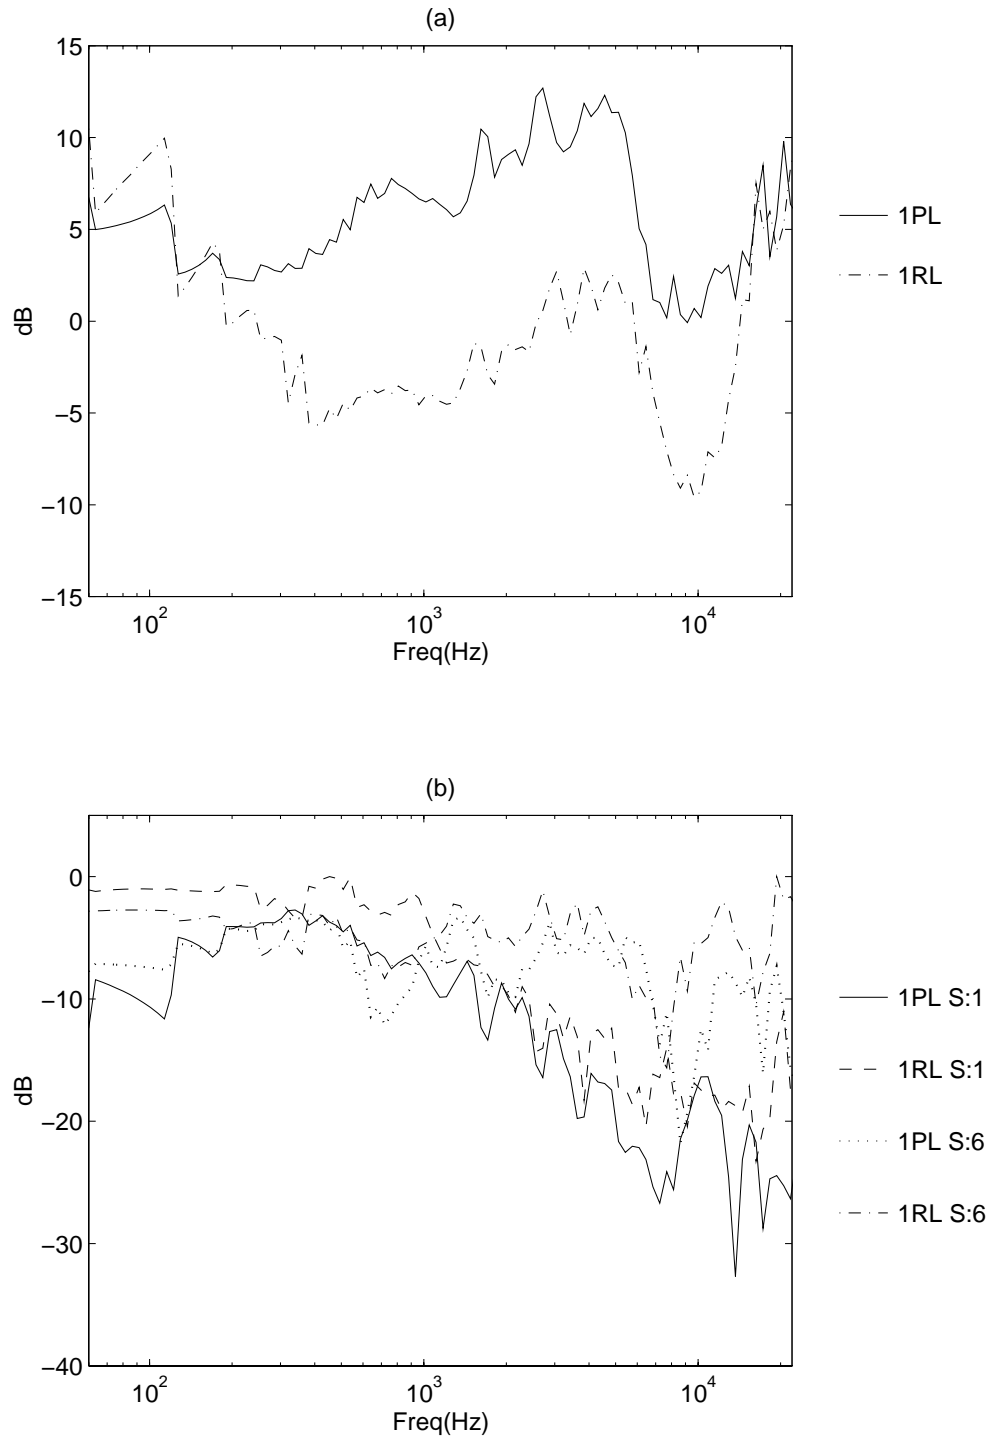


Figure 4.3: Diffuse field equalization examples. **(a)** Diffuse field equalization functions for 1PL and 1RL. **(b)** Diffuse field equalized data showing sources front (WPAFB Source:6) and right (WPAFB Source:1) for 1L.

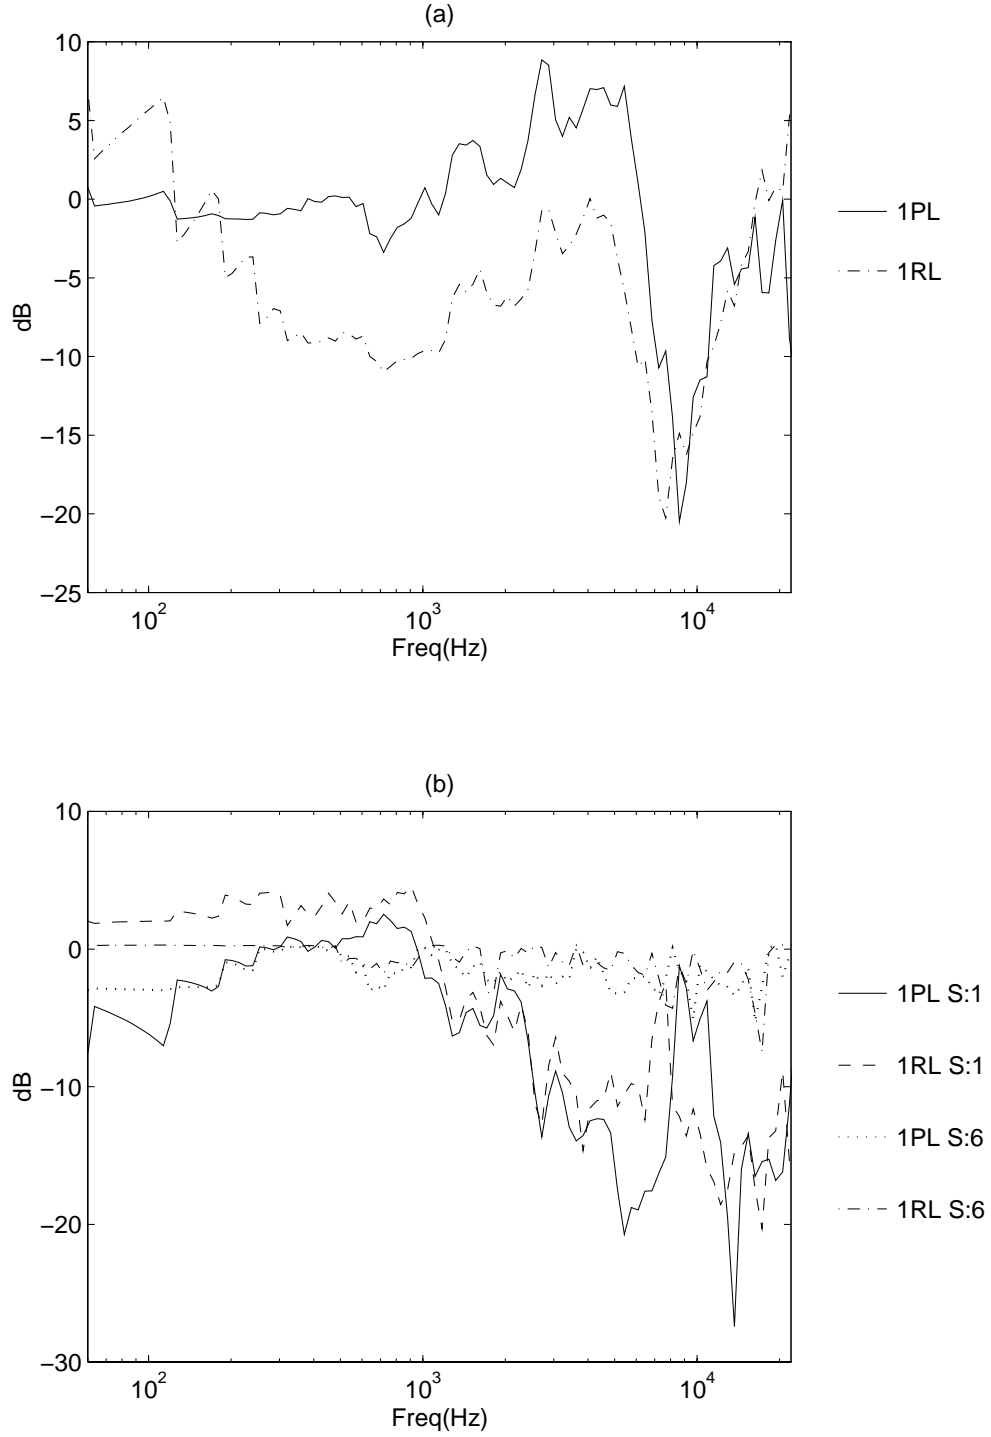


Figure 4.4: Front-source equalization examples. **(a)** Front-source equalization functions for 1PL and 1RL. **(b)** Front-source equalized data showing sources front (WPAFB Source:6) and right (WPAFB Source:1) for 1L.

4.2.1. Apparatus

The measurement apparatus consisted of a spherical array of sources, a pair of receivers, and appropriate hardware. The source array was a dome frame, 4.27 meters in diameter, containing 272 speakers with 10 cm (4 in) diameter. The speakers were located at 15° increments. The dome was located inside an anechoic chamber on a suspended floor. The tubular frame construction of the dome was covered with foam to reduce reflections. Sources were activated using a PC controlled relay unit.


A pair of receiving electret microphones, one for each ear, were each mounted on a small plastic stick. Foam ear plugs were modified to contain a tube in the center. When the foam plugs were inserted into the ear canal, the microphone stick was placed in the tube, thereby mounting the microphone in the center of the ear canal, flush with the end of the foam plug providing for a blocked ear canal measurement (see Fig. 4.5).

The data acquisition was performed using an HP spectrum analyzer to compute the power spectrum of each receiver for any single active source. The sources were controlled using a PC, which controlled the source selection relay system and the analyzer via a GPIB interface. The noise source used was a random white noise signal, and data was taken with 100 averages. The resulting data consisted of a left and right ear power spectrum for each source position, and an interaural time delay (ITD) data set which contained the frequency averaged time delay between the two ears for each source position.

Subsequent to acquisition of the HRTF a calibration measurement was made. This measurement used a rotating apparatus which was placed in the position previously occupied by the head (or subject). The same two sensor microphones were placed on a rotating boom. Each microphone was directed towards each individual source on the sphere, and the transfer function was measured. This transfer function, when used as a calibration function, removed the effects of the electronics, source, and receiver. This resulted in a transfer function which contained the effect of just the presence of the head.

4.2.2. Individual HRTF

For the individual HRTF measurement, the subjects were seated inside the dome. The head was positioned so that the center of the interaural axis was located at the center of the sphere. This was accomplished using two low powered lasers to align the head correctly. The subject's location was stabilized using a chin rest (see Fig. 4.6), thereby keeping the head reasonably immobile during the measurement. This rest was covered with acoustically absorptive foam in order to reduce any effects to the HRTF. Allowing for settling time and data averaging, this procedure lasted about 45 minutes. This time included several rest periods for the subjects. Due to the time involved in this measurement, only a subset of the entire source array was used. Out of the 272, only 70 source positions were measured. These positions consisted of the horizontal plane, the median plane, and three points in the middle of each of the eight octants. This provided ample data for comparison.



a. WPAFB microphone with attached stick and foam plug with mounting tube.



b. Microphone system in use on subject.

Figure 4.5: WPAFB microphone design and placement.




Figure 4.6: WPAFB spherical source array with HRTF measurement subject in position.

4.2.3. Pinna Replica HRTF

Due to the long measurement time for an individual HRTF measurement, an alternate method was employed using pinna replicas made from the individual subjects and placed onto a mannequin head. This method is hoped to provide a sufficient HRTF measurement for use in localization research.

Use of this method requires several assumptions. As the mannequin head is the same for all subjects, it is assumed that the individual specific information required for localization is contained solely in the effects due to the pinna. Any differences in head shape are ignored.[†] It is also assumed that the surface impedance of the pinna replica and mannequin head are either comparable to an actual subject or are of no consequence to the HRTF. The validity of these assumptions is dealt with both through comparisons between the measured data sets and in varying the impedances of the computational model.

4.2.3.1. BAMAR Head

The mannequin used for the pinna replica measurements was the Bio-acoustic Anthropomorphic Manikin for Auditory Research (BAMAR) head built at WPAFB. Material composition of the head consisted of an epoxy fiberglass mixture which resulted in a very rigid and solid structure. The head was comprised of 3 sections, allowing for the interaural spacing (distance between the ears) to be adjusted to that of the subject whose pinna are being used (see Fig. 4.7). This allowed for the interaural time delay (ITD) to be correct for the HRTF measurement. The gaps between the sections were filled with clay to create an acoustically closed surface. Finally, the head is placed onto a torso which simulates the effect of the human torso in terms of reflections from the chest, back, and shoulders. The chin rest from the individual HRTF measurements was removed, as was the chair used by the subject.

4.2.3.2. Pinna Replica Construction

Pinna replicas were built by taking a casting of the subject's pinna and surrounding area. This was done using a diluted solution of dental molding. The ear canals were first blocked with a cotton plug. Then the solution was injected into the ear canal, pinna and surrounding area so that all cavities were filled. The solution was let to rest till firm, about 10 minutes, and was then carefully removed. From this negative mold, a positive mold of the same substance was created. Using the positive a plaster negative was made. The reason for this was that the molding solution used has a limited life span before it deteriorates. From this plaster negative, pinna molds can be made for some time. The cast can also be kept for many years for future study. The final pinna molds were composed of 40% Dow Corning 3110 RTV compound using Type I catalyst, and 60% Dow Corning Silicone Fluid Type 200 [viscosity 10000 cs]. An example of the final pinna mold is shown in Fig. 4.8(a).

[†] The mannequin head used allows for the adjustment of the head width. See section 4.2.3.1.




Figure 4.7: BAMAR mannequin head showing variable width feature.

4.2.3.3. Measurement

Having created the pinna molds, they were then placed into position on the dummy head. The pinna molds were attached using thin velcro straps as shown in Fig. 4.8(b). These were also assumed to have no acoustic effect. The microphones were placed in the pinna molds in the same manner and position as for the individual HRTF measurements. With the mannequin in place in the measurement chamber, the transfer function was measured for all 272 positions, resulting in the completely populated HRTF. As there is no human subject in the test chamber, the signal levels can be greatly increased from the real subject HRTF measurement, providing a much greater signal to noise ratio.

4.2.4. WPAFB Results

Comparisons of the two different measurement techniques can help show the effects of some of the variations due to assumptions made in the BAMAR head/pinna mold construction. These assumptions can be summarized as follows; individual specific characteristics are solely contained in the pinna and the width of the head, head shape (except for width) can be generalized, rigid impedance is an adequate representation, and only the upper torso is required. Though it is difficult to isolate the effects of each of these individual assumptions, comparisons of the two data sets and data between various subject measurements does provide for some information.

Individual symmetry is an initial point of interest, both in the real and the pinna mold measurements. Testing symmetry, as compared between real and pinna mold measurements for the same subject, also shows any variations between the measurement techniques. Analysis of these variations requires reviewing a large amount of data. Several



a. Pinna mold with microphone in place.



b. BAMAR head with pinna molds attached.

Figure 4.8: Pinna mold experimental setup.

source positions and measurement options are shown in Fig. 4.9–Fig. 4.12 for subject 1 and Fig. 4.13–Fig. 4.16 for subject 2. These data have been diffuse field equalized. Details of the equalization and labelling schemes are given in section 4.1 while coordinates for the source locations are given in section D.1.

Variations of the symmetry properties can be seen in several of these examples. In Fig. 4.9 the major null in $1RR$ is at ~ 6 kHz while in $1RL$ and $1P$ it is at ~ 8 kHz. In addition, the agreement between $1RR$ and $1RL$ decreases above 10 kHz, while $1PR$ and $1PL$ agree well into the high frequency range. This was true for most of the subject 1 results. Finally, regarding subject 1 symmetry, in Fig. 4.12 the nulls in $1P$ are at the same frequencies, but the depth varies greatly between L and R , while the first nulls present in $1P$ do not even exist in $1R$. With regards to subject 2, some of these discrepancies are switched. In Fig. 4.13 the major null in $2PL$ is at a higher frequency than $2PR$, while in $2R$ the nulls coincide in frequency. As with subject 1, in Fig. 4.16 the nulls in $2PR$ and $2PL$ do not all coincide and the first nulls in $2PR$ are not present in $2R$.

It is apparent that there is a general symmetry which exists within the HRTF, but there are certain points of interest where they do not agree. In addition, this agreement can not be attributed to the measurement technique (real *vs.* pinna mold). To further examine this question of symmetry, and the variations between real and pinna mold HRTF measurements, it is advantageous to analyze the results for sources in the median plane, 0° azimuth, with varying elevations. Sources in the median plane are considered the most difficult to localize as the inter-aural time delay is zero and, assuming perfect symmetry, the transfer function to each ear should be the same. Examples of two median plane sources, in the frontal direction at $\sim \pm 45^\circ$ from level, are shown for subjects 1 and 2 in Fig. 4.17 and Fig. 4.18 respectively.

From the median plane results it is apparent that there is more agreement within measurement method than between methods, *i.e.* $1PL$ agrees more with $1PR$ than with $1RL$. Fluctuations in null positions exist for both cases. This could be a function of the diffuse equalization method, experimental variations, or the non-symmetric properties of the real subject. To determine the effect of the equalization method, it is useful to use the alternate front-source equalization method as described in section 4.1. The same measurements as shown in Fig. 4.17 and Fig. 4.18 are shown in Fig. 4.19 and Fig. 4.20, although using the front source equalization method instead of the diffuse field equalization method.

The front source equalized data shows much more agreement at lower frequencies between all the measurements and the agreement continues higher in frequency than with the diffuse field equalization. What does not change is the frequency null mismatching that occurs between certain pairs, though the number of mismatches is reduced. The interpretation of the diffuse field *vs.* front source equalized data is important, and is described in detail in section 4.1.

A final numerical analysis considers the variations among subjects for a single measurement technique. This should highlight non-method specific variations, *i.e.* actual subject variations. Pinna mold measurements were made of the greatest number of subjects and therefore are a good data set to use. Secondary to this is the fact that using the pinna




Figure 4.9: Diffuse field equalized HRTF for subject 1 using sources 6 and 7, thereby through symmetry, showing the equivalent same front source for each ear.




Figure 4.10: Diffuse field equalized HRTF for subject 1 using sources 18 and 19, thereby through symmetry, showing the equivalent same rear source for each ear.




Figure 4.11: Diffuse field equalized HRTF for subject 1 using sources 1 and 24 for the left ear and sources 12 and 13 for the right ear, thereby through symmetry, showing the equivalent same opposite side source for each ear.




Figure 4.12: Diffuse field equalized HRTF for subject 1 using sources 12 and 13 for the left ear and sources 1 and 24 for the right ear, thereby through symmetry, showing the equivalent same ear side source for each ear.




Figure 4.13: Diffuse field equalized HRTF for subject 2 using sources 6 and 7, thereby through symmetry, showing the equivalent same front source for each ear.




Figure 4.14: Diffuse field equalized HRTF for subject 2 using sources 18 and 19, thereby through symmetry, showing the equivalent same rear source for each ear.




Figure 4.15: Diffuse field equalized HRTF for subject 2 using sources 1 and 24 for the left ear and sources 12 and 13 for the right ear, thereby through symmetry, showing the equivalent same opposite side source for each ear.




Figure 4.16: Diffuse field equalized HRTF for subject 2 using sources 12 and 13 for the left ear and sources 1 and 24 for the right ear, thereby through symmetry, showing the equivalent same ear side source for each ear.




Figure 4.17: Diffuse field equalized HRTF for subject 1 using source 52 and 57, located in the front median plane, showing the same source for each ear.

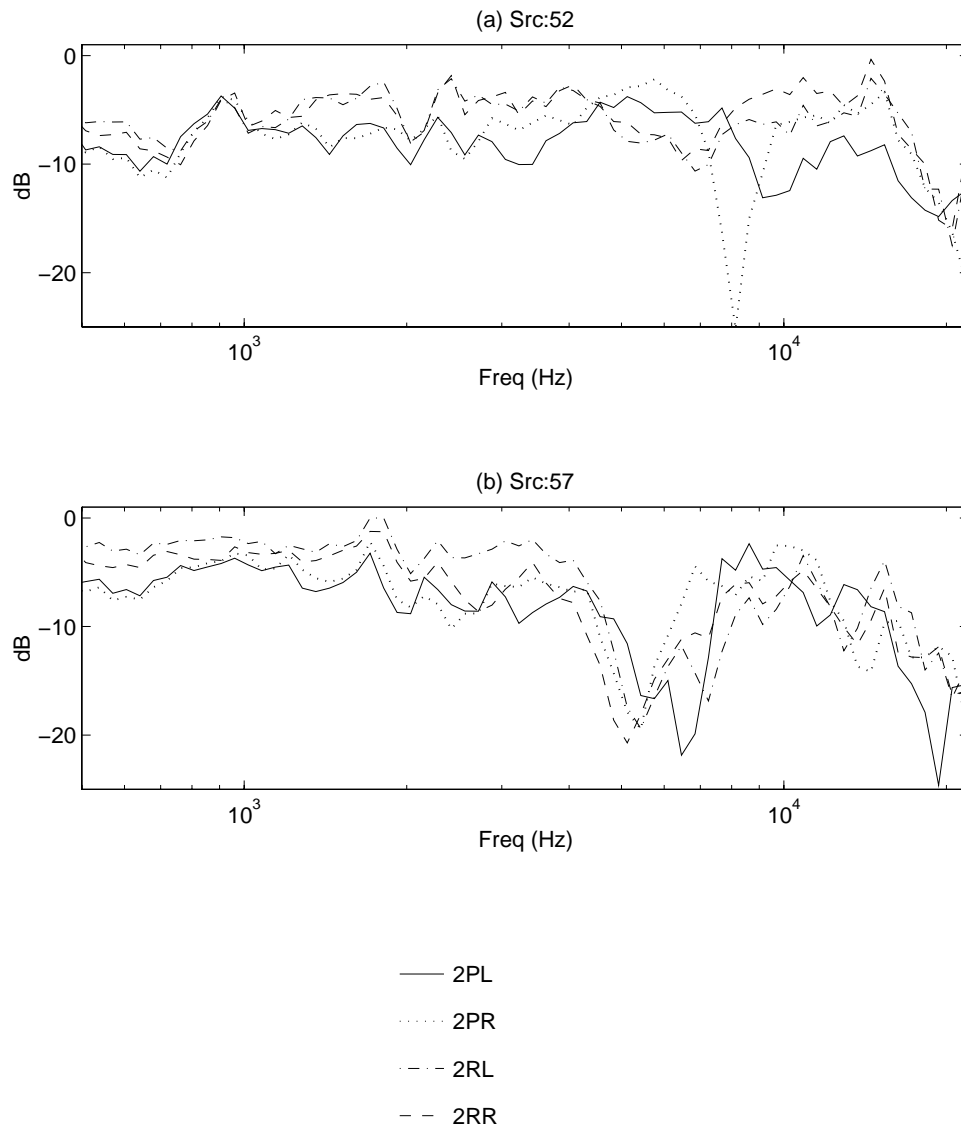


Figure 4.18: Diffuse field equalized HRTF for subject 2 using source 52 and 57, located in the front median plane, showing the same source for each ear.




Figure 4.19: Front source equalized HRTF for subject 1 using sources 52 and 57, located in the front median plane, showing the same source for each ear under real and pinna mold measurement conditions.




Figure 4.20: Front source equalized HRTF for subject 2 using sources 52 and 57, located in the front median plane, showing the same source for each ear under real and pinna mold measurement conditions.

mold measurements for this comparison, it is possible to compare variations between individuals due to the pinna only, as the same mannequin head (allowing for minor width variations which are non-negligible, but common between subjects 1, 3, and 6) was used for all measurements. An example of this comparison is given in Fig. 4.21, using both diffuse field and front-source equalization. These data are for the right ear listening with the source directly to the right. The effect of the pinna should be near its minima for this situation.

From the diffuse field equalized data, the individual data sets agree up to about 1 kHz where variations of 2–3 dB exist. Above 6 kHz the variations increase dramatically, exceeding 15 dB in some instances. The front-source equalized set shows agreement up to 6 kHz with variations of 1–2 dB, at which point similar deviations of 15 dB occur. These results show that for the region up to 6 kHz, even though variations between individuals exists, they are minimized. Once beyond these frequencies though, the individual characteristics based on pinna-only effects vary greatly.

4.3. Session CRE

In order to provide additional comparisons between actual and simulated HRTF measurements, another session for individual HRTF measurements was performed. The equipment used for this measurement was provided and operated by Crystal River Engineering, a subsidiary of Aureal Superconductor Inc., in Fremont California. Aureal is a manufacturer and designer of virtual acoustic presentation hardware and software. Primary interest is in the retail market for auralization systems for personal computers and home theater.

4.3.1. Apparatus

The measurement apparatus, SnapShot, consisted of a single 5 cm (2 in) diameter speaker, a pair of receivers, and appropriate hardware. The source was fixed to a vertical pole, allowing it to be positioned at various vertical positions and tilted towards the subject. The subject was placed on a rotating stool 76 cm (2.5 ft) from the source pole. Using signal processing techniques (described below) the apparatus functions properly in a typical office space, without the need for an anechoic environment. The source position is controlled manually and the data acquisition is performed using a PC and MatLab software. Measurements are performed for six elevations with azimuthal intervals of 30°. Coordinates of the data points is given in section D.2. These data would then be warped to a sphere for use in the auralization hardware, but for data comparisons here was left in its raw cylindrical unmodified form.

A pair of receiving cylindrical electret microphones, one for each ear, were mounted at the entrance of the ear canal using a foam ring to secure the microphones in place. This provided for a blocked ear canal measurement similar to that used for the previous measurement method.

Data acquisition utilized Golay code pairs as the source signal. [48,63] Golay codes are a pair of predetermined digital sequences of short duration which can be used to measure

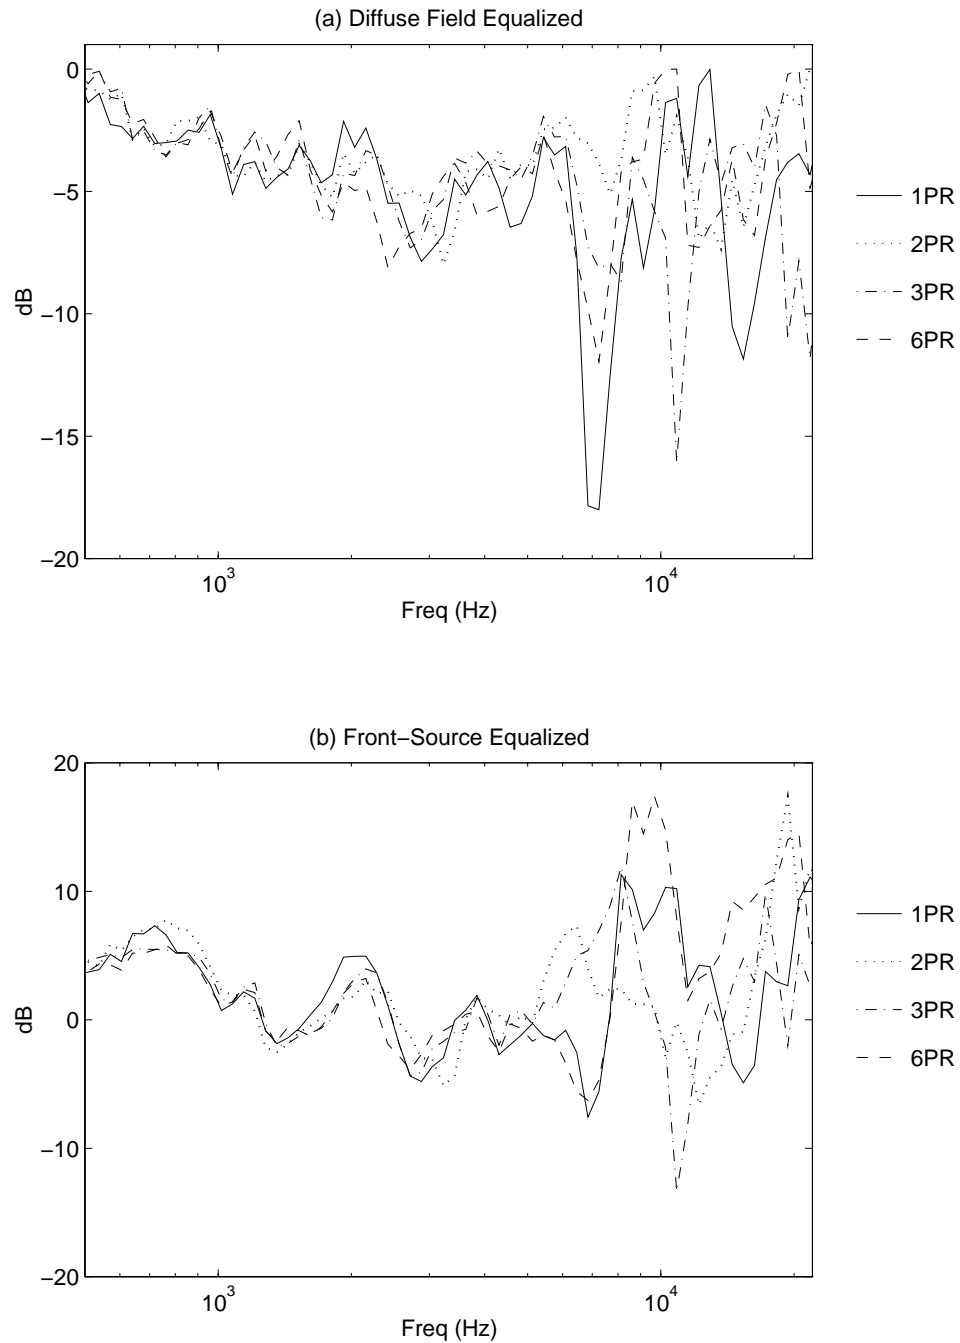


Figure 4.21: Variations between pinna mold measured HRTFs between subjects for the right ear HRTF and source at the right. Results are shown for both diffuse field and front-source equalization methods. Source position is directly right.

the impulse response of a system quickly. This was in contrast to the WPAFB apparatus which used random noise and averaging to calculate the HRTF. The basic property of a Golay code pair is that the sum of their autocorrelation is zero for all τ except for $\tau = 0$, or zero time delay. They are employed in such a fashion that one transfer function using one of the pair is measured, then the second is used to measure a second transfer function. For example, begin with a code pair a_n and b_n . Transforming from the time domain to the frequency domain, the resulting direct measurements would be $F(a_n)H(\omega)$ and $F(b_n)H(\omega)$ where $H(\omega)$ is the frequency response of the system in question and F is the Fourier transform operation. Determining the response of the system would follow Eq. (4.2), where F^* is the complex conjugate of said transform and the codes are each of length L .

$$\begin{aligned} [F(a_n)H(\omega)]F^*(a_n) + [F(b_n)H(\omega)]F^*(b_n) &= [F(a_n)F^*(a_n) + F(b_n)F^*(b_n)]H(\omega) \\ &= 2LH(\omega) \end{aligned} \quad (4.2)$$

The results of this measurement technique are an increase in signal to noise ratio over standard impulse response measurements using true impulses and a measurement time much shorter than standard random noise averaging. For a Golay code pair of length $L = 512$, the gain over a single impulse is on the order of 30 dB. Using a sampling rate of 50 kHz, a single 512 point sequence lasts approximately 0.01 sec. Examining the measured impulse response and limiting the results to the direct sound field, it is possible to calculate an anechoic impulse response in a non-anechoic space, given that there is enough time between the direct and reflected sound. Allowing for settling time of the transducer and decay of the room response, measurements can easily be done in 1 sec. per data point. The limitation of the system, due to such short pulses and the size of the speaker used is in the low frequency region, $f < 1$ kHz. Below this frequency poor results were expected.

4.3.2. Individual HRTF

Comparisons between symmetry and individual differences using this experimental method can be used to help validate the conclusions of the previous WPAFB results. Fig. 4.22–Fig. 4.25 show front equalized data for several subjects for source positions comparable to those in section 4.2.4, including the median plane data. As with the WPAFB results, there are some measurements where nulls and peaks coincide in the symmetry assessment and other times where they do not. In addition, the symmetry measurements agree well up to several kHz, then divergence increases with increasing frequency. It is therefore a reasonably justified conclusion that the variations seen and comments made in the previous section regarding the general trends in HRTF measurements are valid and not measurement type specific.

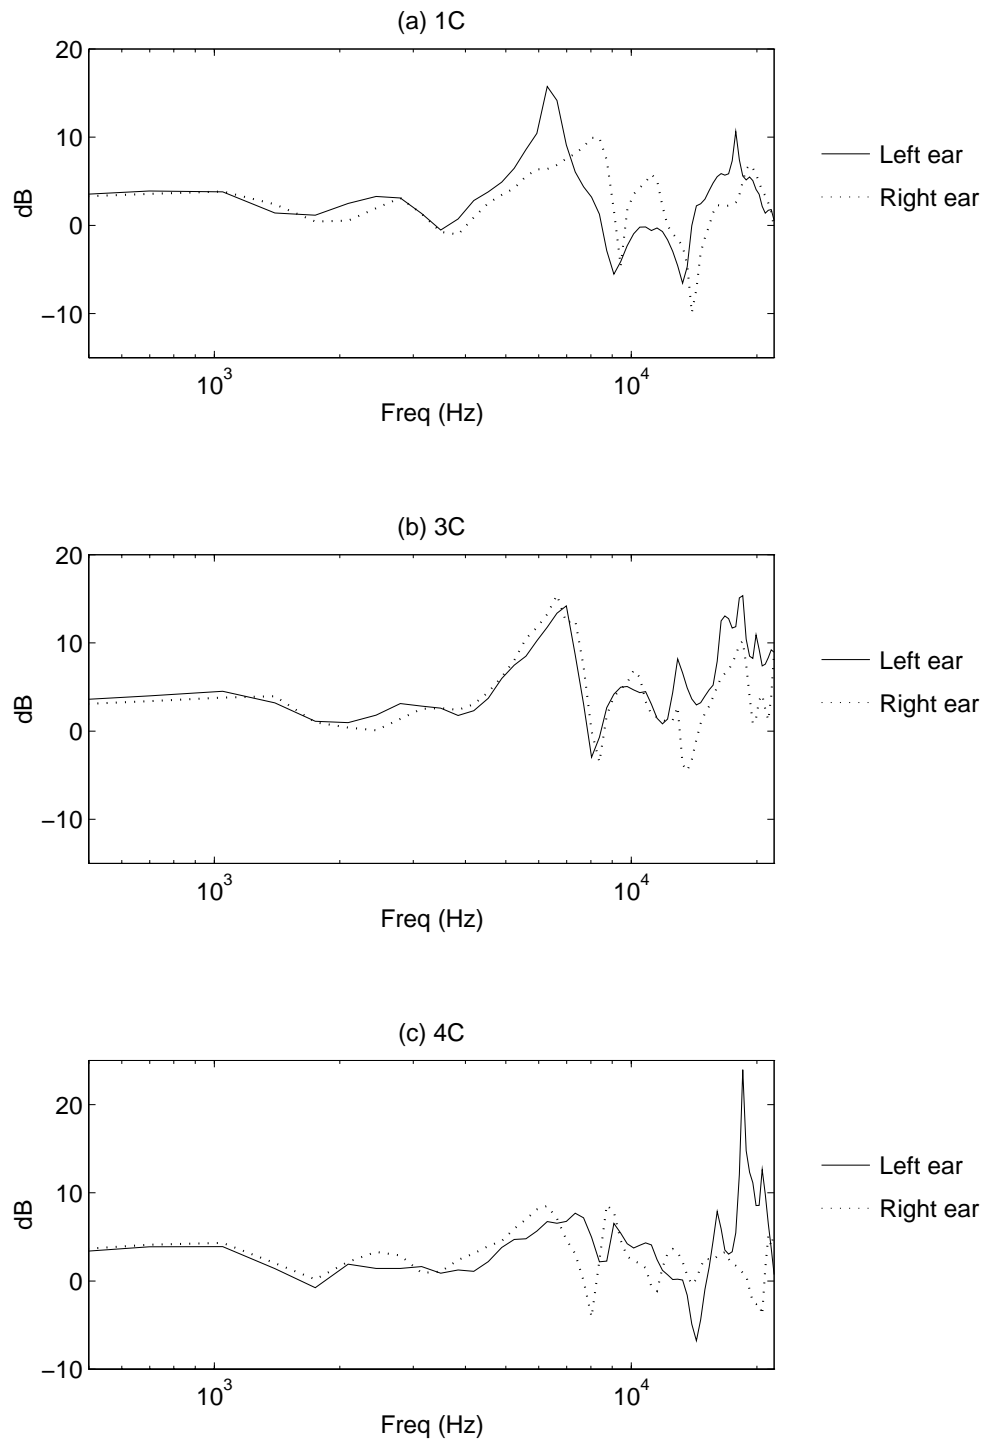


Figure 4.22: Front source equalized HRTF for CRE data for three subjects. Source located on the listening ear side, source 22 for left ear and 58 for right ear.

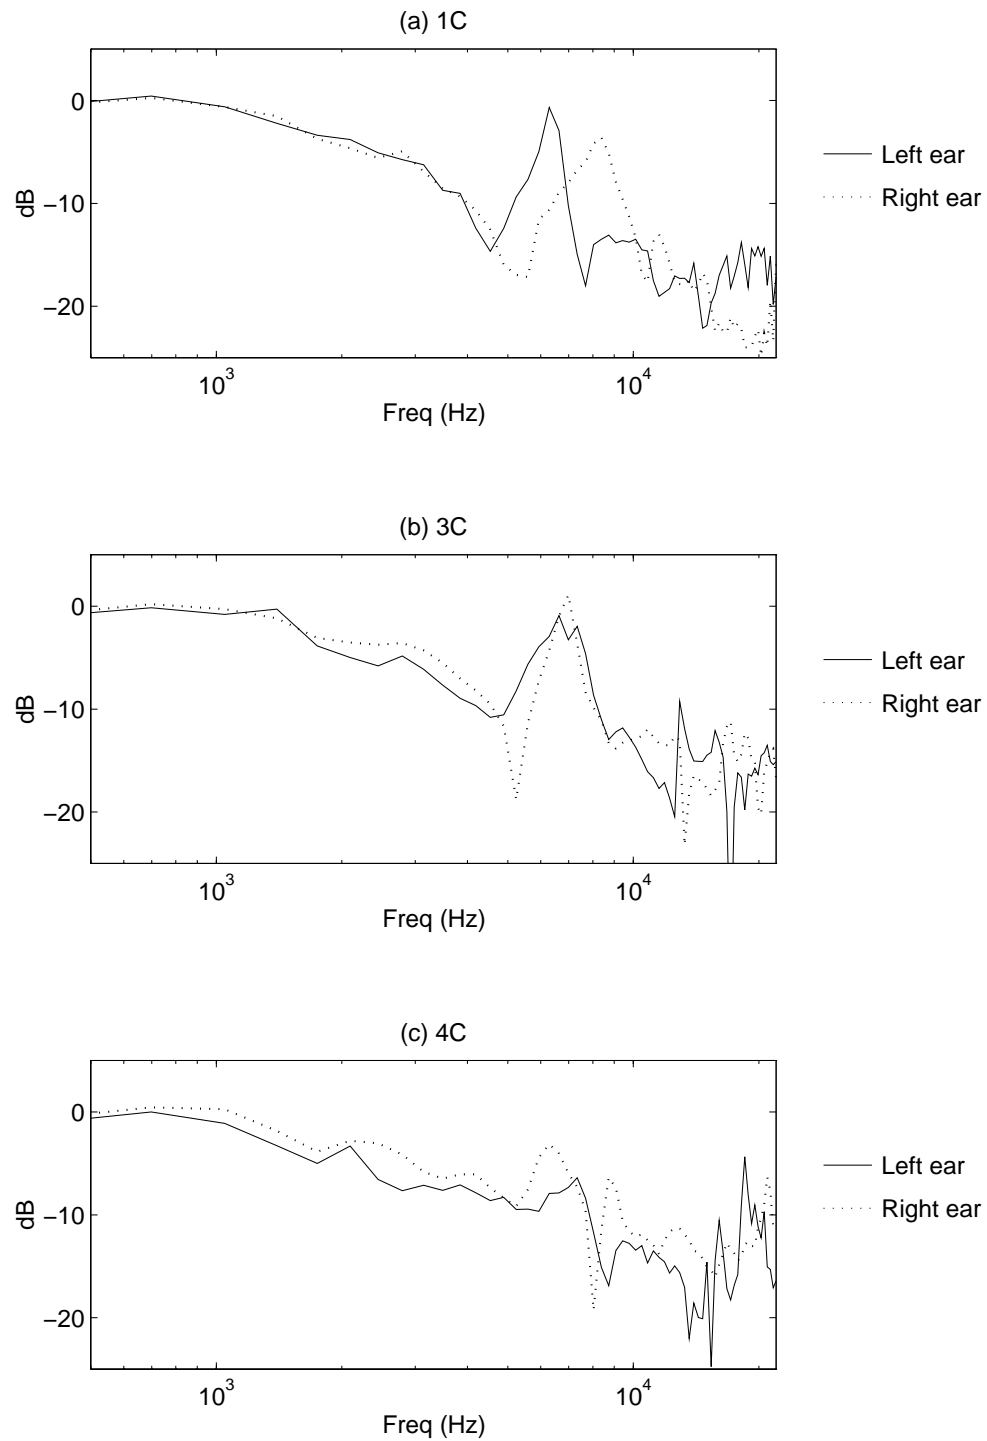


Figure 4.23: Front source equalized HRTF for CRE data for three subjects. Source located on the opposite side of the listening ear, source 58 for left ear and 22 for right ear.

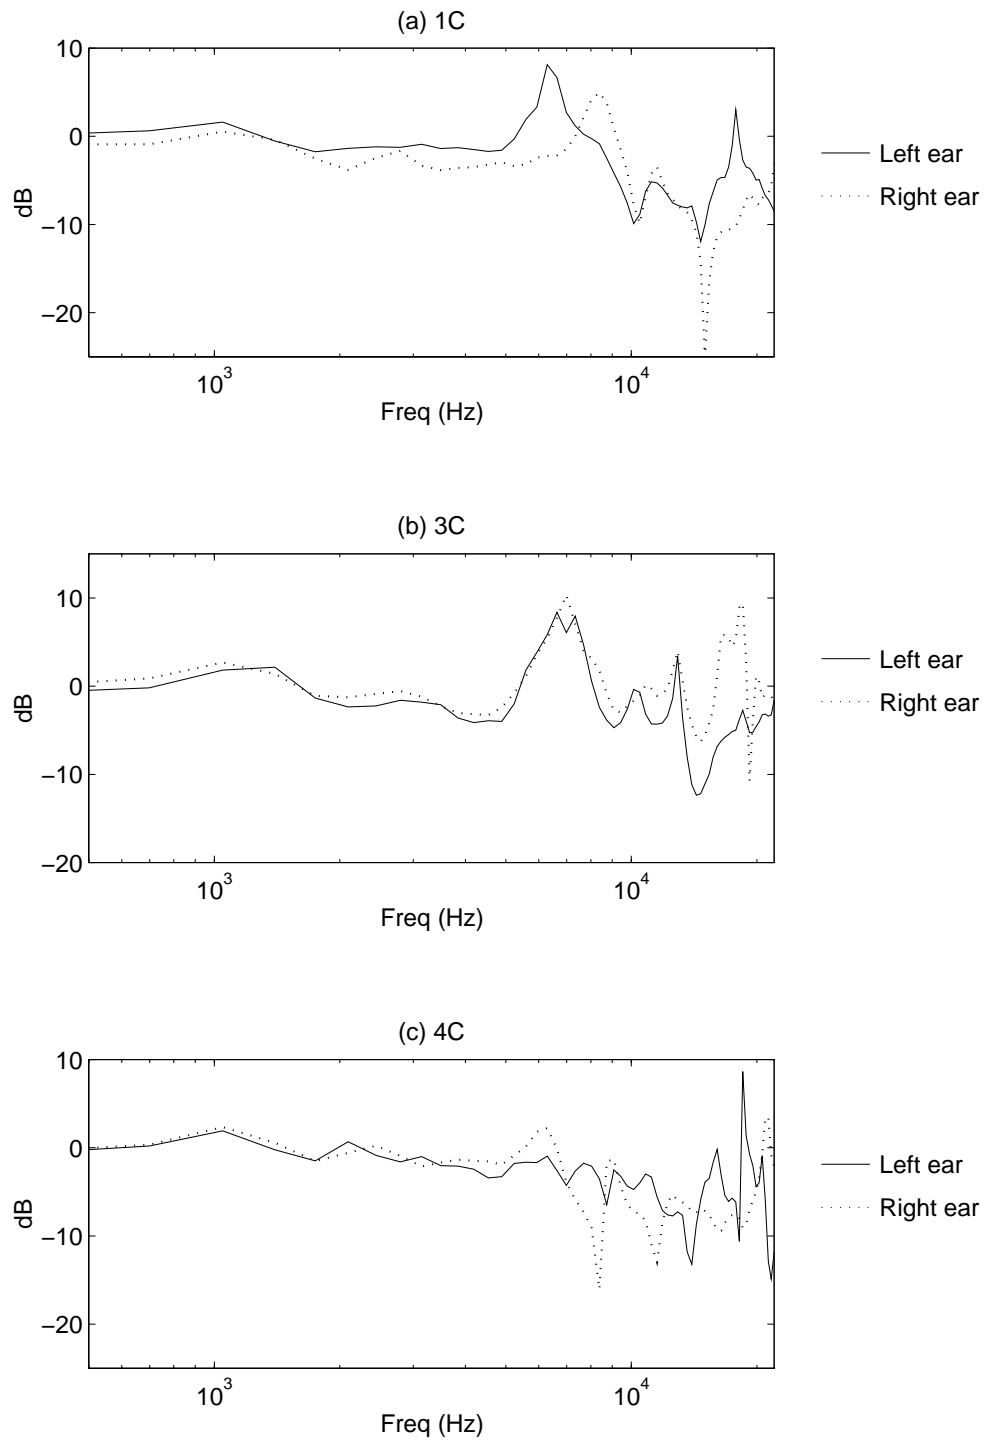


Figure 4.24: Front source equalized HRTF for CRE data for three subjects. Source located at the rear, using source 4.




Figure 4.25: Front source equalized HRTF for CRE data for subject 1. Sources located in the frontal median plane.

4.4. Measured HRTF Data Comparison

Cross comparisons between the three different measurement techniques is the final step in analyzing the experimentally measured HRTFs. Of initial interest is an examination of the resulting equalization functions. These functions, for the left ear of subject 1, *1L*, are shown in Fig. 4.26. The free-field equalized transfer function of the front source, used in the front-source equalization method, is shown in Fig. 4.26(a). The *1PL* and *1RL* have similar shape, differing by approximately 5 dB below 8 kHz. Above 8 kHz the difference between *1PL* and *1RL* is less. This difference may be due to the lower signal power, and hence signal to noise ratio, of the real subject measurements as described in section 4.2.3.3. The *1CL* has a noticeably different shape, and is on the order of 20 dB lower in the low frequency range. This could be an artifact of both the measurement method (of using short pulses) and the use of very small speakers used which have poor low frequency performance. The region of 4 kHz to 10 kHz of *1CL* agrees reasonably well with the shape and relative level of the WPAFB data sets. At higher frequencies the shape is dramatically different. The reason behind this difference is unclear. For the calculated diffuse fields shown in Fig. 4.26(b) similar observations can be made. Sets *1PL* and *1RL* have roughly the same shape, differing in level, though the difference persists up to approximately 15 kHz. Again, the *1CL* set has a very different shape. It has a much lower low frequency level. At higher frequencies, though it does not match in shape, it does approach the same level range as the other data sets.

Several source positions are analyzed for comparison. Data for source position left, on the listening ear side, is given in Fig. 4.27. For the diffuse field equalized results there seems to be general agreement between the WPAFB data sets. The CRE data set agrees very well at some frequencies, but diverges greatly at other frequencies. In contrast, the front-source equalized data given in Fig. 4.27(b) shows that, within about 5 dB, data set *1CL* follows either *1PL* or *1RL* in magnitude and shape. Below approximately 15 kHz *1CL* agrees best with *1RL* but then follows, in general, *1PL*. There is one major null in *1CL* which is not in the other sets, but the second major null in *1CL* agrees well with the null in *1RL*. Results for source position right is given in Fig. 4.28. Again, good agreement is found between the WPAFB data sets in the diffuse field equalized case, but little agreement is found with the CRE data set. The front-source equalized data shows similar results again for all three data sets, but with peaks occurring at different frequencies.

Considering the rear source position, results are shown in Fig. 4.29. As has been continually observed, for the diffuse field data there is good agreement between *1PL* and *1RL*. Data set *1CL* seems to be a poor match in this case. Considering the front-source equalized data, there is general agreement between all three sets up to approximately 2 kHz, at which point all three sets diverge and there is very little agreement. There is a major peak which occurs in all three, though at slightly different frequencies. Above 10 kHz the data all remain in a 5-10 dB window, following somewhat similar trends.

The final source position analyzed is the source directly right, on the opposite side of the listening ear. Results are similar to the rear source position above. For the diffuse equalized data set there is general agreement between *1PL* and *1RL* and very little agreement of these with *1CL*. In this source position, *1CL* diverges a great deal above 2 kHz. Front-source localized data shows general agreement again between all sets up to about 5 kHz.

The major null exists, again at different frequencies for each set. In the region of 8 kHz and above, $1CL$ agrees quite well with $1RL$ and $1PL$, sometimes better with one than the other.

From the results of these experimental measurement methods of HRTFs several comments can be made. Using a given measurement method, the variations between subjects is definite. The basis of these differences can be analyzed by using a pinna mold measurement technique which removes variations due to body shape and head shape variations, except for head width which was adjusted to match the original subject. Comparing both individual source location measurements and the diffuse fields for these variations shows that the system is very sensitive to changes. Even though the diffuse field is assumed to have no directional characteristics, it is not subject independent, even when using pinna molds. Therefore, it can be said that even though direction dependent information is not present (removed using equalization), there is still frequency filtering occurring. This remaining filtering on the basis of left/right side could be due to varying amounts of asymmetry between the pinna. These variations in the HRTF, even though direction independent, are side dependent and therefore contribute useful information which the brain will process and potentially use in localization.

In comparing the different experimental methods it is difficult to determine if some of the variations are due to experimental method or experimental error. The fact that frequency peaks and nulls do not always coincide, and that for a given pair of subjects and a variety of the measurements there is no pattern, leads to the conclusion that there is more in experimental variation error than is described in the literature. By comparing different methods it has been shown that the diffuse field equalization method does not seem to retain enough information which is portable between different measurement techniques and data sets. The front-source equalization method shows better agreement between methods and subjects. In direct interpretation of the front-source equalized data it is necessary to conceptualize the exact meaning of the transfer function. Whether or not this is necessary depends on the goals of the analysis. The interpretation used here for the front-source equalization method is that the results identify changes in the transfer function as one moves away from the front direction. This should highlight any variations in the HRTF with respect to position. In addition, when comparing different measurement schemes the results focus on the direction-based variations of the HRTF and not on global trends throughout the data, which could be caused by particularities of the experimental method.

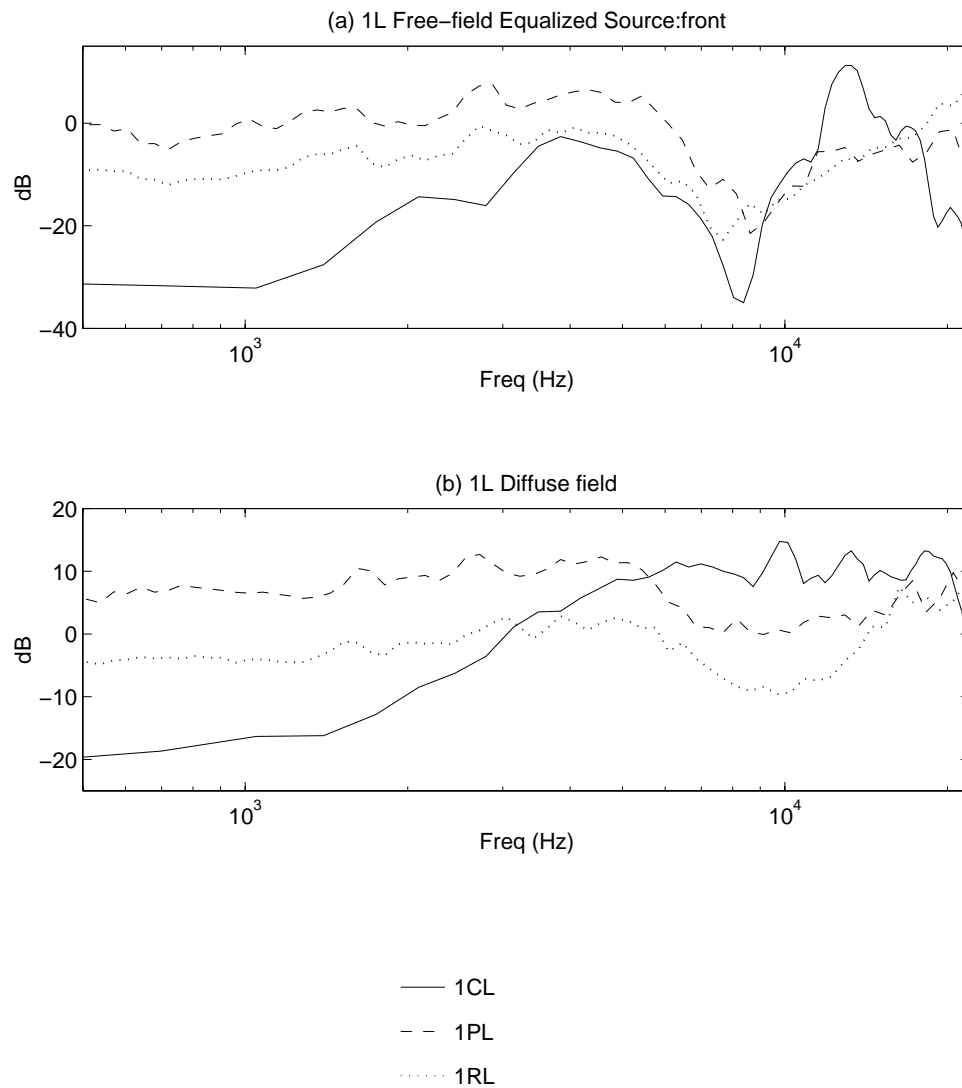


Figure 4.26: Equalization functions for the three experimental method data sets for 1L. **(a)** Free-field transfer function for front-source, used in front-source equalization method. **(b)** Diffuse field calculations for diffuse field equalization.

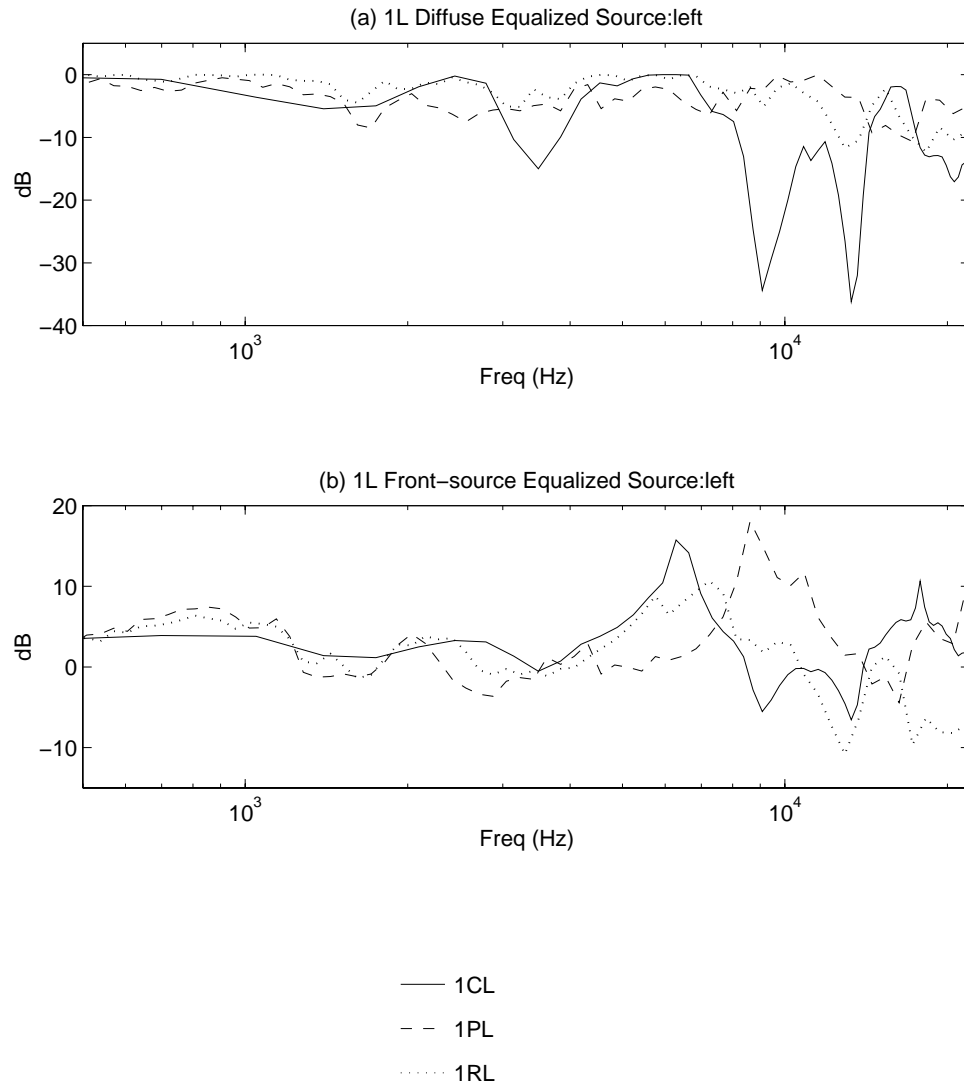


Figure 4.27: Experimentally measured HRTF for source left with data sets $1L$ for each experimental and equalization method. Sources were $1CL : 22$, $1PL : 1$, and $1RL : 1$.

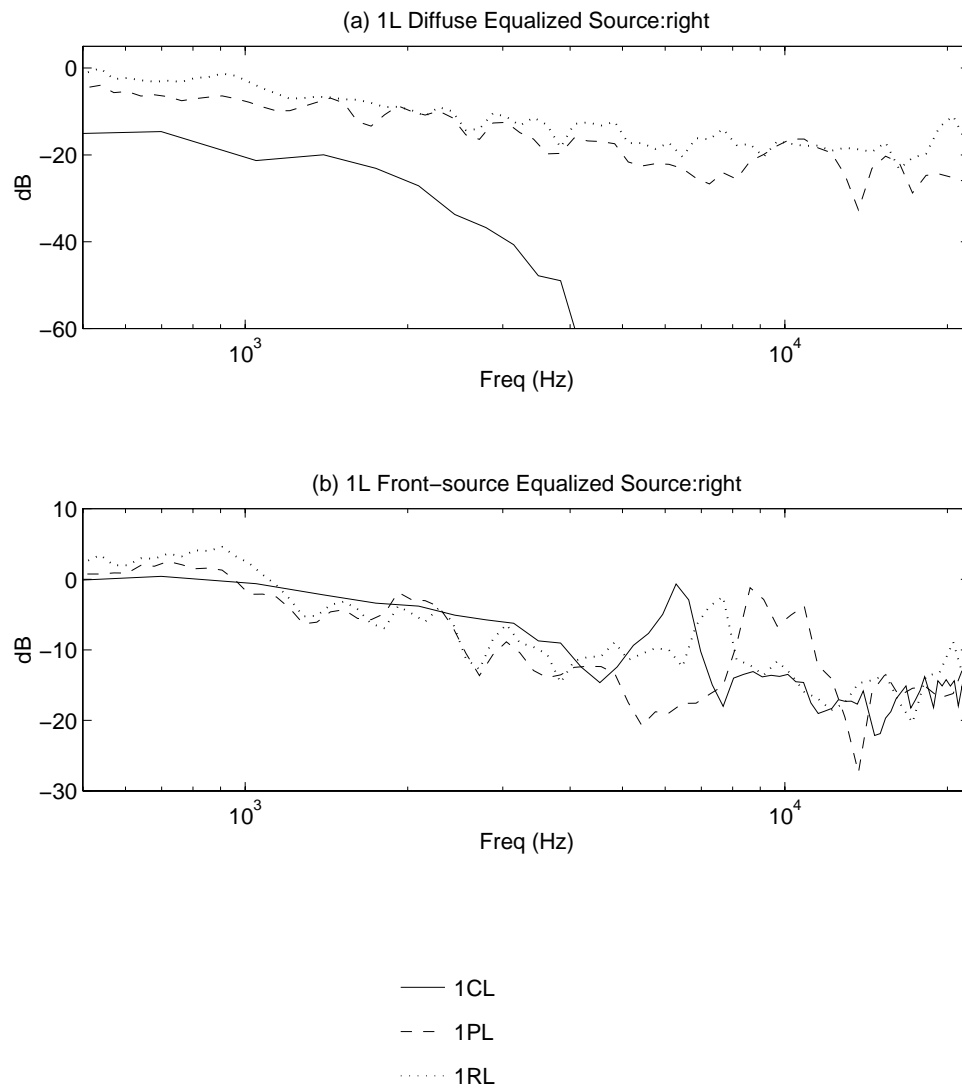


Figure 4.28: Experimentally measured HRTF for source right with data sets $1L$ for each experimental and equalization method. Sources were $1CL$: 58, $1PL$: 12, and $1RL$: 12.

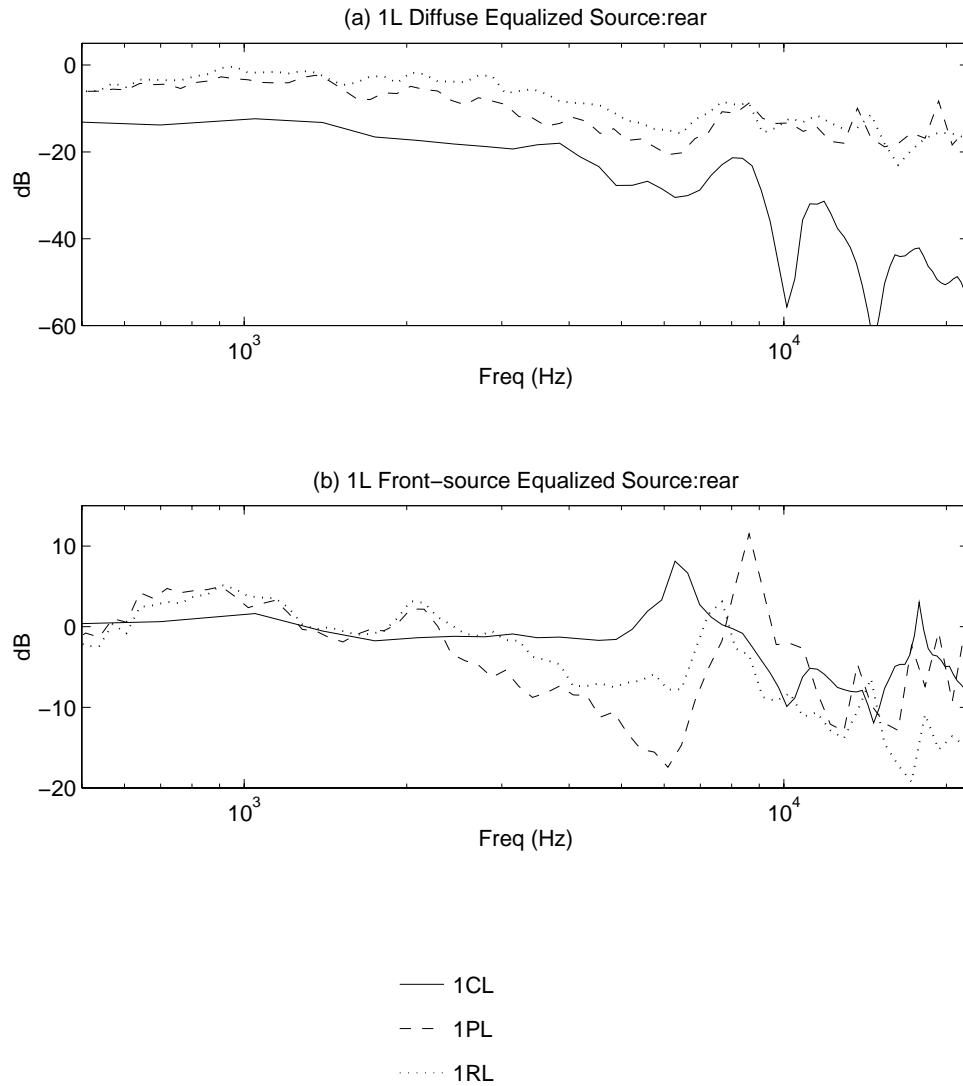


Figure 4.29: Experimentally measured HRTF for source rear with data sets $1L$ for each experimental and equalization method. Sources were $1CL$: 4, $1PL$: 18, and $1RL$: 18.

Chapter 5.

HRTF Calculation

The central work of this research consisted of the computational generation of an HRTF involving the construction of a computer model of an individual from which a simulation of the actual measurements performed in the previous chapter could be made. The method used in this research was to consider the human head as a surface only, with respect to acoustical propagation, and that this surface could be represented by a boundary element model. From this boundary element model, the transfer function can be computed for any given receiver point. In this way, given a boundary surface model of the individual, the HRTF can be computed for any set of source locations.

5.1. Boundary Element Method Approach

Calculation of the HRTF is made on the assumption that only the surface characteristics of the head are pertinent, propagation through the head structure is ignored. In determining a method for calculation, several options are available which all vary in complexity and computation requirements. The method chosen here, which is potentially the most efficient manner for this type of problem (*i.e.* scattering) is the *Boundary Element Method* or BEM . The BEM represents a surface as a mesh of discrete elements. These elements, and their size, limit the interpretation of the results and frequency range of validity.

The act of discretization converts a smooth continuous surface to a set of small planar elements (in this scheme). The size of these elements directly relates to the upper frequency bound for which the discretized surface mesh is valid. Generally it is suggested that six elements are required to accurately represent a period of an acoustic wave, though in the limit four is possible. The frequency bound is then determined by the largest element such that six (or four) element edges would equal one wavelength; $f_{max} = c/(6 \times edge_{max})$, where c is the speed of sound.

There are two different computation methods which are employed in the BEM, direct (or collocation) and indirect (or variational). The direct method uses pressure and normal velocity on the boundary surface as the unknowns. The indirect method uses a representation of the pressure and velocity on each side of the boundary, pressure discontinuity (pressure jumps, double layer potentials) and normal velocity (velocity jumps, single layer potential), as the unknowns. The direct technique uses a non-symmetric matrix of equations while the indirect technique uses a symmetric set. Determining which method to use is a function of problem size, computing power, and computing hardware. These issues, and their dependence on solution method are described in section 5.3.1 where the justification is made for using the indirect approach.

The indirect approach incorporates the interior and exterior of the boundary mesh. The indirect method can be described in the following way.[†] Using the exterior side normal, n , to define direction one can define p^+ as the pressure on the exterior side of the surface and p^- as the pressure on the interior side. The normal derivatives of the respective pressures are $\delta p^+/\delta n$ and $\delta p^-/\delta n$. As mentioned above, the indirect method utilizes pressure and velocity jumps. It is therefore necessary to make the following definitions where μ is the double layer potential and σ is the single layer potential.

$$\mu = p^+ - p^- \quad (5.1)$$

$$\sigma = \frac{\delta p^+}{\delta n} - \frac{\delta p^-}{\delta n} = -i\rho\omega(u^+ - u^-) \quad (5.2)$$

In order to solve the set of equations which result in the values for the layer potential sources, the boundary conditions on the surface must be defined. The boundary condition can be defined by describing the pressure on the surface, normal velocity on the surface, or normal impedance on the surface. The final option is the method used in the HRTF calculations as the desire is to determine the scattering effect around the head and the effects of skin and hair impedance.

The discretization of the surface and the layer potentials can be regarded as a monopole (single layer) and dipole (double layer) distribution which, when given the appropriate values reproduces the scattered field by the boundary ‘surface.’ The acoustical variables in the fluid (*i.e.* pressure, velocity, intensity) can be determined by summing, through superposition, the solution set of representative sources. The pressure at a point away from the surface can be calculated using Eq. (5.3), where \oint_s is the surface integral over the boundary surface s , \bar{x} is the field point, \bar{x}_s is the source location, and $G(\bar{x}|\bar{x}_s)$ is the free space Green’s function.

$$\hat{p}(\bar{x}) = \oint_s \left[\mu(\bar{x}_s) \frac{\delta G(\bar{x}|\bar{x}_s)}{\delta n} - \sigma(\bar{x}_s) G(\bar{x}|\bar{x}_s) \right] ds \quad (5.3)$$

In the BEM there exists the potential for errors regarding the exterior problem where the solution can be non-unique at some frequencies. The error arises at resonance frequencies of the internal problem. These frequencies, in the BEM, are referred to as irregular frequencies in the solution. To reduce the contribution of the internal resonances the methodology is to prescribe finite impedances on points or elements (depending in solution option) within the interior of the surface. For the indirect method these are referred to as overdetermination elements (overdetermination points for the direct method). The elements are intended to suppress the internal resonances of the model. The number and placement of these elements is a function of the geometry of the mesh. An unsymmetrical mesh requires few overdetermination elements to suppress the internal resonances, while a very symmetric mesh may require many internal elements in order to suppress all the of these irregular frequencies.

[†] The description of the BEM method presented here in this section is primarily derived from the Sysnoise Manual [50].

5.2. BEM Model Generation

The primary component of the calculation is the computational mesh. This mesh represents the physical surface in the calculation. The goal is to obtain a mesh which resembles the physical structure to be modelled as closely as possible, while still being of a manageable size. The size of the mesh determines the speed of the calculation and other hardware requirements such as memory and available disk space. The structure of the mesh, element size, also determines the upper frequency for which the subsequent calculation is valid. The result of the BEM HRTF calculation is a balancing of these factors: mesh accuracy, size, homogenous discretization.

5.2.1. Head Geometry Acquisition

Obtaining the geometrical data of an individual was the first step in designing the boundary element model to be used in the calculation. There are a number of different methods which exist for gathering spatial data for use on a computer. A 3D mouse or pointer, which functions using either a magnetic or radio transmitter/receiver pair, can be used to input any selection of 3D data points which can then be constructed into a surface. This method works well for simple structures, where few data points are needed; otherwise the procedure is very tedious and difficult. A solid cast could be made of the object which would then be sliced into very thin layers, each layer being scanned with a 2D optical scanner. The scan would then be processed to a set of points only on the boundary. Then all the layers would be reassembled into the boundary surface. This is also a time intensive method for gathering data when fine resolution is desired. In addition, the construction of the surface from the various layers is not trivial. This method is used, in theory, by many medical scanners. These scanners take layered images of subjects, processed as optical images, which are then converted to interior surfaces. This procedure is rather elegant but requires very specific equipment and personnel; it is not necessarily trivial. The advantage of this method is that the result is a very detailed 3D image of the individual which includes any obstructed, difficult to reach, or hard to see areas. It can also include the more interior structures of the ear canal which would be extremely difficult to map with a physical or optical probe.

A fourth method, and the one used here, involves using a laser scanner to obtain 3D data from the surface of the object. For the data collected, a Cyberware Head Scanner (Cyberware Color 3D Scanner #4020 RGB/PS-D) was used, courtesy of Wright-Patterson Air Force Base. The system operates using a vertical linear transmitter/receiver array which is rotated about the subject, resulting in a cylindrical data set containing the spatial information of the subject (in addition to a color map which was discarded for this research). There are a number of limitations with this system. The primary limitation of the system was that it utilized an optical laser. This requires the object to be somewhat reflective. The major problem area for this was hair. It was necessary for the subjects to cover their hair with a rubber cap, providing for a reflective surface. Due to the linear design of the system, the data was restricted to line-of-sight data in a radial direction about the head. Because of this, any occluded areas were not able to be measured. The

result of this was that areas such as the space behind the ears and detailed folds of the ears were seen as solid. This effect can be seen quite clearly in Fig. 5.1 and Fig. 5.9.

The ear canal was also seen as blocked, as no light reflected from it. As the system relies on reflections back towards the array, surfaces that were almost perpendicular to the array are difficult to view. The primary surface where this occurred was the top of the head, which is left open, as is the bottom of the scan as the subject continues to exist beyond the range of the system. Due to the size of the array, only the head and neck of the subject were able to be scanned. This is also a limitation as it is believed by many that the shoulders and torso (and even legs when seated) all affect the HRTF to some noticeable extent. Due to the enormous data set such a model would include, and the time necessary for such a computation (see section 5.3.1), these effects are not considered at this stage of the research. It was acknowledged that this is a possible deficiency but was a required assumption. Finally, the typical use for this system is for graphical applications. For this reason, most of the data processing software was designed to optimize graphical rendering, *i.e.* fewer elements of greater size without need for strict element connectivity. These issues required the need for complex modifications to the mesh to be performed before any calculations could be done, as described in section 5.2.2.

The result of the optical data scan was a cylindrical shell, open at both ends, which was a surface model of the subject's head and neck. The data set contained approximately 300,000 triangular elements with edges approximately 1 mm in length on average. The resolution in the vertical direction was also much greater than the horizontal direction, making the elements typically elongated. Allowing for six nodes per wavelength, this model would be valid up to approximately 57 kHz, well beyond the audible frequency range. Solving this problem would also take an enormous amount of computer power, storage, and time (see section 5.3.1).

5.2.2. Mesh Modification

Several modifications were necessary to make the mesh suitable for the HRTF calculation. To make the mesh usable by the boundary element software, it was necessary to have a closed mesh. This required closing the top and bottom of the cylindrical scan mesh and sealing any holes that may have been present due to the lack of element connectivity, *i.e.* two adjacent elements not being connected as duplicate coincident nodes exist, one for each element. Resolution reduction of the mesh was necessary to make the frequency range limited to the audible spectrum and the problem realistic in size. The mesh also needed to be refined in some areas where large elements existed due to the apparent elongation of many elements due to the scanning method. This also led to the goal of creating elements more equilateral in shape, which is advantageous to the boundary element solver. As no software existed to perform the changes necessary to the mesh, several programs were written to deal with these specific tasks. The code for these programs is given in Appendix B. The primary code used was `decimate.c` (see section B.1), which contained a large number of mesh modification subroutines which were executed in various order to obtain the final mesh state.




Figure 5.1: Cylindrical scanner omission effect. Surfaces for which the cylindrical scanner can not obtain a reflection from, or around, are closed or filled according to line-of-sight towards the center of the head. The effect is most prominent for the area behind the ear, which is interpreted as being filled solid.

5.2.2.1. Data Format

The raw mesh format was designed for graphical representation of the data, providing for many more parameters than were necessary for this research and the boundary element solver. These parameters included element normals, texture, and lighting conditions. The data format used by the solver contained a good deal of extraneous and redundant information, such as multiple numbering schemes and element type identifiers. When this format was used it greatly increased the data read and write time when performing the mesh modifications due to the size of the data mesh. For these reasons, and for additional bookkeeping and error checking, an intermediate data format was used during the mesh modification procedure. This explains the various file read and write formats presented in `decimate.c`.

5.2.2.2. Closure

Creating a closed mesh (a solid surface with no holes) from the raw data required several procedures. The first requirement was to remove any coincident nodes which from visual inspection appear to be the same node, but structurally are a hole in the mesh. There was no clear method to determine which nodes were coincident with other nodes other than a brute force approach. Therefore, the distance from every node to every other node was evaluated. If the distance was too small (determined to be 0.015 cm) then the two nodes would be merged. This procedure (`decimate.c: check_nodes`) also eliminated any elements which were very small, *i.e.* where an edge falls within the determined tolerance. It was important during this and subsequent procedures to keep very good track of which nodes are connected to which elements, and vice versa. A data array with this information must be maintained and updated for every modification made, otherwise the mesh can easily be corrupted. Corruption of a mesh could be the creation of more holes, overlapping of elements, or distortion of the mesh shape. Connectivity of nodes to elements is contained within the standard data format, but the reverse is not. This reverse connectivity is extremely useful in making changes to the mesh, especially when trying to determine neighboring elements for selected elements. Without this list, it would be necessary to search through the entire mesh every time a neighboring element was needed. Using `decimate.c`, this connectivity list was generated when the mesh was initially read in.

After resolving coincident nodes, final closure involved closing the top of the head and the neck. This was accomplished by choosing cutoff planes for both the top and the bottom at which the closure would be created. From this plane, all points either above or below were brought to the cutoff plane and then connected together. This created a plane with very long narrow elements all directed towards the center (arbitrarily chosen point). To maintain the frequency range of the model it was necessary to refine these large elements into smaller, more equilateral elements.

5.2.2.3. Coarsen Mesh

Coarsening of the mesh is the process of reducing the number of elements (and nodes) in the mesh to make the problem a manageable size, and to remove excessive resolution. Due to the nature of finite element solutions, the frequency limitation of the results is a function of the size of the individual elements. Particularly, the results are not valid when the largest element edge is a quarter wavelength at maximum. It is typically suggested that the edge be at least $1/6$ wavelength. The raw mesh is valid (not including the closure areas at the top and bottom) up to approximately 57 kHz and has in excess of 300,000 elements. It was desired to reduce the mesh to no larger than 30,000 elements. The size of an element was evaluated by determining the length of each side. If any side was smaller than the selected tolerance, that element was removed. The removal of an element is not a trivial matter in terms of connectivity and adjacent elements. It was desired that upon removing an element, there would be the minimal amount of change to the mesh in general. Two methods were considered in removing an element. The first method was to collapse the element, moving all three vertices to the center. This method removes the element in question, and the three elements which shared sides with the element in question. It was decided that this method distorted the mesh too much. The second method, the method chosen, collapses the edge in question, moving the two end points to a single midpoint. The result of this is the removal of the element in question and the single element which shared the small edge.




Figure 5.2: Mesh coarsening methods. **(a)** Raw mesh with selected element and edge to eliminate. **(b)** Element collapsed to center. **(c)** Edge collapsed to midpoint.

It was observed that if the coarsening procedure progressed through the elements sequentially, *i.e.* neighboring elements were consecutively removed in the early stages of the procedure, the mesh would be corrupted and greatly distorted. This was due to the fact that early in the procedure so many elements needed to be removed. The result was a mesh that very much was a function of the numbering scheme, which should not be a factor in the mesh shape. The solution was to loop through the elements several times, skipping several elements while progressing through the mesh, each time through the loop skipping

fewer elements until the final pass in which no elements were skipped. This solution solved the problem and worked quite well. This method was used for many of the procedures which affected large numbers of elements.

Since the pinna are very complicated and probably the most crucial structures in the HRTF mechanism, they were excluded from the majority of the coarsening procedures. Coarsening was typically done in several stages, with the minimum size increasing each time, and the model checked periodically to ensure no distortion was occurring. The resulting mesh therefore contains a much greater amount of detail and number of elements in and around the pinna region. This can be seen in the final mesh shown in Fig. 5.8–Fig. 5.9.

5.2.2.4. Refine Mesh

The closure of the top and bottom of the head created a number of large elements. The elements reduce the maximum frequency for which the model is valid in the solver. It was therefore necessary to refine these elements into more numerous smaller elements. A search procedure similar to the one used for coarsening the mesh was used to determine which elements needed to be refined. The length of each side of every element was calculated, and if it exceeded the maximum allowed length the element was refined. Setting the maximum length close to the minimum length used for the coarsening procedure results in the mesh being composed of somewhat regular, more equilateral, triangles. This is advantageous for the solver, in that the elements are more regularly spaced.

The method by which an element was broken into smaller elements was a matter of consideration as it was with element removal. Two possible methods for refining an element were evaluated. The first method determined the center of the element triangle. From this point, the triangle was divided into three new triangles, using the center point as a new connecting node. The difficulty with the method was that a long skinny triangle when refined in this way results in more long and skinny triangles. This is the situation that was undesired. The second method for element refining was much the same as the selected coarsening procedure. The method by which an element was refined was to first determine the midpoint of the oversized edge in question. The neighboring element which contained the same edge was then determined. Using the midpoint as a new vertex, these two elements were divided into four. The new node and elements were then added to the connectivity lists. This method ensured that the elements were made progressively shorter and more regular, instead of thinner while remaining long.

5.2.3. Field Mesh

The field mesh was the set of points at which the transfer function was computed. For comparison purposes, the measurement mesh was generated to duplicate the measurement points for the experimentally measured HRTF. A spherical mesh was constructed which corresponded, in dimension and numbering, to the sphere used at WPAFB. this mesh is shown in Fig. 5.4. The BEM solution calculates the total field at a given point due to all




Figure 5.3: Mesh refining methods. **(a)** Raw mesh with selected element and edge to eliminate. **(b)** Element broken at center. Note: The selected edge still exists. **(c)** Edge cut at midpoint and neighboring elements cut in half.

sources. If the WPAFB setup was reproduced verbatim, there would be a separate calculation for each source location on the sphere (see section 4.2.1). Instead, if the properties of reciprocity are utilized, the source/receiver positions can be interchanged. The result of this switch is only two sources (*i.e.* two calculations) and 272 receivers. There is still some calculation due to each receiver, but these calculations are much faster than the source calculation. The initial calculation comprises solving for the potentials for all the elements such that the source and boundary conditions are met. Then, using the results of these potentials, the pressure field is calculated at each field point through superposition of all the contributions.

This is different than the experimental setup which requires a separate measurement for each direction. It would be theoretically possible to perform the same experimental measurement as the computational measurement except for the proximity of a high powered source to a subject's eardrum, a problem for most human experimentation policies. In addition, there would be the technical challenge of designing a very small yet powerful source which would fit within the ear canal.




Figure 5.4: Field mesh used in BEM calculations which corresponds to the WPAFB experimental setup. Element connectivity is not important, only the nodal positions. Also shown in the head mesh in place.

5.3. Calculation

5.3.1. Calculation Speed

Boundary element computations are very powerful, but they are also quite computationally intensive. The nature of the BEM is the solution of a set of simultaneous equations. Each equation contains a component for every element in the mesh, and there is an equation for each element. Therefore, as the size of the mesh grows, the size of the calculation grows as the square of the number of elements. There are some variations to the solution methods which help speed up the calculations, and each commercial package has their own tricks, but the general rule still applies. Sysnoise was used as the BEM solver for all the HRTF calculations here.[†]


When discussing the ‘size’ of the calculation there are really two separate issues involved. First is the amount of calculation that needs to be performed, a larger problem would need more time. The second consideration is the actual size of the set of simultaneous equations, as it refers to computer memory and disk storage space. As can be expected, as the mesh gets larger, the amount of computational space increases as well as the amount of computational time. The two variations on the BEM solution optimize each of these, but not both. (For details regarding the BEM options see 5.1.)

For problems of any real size the direct method is optimized for speed. It solves a set of full simultaneous equations and in doing so is rather fast. But, in creating such a large set of equations, more storage space for the equations is required. In contrast, the indirect approach utilizes a symmetrical set of equations. The indirect method is slower than the direct method, but requires less memory (either RAM or temporary disk storage) than the direct method. [50]


In preparing for the HRTF calculation procedure using the BEM solution, several tests were made to determine the expected amount of storage space and computational time that would be required for such a problem. The required temporary disk storage space for the calculation (*not* for storage of the results) was examined as a function of mesh size for both the direct and indirect method. The results of this comparisons are shown in Fig. 5.5(a). This storage space can be accommodated either in RAM or on disk. As few machines have the large amounts of memory required for these calculation the results are shown for disk space requirements. Several different machines were available which provided for a variety of calculation speed comparisons. The results of some of these tests are given in Fig. 5.5(b). The direct method could not be tested using all the sample meshes as once the mesh size increased beyond approximately 12,000 elements, the disk space requirement for the direct calculation exceeded the free space available (1.4 Gigabytes).

It can be seen quite clearly that the computational speed quickly becomes impractical with increasing mesh size. Ideally, the computation would calculate the HRTF over the entire audible frequency range. Calculating up to 20 kHz at a high frequency resolution,

[†] Sysnoise, a product of LMS Numerical Technologies (LMS-NIT), versions 5.1–5.3 were used during this research. Sysnoise for the SGI was graciously donated for this research project and the author would like to acknowledge the kind contributions of LMS-NIT in aiding this extremely computationally intensive work.



a. Required space for BEM calculations as a function of mesh size for direct and indirect method.



b. Calculation speed comparison for two BEM transfer functions schemes, on several machines, as a function of mesh size.

Figure 5.5: BEM calculation space and time requirements.

for example in 25 Hz increments, would equate to 800 frequencies. The largest mesh attempted, again due to storage limitations, was 22,000 elements. This mesh, shown in Fig. 5.6, was only valid up to 5.4 kHz. This calculation took approximately 28 hours of CPU time for a single frequency. To calculate the ideal number of frequencies at this speed, neglecting the issue that at higher frequency resolution the computation times increases, would take approximately 2.6 years of CPU time. In addition, it is necessary to separately calculate the HRTF for each ear, doubling the calculation time to over 5 years. It was therefore necessary to balance the desired results with a plausible time frame. A calculation speed of one frequency/day was decided as acceptable, limiting the frequency range to 5.4 kHz. With this limited frequency range it would not be plausible to perform any subjective testing as the audio quality and information contained over this range would not be sufficient. As there is minimal to no frequency relevant information in the HRTF below 1 kHz (see section 2.1), the span of the calculation can be limited to the range of 1–5.4 kHz. With the removal of subjective testing, the calculation can then be limited to a single ear. This quickly reduced the computation time in half. Finally, sufficient frequency resolution was determined to be at 100 Hz intervals maximum. For some of the mesh variation conditions this spacing was increased up to 400 Hz. The result was a computation of approximately 50 days of CPU time. Depending on the load of the system due to other users, the actual time will increase by an unpredictable amount, though doubling is not unquestionable. The inclusion of hair impedance on the external surface also increased the computational time from that of a purely rigid model.

5.3.2. Various Meshes

The primary mesh for the HRTF calculation is shown in Fig. 5.1 and Fig. 5.6–Fig. 5.10. The first test condition was for a rigid model in which all the elements are defined as having infinite (*i.e.* perfectly rigid) impedance. Additional modifications were made to this mesh to test various other conditions and determine the validity of this approach. Example Sysnoise setup and command files for some of these meshes are given in Appendix C.

Using the impedance measurements of hair from section 3.3.4, the boundary conditions of some elements were modified from purely rigid to the measured impedance values as a function of frequency. As the BEM is designed around surfaces, the effects of sound propagation *through* hair is not feasible without defining a secondary fluid over part of the head having some defined impedance. Using the method here, each element in the defined region is assigned the impedance condition and it is thought that as the sound diffracts around the head, it will interact with the impedance effects of the elements it propagates over. The BEM solver used only provides for the definition of normal acoustic impedance, not grazing incidence values, and considers each element as a locally reacting surface not capable of imparting vibrations to adjoining elements. This, in combination with the boundary definition itself, makes this method for inclusion of hair impedance a rough approximation. For these reasons, the hair region is only simply defined, not precisely defined by the geometry of the original subject. The set of elements defined with the hair response can be seen in Fig. 5.11.

As a test to see the effects of geometric changes, a large modification was made to the measurement mesh. This modification was the entire removal of the pinna, something which would be extremely difficult with a real subject. Using `decimate.c: remove_ears` the pinnae were basically sliced off using a plane which was roughly normal to the head surface over the pinna region. The resulting mesh can be seen in Fig. 5.12.

One final mesh geometry was used in the HRTF calculation comparisons. Two sized spherical meshes were created. The first mesh had the same diameter as the inter-aural spacing (IAS) between the the entrance of the two ear canals. The mesh can be seen, in its relative position to the head mesh, in Fig. 5.16 and is referred to as either the “small sphere” or “IAS Sphere.” The same ear canal point (relative to the placement of the measurement mesh) was used and the transfer function was calculated. The second mesh, shown in Fig. 5.17, has the same volume as the head mesh and is referred to as the “Eqvol Sphere.” The mesh is orientated such that the ear canal positions are on the same axis as the head mesh. Due to the symmetry of these models more overdetermination elements were required, as discussed in 5.4.1.




Figure 5.6: The full BEM mesh used. Frequency limitation of 5.4 kHz for 4 nodes/wavelength. Note the increased resolution in the pinna region. Isometric view.

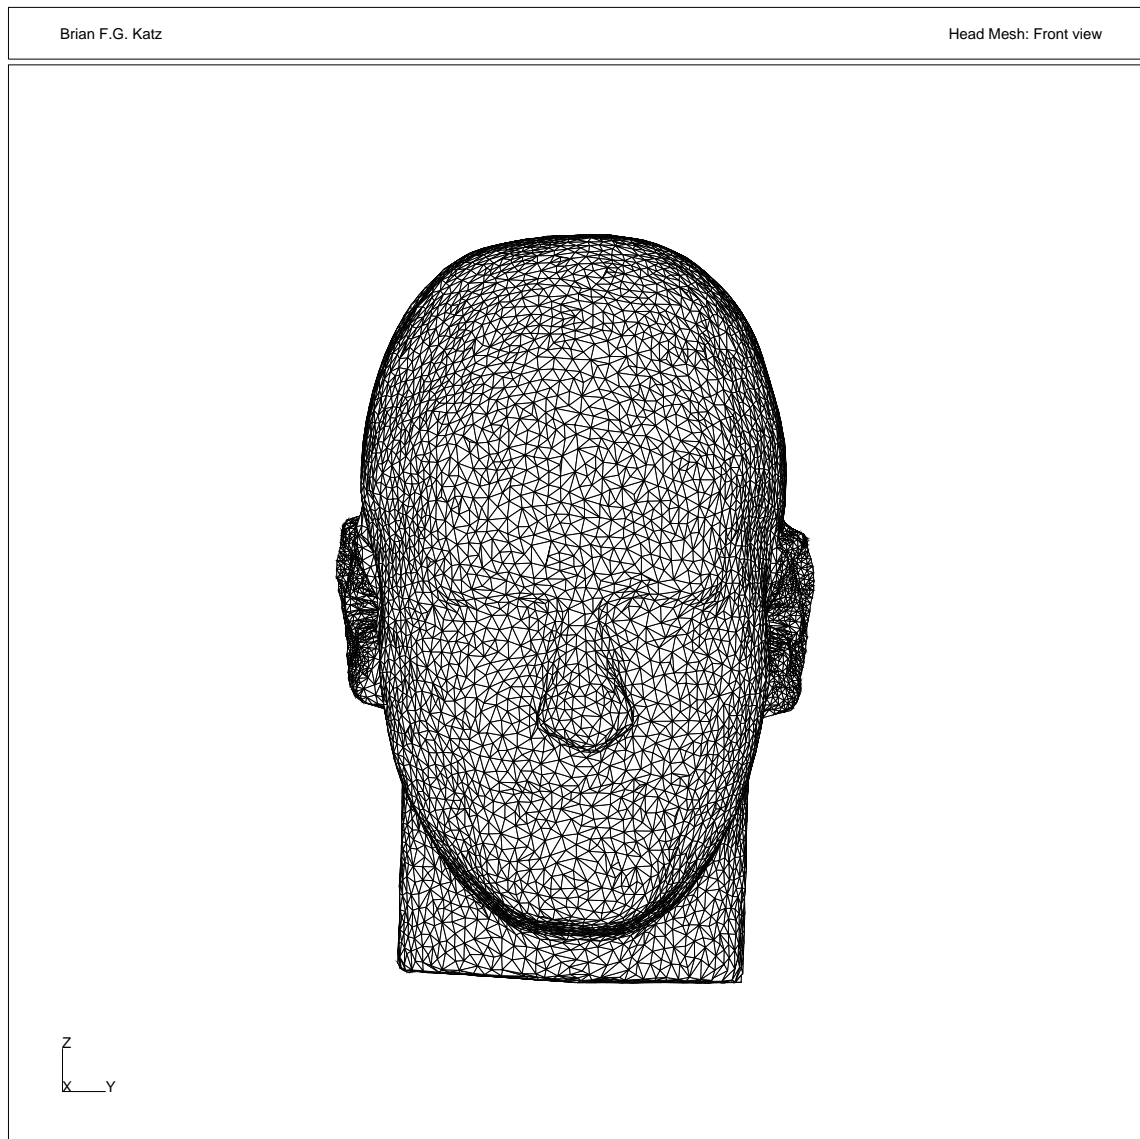


Figure 5.7: The full BEM mesh used. Frequency limitation of 5.4 kHz for 4 nodes/wavelength. Note the increased resolution in the pinna region. Front view.




Figure 5.8: The full BEM mesh used. Frequency limitation of 5.4 kHz for 4 nodes/wavelength. Note the increased resolution in the pinna region. Side view.

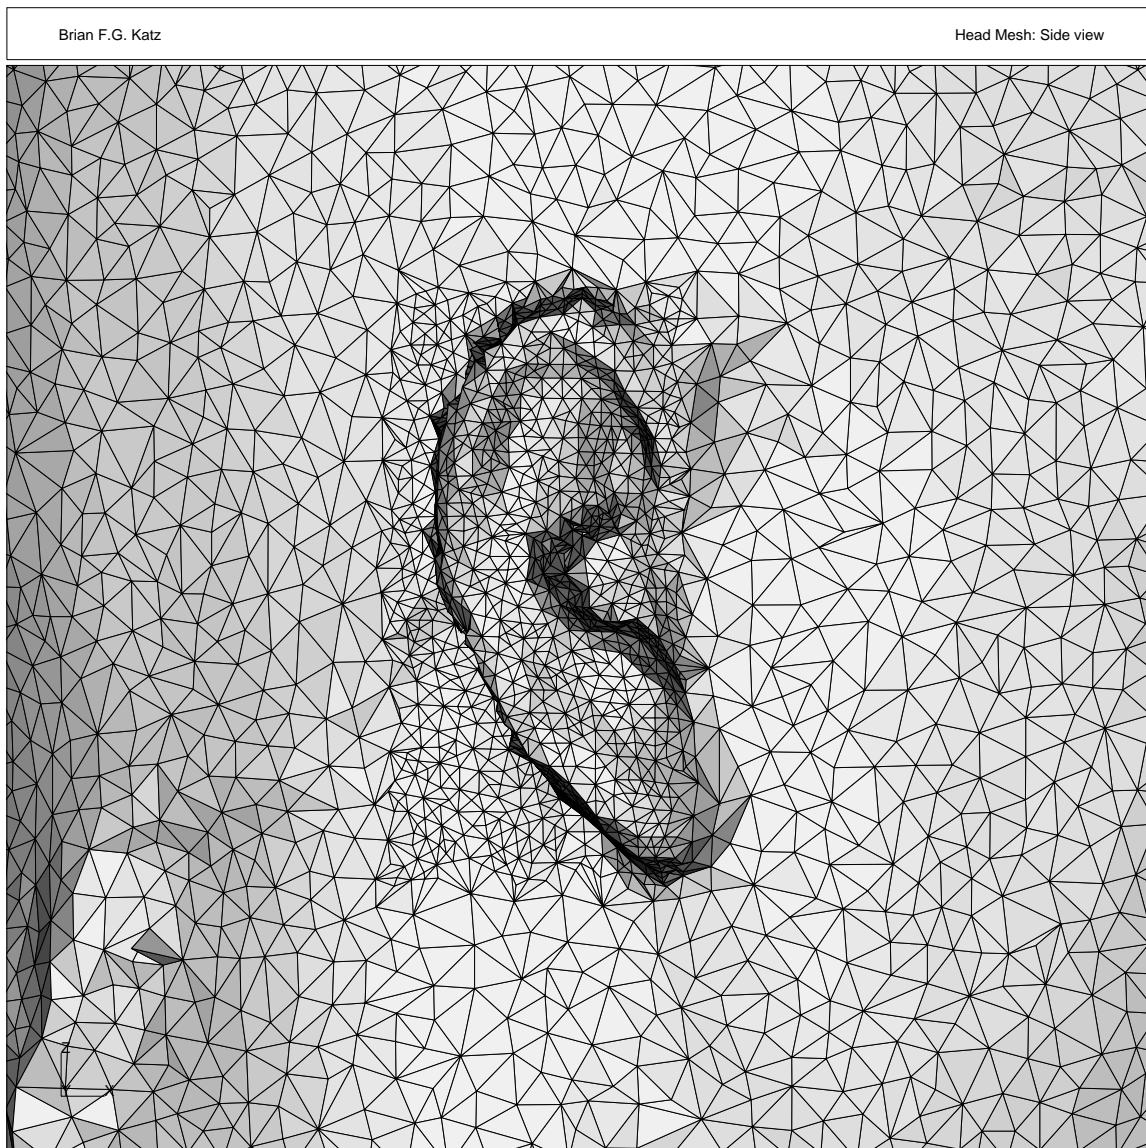


Figure 5.9: The full BEM mesh used. Frequency limitation of 5.4 kHz for 4 nodes/wavelength. Note the increased resolution in the pinna region. Right side ear close-up.




Figure 5.10: The full BEM mesh used. Frequency limitation of 5.4 kHz for 4 nodes/wavelength. Note the increased resolution in the pinna region. Underside view.




Figure 5.11: Mesh definition including the hair subset for which the measured values of hair impedance were defined over.




Figure 5.12: BEM mesh used for evaluation of effect of the removal of the pinna contribution. Isometric view.




Figure 5.13: BEM mesh used for evaluation of effect of the removal of the pinna contribution. Front view.




Figure 5.14: BEM mesh used for evaluation of effect of the removal of the pinna contribution. Side view.




Figure 5.15: BEM mesh used for evaluation of effect of the removal of the pinna contribution. Top view.

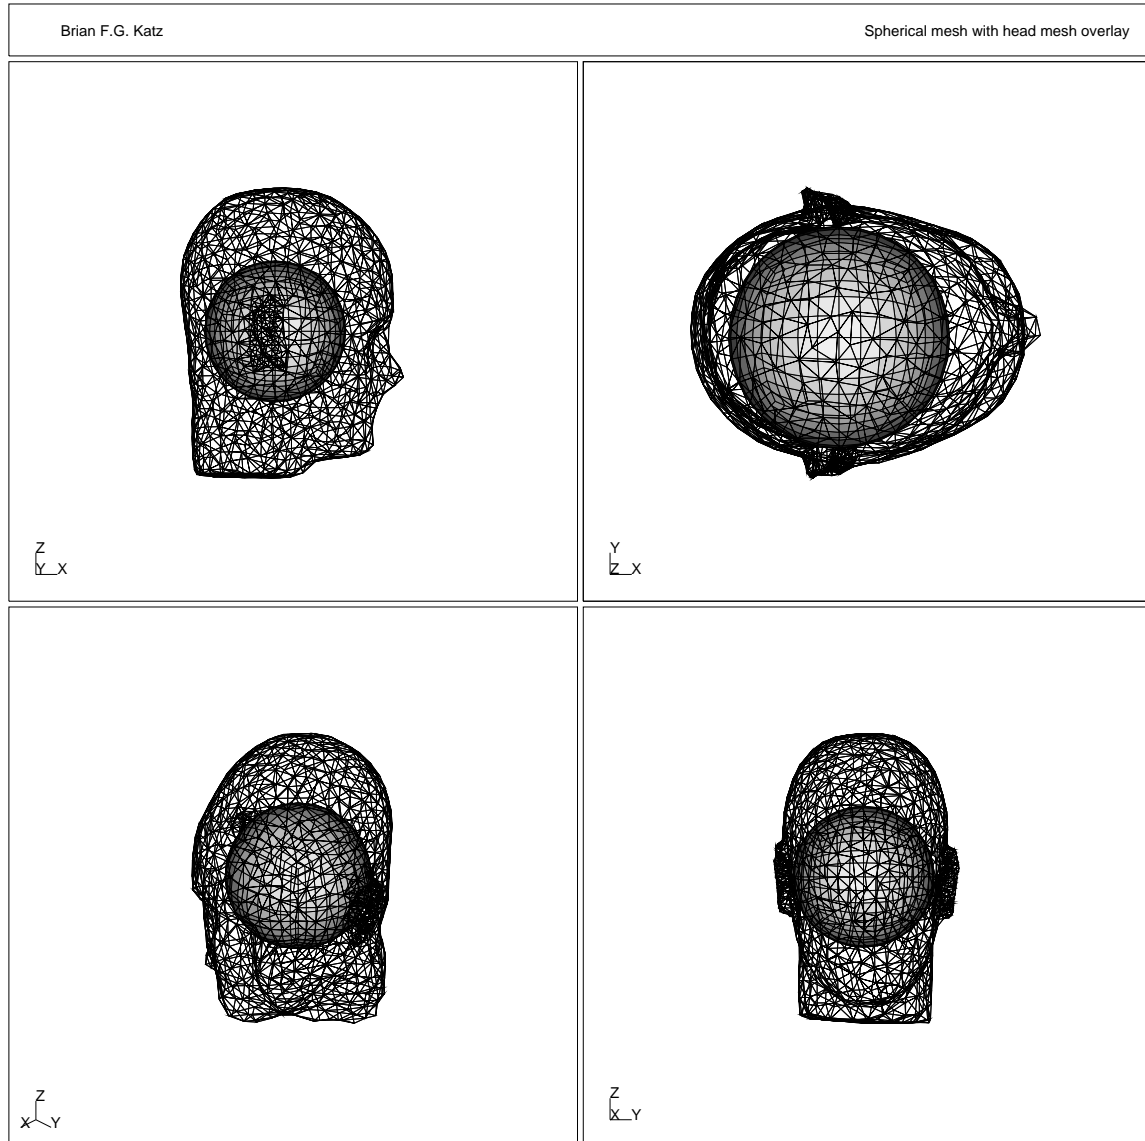


Figure 5.16: BEM IAS spherical mesh and head mesh overlay for dimensional and placement reference. Sphere size designed such that the diameter of the sphere is equal to the inter-aural spacing *i.e.* the sphere is coincident with the entrances to the ear canals.

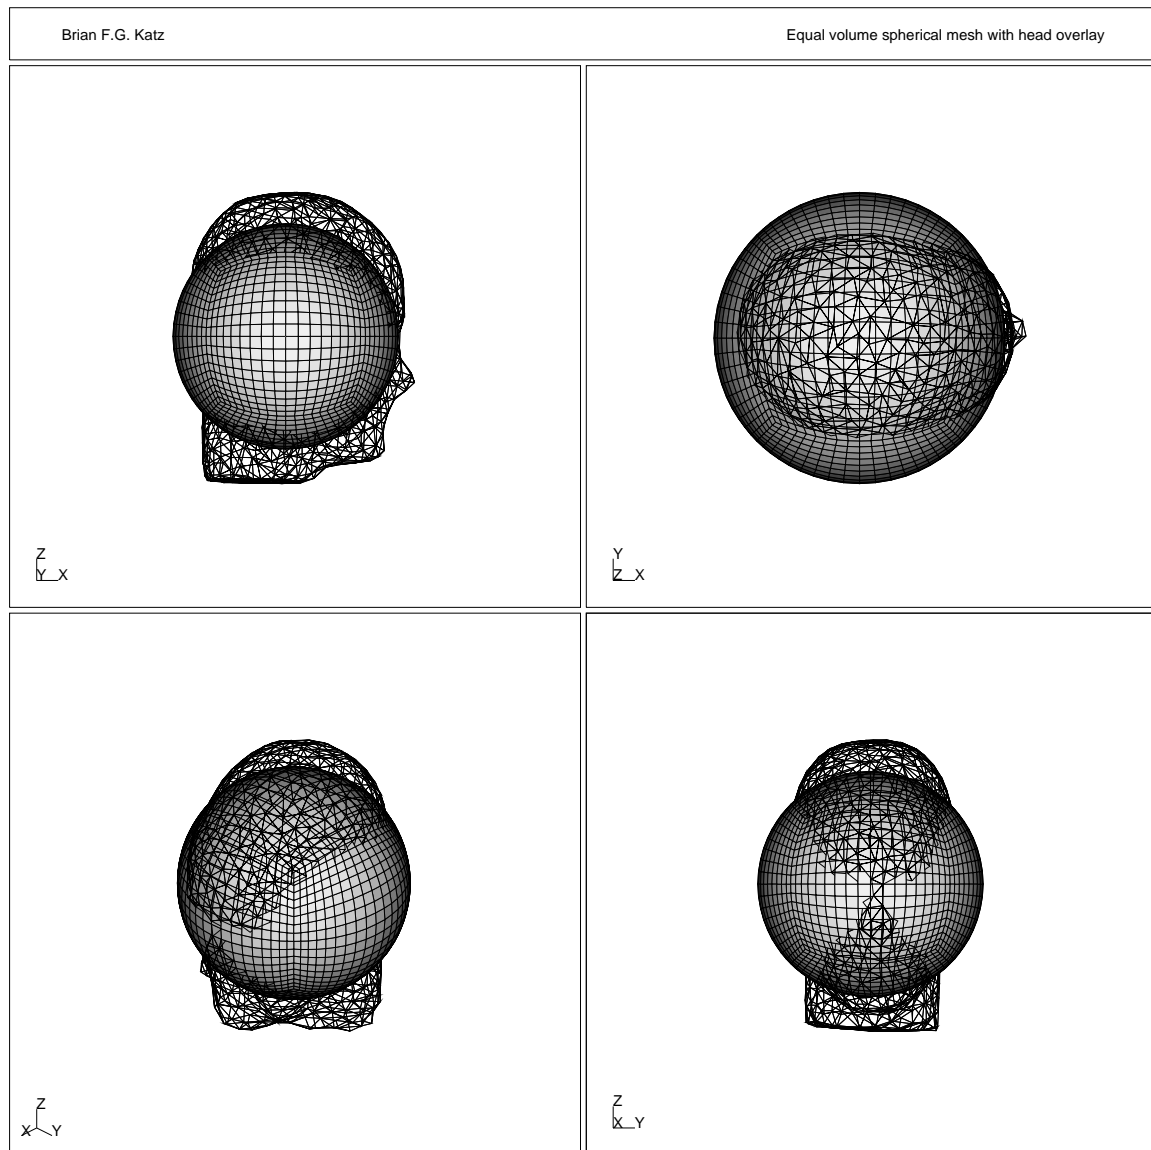


Figure 5.17: BEM Eqvol spherical mesh and head mesh overlay for dimensional and placement reference. Sphere size designed such that the sphere and head mesh have equal volumes.

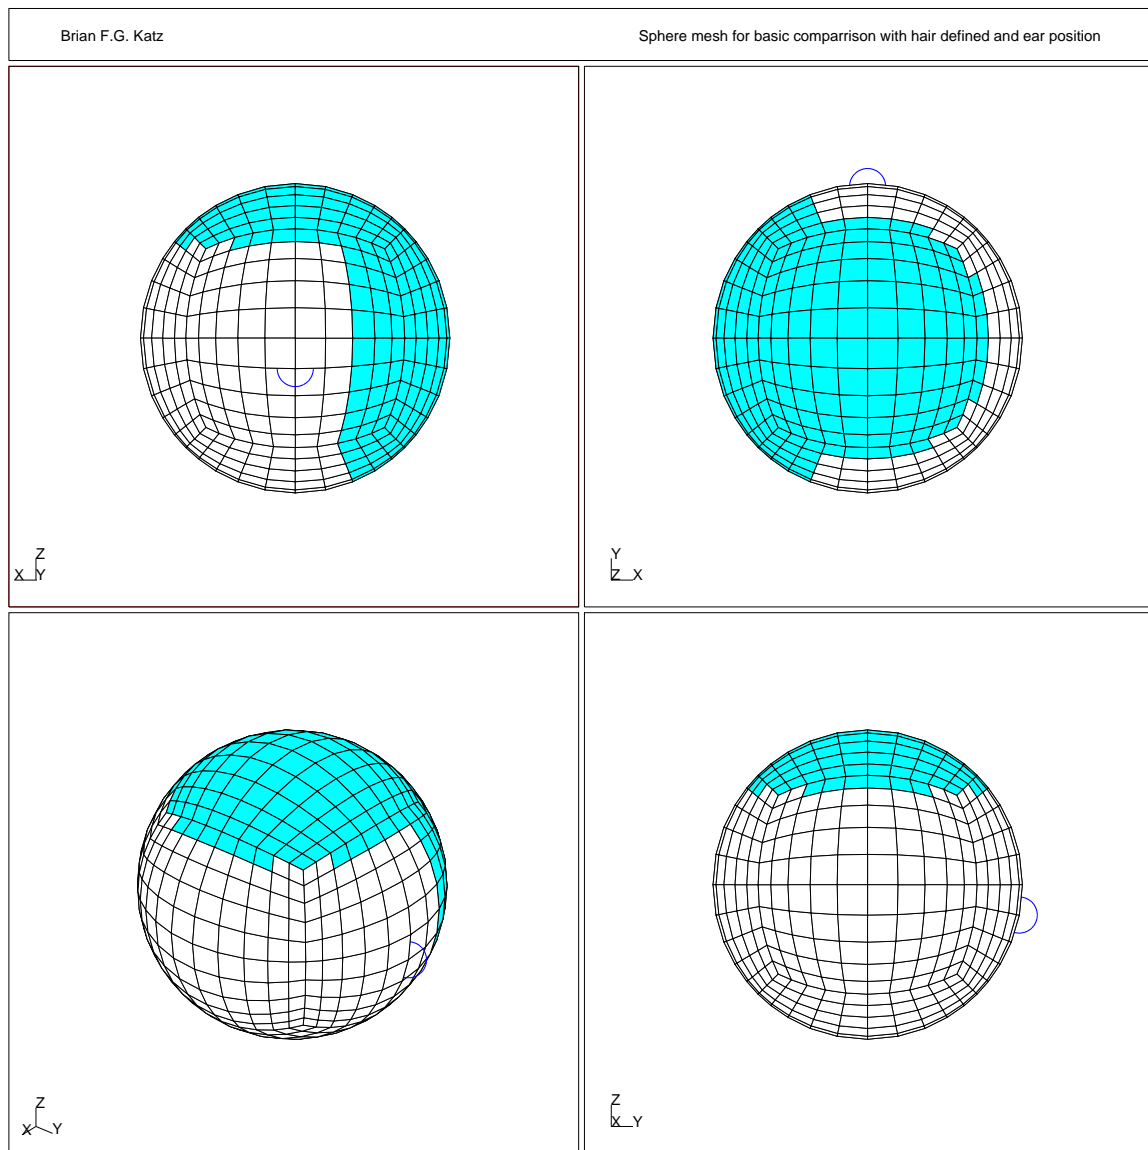


Figure 5.18: BEM spherical mesh with hair element set defined.

5.4. Results

5.4.1. Spherical model

In order to have some level of confidence in the model solution, and to provide for a simple comparison situation, the smaller spherical mesh was utilized in a basic evaluation between theory and a BEM model solution. The problem, defined as the radiation from a point source on a sphere, has a well defined analytical solution. [31,47] The effect of the sphere can be described using Eq. (5.4) where $\mu = kr$, r being the radius of the sphere, P_n is the Legendre polynomial of n^{th} order and B_n and δ_n are defined in Eq. (5.5) from the spherical Hankel function h_n .

$$F_r(\theta) = \frac{1}{\mu^4} \sum_{\substack{m=0 \\ n=0}}^{\infty} \frac{(2m+1)(2n+1)}{B_m B_n} P_m(\cos \theta) P_n(\cos \theta) \cos[\delta_m - \delta_n + \frac{\pi}{2}(m-n)] \quad (5.4)$$

$$\frac{d}{dz} h_n(z) = i B_n(z) e^{i\delta_n(z)} \quad (5.5)$$

Several approximation functions for B_n and δ_n are used, as defined in Eq. (5.6) and Eq. (5.7), to simplify calculations. [31,47]

$$B_n \simeq \begin{cases} \frac{1}{(kr)^2} & ; m = 0 \\ \frac{1 \cdot 3 \cdot 5 \cdots (2n-1)(n+1)}{(kr)^{n+2}} & ; kr \ll n + \frac{1}{2} \\ \frac{1}{kr} & ; kr \gg n + \frac{1}{2} \end{cases} \quad (5.6)$$

$$\delta_n \simeq \begin{cases} \frac{1}{3}(kr)^3 & ; n = 0 \\ \frac{-n(kr)^{2n+1}}{1^2 \cdot 3^2 \cdot 5^2 \cdots (2n-1)^2 (2n+1)(n+1)} & ; kr \ll n + \frac{1}{2} \\ kr - \frac{\pi}{2}(m+1) & ; kr \gg n + \frac{1}{2} \end{cases} \quad (5.7)$$

Comparison of the theoretical approximation function and a BEM solution, using a coarse mesh and small frequency steps, results in Fig. 5.19. The number of terms in the summation was varied from 10 to 30 with little to no change in the results. The theoretical approximation functions result in a step approximation of the true function, and therefore result in oscillations of the solution as seen in the plot. A simple curve fit of the approximation result is therefore included and used. The BEM solution and theory for the point source on a sphere follow the same trend, though there is a slight level reduction of about 0.5–1 dB in the BEM solution which is considered negligible. With the relatively close match between theory and calculation, there can be a good degree of confidence in the calculations of the more complex models of the head.

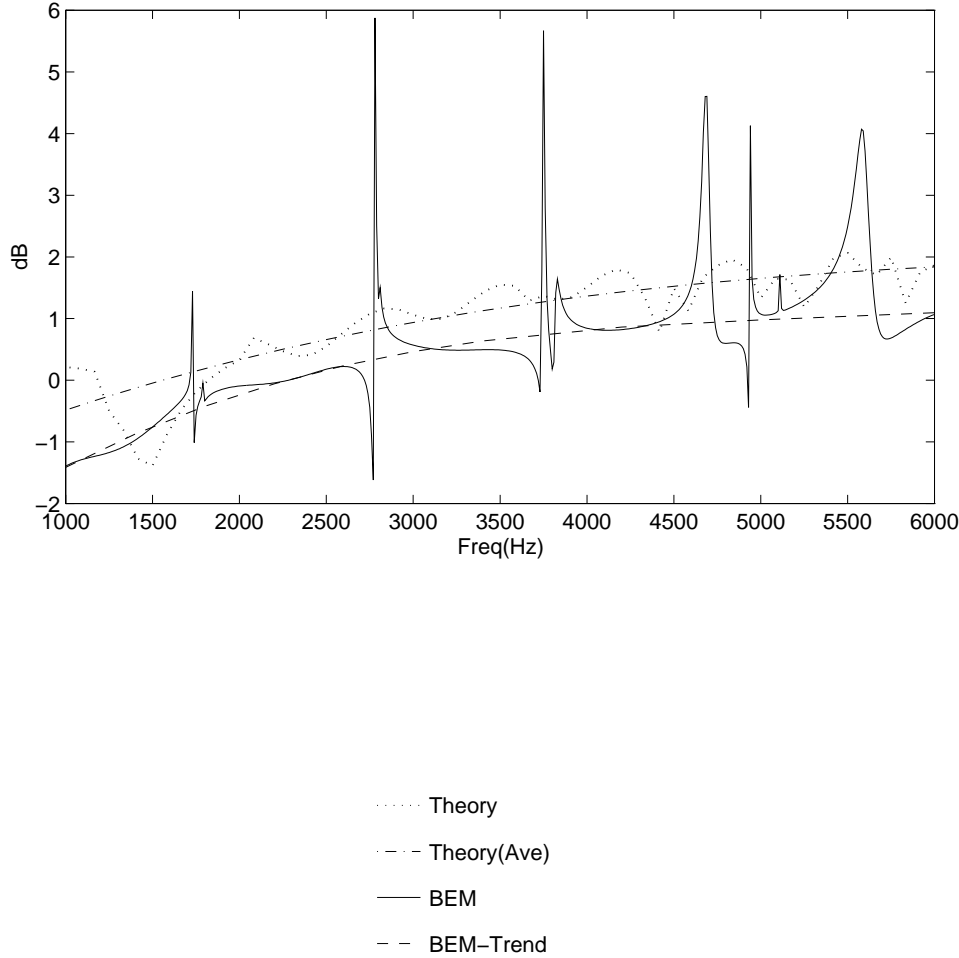


Figure 5.19: Comparison of BEM solution (free field equalized) of point source on a sphere (left ear with source left) to the theoretical approximation. In addition, a smooth fit of the theory is shown with a smooth fit of the BEM calculation disregarding computational anomalies due to the effects of the internal resonance solutions of the sphere.

Of interest in this result are the spikes in the BEM solution. These are artifacts of the boundary element approach and correspond to internal pressure release modes of the sphere (see section 5.1). The mesh used is valid up to 4.5–6 kHz depending on whether six or four elements per wavelength are prescribed. As the solution goes higher in frequency the number of error spikes increases, as would be expected with the increase in internal resonance modes with frequency. One reason for such an effect of these modal solutions is that the sphere has a great deal of symmetry. These effects were not seen in the head solutions. Some of this may be due to the placement of the overdetermination planes, which is more of an art than a science, and the fact that the head is much less symmetrical than the sphere.

In an effort to understand the effects of the overdetermination points better, the same spherical calculation was performed using a different and larger set of overdetermination points. The results of this comparison are shown in Fig. 5.20. Two distributions for the overdetermination elements are used. The first set uses a single plane of elements, divided into a 10×10 grid, placed according to the setup file given in C.1. This is the set used for the head mesh calculations. The second set uses three planes, divided in the same manner, placed according to the setup file given in C.3. For this setup, the elements are larger than the previous set, and with three different planes there is a larger distribution in the interior of the mesh. This set was used for the spherical model calculations.




Figure 5.20: Effect of varying location and number of overdetermination elements in BEM indirect solution. Solid line represents overdetermination locations defined in Appendix C.1. Dashed line represents element set defined in C.3. Circles represent data with 400 Hz spacing.


Results of this comparison show that an increase in the number, size, and distribution of overdetermination elements raises the frequency of the first resonance and shifts the frequencies at which subsequent irregular frequencies occur in the solution. This procedure does not remove these errors altogether, but as they are easily identifiable they can be ignored in the presented results. In subsequent plots, it is important to remember the occurrence of these calculation errors when sharp peaks or nulls are present.

5.4.2. Rigid models


Variations in the geometry of the mesh highlight contributions of the shape of the head to the HRTF. The four mesh geometries used here –full head, head without pinna, IAS sphere, eqvol sphere– show several major points. As the HRTF is a function of space as well as frequency, different source locations may be affected differently by geometrical changes. For these results, the data was free-field equalized. Subsequently, in the comparisons with the measured data (see chapter 6) the front-source equalization method was used.

As described earlier, all calculations are performed using only the left ear. The first position for consideration is directly left. Results for all four geometries are given in Fig. 5.21(a). From this result, it can be seen that the pinnaless head follows the same response as the intact head up to approximately 3 kHz, at which point it follows the trend of the eqvol sphere. The IAS sphere has roughly the same shape as the eqvol sphere but with a lower magnitude. At higher frequencies the two spheres, and the pinnaless head, are approaching similar values. As described in the previous section, there are frequencies at which the internal modal resonances of the mesh induce error and affect the exterior solution. For this reason, it is important to ignore sharp spikes in the BEM solution, such as the spikes in the IAS spherical solution in Fig. 5.21(a). The source position is moved clockwise in the horizontal plane to position 9, midpoint between left and front. For this position (results shown in Fig. 5.21(b)), deviation between the pinna and pinnaless model occurs at 2 kHz, and even though the pinnaless and sphere models follow the same trend, the pinnaless model is offset with approximately a 2 dB increase from the IAS sphere and 1 dB from the eqvol sphere.

With the source position directly in front, the variations between the pinna and pinnaless meshes increases, as shown in Fig. 5.22(b). Trend agreement only exists up to approximately 1.5 kHz, with a 1 dB offset. The pinnaless solution follows the rough trend of the spherical solutions up to about 4.5 kHz, at which point the solutions diverge. Continuation of source rotation to the opposite side of the listening ear, directly right, is shown in Fig. 5.22(a). It is quite evident that the pinna, pinnaless, and equal volume mesh solutions agree well over most of the frequency range. There are some deviations on the order of 3 dB, but these are at locations where the transfer function result is already on the order of -15 dB. The pinna and pinnaless model both have a ‘step’ as a function of frequency starting in the neighborhood of 3–3.5 kHz. This is not present in the eqvol sphere. The total reduction is much less than for the IAS spherical model. The final position in the horizontal plane under consideration is the rear position, shown in Fig. 5.23(a). As with the front source (Fig. 5.22(b)), the pinna and pinnaless models agree only up to about 2 kHz, at which point the pinnaless model tends more towards the spherical models. The eqvol sphere matches better with the pinnaless head than the IAS sphere. Also of interest is the radical difference between the front and rear source positions for the intact model, and how the pinnaless and spherical models are not dramatically different. Two median plane sources are shown in Fig. 5.24. In general the pinnaless model still conforms more to the spherical models than the intact pinna model. Finally, a point not on either the horizontal or median plane is shown. In Fig. 5.23(b) the results for a source in the front sphere, in the center of the lower right octant, is shown. What is useful in reviewing these positions is the radical difference between all four models.




a. Source WPAFB position 12, left.




b. Source WPAFB position 9, front left.

Figure 5.21: Free-field equalized rigid BEM transfer functions for source positions WPAFB 9 and 12.




a. Source WPAFB position 1, right.




b. Source WPAFB position 6, front.

Figure 5.22: Free-field equalized rigid BEM transfer functions for source positions WPAFB 1 and 6.




a. Source WPAFB position 18, rear.




b. Source WPAFB position 86, center of front lower right octant.

Figure 5.23: Free-field equalized rigid BEM transfer functions for source positions WPAFB 18 and 86.



a. Source WPAFB position 52, median plane -45° .



b. Source WPAFB position 57, median plane $+45^\circ$.

Figure 5.24: Free-field equalized rigid BEM transfer functions for source positions WPAFB 52 and 57.

It can be concluded from these results that the shape of the head in general is important to the low frequency region of the HRTF. Below approximately 2 kHz the effect of the pinna is minimal. For sources on the opposite listening side, the pinna effect is greatly reduced. This is most evident for the source right position, where there was minimal difference between the pinna and pinnaless model. Approximating the head with a sphere of equal volume yields very similar results to a head *without* pinna, but does not match the effects of a head with pinna except in the case of the source on the opposite side of the listening ear.

It is important to remember that a number of assumptions were made in the construction of the mesh which differ from the true individual's head and body geometry. Primarily, these assumptions are the lack of any body below the neck. In addition was the effect of the cylindrical scanner, as described in section 5.2.1.


5.4.3. Impedance Condition

Using the inclusion of measured hair impedances (section 3.3.4) defined for the regions shown in Fig. 5.11 and Fig. 5.18, the HRTF calculation procedure was performed. These results show the effects of an absorptive region on the top and back of the head to the HRTF, and may help in determining whether or not consideration of such an effect should be included in localization and HRTF theory.


Results are given in pairs, showing the effect of hair on the intact head and both spheres. In this manner, trends can be seen, and possible anomalies can be considered if the trends do not agree. In addition, it is important to remember the effects of the internal modal solution and the anomalies which they can cause (as described in 5.4.1 and 5.1). Source locations used are the same as those considered in the previous section.

Placing the source directly to the left, still using the left ear as the receiver, results in the data given in Fig. 5.25(a). As can be seen, in the region above 3 kHz there is approximately a 0.5–1 dB decrease in the transfer function due to the presence of hair for both the head and IAS sphere solution. The eqvol sphere shows little change over the same region (taking into account the two spikes which are attributed to internal resonances). Above 4 kHz the eqvol rigid sphere continues its trend while the eqvol sphere with hair follows the IAS spheres trends. Below 3 kHz there is a general increase of roughly the same order, or greater (1.5 dB) for the head and IAS sphere with the opposite trend in the eqvol sphere. Variation of such a small degree, though interesting, are difficult to attribute either a great deal of meaning to or consider as generally perceivable to the listener.

As the source is again moved clockwise in the horizontal plane, towards the front, the effect of the hair changes slightly as shown in Fig. 5.25(b). For the head model, at frequencies above 4.5 kHz there is an increase due to the hair of 1 dB. In the range below this there is little contribution due to the hair as was seen with the eqvol sphere above. For the IAS sphere, this difference in response is not found, and the variations on the order of 0.5 dB still exist in the same regions. The eqvol sphere shows little effect below 3 kHz while above 3 kHz the inclusion of hair decreases the response by approximately 1 dB. Much the same hair variation is found when the source is at the front of the head




a. Source WPAFB position 12, left.




b. Source WPAFB position 9, front left.

Figure 5.25: Free-field equalized BEM transfer functions with defined hair impedance for source positions WPAFB 9 and 12 using intact head and two sphere meshes.




a. Source WPAFB position 1, right.




b. Source WPAFB position 6, front.

Figure 5.26: Free-field equalized BEM transfer functions with defined hair impedance for source positions WPAFB 1 and 6 using intact head and two sphere meshes.




a. Source WPAFB position 18, rear.




b. Source WPAFB position 86, center of front lower right octant.

Figure 5.27: Free-field equalized BEM transfer functions with defined hair impedance for source positions WPAFB 18 and 86 using intact head and two sphere meshes.



a. Source WPAFB position 52, median plane -45° .



b. Source WPAFB position 57, median plane $+45^\circ$.

Figure 5.28: Free-field equalized BEM transfer functions with defined hair impedance for source positions WPAFB 52 and 57 using intact head and two sphere meshes.

model (Fig. 5.26(b)). In this position, the eqvol sphere shows little variation due to the inclusion of hair.

Moving the source to the opposite side of the listening ear shows some real effect for the head model. Variations in the region above 3.5 kHz begin at about 1 dB and increase to over 6 dB, where the response due to hair is greater than without. Below this frequency, the hair has the effect of reducing the transfer function by up to 3 dB. With the eqvol sphere there is a dramatic increase in the response for the hair model over the rigid model which is similar to that found in the upper frequencies of the head model. This is not the case in the IAS spherical model. A possible explanation for this could be that the effect of destructive interference, due to the sound waves taking different path around the head (from the opposite side) and then converging and combining at the far side of the head, is reduced due to the reduction in level and/or phase of the sound which propagates over the complex hair impedance. The interference effect depends upon frequency and path length. The eqvol sphere and head have similar dimensions and therefore have similar path lengths.

The rear source position shows similar results for the head model, though not to the same degree (see Fig. 5.27(a)). The head response shows an increase for the hair model beginning at approximately 3 kHz, though only reaching a maximum deviation of 1 dB. For the spherical models, little to no effect was again seen. This highlights the fact that the HRTF and its contributing geometrical parameters are greatly a function of angle. A change with assists the model in one location may have little, or the opposite, effect at other incident angles.

The median plane source positions as well as the lower left position (shown in Fig. 5.28 and Fig. 5.27(b)) show similar results as seen for the other front sources. Variations on the order of 1 dB or less exist. For the low median source there is little effect in the sphere case. For the head mesh below approximately 5 kHz the hair response is lower, but above 5 kHz the hair response is greater by the same degree. This is interesting, though with the lack of a torso in the model, it is difficult to attribute much meaning to this occurrence. The torso provides additional reflections from the chest, back, and shoulders. In addition, it provides a kind of baffle which blocks sound from going under the head, which can occur in the BEM model.

From these calculations it can be seen that the inclusion of hair effects does result in some degree of variation in the HRTF. These variations range from 0.5 dB to 6 dB for the source positions shown. In general the effects seen in the eqvol sphere follow the same trends as exhibited by the head model. The effect for almost all cases shown for the IAS spherical model was much less than for the head model and eqvol sphere. This can be most likely attributed to the difference in size and the larger surface with defined hair impedance in the head and eqvol sphere mesh as compared to the IAS sphere mesh. Whether or not this inclusion of hair improves the model as compared to measured data is discussed in chapter 6.

Chapter 6.

HRTF Comparisons

Final analysis of the work presented here consists of comparing the HRTFs obtained through both experimental methods and calculations, both with and without hair. Before the various source positions are shown, the unequalized data for the front source is presented in Fig. 6.1. It can be seen that the calculated HRTF is in the same range as $1R$ and $1P$, though two distinct peak features are missing. Data for $1C$ is different, though somewhat expected, as the data acquisition method and some diffuse equalization post processing in the data acquisition system exist and do not compare directly for raw data between the other methods. The variation between the rigid and hair calculation for this point is minimal.




Figure 6.1: Front source equalization function used for the three experimental HRTF measurement methods and for the BEM head calculations with and without the inclusion of hair impedance values.


Having presented the front-source equalization functions the various other source position data, having been front-source equalized, is present in Fig. 6.2–Fig. 6.4. Initial consideration is given to the source-left position as given in Fig. 6.2(a) which shows good agreement between the computational models and measurement data $1C$ below 1.8 kHz and $1R$ above 2.5 kHz. There is a prominent peak in the experimental data sets in the

region of 2 kHz which is not present in the calculations. This peak could very well be due to reflections off of the left shoulder, a limitation of the calculations model which considered only the head. In addition, the inclusion of hair seems to improve the calculation from 3–4 kHz, though above this the hair model tends more towards $1P$, the rigid head apparatus, than the real measurements.


Examination of the source-right position, where the inclusion of hair had the most dramatic effect, provides for somewhat different results (Fig. 6.2(b)). The rigid head model appears as a good average value compared to the experimental data up to 2.5 kHz. The hair model matches the null values of the experimental data over this region. After this point the rigid model level decreases substantially, while the hair model maintains a higher level. This is similar to the comparison between $1P$ and $1C$, where the real subject data (with hair) is noticeably higher in level as compared to the rigid head pinna mold measurement.

Analysis of the median plane sources is shown in Fig. 6.3. Both calculated positions show little variation. To some extent $1C$ also has less variation when compared to $1P$ and $1R$. In reference to these last two measurement techniques, one could explain the response of the calculations in that this equalization method shows the variation of a position with respect to the frontal source position of the horizontal plane. The computational model mesh does not accurately represent the system in this region as the torso has the effect of being a rather large baffle in this frequency region. It is not surprising that there is a dramatic difference between the results. But, as the $1C$ measurement does not show the same variations, and more resembles the calculated response in shape, it is difficult to draw a discrete conclusion.

The final position considered is source-rear. For this location the calculated responses follow the trend of $1C$ below 2.5 kHz and $1R$ above 3.5 kHz. The inclusion of hair does seem to improve the calculation in the low frequency region, but aids little in the region above this.




a. Full HRTF set comparisons for source left.




b. Full HRTF set comparisons for source right.

Figure 6.2: Comparisons of measured and calculated HRTFs for source positions located directly left and right. Source positions used were WPAFB Source:12 & Source:1 and CRE Source:22 & Source:58 respectively for left and right positions.



a. Full HRTF set comparisons for high frontal median plane source.



b. Full HRTF set comparisons for low median plane frontal source.

Figure 6.3: Comparisons of measured and calculated HRTFs for frontal median plane source positions. Source positions used were WPAFB Source:58 & Source:53 and CRE Source:38 & Source:42 respectively for high and low positions.




Figure 6.4: Comparisons of measured and calculated HRTFs for directly rear source positions. Source positions used were WPAFB Source:18 and CRE Source:4.

Chapter 7.

Conclusion

Human ability to locate sound sources is a remarkably complex process. For it to function it requires three distinct anatomical components of the body to work together. The pinna or outer ear, along with the head, alter the incident sound waves resulting in frequency modified and time delayed signals to the inner ear. The inner ear, responsible for converting the sound waves from the pinna to electrical neural signal for the brain, increases level of these changes and performs signal processing on each ear signal, and on the variations between the two. The results of these processes are all fed to the brain, which calculates more complex relations between the input signals and compares the results to a vast memory of results to determine where the sound originated from.

This work has focussed primarily on the first part of this complex procedure, the Head-Related Transfer Function. It has been shown that measurement of the HRTF is not a trivial matter. Variations between measurement methods persists and must be taken into consideration when comparing measurements and making statements regarding the physics of the system. The use of replacement individuals, utilizing personal pinna replicas in a generically shaped head, has shown relatively good agreement with measurements made on the respective real individual. Future work involving the detailed analysis of the contribution of measurement technique, experimental error, and repeatability is recommended to fully understand the variations observed. In addition, the use of truly symmetrical pinna for an individual would be an interesting experimental investigation.

A computational model based on the geometric shape of an individual was used to calculate a limited frequency HRTF. This model showed good agreement with the various measured data, within the observed variations of the experimental results. The success of this work now provides a useful tool in examining the contribution of head and pinna shape to the properties of the HRTF. Continuing work in this area could focus on the details of the pinna. In addition, with faster and larger computers available all the time the detail and frequency span of the calculation can be increased so that possibly a full frequency band HRTF spanning the entire audible frequency range could be computed.

The contribution of the pinna to the characteristics of the HRTF was distinctly shown by utilizing the computer model. With this technique the pinna of an individual subject were totally removed, leaving a plane flat surface on the sides of the head. The calculated HRTFs from this computational experiment clearly identified the contribution of the pinna to the HRTF as a function of frequency and direction. For sources towards the front and rear of the listener the pinna effects begin at frequencies in the region of 1.5–2 kHz. For sound sources directed towards the ear pinna contributions begin at frequencies above 3.5 kHz. This coincides with the region reported in the literature where localization ability exists not as a function phase difference between the ears but as a function of frequency variations due to incident angle.

The contribution of acoustical parameters of the individual, the impedance of skin and hair, was also investigated. An qualitative analysis of the errors induced by variations in the estimated positions of the microphones during the measurements was performed. This analysis showed that the ability to both determine these distances and perform the measurements according ANSI standards is not possible for rigid terminations. Taking this into account, it was still possible to show that for all practical purposes skin is acoustically rigid. But, hair was shown to have significant absorptive properties, increasing as a function of frequency. The effects of this absorption on the HRTF were examined physically using an experimental rigid head with replica pinna and numerically by introducing the measured impedance data into the HRTF computation. Both procedures showed perceptible variations in the HRTF. The effects were more pronounced, though not limited to, sources originating at or behind the ears in direction. The most prominent effect was for the sound source directed at the ear on the opposite side of the head. Results showed differences on the order of 6 dB due to the inclusion of hair in the head model. Future work in this area could involve the extension of the frequency range of the impedance measurements. In addition, the inclusion of hair in the computational model could be approached differently. By modelling the hair as a absorptive medium, though which the sound must propagate, instead of the current boundary surface representation, the results may be a more accurate representation of the actual physical phenomenon.

In the continuing path of research into human localization ability, it is hoped that this research has provided some new information, and may lead to further investigations using the same techniques. The history of this field is long, and yet, there is still much to be learned in what some perceive as a relatively new field. With advances in hardware miniaturization, computational speed, and signal processing techniques, the possibilities of research in this area continue to grow rapidly.

References

1. Angell, James and Warner Fite. "The monaural localization of sound.", *The Psychological Review* **VIII**, n. 3, 225-246 (May 1901).
2. Angell, James and Warner Fite. "Further observations on the monaural localization of sound." *The Psychological Review* **VIII**, n. 5, 449-458 (September 1901).
3. ASTM E 1050 – 90. Standard Test Method for Impedance and Absorption of Acoustical Materials Using a Tube, Two Microphones, and a Digital Frequency Analysis System.
4. Banister, H. "A suggestion towards a new hypothesis regarding the localization of sound." *Journal of Psychology* **XVII**, 2, 142-153 (October 1926).
5. Batteau. "Role of the pinnae in human localization." *Proceedings of the Royal Society of London* **B168**, 158-180 (1967).
6. Begault, Durand R. *3-D Sound for Virtual Reality and Multimedia*. Academic Press, New York, 1994.
7. Blauert, Jens. *Spatial Hearing*, 3rd ed. The MIT Press, Massachusetts, 1997.
8. Bodén, Hans and Mats Åbom. "Influence of errors on the two-microphone method for measuring acoustic properties in ducts." *Journal of the Acoustical Society of America* **79(2)**, 541-549 (1986).
9. Butler, R. A. and K. Belendiuk. "Spectral cues utilized in the localization of sound in the median sagittal plane." *Journal of the Acoustical Society of America* **61**, 1264-1269 (1977).
10. Chu, W. T. "Transfer function technique for impedance and absorption measurements in an impedance tube using a single microphone." *Journal of the Acoustical Society of America* **80(2)**, 555-60 (1986).
11. Chung, J. Y. and D. A. Blaser. "Transfer function method of measuring in-duct acoustic properties. I. Theory." *Journal of the Acoustical Society of America* **68(3)**, 907-913 (1980).
12. Chung, J. Y. and D. A. Blaser. "Transfer function method of measuring in-duct acoustic properties. II. Experiment." *Journal of the Acoustical Society of America* **68(3)**, 914-921 (1980).
13. Clark, B. and A. Graybiel. "The effect of angular acceleration on sound localization: The auditory illusion." *Journal of Psychology* **28**, 235-244 (1949).
14. Demaske, P. and B. Wagener. "Richtungshörversuche über einen nachgebildeten Kopf [Investigations of directional hearing using a dummy head]." *Acoustica* **22**, 191-204 (1969).
15. Durrant, John and Jean Lovrinic. *Bases of Hearing Science*, 3rd Edition. Williams & Wilkins, Baltimore, 1995.
16. Farwell, Robert William. *Acoustic Absorption of Haired and Hairless Mice*. Masters thesis. The Pennsylvania State University, 1955.

17. Franke, Earnest K. *Mechinical impedance measurement of the human body surface*. Air Force Technical Report No. 6469, Aero medical laboratory, Air Material Command, Wright-Patterson Air Force Base, Dayton Ohio, 1951.
18. Gardner, M. B. and R. S. Gardner. "Problem of localization in the median plane: Effect of pinnae cavity occlusion." *Journal of the Acoustical Society of America* **53**, 400-408 (1973).
19. Gerber, Sanford E. *Introductory Hearing Science*. W. B. Saunders Company, Philadelphia, 1974.
20. von Gierke, Henneng E. *Measurement of the acoustic impedance and the acoustic absorption coefficient of the surface of the human body*. Air Force Technical Report No. 6010, Aero medical laboratory, Air Material Command, Wright-Patterson Air Force Base, Dayton Ohio, 1950.
21. Hammershoi, Dorte and Henrik Moller. "Sound transmission to and within the human ear canal." *Journal of the Acoustical Society of America* **100(1)**, 408-427 (1996).
22. Hebrank, J. and D. Wright. "Spectral cues used in the localization of sound sources in the median plane." *Journal of the Acoustical Society of America* **56**, 1829-1834 (1974).
23. Hochberg, Irving. *Auditory Localization of Speech and Its Presumed Underlying Factors*. Doctoral thesis, The Pennsylvania State University, 1962.
24. Katz, Brian F.G. "New Approach for Obtaining Individualized Head-Related Transfer Functions," *Journal of the Acoustical Society of America* **100(4)**, 2609(1996).
25. Kimitoshi, Fukudome and Masaki Yamada. "Influence of the shape and size of a dummyhead upon Thévenin acoustic impedance and Thévenin pressure." *Journal of the Acoustical Society of Japan* **(E) 10**, 1, 11-22 (1989).
26. Kinsler, Lawrence E. and Austin R. Frey. *Fundamentals of Acoustics*, John Wiley & Sons, New York, 1950.
27. Kinsler, Lawrence E., Austin R. Frey, Alan B. Coppens, and James V. Sanders. *Fundamental of Acoustics*, 3rd Ed., John Wiley & Sons, New York, 1982.
28. Levin, Harold and Julian Schwinger. "On the radiation of sound from an unflanged circular pipe." *Physical Review* **73(4)**, 383-406 (1948).
29. More, Louis. "On the appreciation of difference of phase of sound waves." *Philosophical Magazine* **s. 6 XIII**, n. 76, 452-459 (April 1907).
30. More, Louis. "On the localization of the direction of sound." *Philosophical Magazine* **s. 6 XVIII**, n. 104, 308-319 (August 1909).
31. Morse, Philip M. and K. Uno Ingard. *Theoretical Acoustics*, Princeton University Press, Princeton NJ, 1986.
32. Münsterberg, Hugo and Arthur Pierce. "Studies from the Harvard Psychological Laboratory (II): The Localization of Sound." *The Psychological Review* **I**, n. 5, 461-476 (September 1894).
33. Newby, Hayes A. *Audiology*. Appelton-Century-Crofts, Inc., New York, 1958.

34. *Noise and Vibration Control Engineering* (Ed. Leo Beranek and István Vér). John Wiley & Sons, Inc., New York, 1992.
35. Pierce, Allan D. *Acoustics*. Acoustical Society of America, New York, 1991.
36. Rayleigh, Lord. "Our perception of the direction of sound." *Nature* **XIV**, 32-33 (1876).
37. Rayleigh, Lord. "Acoustical observations. I: Perception of the direction of sound." *Philosophical Magazine* **III**, 456-464 (1877).
38. Rayleigh, Lord. "Acoustical observations. IV: Estimation of the direction of sound with one ear." *Philosophical Magazine* **XIII**, 340-347 (1882).
39. Rayleigh, Lord. "Acoustical notes. VI: Beats of sound led to the two ears separately." *Philosophical Magazine* **II**, 280-285 (1901).
40. Rayleigh, Lord. "Acoustical notes. VII: On our perception of sound direction." *Philosophical Magazine* **XIII**, 214-232 (1907).
41. Rayleigh, Lord. "Acoustical notes. VIII: Discrimination between sounds directly in front and directly behind the observer." *Philosophical Magazine* **XVI**, 235-246 (1908).
42. Roffler, S. K. and R. A. Butler. "Factors that influence the localization of sound in the vertical plane." *Journal of the Acoustical Society of America* **43**, 1255-1259 (1968).
43. Searle, C. L., L. D. Briada, D. R. Cuddy, and M. F. Davis. "Binaural pinna disparity: Another auditory localization cue." *Journal of the Acoustical Society of America* **57**, 448-455 (1975).
44. Seybert, A. F. and D. F. Ross. "Experimental determination of acoustic properties using a two-microphone random excitation technique." *Journal of the Acoustical Society of America* **61(5)**, 1362-70 (1977).
45. Shaw, E. A. G. and R. Teranshi. "Sound pressure generated in an external-ear replica and real human ears by a nearby point source." *Journal of the Acoustical Society of America* **44(1)**, 240-249 (1968).
46. Shaw, E. A. G. and M. M. Vaillancourt. "Transformation of the sound-pressure from the free field to the eardrum presented in numerical form." *Journal of the Acoustical Society of America* **78(3)**, 1120-1123 (1985).
47. Skudrzyk, Eugen. *The Foundations of Acoustics*, Springer-Verlag, New York-Wein, 1971.
48. *SnapShot User's Manual*, Crystal River Engineering, Aureal Semiconductor, Fremont CA, 1986.
49. Stevens, S. and E. Newman. "The localization of actual sources of sound." *The American Journal of Psychology* **XLVIII**, n. 2, 297-306 (April 1936).
50. *Sysnoise User's Manual: Revision 5.3*. LMS Numerical Technologies, Leuven, Belgium, 1997.
51. Thompson, Silvanus. "On binaural audition." *Philosophical Magazine* **IV**, n. 25, 274-276 (October 1877).

52. Thompson, Silvanus. "On binaural audition- Part II." *Philosophical Magazine* **s. 5 VI**, n. 38, 383-391 (November 1878).
53. Thompson, Silvanus. "On binaural audition- Part III." *Philosophical Magazine* **s. 5 XII**, n. 76, 351-355 (November 1881).
54. Vassallo, Lawrence A. *Comparison of Several Procedures for Measuring Sound Localization Ability*, Masters thesis. The Pennsylvania State University, 1962.
55. WADC Technical Report 57-461. *Sound Absorption at the Surfaces of Small Laboratory Animals*. Wright Air Development Center, June 1957.
56. Wallach, Hans. "The Role of head movements and vestibular and visual cues in sound localization." *Journal of Experimental Psychology* **XXVII**, n. 4, 339-368 (October 1940).
57. Weinrich, Søren Gert. *Sound field calculations around the human head*. Report No. 37, Acoustics Laboratory, Technical University of Denmark, 1984.
58. Wightman, Fred and Doris Kistler. "Headphone simulation of free field listening. I: Stimulus synthesis." *Journal of the Acoustical Society of America*; 85, 858-867(1989).
59. Wightman, Fred and Doris Kistler. "Headphone simulation of free field listening. II: Psychophysical validation." *Journal of the Acoustical Society of America*; 85, 868-878(1989).
60. Wood, Alexander. *Acoustics*, 1st ed. Blackie & Son Limited, Glasgow, 1940 (reprinted 1950).
61. Young, Paul. "Auditory localization with acoustical transposition of the ears." *Journal of Experimental Psychology* **XI**, n. 6, 399-429 (December 1928).
62. Young, Paul. "The rôle of head movement in auditory localization." *Journal of Experimental Psychology* **XIV**, n. 2, 95-124 (April 1931).
63. Zhou, Bin, David M. Green and John C. Middlebrooks. "Characterization of external ear impulse responses using Golay codes." *Journal of the Acoustical Society of America* **92(2)**, 1169-1171 (1992).

Appendix A.

MatLab Scripts

A.1. anspace.m

```
% MatLab script to determine the separation distance between
% the 2 measurement microphones in the impedance tube using
% an anechoic termination and determining the distance
% through the phase of the transfer function.
% Includes calibration section as well.
%
% Brian Katz
% 15-May-1997
%
% The basic premise is from Boden and Abom.
%
% Data variables to use from hp converted '*.mat' (format)
fprintf(1,'\n Checking existence of Hcal1 [Regular position: Calibration Data file] ');
if exist('Hcal1') == 0
    dname = input('data file header name: ','s') ;
    eval(['load ' dname]);
    freq = x ;
    Hcal1 = y ;
end
fprintf(1,'\n Checking existence of Hcal2 [Switched position: Calibration Data file] ');
if exist('Hcal2') == 0
    dname = input('data file header name: ','s') ;
    eval(['load ' dname]);
    if freq ~= x fprintf(1,'Data files do not match!!\n'); end
    Hcal2 = y ;
end
fprintf(1,'\n Checking existence of Hanech [Data file] ');
if exist('Hanech') == 0
    dname = input('data file header name: ','s') ;
    eval(['load ' dname]);
    if freq ~= x fprintf(1,'Data files do not match!!\n'); end
    Hanech = y ;
end
% Determine sound speed due to weather (call whatc.m)
fprintf(1,'\n Using variables [temp,RH,p_mbar] for weather conditions \n')
whatc
%
% Calibration
%
dfreq = freq(2)-freq(1) ;
k = 2*pi*freq/c ;
% avoid divid by zero error
if freq(1) == 0 freq(1)=0.000001; end
% Take geometric mean of the calibration transfer functions
Hcal = sqrt(Hcal1.*Hcal2) ;
```

```

H = Hanech ./ Hcal          ;
%
% Determine spacing
%
% Spacing, as a function of frequency, normalized by frequency
s = -angle(H) ./ k          ;
plot(freq,s,'w')
axis([100 7000 min( s(20:length(s))) max( s(20:length(s))) ]) ;
xlabel('Freq (Hz)')
ylabel('s (m)')
title('Microphone Separation Distance')
% Perform running average
clear s_runav
bar = 25                     ; % Half-Span of average in freq bins
for bin = 2*bar:length(freq)-(2*bar)
    s_runav(bin,:) = [ freq(bin) , mean(s(bin-bar:bin+bar)) ] ;
end
hold on; aveplot = plot(s_runav(:,1),s_runav(:,2),'w-'); hold off
set(aveplot,'LineWidth',1.5)
% User select computed average start/stop
strtf = input('Start s average computation at (Hz): ');
stopf = input('Stop s average computation at (Hz) : ');
if (stopf > max(freq) ) stopf = max(freq); end
% determine bin
strtav = round( (( strtf-freq(1) )/dfreq) + 1);
stopav = round( (( stopf-freq(1) )/dfreq) + 1);
if (stopav > length(freq) ) stopav = length(freq); end
s_av = mean(s(strtav:stopav)) ;
s_max = max(s(strtav:stopav)) ;
s_min = min(s(strtav:stopav)) ;
fprintf(1,'\n Over the selected range \n')
fprintf(1,' For calibrated data: min(s): %0.4f, max(s): %0.4f, mean(s): %0.4f \n',s_min,s_max,s_av);
if (stopav > length(s_runav) ) stopav = length(s_runav); end
sr_av = mean(s_runav(strtav:stopav,2)) ;
sr_max = max(s_runav(strtav:stopav,2)) ;
sr_min = min(s_runav(strtav:stopav,2)) ;
fprintf(1,' For running average: min(s): %0.4f, max(s): %0.4f, mean(s): %0.4f \n',sr_min,sr_max,sr_av);
avline = line([0 max(freq)], [s_av s_av]); set(avline,'LineStyle',':'); set(avline,'Color','w')
avline2 = line([strtf stopf], [s_av s_av]); set(avline2,'LineStyle','-'); set(avline2,'LineWidth',1.5);
set(avline2,'Color','r')
% Optimum frequency for calculated spacing (minimum error)
qline = line([c/(4*s_av) c/(4*s_av)], [0 1]); set(qline,'LineStyle','-'); set(qline,'Color','w')

```

A.2. polynul2.m

```

% MatLab script to interpolate pressure curve from impedance end calibration
% measurement to determine exact frequency of nulls. This is then used to
% determine the exact distance to the end of the duct from the last mic to
% a rigid end.
%
% Brian Katz
% 10-March-1997
%
% Use:  graphically click on the null of choice

```

```

%      script will grab points from left and right and use
%      a poly fit to interpolate
%      The result of this interpolation is then used as the
%      initial guess for a second interpolation.
% Data variables to use from hp converted '*.mat' (format)
fprintf(1,'Checking existance of x [Data file] \n');
if exist('x')== 0
    dname = input('data file header name: ','s')    ;
    eval(['load ' dname]);
    y = abs(y);
end
% determine sound speed due to weather (call whatc.m)
fprintf(1,'Using variables [temp,RH,p_mbar] for weather conditions \n')
whatc
% Configuration variables
miss = 15      ; % points to the left and right of null which are NOT included
incl = 60      ; % points to include in each interpolation
order = 2      ; % polynomial order for the fit
diffy_tol = 0.000001 ; % tolerance for null position determination
% Get the null
clf
plot(y,'w')
axis;axis([ ans(1) length(x) ans(3)-.1 ans(4) ])
title(dname)
fprintf(1,'Click the null you wish to use... \n')
xlabel('Click the null you wish to use')
null = ginput(1)      ; % gets [x,y] coordinate from plot
%null(1) = tmp;
title('Guess')
hold on; plot( null(1),y(null(1)),'w*'); hold off
null(1) = round( null(1) )      ;
fprintf(1,' User guess selected null at %6.2f Hz as the null point guess\n',x(null(1)) )
subplot(1,2,1)
plot(x,y,'w')
xlabel('Freq (Hz)')
ylabel('Normalized Linear pressure')
% Determine partial data to fit from user null guess
leftx = x( null(1)-miss-incl : null(1)-miss )      ;
lefty = y( null(1)-miss-incl : null(1)-miss )      ;
ritex = x( null(1)+miss : null(1)+miss+incl )      ;
ritey = y( null(1)+miss : null(1)+miss+incl )      ;
[leftp,leftS] = polyfit(leftx,lefty,order)      ;
leftnewx = min(leftx)-3*x(miss):max(leftx)+3*x(miss)      ;
leftnewy = polyval(leftp,leftnewx)      ;
[ritep,riteS] = polyfit(ritex,ritey,order)      ;
ritenewx = min(ritex)-3*x(miss):max(ritex)+3*x(miss)      ;
ritenewy = polyval(ritep,ritenewx)      ;
axis;axis([ ans(1) max(x) -.1 ans(4) ])
% Plot
hold on; plot(leftnewx,leftnewy,'r:'); hold off
hold on; plot(ritenewx,ritenewy,'r:'); hold off
title(dname)
% Zoom plot
subplot(1,2,2)
plot(x,y,'w.')

```

```

hold on; plot( x(null(1)),y(null(1)), 'wo'); hold off      % user guess
hold on; plot(leftnewx,leftnewy, 'r:'); hold off
hold on; plot(ritenewx,ritenewy, 'r:'); hold off
title('Zoom')
% Determine intersection
step = 0          ;
stepsize = 1      ; % start freq step size at 1 Hz
flipflop = 0      ; % used to see when search goes to far
oldsign = 0       ;
diffy = 1         ;
f = x(null(1))-5  ; % start search at guess null point
while abs(diffy) > diffy_tol,
    f = f + step      ;
    diffy = polyval(leftp,f) - polyval(ritep,f)      ;
    if ( oldsign ~= sign(diffy) ),
        flipflop = flipflop + 1 ;
        if (flipflop > 4),
            stepsize = stepsize/10 ;
        end
    end
    oldsign = sign(diffy)      ;
    step = sign(diffy)*stepsize ;
end
fprintf(1, ' Estimated pressure null at  %6.2f Hz\n',f)
hold on; plot(f,polyval(leftp,f), 'r+'); hold off
% zoom axis
axis([ f-x(miss/2) f+x(miss/2) polyval(leftp,f)-.1 polyval(leftp,f)+.1 ])
%
% Second go around
%
fprintf(1, ' Second estimation...\n')
null(1) = round( (f-x(1)) / (x(2)-x(1)) );
subplot(1,2,1)
% Determine partial data to fit from user null guess
leftx = x( null(1)-miss-incl : null(1)-miss )      ;
lefty = y( null(1)-miss-incl : null(1)-miss )      ;
ritex = x( null(1)+miss : null(1)+miss+incl )      ;
ritey = y( null(1)+miss : null(1)+miss+incl )      ;
[leftp,leftS] = polyfit(leftx,lefty,order)          ;
leftnewx = min(leftx)-3*x(miss):max(leftx)+3*x(miss) ;
leftnewy = polyval(leftp,leftnewx)                  ;
[ritep,riteS] = polyfit(ritex,ritey,order)          ;
ritenewx = min(ritex)-3*x(miss):max(ritex)+3*x(miss) ;
ritenewy = polyval(ritep,ritenewx)                  ;
% Plot
hold on; plot(leftnewx,leftnewy, 'c'); hold off
hold on; plot(ritenewx,ritenewy, 'c'); hold off
% Zoom plot
subplot(1,2,2)
hold on; plot(leftnewx,leftnewy, 'c'); hold off
hold on; plot(ritenewx,ritenewy, 'c'); hold off
% hold on; plot( x(null(1)),y(null(1)), 'ro'); hold off
% Determine intersection
step = 0          ;
stepsize = 1      ; % start freq step size at 1 Hz

```

```

oldsign = 0      ;
diffy = 1        ;
f = x(null(1))-5 ; % start search at guess null point
while abs(diffy) > diffy_tol,
    f = f + step      ;
    diffy = polyval(leftp,f) - polyval(ritep,f)      ;
    if ( oldsign ~= sign(diffy) ),
        stepsize = stepsize/10 ;
    end
    oldsign = sign(diffy)      ;
%fprintf(1,'here I am\n')
    step = sign(diffy)*stepsize      ;
end
fprintf(1,' Determined pressure null at %6.2f Hz\n',f)
hold on; plot(f,polyval(leftp,f),'co'); hold off
%
% Determine distance from freq
%
% Evaluate for each null position possibility
d1 = (c_dry/f)*(1/4)*100;
d1w = (c_wet/f)*(1/4)*100;
fprintf(1,'1/4 wavelength distance: %8.3f cm (%8.3f cm w/out humid effect)\n',d1w,d1)
d2 = (c_dry/f)*(3/4)*100;
d2w = (c_wet/f)*(3/4)*100;
fprintf(1,'3/4 wavelength distance: %8.3f cm (%8.3f cm w/out humid effect)\n',d2w,d2)
d3 = (c_dry/f)*(5/4)*100;
d3w = (c_wet/f)*(5/4)*100;
fprintf(1,'5/4 wavelength distance: %8.3f cm (%8.3f cm w/out humid effect)\n',d3w,d3)

```

A.3. whatc.m

```

%
% Determine speed of sound due to weather
%
% input parameters
if exist('temp') ~= 1
    temp = input('Tempurare (C): ')      ;
end
if exist('RH') ~= 1
    RH = input('Relative Humidity (%RH): ')      ;
end
if exist('p_mbar') ~= 1
    p_mbar = input('Barometric pressure (mbar): ')      ;
end
%
vp_p = [0.0658    0.1445    59.3471  560.5460]      ; % polyfit of vapor pressure of water
                                                % at T(C) [Pa]
                                                % Determined from Pierce data p.555

pvpT = polyval(vp_p,temp)      ;
p = p_mbar*100      ; % convert mbar to Pa
h = 10^(-2)*RH*pvpT/p      ;
c_dry = 331 + 0.6*temp      ; % variation in speed of sound due to temperature
c_wet = (1 + 0.16*h)*c_dry      ; % variation in speed of sound due to humidity
c = c_wet;

```

```
fprintf(1,' Speed of sound (c) due to weather conditions: %3.2f m/s \n',c);
```

Appendix B.

C Code

B.1. decimate.c

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
struct node_struct {
    int index1;
    int index2;
    float x,y,z;
};
struct elem_struct {
    int index1;
    int index2;
    int vertex[3];
    int valid;
};
struct node_list_struct {
    int *elem;
    int valid;
    int alloc;
};
#define MIN_LENGTH ( 0.000008 ) /* length is a squared value */
#define MAX_LENGTH ( 0.00004 ) /* length is a squared value */
#define Ear_Top ( 0.062 )/* parameters which to avoid data reduction */
#define Ear_Bottom ( -0.009 )
#define Ear_Front ( 0.02 )
#define Ear_Back ( -0.015 )
#define REPORT_INTERVAL ( 20000 )
#define NL_INIT ( 10 ) /* the initial number of elements for each node in the node list */
struct node_struct *n;
struct elem_struct *elem;
struct node_list_struct *node_list;
int numelem, numnode;
void add_to_node_list( int i ) {
    if( node_list[i].valid >= ( node_list[i].alloc ) ) {
        node_list[i].elem = (int *) realloc( (void *)node_list[i].elem, (node_list[i].alloc + 10)*sizeof(int)
    );
        if( NULL == node_list[i].elem ) {
            printf(" Could not allocate node list element %d storage (attempted list of size %d)\n",i,node_list[i].alloc );
            exit ( 1 );
        }
        node_list[i].alloc += 10;
    }
}
void setup_memory() {
    int i;
    n = (struct node_struct *) malloc( (numnode*2)* sizeof( struct node_struct ) );
    if( NULL == n ) {
```



```

    printf(" Could not allocate node storage\n");
    exit ( 1 );
}
node_list = ( struct node_list_struct * ) malloc( (numnode*2)* sizeof( struct node_list_struct
) );
if( NULL == node_list ) {
    printf(" Could not allocate node list storage\n");
    exit ( 1 );
}
for( i = 1; i < (numnode*2) ; i++ ) {
    node_list[i].elem = (int *) malloc( NL_INIT*sizeof(int) );
    if( NULL == node_list[i].elem ) {
        printf(" Could not allocate node list element storage\n");
        exit ( 1 );
    }
    node_list[i].valid = 0;
    node_list[i].alloc = NL_INIT;
}
elem = (struct elem_struct *) malloc( (numelem*2) * sizeof( struct elem_struct ) );
if( NULL == elem ) {
    printf(" Counld not allocate element storage\n" );
    exit( 1 );
}
}
void read_tim_mesh() {
    FILE *infile;
    int i,j;
    int junk;
    int valid;
    infile = fopen( "dec_4.0.tim", "r" );
    if( infile == NULL )
    {
        printf("cannot open input file\n" );
        exit( 0 );
    };
    printf("Decimate: openning input file [format .tim]\n");
    fscanf( infile, "%d %d", &numnode, &numelem );
    setup_memory();
    for( i = 1; i < (numnode+1); i++ ) {
        fscanf( infile, " N%d %f %f %f %d ", &junk, &n[i].x, &n[i].y, &n[i].z, &valid );
        node_list[i].valid = valid;
    }
    for( i= 1; i< (numelem+1); i++ ) {
        fscanf( infile, " E%d %d %d %d %d", &junk,
            &(elem[i].vertex[2]),
            &(elem[i].vertex[1]),
            &(elem[i].vertex[0]), &valid );
        elem[i].valid = valid;
        if( valid > 0 ) {
            add_to_node_list( elem[i].vertex[0] );
            node_list[ elem[i].vertex[0] ].elem[ node_list[ elem[i].vertex[0] ].valid++ ] = i;
            add_to_node_list( elem[i].vertex[1] );
            node_list[ elem[i].vertex[1] ].elem[ node_list[ elem[i].vertex[1] ].valid++ ] = i;
            add_to_node_list( elem[i].vertex[2] );
            node_list[ elem[i].vertex[2] ].elem[ node_list[ elem[i].vertex[2] ].valid++ ] = i;
        }
    }
}

```

```

    }
    printf("READ %d nodes and %d elements\n", numnode, numelem );
}
void read_sys_tria_mesh() {
    FILE *infile;
    int i,j;
    int junk;
    char junkchar[255];
    int valid;
    int mnoel;
    char oneline[255];
    int a,b;
    float c,d,e;
    infile = fopen( "out.free", "r" );
    if( infile == NULL )
    { printf("cannot open input file\n" );
      exit( 0 );
    };
    printf("Decimate: openning .free input file\n");
    fgets( oneline, 255, infile );
    sscanf(oneline,"%s",junkchar); /* Read in header lines as junk */
    printf("first line: %s \n",junkchar);
    fgets( oneline, 255, infile );
    sscanf(oneline,"%s",junkchar); /* Read in header lines as junk */
    printf("second line: %s \n",junkchar);
    fgets( oneline, 255, infile );
    sscanf(oneline,"%s",junkchar); /* Read in header lines as junk */
    printf("third line: %s \n",junkchar);
    fgets( oneline, 255, infile );
    sscanf(oneline,"%s",junkchar); /* Read in header lines as junk */
    printf("fourth line: %s \n",junkchar);
    fscanf( infile, "%d %d %d",&numnode,&numelem,&nnoel );
    printf("numnode:%d numelem:%d \n",numnode,numelem);
    setup_memory();
    fgets( oneline, 255, infile );    sscanf(oneline,"%s",junkchar); /* Read 'NODES' */
    fgets( oneline, 255, infile );
    sscanf(oneline,"%s",junkchar); /* Read 'NODES' */
    printf("node header line: %s \n",junkchar);
    for( i=1; i<numnode+1; i++ ) { /* read the node line. "i" is the internal node number */
        fscanf( infile, "%d %d %f %f %f",&n[i].index1,&n[i].index2,&n[i].x,&n[i].y,&n[i].z );
        node_list[i].valid = 1;
        if( n[i].index1 != n[i].index2 ) {
            printf("Decimate: internal and external numbers must be the same \n");
            printf(" Use renumber-s to fix [on sabine] \n");
            exit(1);
        }
    }
    /* printf("node %d: %d %d %f %f %f \n",i,n[i].index1,n[i].index2,n[i].x,n[i].y,n[i].z ); */
}
    fgets( oneline, 255, infile );    sscanf(oneline,"%s",junkchar); /* Read 'ELEMENTS' */
    fgets( oneline, 255, infile );
    sscanf(oneline,"%s",junkchar); /* Read 'ELEMENTS' */
    printf("element header line: %s \n",junkchar);
    for( i=1; i<numelem+1; i++ ) { /* read the element line */
        fscanf( infile, "%d %d %d %d %d %d",&elem[i].index1,&elem[i].index2,&junk,&junk,
            &elem[i].vertex[0],

```

```

        &elem[i].vertex[1],
        &elem[i].vertex[2] );
    elem[i].valid = 1;
    add_to_node_list( elem[i].vertex[0] );
    node_list[ elem[i].vertex[0] ].elem[ node_list[ elem[i].vertex[0] ].valid++ ] = i;
    add_to_node_list( elem[i].vertex[1] );
    node_list[ elem[i].vertex[1] ].elem[ node_list[ elem[i].vertex[1] ].valid++ ] = i;
    add_to_node_list( elem[i].vertex[2] );
    node_list[ elem[i].vertex[2] ].elem[ node_list[ elem[i].vertex[2] ].valid++ ] = i;
}
printf("READ %d nodes and %d elements\n",numnode,numelem );
}
void read_mesh() {
FILE *infile;
static char oneline[255];
int junk;
int i,j;
float x[4],y[4],z[4];
int n1[3], n2[3], n3[3];
int count;
count = 1;
infile = fopen( "psu-rkxx.g", "r" );
if( infile == NULL )
{   printf("cannot open input file\n" );
    exit( 0 );
};
fgets( oneline, 255, infile );
sscanf( oneline, "%d %d %d %d", &junk, &numnode, &numelem, &junk );
fgets( oneline, 255, infile );
sscanf( oneline, "%d %d", &junk, &junk );
setup_memory();
for( i= 0; i<(numnode/4-0.1); i++ ) {
    fgets( oneline, 255, infile );
    junk = sscanf( oneline, "%f %f %f %f %f %f %f %f %f %f",
        &x[0], &y[0], &z[0],
        &x[1], &y[1], &z[1],
        &x[2], &y[2], &z[2],
        &x[3], &y[3], &z[3]
    );
    for( j = 0; j < (junk/3); j++ ) {
        n[count].x = x[j];
        n[count].y = y[j];
        n[count].z = z[j];
        node_list[count].valid = 0;
        count++;
        if( count % REPORT_INTERVAL == 0 ) printf("Read node %d\n",count);
    }
}
count = 1;
for( i= 0; i<(numelem/3 + 0.9); i++ ) {
    fgets( oneline, 255, infile );
    junk = sscanf( oneline, "%d %d %d %d %d %d %d %d %d",
        &n1[0], &n2[0], &n3[0],
        &n1[1], &n2[1], &n3[1],
        &n1[2], &n2[2], &n3[2]

```

```

    );
    for( j = 0; j < (junk/3); j++ ) {
        elem[count].vertex[0] = n1[j];
        elem[count].vertex[1] = n2[j];
        n3[j] = -n3[j];
        elem[count].vertex[2] = n3[j];
        elem[count].valid = 1;
        add_to_node_list( n1[j] );
        node_list[ n1[j] ].elem[ node_list[ n1[j] ].valid++ ] = count;
        add_to_node_list( n2[j] );
        node_list[ n2[j] ].elem[ node_list[ n2[j] ].valid++ ] = count;
        add_to_node_list( n3[j] );
        node_list[ n3[j] ].elem[ node_list[ n3[j] ].valid++ ] = count;
        count++;
        if( count % REPORT_INTERVAL == 0 ) printf("Read element %d\n",count);
    }
}

printf("READ %d nodes and %d elements\n", numnode, numelem );
fclose(infile);
}

void write_tim_mesh() {
    FILE *outfile;
    int i;
    outfile = fopen( "after_brian.tim", "w" );
    fprintf( outfile, "%d %d\n", numnode, numelem );
    for( i = 1; i < (numnode+1); i++ ) {
        /*      if ( node_list[ i ].valid < 5 ) { */
        fprintf( outfile, "N%d %.3f %.3f %.3f NL%d \n", i, n[i].x, n[i].y, n[i].z, node_list[i].valid
);
        /*      } else
        fprintf( outfile, "N%d %.3f %.3f %.3f -1 \n", i, n[i].x, n[i].y, n[i].z ); */
    }
    for( i = 1; i < (numelem+1); i++ ) {
        if ( elem[ i ].valid == 0 ) {
            fprintf( outfile, "E%d %d %d %d 0 \n", i,
                    elem[i].vertex[2],
                    elem[i].vertex[1],
                    elem[i].vertex[0] );
        } else
            fprintf( outfile, "E%d %d %d %d 1 \n", i,
                    elem[i].vertex[2],
                    elem[i].vertex[1],
                    elem[i].vertex[0] );
    }
    fclose( outfile );
    printf("WROTE %d nodes and %d elements\n", numnode, numelem );
}

void write_mesh() {
    FILE *outfile;
    int i;
    outfile = fopen( "after_brian.iv", "w" );
    fprintf( outfile, "#Inventor V1.0 ascii\n\n");
    fprintf( outfile, "Separator {\n" );
    fprintf( outfile, "    Coordinate3 {\n" );
    fprintf( outfile, "        point [ \n" );

```

```

for( i = 1; i < (numnode+1); i++ ) {
    fprintf( outfile, "          %.3f %.3f %.3f,\n", n[i].x, n[i].y, n[i].z );
}
fprintf( outfile, "      ]\n" );
fprintf( outfile, "    }\n\n\n" );
fprintf( outfile, "    IndexedFaceSet {\n" );
fprintf( outfile, "        coordIndex    [\n" );
for( i = 1; i < (numelem+1); i++ ) {
    if( elem[i].valid == 1 )
        fprintf( outfile, "            %d, %d, %d, -1,\n",      elem[i].vertex[2]-1,
                                elem[i].vertex[1]-1,
                                elem[i].vertex[0]-1 );
}
fprintf( outfile, "        ]\n" );
fprintf( outfile, "    }\n\n\n" );
fprintf( outfile, " }\n\n\n" );
fclose(outfile);
}

void write_sys_mesh() {
FILE *outfile;
int i;
int numnode_sys, numelem_sys, nnoel, itype, nnode;
nnoel = 3;
itype = 4;
nnode = 3;
numnode_sys = 0;
numelem_sys = 0;
/* Count the number of valid nodes and elements */
for( i = 1; i < (numnode+1); i++ ) {
    if( node_list[i].valid > 0 ) numnode_sys++;
}
for( i = 1; i < (numelem+1); i++ ) {
    if( elem[i].valid > 0 ) numelem_sys++;
}

    outfile = fopen( "after_brian.free", "w" );
    printf("Writing sysnoise output file...");
    fprintf( outfile, "SYSNOISE MESH FILE\n" );
    fprintf( outfile, "Rev 5.1 HP-UX/7xx  30-SEP-94\n" );
    fprintf( outfile, "Possibly fixed head model\n" );
    fprintf( outfile, " 1-JAN-2000 12:00:00\n" );
    fprintf( outfile, "%10d%10d%10d\n", numnode_sys, numelem_sys, nnoel );
fprintf( outfile, "NODES\n" );
for( i = 1; i < (numnode+1); i++ ) {
    if( node_list[i].valid > 0 ) {
        fprintf( outfile, "%10d%10d%20.08E%20.08E%20.08E\n", i, i, n[i].x, n[i].y, n[i].z);
    }
}
fprintf( outfile, "ELEMENTS\n" );
for( i = 1; i < (numelem+1); i++ ) {
    if( elem[i].valid > 0 ) {
        fprintf( outfile, "%10d%10d%10d%10d%10d%10d\n", i, i, itype, nnode,
                                elem[i].vertex[0], elem[i].vertex[1], elem[i].vertex[2] );
    }
}
printf("Wrote %d Nodes and %d Elements\n", numnode_sys, numelem_sys);

```

```

    fclose(outfile);
}
float get_length( int p1, int p2 ) {
    static float dx,dy,dz;
    static float answer;
    if( ( p1 < 1 ) || ( p2 < 1 ) || ( p1 > numnode ) || ( p2 > numnode ) || ( p1 == p2 ) ) return(
1000 );
    dx = n[p1].x - n[p2].x;
    dy = n[p1].y - n[p2].y;
    dz = n[p1].z - n[p2].z;
    answer = (dx*dx + dy*dy + dz*dz);
    return( answer );
}
int check_for_redundacy( int elem_n ) {
    int v1,v2,v3;
    if( elem[ elem_n ].valid == 0 ) return( 0 );
    v1 = elem[ elem_n ].vertex[0];
    v2 = elem[ elem_n ].vertex[1];
    v3 = elem[ elem_n ].vertex[2];
    if( ( v1 == v2 ) || ( v2 == v3 ) || ( v3 == v1 ) ||
        (node_list[v1].valid == 0) || (node_list[v3].valid == 0) || (node_list[v2].valid == 0) ) {
        elem[ elem_n ].valid = 0;
        return( 1 );
    } else {
        return( 0 );
    };
}
void rereference( int oldv1, int oldv2, int newv ) {
    int i,j;
    int current_elem;
    /* first change the references */
    for( i = 0; i < node_list[ oldv1 ].valid; i++ ) {
        current_elem = node_list[ oldv1 ].elem[i];
        for( j = 0; j < 3; j++ )
            if( elem[current_elem].vertex[j] == oldv1 ) elem[current_elem].vertex[j] = newv;
    }
    if( oldv1 == oldv2 ) goto next_step;
    for( i = 0; i < node_list[ oldv2 ].valid; i++ ) {
        current_elem = node_list[ oldv2 ].elem[i];
        for( j = 0; j < 3; j++ )
            if( elem[current_elem].vertex[j] == oldv2 ) elem[current_elem].vertex[j] = newv;
    }
    /* then check any elements for redundant verticies */
next_step:
    for( i = 0; i < node_list[ oldv1 ].valid; i++ ) check_for_redundacy( node_list[ oldv1 ].elem[i]
);
    if( oldv1 == oldv2 ) goto next_step2;
    for( i = 0; i < node_list[ oldv2 ].valid; i++ ) check_for_redundacy( node_list[ oldv2 ].elem[i]
);
next_step2:
    node_list[ oldv1 ].valid = 0;
    node_list[ oldv2 ].valid = 0;
}
void rereference2( int oldv, int newv ) {
    /* routine will change references of node oldv to newv and "remove" oldv */

```

```

/* this is a rewrite of 'rereference' which is missing some features */
int i,j,k,flag;
int current_elem;
/* first change the element references from oldv to newv */
for( i = 0; i < node_list[ oldv ].valid; i++ ) {
    current_elem = node_list[ oldv ].elem[i];
    for( j = 0; j < 3; j++ ) {
        if( elem[current_elem].vertex[j] == oldv ) elem[current_elem].vertex[j] = newv;
    }
    flag = 0;
    for( k = 0; k < node_list[ newv ].valid; k++ ) {
        if( node_list[ newv ].elem[ k ] == current_elem ) flag = 1;
    }
    if ( flag == 0 ) {
        add_to_node_list( newv );
        node_list[ newv ].elem[ node_list[ newv ].valid++ ] = current_elem;
    }
}
/* check any elements using newv for redundant vertices (ie. collapse of element) */
for( i = 0; i < node_list[ newv ].valid; i++ ) check_for_redundacy( node_list[ newv ].elem[i]
);
/* clear nodelist of oldv */
node_list[ oldv ].valid = 0;
}

void remove_element( int elem_n ) {
    float tempx, tempy, tempz;
    int v1,v2,v3;
    v1 = elem[elem_n].vertex[0];
    v2 = elem[elem_n].vertex[1];
    v3 = elem[elem_n].vertex[2];
    /* collapse the three vertices to the middle of the element */
    tempx = ( n[v1].x + n[v2].x + n[v3].x ) /3.0;
    tempy = ( n[v1].y + n[v2].y + n[v3].y ) /3.0;
    tempz = ( n[v1].z + n[v2].z + n[v3].z ) /3.0;
    n[v1].x = n[v2].x = n[v3].x =tempx;
    n[v1].y = n[v2].y = n[v3].y =tempy;
    n[v1].z = n[v2].z = n[v3].z =tempz;
    /* change all references to v2 and v3 to v1 */
    rereference2( v2, v1 );
    rereference2( v3, v1 );
    elem[elem_n].valid = 0;
}

void remove_edge( node1, node2, elem_n ) {
    float newx, newy, newz;
    /* Collapse 2 nodes into 1 (midpoint) thereby removing 2 elements after validate is performed */
    newx = ( n[node1].x + n[node2].x ) / 2.0;
    newy = ( n[node1].y + n[node2].y ) / 2.0;
    newz = ( n[node1].z + n[node2].z ) / 2.0;
    n[node2].x = newx;
    n[node2].y = newy;
    n[node2].z = newz;
    rereference2( node1, node2 );
    elem[ elem_n ].valid = 0;
}

void decimate_mesh() {

```

```

/* This routine is used to remove small elements */
int i,j;
float temp12, temp23, temp13;
int n1,n2,n3;
int killed_edges, tot_killed_edges;
    printf("Decimate: decimate_mesh\n");
    tot_killed_edges = 0;
    for( j = 5; j > 0 ; j-- ) {
        printf("Decimation: Pass %d\n", (6-j));
        killed_edges = 0;
        for( i = 1; i< (numelem+1) ; i=i+j ) {
/* This incrementation is very clever */
            if( i % REPORT_INTERVAL == 0 ) printf("decimating %d\n",i);
            if( elem[i].valid == 1 ) {
                n1 = elem[i].vertex[0];
                n2 = elem[i].vertex[1];
                n3 = elem[i].vertex[2];
/*
                if( ( (n[n1].y > Ear_Top) || (n[n1].y < Ear_Bottom) ) ||
                    ( (n[n1].z > Ear_Front) || n[n1].z < Ear_Back) )
*/
                {
                    temp12 = get_length( n1, n2 );
                    temp23 = get_length( n2, n3 );
                    temp13 = get_length( n1, n3 );
                    if( temp12 < MIN_LENGTH/j ) {
                        remove_edge( n1, n2, i );
                        killed_edges++;
                        goto NEXT_ELEM;
                    }
                    if( temp23 < MIN_LENGTH/j ) {
                        remove_edge( n2, n3, i );
                        killed_edges++;
                        goto NEXT_ELEM;
                    }
                    if( temp13 < MIN_LENGTH/j ) {
                        remove_edge( n1, n3, i );
                        killed_edges++;
                        goto NEXT_ELEM;
                    }
                    } NEXT_ELEM: ;
                }
            }
        }
        tot_killed_edges = tot_killed_edges + killed_edges;
        printf("Removed %d edges\n", killed_edges );
    }
    printf("Decimation: Removed %d edges\n", tot_killed_edges );
}

void validate_mesh() {
    int i;
    printf("Decimate: validate_mesh\n");
    for( i= 1; i< (numelem+1) ; i++ ) {
        if( i % REPORT_INTERVAL == 0 ) printf("validating %d\n",i);
        if( check_for_redundacy(i) == 1 ) printf(" Invalid element removed (%d)\n", i );
    }
}

```



```

void squish() {
    int i,j;
    int v1, v2, v3;
    float tempx, tempy, tempz;
    float dx,dy,dz;
    for( i= 1; i< (numelem+1) ; i++ ) {
        if( i % REPORT_INTERVAL == 0 ) printf("squishing %d\n",i);
        v1 = elem[i].vertex[0];
        v2 = elem[i].vertex[1];
        v3 = elem[i].vertex[2];
        /* the center of the element */
        tempx = ( n[v1].x + n[v2].x + n[v3].x ) /3.0;
        tempy = ( n[v1].y + n[v2].y + n[v3].y ) /3.0;
        tempz = ( n[v1].z + n[v2].z + n[v3].z ) /3.0;
        n[v1].x = ( tempx + n[v1].x ) /2.0;
        n[v2].x = ( tempx + n[v2].x ) /2.0;
        n[v3].x = ( tempx + n[v3].x ) /2.0;
        n[v1].y = ( tempy + n[v1].y ) /2.0;
        n[v2].y = ( tempy + n[v2].y ) /2.0;
        n[v3].y = ( tempy + n[v3].y ) /2.0;
        n[v1].z = ( tempz + n[v1].z ) /2.0;
        n[v2].z = ( tempz + n[v2].z ) /2.0;
        n[v3].z = ( tempz + n[v3].z ) /2.0;
    }
}

void check_nodes() {
    /* Merges superimposed nodes */
    int i, j,k;
    printf("Decimate: check_nodes\n");
    for(k = 0; k < 3; k++ ) {
        for(i = 1; i < numnode+1; i ++ ) {
            if( i % REPORT_INTERVAL == 0 ) printf("PASS %d: checking %d\n",k+1,i);
            for( j = i+1; j < numnode+1; j++ ) {
                if( get_length( i,j ) < 0.015 ) {
                    rereference2( i, j );
                    printf("merging node %d to node %d\n",i,j );
                    goto NEXT_NODE;
                }
            }
        }
        NEXT_NODE: ;
    }
}

void special_merge() {
    /* rereference( 176, 176, 2 ); */
    int i,j,k,count;
    count = 0;
    for(i = 1; i < numnode+1; i ++ ) {
        /* Top of the head node */
        if( (i != 89088) && (n[i].x < 9.6 ) && (n[i].z < 20.1) && (n[i].y < 163.8 )
            && (n[i].x > 9.4 ) && (n[i].z > 19.9) && (n[i].y > 163.5 ) ) {
            rereference2( i, 89088);
            count++;
            printf( "Merged node %d. Merge #%d. node_list[89088].valid = %d\n",
                i,count,node_list[89088].valid);
        }
    }
}

```

```

    }
/* Bottom of the head node */
    if( (i != 349) && (n[i].x < 0.2 ) && (n[i].z < 20.1) && (n[i].y < -106.7 )
        && (n[i].x > 0.0 ) && (n[i].z > 19.9) && (n[i].y > -106.5 ) ) {
        rereference2( i, 349);
        count++;
        printf( "Merged node %d. Merge #%d. node_list[89088].valid = %d\n",
            i,count,node_list[89088].valid);
    }
}
count = 0;
for(i = 1; i < numnode+1; i ++ ) {
    if( (i != 349) && (n[i].x < 0.2 ) && (n[i].z < 20.1) && (n[i].y < -106.7 )
        && (n[i].x > 0.0 ) && (n[i].z > 19.9) && (n[i].y > -106.5 ) ) {
        rereference2( i, 349);
        count++;
        printf( "Merged node %d. Merge #%d. node_list[89088].valid = %d\n",
            i,count,node_list[89088].valid);
    }
}
printf( "\n Merged specific nodes \n" );
}

void cut_elem( elem_n, edge_end1, edge_end2, elem_node ) {
int i,j;
float newx, newy, newz;
int other_elem, other_elem_node;
    other_elem = -1;
    other_elem_node = -1;
/* Create new node at midpoint of selected edge */
    newx = ( n[edge_end1].x + n[edge_end2].x ) / 2.0;
    newy = ( n[edge_end1].y + n[edge_end2].y ) / 2.0;
    newz = ( n[edge_end1].z + n[edge_end2].z ) / 2.0;
    numnode++;
    n[numnode].x = newx;
    n[numnode].y = newy;
    n[numnode].z = newz;
/* Determine the element sharing the edge with elem_n */
    for( i = 0; i < node_list[ edge_end1 ].valid; i++ ) {
        for( j = 0; j < node_list[ edge_end2 ].valid; j++ ) {
            if( (node_list[edge_end1].elem[i] == node_list[edge_end2].elem[j] ) &&
                (elem[node_list[edge_end1].elem[i]].valid == 1) &&
                (node_list[edge_end1].elem[i] != elem_n) ) {
                other_elem = node_list[edge_end1].elem[i];
            }
        }
    }
    if( other_elem == -1 ) {
        printf("Could not determine other element for E%d with N%d and N%d\n",
            elem_n,edge_end1,edge_end2);
        goto ERROR;
    }
/* Determine the node on other_elem which is not on the edge */
    for( i = 0; i < 3; i++ ) {
        if( ( elem[ other_elem ].vertex[i] == edge_end1 ) ||
            ( elem[ other_elem ].vertex[i] == edge_end2 ) ){

```

```

        } else
            other_elem_node = elem[ other_elem ].vertex[i];
    }
    if( other_elem_node == -1 ) {
        printf("Could not determine other element node for E%d with E%d\n",other_elem,elem_n);
        goto ERROR;
    }
/* Assign new node connections for the 2 old elements and create 2 new elements */
    elem[elem_n].vertex[0] = numnode;
    elem[elem_n].vertex[1] = edge_end1;
    elem[elem_n].vertex[2] = elem_node;
    elem[other_elem].vertex[0] = numnode;
    elem[other_elem].vertex[1] = other_elem_node;
    elem[other_elem].vertex[2] = edge_end1;
    numelem++;
    elem[numelem].vertex[0] = numnode;
    elem[numelem].vertex[1] = edge_end2;
    elem[numelem].vertex[2] = elem_node;
    elem[numelem].valid = 1;
    numelem++;
    elem[numelem].vertex[0] = numnode;
    elem[numelem].vertex[1] = other_elem_node;
    elem[numelem].vertex[2] = edge_end2;
    elem[numelem].valid = 1;
/* Update node_list for old and new nodes */
    node_list[numnode].valid = 4;
    node_list[numnode].elem[0] = elem_n;
    node_list[numnode].elem[1] = other_elem;
    node_list[numnode].elem[2] = (numelem-1);
    node_list[numnode].elem[3] = numelem;
    for( i = 0; i < node_list[edge_end2].valid; i++ ) {
        if( node_list[edge_end2].elem[i] == elem_n ) {
            node_list[ edge_end2 ].elem[ i ] = (numelem-1);
        }
        if( node_list[edge_end2].elem[i] == other_elem ) {
            node_list[ edge_end2 ].elem[ i ] = numelem;
        }
    }
    add_to_node_list( elem_node );
    node_list[ elem_node ].elem[ node_list[ elem_node ].valid++ ] = (numelem-1);
    add_to_node_list( other_elem_node );
    node_list[ other_elem_node ].elem[ node_list[ other_elem_node ].valid++ ] = numelem;
/*
    printf("Created N%d, E%d and E%d from E%d and E%d\n",
        numnode, (numelem-1), numelem, elem_n, other_elem);
*/
    ERROR;;
}
void refine() {
/* This routine is used to break up large elements */
    int i,j;
    float temp1, temp2, temp3;
    int n1,n2,n3;
    int killed_edges,tot_killed_edges,numelem_old;
    printf("Decimate: refine\n");

```

```

    killed_edges = 999;
    tot_killed_edges = 0;
    while( killed_edges > 1 ) {
printf("Killed %d\n",killed_edges);
        for( j = 7; j > 0; j-- ) {
            killed_edges = 0;
            numelem_old = numelem;
            for( i= 1; i< (numelem_old) ; i=i+j ) {
                if( elem[i].valid == 1 ) {
                    n1 = elem[i].vertex[0];
                    n2 = elem[i].vertex[1];
                    n3 = elem[i].vertex[2];
                    temp1 = get_length( n1, n2 );
                    temp2 = get_length( n2, n3 );
                    temp3 = get_length( n1, n3 );
                    if( (temp1 > MAX_LENGTH*j) && (temp1 > temp2) && (temp1 > temp3) ) {
                        cut_elem( i, n1, n2, n3 );
                        killed_edges++;
                        goto NEXT_ELEM;
                    }
                    if( (temp2 > MAX_LENGTH*j) && (temp2 > temp1) && (temp2 > temp3) ) {
                        cut_elem( i, n2, n3, n1 );
                        killed_edges++;
                        goto NEXT_ELEM;
                    }
                    if( (temp3 > MAX_LENGTH*j) && (temp3 > temp1) && (temp3 > temp2) ) {
                        cut_elem( i, n1, n3, n2 );
                        killed_edges++;
                        goto NEXT_ELEM;
                    }
                }
                NEXT_ELEM: ;
            }
            tot_killed_edges = tot_killed_edges + killed_edges;
        }
    }
    printf("Cut %d edges\n", tot_killed_edges );
}

void fix_node_list(node_n) {
    int i,j;
    int temp_list[ 500 ];
    int temp_valid;
    for( i = 0; i < 500; i++ ) temp_list[i] = 0;
    temp_valid = 0;
    /* Check validity of each element in node_list and remove invalids */
    for( i = 0; i < node_list[node_n].valid; i++ ) {
        if( elem[ node_list[node_n].elem[i] ].valid > 0 ) {
            temp_list[ temp_valid ] = node_list[node_n].elem[i];
            temp_valid++;
        }
        node_list[node_n].elem[i] = 0;
    }
    for( i = 0; i < temp_valid; i++ ) node_list[node_n].elem[i] = temp_list[i];
    node_list[node_n].valid = temp_valid;
}

```

```

void fix_node_elem(node_n) {
/* If a node has < 3 elements connected remove all the elements as bogus */
int i;
if( node_list[node_n].valid < 3 ) {
    for( i = 0; i < node_list[node_n].valid; i++ ) {
        elem[ node_list[node_n].elem[i] ].valid = 0;
        printf("Fixed node %d\n",node_n);
    }
    node_list[node_n].valid = 0;
}
}

void check_node_list () {
int i;
printf("Decimate: check_node_list\n");
for( i = 0; i < numnode; i++ ) {
    fix_node_list( i );
    fix_node_elem( i );
}
}

void flatten_ears () {
/* flatten ear section to common y value */
int i,j;
int n1,n2,n3;
j = 0;
printf("Decimate: flatten_ears \n");
for( i = 1; i < (numelem+1) ; i++ ) {
    if( elem[i].valid == 1 ) {
        n1 = elem[i].vertex[0];
        n2 = elem[i].vertex[1];
        n3 = elem[i].vertex[2];
        if( ( n[n1].z < Ear_Top) && (n[n1].z > Ear_Bottom) ) &&
            ( (n[n1].x < Ear_Front) && n[n1].x > Ear_Back) ) {
            if( n[n1].y > 0 ) {
                n[n1].y = 0.0747 - (0.062 - n[n1].z)*(0.0131/0.071);
                n[n2].y = 0.0747 - (0.062 - n[n2].z)*(0.0131/0.071);
                n[n3].y = 0.0747 - (0.062 - n[n3].z)*(0.0131/0.071);
                j++;
            }
            if( n[n1].y < 0 ) {
                n[n1].y = -0.0742 + (0.0634 - n[n1].z)*(0.0108/0.071) ;
                n[n2].y = -0.0742 + (0.0634 - n[n2].z)*(0.0108/0.071) ;
                n[n3].y = -0.0742 + (0.0634 - n[n2].z)*(0.0108/0.071) ;
                j++;
            }
        }
    }
}
printf(" Flattened %d nodes \n",j);
}

void main() {
    read_sys_tria_mesh();
/*   read_tim_mesh();   */
/*   read_mesh();       */
/*   check_nodes();     */
/*   special_merge();    */
}

```

```

    decimate_mesh();
    flatten_ears();
    decimate_mesh();
/*    squish();        */
/*    refine();        */
/*    check_node_list();    */
/*    decimate_mesh();    */
/*    validate_mesh();    */
/*    check_node_list();
/*    write_mesh();        */
/*    write_tim_mesh();    */
    write_sys_mesh();
}

```

B.2. edgealign.c

```

/*
** Purpose
**     Program to detect all the edges of a free format mesh, and align them to a
**     common plane
** History
** File : edgealign.c
** Brian Katz, Acoustics Dept, The Pennsylvania State University
** Last modified November 20, 1995
**     File : mesh2free.c
** Limitations/Assumptions
** The length of a single source line is limited to 255 characters
**     The number of verticies and faces are statically allocated.
**     Each face is assumed to be a triangle
** All vertex numbers will be positive
** Usage
**
** Invoke the program with three command line arguments as follows
**
**     edgealign infile outfile scale_factor edgefile
**
**     The file "infile" will be processed and the file "outfile" will be produced.
** Note the output file will overwrite an existing file without prompt or warning.
** scale_factor is a floating point number which is used to scale all the coordinates
** of each vertex.  (e.g. a scale_factor of 10 will make the object 10 times larger.)
**
** Terminology
** In switching the code to Sysnoise format some terms have been added with the
** following identitities
**     node = vertex
**     element = face
**
*/
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
/* the maximum number of verxticies and faces */
#define MAXNODE 10000
#define MAXELEM 100000

```

```

/* a node structure */
struct node_struct {
    int index1;      /* the given internal index number of the node */
    int index2;      /* the given external index number of the node */
    int valid;       /* an integer indicating if the node has valid data */
    float x,y,z;     /* the three coordinates of the node */
};

/* the structure of an element */
struct elem_struct {
    int index1;      /* the internal element number */
    int index2;      /* the external element number */
    int vertex[3];   /* an array of the three node indexes that define the element */
    struct node_struct normal; /* a vector defining the normal of the face */
};

void main( int argc, char **argv ) {
    struct node_struct n[MAXNODE]; /* an array of nodes */
    struct elem_struct elem[MAXELEM]; /* an array of all the elements */
    FILE *infile, *outfile, *edgefile; /* the input and output files */
    char oneline[255]; /* one text line from the file */
    float scale_factor; /* the scale factor */
    struct node_struct A,B; /* the vectors used to calculate element Normals */
    int n1,n2,n3; /* the three node indicies that define an element */
    int i=0,j=1,count=1; /* counters */
    float x,y,z; /* temporary variable to store x,y,z values */
    float normal_length; /* the length of the normal vector */
    int numelem=0, numnode=0; /* the total number of elements and nodes */
    int i1=0, i2=0; /* internal and external indexes */
    int badnode; /* stroage of a bad vertex reference in a face */
    int mnoel=3; /* max number of nodes/element (given data = 3) */
    int itype=4, nnode=3; /* element type and number of connecting nodes (given data) */
    char junk[255]; /* junk string */
    int junknum=0; /* junk number string */
    int edgecount[MAXNODE]; /* array to count which nodes on edges */
    int edgelist[MAXNODE]; /* list of nodes on edge */
    /* if the program is not used properly report the correct usage */
    if( argc < 5 ) {
        printf("USAGE : convert infile outfile scale-factor edgefile\n");
        return;
    }
    sscanf(argv[3],"%f",&scale_factor); /* set the scale factor */
    /* initialize the vertex data */
    for( i=0; i< MAXNODE; i++ ) n[i].valid = 0;
    /* open the input file */
    infile = fopen(argv[1], "r" );
    if( NULL == infile ) {
        printf("Error opening input file %s\n",argv[1]);
        return;
    }
    /* open the output file */
    outfile = fopen(argv[2], "w" );
    if( NULL == outfile ) {
        printf("Error opening output file %s\n",argv[2]);
        return;
    }
    /* open the edgefile */

```

```

edgefile = fopen(argv[4], "w" );
if( NULL == edgefile ) {
    printf("Error opening edgefile %s\n",argv[4]);
    return;
}
/* Read input file */
fgets( oneline, 255, infile );
sscanf(oneline,"%s",&junk);    /* Read in header lines as junk */
fgets( oneline, 255, infile );
sscanf(oneline,"%s",&junk);
fgets( oneline, 255, infile );
sscanf(oneline,"%s",&junk);
fgets( oneline, 255, infile );
sscanf(oneline,"%s",&junk);
fgets( oneline, 255, infile );
sscanf(oneline,"%d %d %d",&numnode,&numelem,&nnoel );
fgets( oneline, 255, infile );
sscanf(oneline,"%s",&junk);    /* Read 'NODES' */
for( i=1; i<numnode+1; i++ ) {
    /* read the node line. "i" is the node number */
    fgets( oneline, 255, infile );
    sscanf(oneline,"%d %d %f %f %f",&i1,&i2,&x,&y,&z );
    n[i].x = x;        /* x value */
    n[i].y = y;        /* y value */
    n[i].z = z;        /* z value */
    n[i].valid = 1;    /* mark as valid node */
    n[i].index1 = i1;   /* node index1 number */
    n[i].index2 = i2;   /* node index2 number */
}
fgets( oneline, 255, infile );
sscanf(oneline,"%s",&junk);    /* Read 'ELEMENTS' */
for ( j=1; j<numelem+1; j++ ) {
    /* read the face line */
    fgets( oneline, 255, infile );
    sscanf(oneline,"%d %d %d %d %d %d",&i1,&i2,&junknum,&junknum,&n1,&n2,&n3 );
    elem[j].index1 = i1;
    elem[j].index2 = i2;
    elem[j].vertex[0] = n1;    /* element vertex #1 */
    elem[j].vertex[1] = n2;    /* element vertex #2 */
    elem[j].vertex[2] = n3;    /* element vertex #3 */
}
/* Determine Edges */
for ( j=1; j<numelem+1; j++ ) {
    edgecount[ elem[j].vertex[0] ]++ ;
    edgecount[ elem[j].vertex[1] ]++ ;
    edgecount[ elem[j].vertex[2] ]++ ;
}
for (j=1; j<MAXNODE; j++ ) {
    if (edgecount[j]>0) {
        fprintf(edgefile,"%d\n",edgecount[j]);
    }
}
/* print Sysnoise Free Format Header */
fprintf(outfile,"SYSNOISE MESH FILE\n");    /* FTYPE File Type */
fprintf(outfile,"Rev 5.1 HP-UX/7xx 30-SEP-94\n");    /* REVIS Software Revision */

```



```

fprintf(outfile,"%s\n",argv[2]);          /* TITLE Model Title */
fprintf(outfile," 1-JAN-2000 12:00:00\n"); /* CDATE Creation Date -bogus currently */
fprintf(outfile,"%10d%10d%10d\n",numnode,numelem,nnoel) ; /* NPOIN,NELEM,NNOEL */
/* print the Sysnoise Free Node/Element format */
fprintf(outfile,"NODES\n") ;
for( i=0; i<MAXNODE; i++ ) {
    if( n[i].valid == 1) {
        fprintf(outfile,"%10d%10d%20.08E%20.08E%20.08E\n",
            i,n[i].index2,n[i].x,n[i].y,n[i].z) ;
    }
}
fprintf(outfile,"ELEMENTS\n") ;
for( j=1; j<(numelem+1); j++ ) {
    fprintf(outfile,"%10d%10d%10d%10d%10d%10d\n",j,j,itype,nnode,
        elem[j].vertex[0],elem[j].vertex[1],elem[j].vertex[2]) ;
}
printf("Processed %d Nodes and %d Elements\n",numnode,numelem);
fclose(infile);
fclose(outfile);
fclose(edgefile);
}

```

Appendix C.

Sysnoise Related Command Files

C.1. Example setup command file for head mesh

Sysnoise command file to set-up a BEM model with the head mesh, hair definitions, overdetermination points, source, and receiver positions.

```
Option BEM Indirect Return
Input Mesh Format Free File 'head.free' Return
{ Check mesh for degenreacy
Check Mesh Set Domain Return
Overdetermine
  Plane 0 0 0 To 0 0.01 0 Divide 10 To 0.01 0 0.01 Divide 10
  Return
Boundary Admittance Real 1000 Imag 0
  Elements Type QUAD4
  Singular
  Return
Fmax return
Extract Sets Return
{ Define hair set (10)
Read 'hair_define.cmd'
{ Load in measured admittance data
Table 1 File 'admdat1.asc'
{ Assign hair admittance to hair set
Boundary Set 10 Admittance Table 1 Return
{ Left ear canal source point
Source Spherical Amplitude 1 Node 7223 Return
{ Measurement sphere of points
Input Point Format Free File 'thesphere.free' Return
Extract Source Return
Extract Summary Return
Extract Points All Return
```

C.2. Hair definition command file

Sysnoise command file to define the hair regions on the meshes.

```
{ Define sets for top and back of head for impedance deinitions
{ Back Hair
set 11 name "back hair pt 1" element between x = -1.41E-2 , -2 return
set 12 name "back hair pt 2" element between z = -4.65E-2 , 2 return
set 13 name "back hair" intersection 11 12 return
{ Top Hair
set 14 name "top hair" element between z = 6.56E-2 , 2 return
{ Hair
set 15 name "possible hair" union 13 14 return
set 10 name "hair" intersection 1 15 return
{ Please define admittance condition on set 10 for hair effects if desired
```

C.3. Example setup command file for sphere mesh

Sysnoise command file to set-up a BEM model with a sphere mesh, hair definitions, overdetermination points, source, and receiver positions. An extra set of overdetermination elements is used for the sphere. This is necessary as the high symmetry of the sphere causes an increase in the errors due to internal modes.

```
Option BEM Indirect Return
Point Sphere 1E-2 -0.0035 2.99e-2 Radius 0.0655 Divide 20 Return
Output Point Format Free File finesphere.free Return
Input Mesh Format Free File finesphere.free Return
Read 'hair_define.cmd'
Overdetermine
  Plane 0 0 0 To 0 .01 0 Divide 10 To .01 0 .01 Divide 10
  Return
Overdetermine
  Plane 0.05 0 0 To 0.05 0.01 0 Divide 10 To 0.06 0 0.01 Divide 10
  Return
Check Mesh Set Return
Boundary Admittance Real 1000 Imag 0
  Elements Set 5
  Elements Set 22
  Singular
  Return
Set 6 Name "left_ear"
  Nodes Near 0.98e-2,6.20e-2,2.30e-2
  Return
Extract Set 6 return
{ Put source here at set 6
{ Measurement sphere of points
Input Point Format Free File 'thesphere.free' Return
Exit Save Journal run40.journal Save Models
```

C.4. Example computation command file

Sysnoise command file to calculate the frequency response (HRTF) for a given model.

```
Open Model 1 File head.sdb Return
Response
  Frequency 1000 to 6000 linstep 400
  Save Potentials Step 1
  Save Displacements Step 1
  Save Results Step 1
  Near 2
  Far 5
  Quadrature 2 2 1
  Output Results Step 1 File 'run23fr' Format Free
  Return
Exit Save Models
```

C.5. Unix sysmon script

This Unix shell script controls the running of a number of Sysnoise processes. Due to the long running time of the calculations it was desired that each frequency be run as a independent self-contained job so that in the event of a system failure, or the process being cancelled in the event that others needed the full machine, the total calculation would not suffer greatly. The script shown monitors the Sysnoise jobs running (on a four processor machine) and in the event that none are currently running proceeds with the next set of frequencies. Usage requires a file which lists Sysnoise command files to be executed (similar to C.4 where only a single frequency for calculation is listed) and another file which keeps track of where in the list file the current run is. Finally, a log file is kept to check on progress. An example log file is shown in C.6. Execution is done by periodically running the script using cron.

```
#!/usr/bin/tcsh
cd /usr/people/bfk/sysruns
set sysmon_current_runs='ps -u bfk | grep -ic sysnoise'
set sysmon_runpos='cat sysmon.runpos'
set sysmon_totruns='cat sysmon.list | grep -ic r'
touch sysmon.log
touch sysout.log.1
touch sysout.log.2
touch sysout.log.3
touch sysout.log.4
date >> sysmon.log
if ( $sysmon_current_runs > 1 ) then
    echo "Sysnoise is running: do nothing" >> sysmon.log
else
    set sysmon_runpos='expr $sysmon_runpos + 1'
    echo $sysmon_runpos > sysmon.runpos
    if ( $sysmon_runpos > $sysmon_totruns ) then
        echo "End of batch file list" >> sysmon.log
        echo "    Resetting sysmon.runpos" >> sysmon.log
        echo "0" > sysmon.runpos
        crontab -r
    else
        rm /usr/old_drive/bfk/tmp/*.TMP
        set sysmon_line='tail +$sysmon_runpos sysmon.list | head -n 1'
        nice -10 sysnoise -nogui -m11 < $sysmon_line >>& sysout.log.1 &
        echo $sysmon_runpos $sysmon_line >> sysmon.log
        set sysmon_runpos='expr $sysmon_runpos + 1'
        set sysmon_line='tail +$sysmon_runpos sysmon.list | head -n 1'
        echo "nice -10 sysnoise -nogui -m11 < $sysmon_line >> sysout.log.2" | at now + 5 min
        echo $sysmon_runpos $sysmon_line >> sysmon.log
        set sysmon_runpos='expr $sysmon_runpos + 1'
        set sysmon_line='tail +$sysmon_runpos sysmon.list | head -n 1'
        echo "nice -10 sysnoise -nogui -m11 < $sysmon_line >> sysout.log.3" | at now + 10 min
        echo $sysmon_runpos $sysmon_line >> sysmon.log
        set sysmon_runpos='expr $sysmon_runpos + 1'
        set sysmon_line='tail +$sysmon_runpos sysmon.list | head -n 1'
        echo "nice -10 sysnoise -nogui -m11 < $sysmon_line >> sysout.log.4" | at now + 15 min
        echo $sysmon_runpos $sysmon_line >> sysmon.log
        echo $sysmon_runpos > sysmon.runpos
    endif
endif
```

```
endif
```

C.6. Example sysmon log file

```
Fri Apr  4 10:25:36 EST 1997
1 head0800
2 head0900
3 head1000
4 head1100
Fri Apr  4 19:00:03 EST 1997
Sysnoise is running: do nothing
Sat Apr  5 00:00:03 EST 1997
Sysnoise is running: do nothing
Sat Apr  5 07:00:03 EST 1997
Sysnoise is running: do nothing
Sat Apr  5 19:00:03 EST 1997
Sysnoise is running: do nothing
Sun Apr  6 00:00:04 EST 1997
Sysnoise is running: do nothing
Sun Apr  6 07:00:04 EDT 1997
Sysnoise is running: do nothing
Sun Apr  6 19:00:04 EDT 1997
Sysnoise is running: do nothing
Mon Apr  7 00:00:02 EDT 1997
5 head1200
6 head1300
7 head1400
8 head1500
Mon Apr  7 07:00:04 EDT 1997
Sysnoise is running: do nothing
Mon Apr  7 19:00:03 EDT 1997
Sysnoise is running: do nothing
Tue Apr  8 00:00:04 EDT 1997
Sysnoise is running: do nothing
Tue Apr  8 07:00:04 EDT 1997
Sysnoise is running: do nothing
Tue Apr  8 19:00:04 EDT 1997
Sysnoise is running: do nothing
Wed Apr  9 00:00:05 EDT 1997
Sysnoise is running: do nothing
Wed Apr  9 07:00:05 EDT 1997
Sysnoise is running: do nothing
Wed Apr  9 19:00:03 EDT 1997
9 head1600
10 head1700
11 head1800
12 head1900
Thu Apr 10 00:00:03 EDT 1997
Sysnoise is running: do nothing
Thu Apr 10 07:00:04 EDT 1997
Sysnoise is running: do nothing
Thu Apr 10 19:00:03 EDT 1997
13 head2000
14 head2100
```

15 head2200
16 head2300
Fri Apr 11 00:00:04 EDT 1997
Sysnoise is running: do nothing
Fri Apr 11 07:00:04 EDT 1997
Sysnoise is running: do nothing
Fri Apr 11 19:00:04 EDT 1997
Sysnoise is running: do nothing
Sat Apr 12 00:00:04 EDT 1997
Sysnoise is running: do nothing
Sat Apr 12 07:00:04 EDT 1997
Sysnoise is running: do nothing
Sat Apr 12 19:00:04 EDT 1997
Sysnoise is running: do nothing
Sun Apr 13 00:00:04 EDT 1997
Sysnoise is running: do nothing
Sun Apr 13 07:00:03 EDT 1997
17 head2400
18 head2500
19 head2600
20 head2700
Sun Apr 13 19:00:04 EDT 1997
Sysnoise is running: do nothing
Mon Apr 14 00:00:03 EDT 1997
Sysnoise is running: do nothing
Mon Apr 14 07:00:04 EDT 1997
Sysnoise is running: do nothing
Mon Apr 14 19:00:03 EDT 1997
21 head2800
22 head2900
23 head3000
24 head3100
Tue Apr 15 00:00:04 EDT 1997
Sysnoise is running: do nothing
Tue Apr 15 07:00:03 EDT 1997
25 head3200
26 head3300
27 head3400
28 head3500
Tue Apr 15 19:00:03 EDT 1997
29 head3600
30 head3700
31 head3800
32 head3900
Wed Apr 16 00:00:03 EDT 1997
33 head4000
34 head4100
35 head4200
36 head4300
Wed Apr 16 07:00:03 EDT 1997
37 head4400
38 head4500
39 head4600
40 head4700
Wed Apr 16 19:00:02 EDT 1997

41 head4800
42 head4900
43 head5000
44 head5100
Thu Apr 17 00:00:04 EDT 1997
Sysnoise is running: do nothing
Thu Apr 17 07:00:05 EDT 1997
Sysnoise is running: do nothing
Thu Apr 17 19:00:04 EDT 1997
Sysnoise is running: do nothing
Fri Apr 18 00:00:05 EDT 1997
Sysnoise is running: do nothing
Fri Apr 18 07:00:06 EDT 1997
Sysnoise is running: do nothing
Fri Apr 18 19:00:05 EDT 1997
Sysnoise is running: do nothing
Sat Apr 19 00:00:04 EDT 1997
Sysnoise is running: do nothing
Sat Apr 19 07:00:03 EDT 1997
45 head5200
46 head5300
47 head5400
48 head5500
Sat Apr 19 19:00:04 EDT 1997
Sysnoise is running: do nothing
Sun Apr 20 00:00:03 EDT 1997
Sysnoise is running: do nothing
Sun Apr 20 07:00:04 EDT 1997
Sysnoise is running: do nothing
Sun Apr 20 19:00:03 EDT 1997
End of batch file list
Resetting sysmon.runpos

Appendix D.

Data

D.1. WPAFB source locations

Source positions for the WPAFB spherical apparatus. Defined are the azimuth and elevation for the 272 source locations used, where 0° azimuth 90° elevation is directly front for the subject. Elevation is measured from 0° being directly under the subject. Azimuth is measured in the clockwise direction.

<u>num</u>	<u>az</u>	<u>el</u>
[1]	79.4	90.0
[2]	58.2	90.0
[3]	45.2	90.0
[4]	33.0	90.0
[5]	21.0	90.0
[6]	7.5	90.0
[7]	352.5	90.0
[8]	339.0	90.0
[9]	327.0	90.0
[10]	314.8	90.0
[11]	301.8	90.0
[12]	280.5	90.0
[13]	259.5	90.0
[14]	238.2	90.0
[15]	225.2	90.0
[16]	213.0	90.0
[17]	201.0	90.0
[18]	187.5	90.0
[19]	172.5	90.0
[20]	159.0	90.0
[21]	147.0	90.0
[22]	134.8	90.0
[23]	121.8	90.0
[24]	100.4	90.0
[25]	90.0	10.6
[26]	90.0	31.7
[27]	90.0	44.7
[28]	90.0	57.2
[29]	90.0	69.1
[30]	90.0	82.4
[31]	90.0	97.6
[32]	90.0	110.9
[33]	90.0	122.8
[34]	90.0	135.3
[35]	90.0	148.3
[36]	90.0	169.4
[37]	270.0	169.4
[38]	270.0	148.3
[39]	270.0	135.3

[40]	270.0	122.8
[41]	270.0	110.9
[42]	270.0	97.6
[43]	270.0	82.4
[44]	270.0	69.1
[45]	270.0	57.2
[46]	270.0	44.7
[47]	270.0	31.7
[48]	270.0	10.6
[49]	0.0	7.6
[50]	0.0	20.9
[51]	0.0	32.8
[52]	0.0	45.3
[53]	0.0	58.3
[54]	0.0	79.4
[55]	0.0	100.6
[56]	0.0	121.7
[57]	0.0	134.7
[58]	0.0	147.2
[59]	0.0	159.1
[60]	0.0	172.4
[61]	180.0	172.4
[62]	180.0	159.1
[63]	180.0	147.2
[64]	180.0	134.7
[65]	180.0	121.7
[66]	180.0	100.6
[67]	180.0	79.4
[68]	180.0	58.3
[69]	180.0	45.3
[70]	180.0	32.8
[71]	180.0	20.9
[72]	180.0	7.6
[73]	69.7	22.4
[74]	69.5	37.6
[75]	74.0	50.2
[76]	77.1	62.9
[77]	79.4	75.3
[78]	35.0	18.0
[79]	49.1	31.6
[80]	55.1	45.8
[81]	60.7	57.6
[82]	65.1	69.9
[83]	68.9	82.4
[84]	23.7	29.8
[85]	38.8	42.6
[86]	45.0	54.7
[87]	49.9	65.8
[88]	54.2	77.7
[89]	18.9	42.4
[90]	30.5	54.0
[91]	36.1	65.6
[92]	40.9	77.8
[93]	15.1	55.2
[94]	21.8	66.7

[95]	27.6	78.5
[96]	8.1	69.1
[97]	14.9	79.8
[98]	290.3	22.4
[99]	290.5	37.6
[100]	286.0	50.2
[101]	282.9	62.9
[102]	280.6	75.3
[103]	325.0	18.0
[104]	310.9	31.6
[105]	304.9	45.8
[106]	299.3	57.6
[107]	294.9	69.9
[108]	291.1	82.4
[109]	336.3	29.8
[110]	321.2	42.6
[111]	315.0	54.7
[112]	310.1	65.8
[113]	305.8	77.7
[114]	341.7	42.4
[115]	329.5	54.0
[116]	323.9	65.6
[117]	319.1	77.8
[118]	344.9	55.2
[119]	338.2	66.7
[120]	332.4	78.5
[121]	351.9	69.1
[122]	345.1	79.8
[123]	249.7	22.4
[124]	249.5	37.6
[125]	254.0	50.2
[126]	257.1	62.9
[127]	259.4	75.3
[128]	215.0	18.0
[129]	229.1	31.6
[130]	235.1	45.8
[131]	240.7	57.6
[132]	245.1	69.9
[133]	248.9	82.4
[134]	203.7	29.8
[135]	218.8	42.6
[136]	225.0	54.7
[137]	229.9	65.8
[138]	234.2	77.7
[139]	198.3	42.4
[140]	210.5	54.0
[141]	216.1	65.6
[142]	220.9	77.8
[143]	195.1	55.2
[144]	201.8	66.7
[145]	207.6	78.5
[146]	188.1	69.1
[147]	194.9	79.8
[148]	110.3	22.4
[149]	110.5	37.6

[150]	116.0	50.2
[151]	102.9	62.9
[152]	100.6	75.3
[153]	145.0	18.0
[154]	130.9	31.6
[155]	124.9	45.8
[156]	119.3	57.6
[157]	114.9	69.9
[158]	111.1	82.4
[159]	156.3	29.8
[160]	141.2	42.6
[161]	135.0	54.7
[162]	130.1	65.8
[163]	125.8	77.7
[164]	161.7	42.4
[165]	149.5	54.0
[166]	143.9	65.6
[167]	139.1	77.8
[168]	164.9	55.2
[169]	158.2	66.7
[170]	152.4	78.5
[171]	171.9	69.1
[172]	165.1	79.8
[173]	79.4	104.7
[174]	77.1	117.1
[175]	74.0	129.8
[176]	69.5	142.4
[177]	69.7	157.6
[178]	68.9	97.6
[179]	65.1	110.1
[180]	60.7	122.4
[181]	55.1	134.2
[182]	49.1	148.4
[183]	35.0	162.0
[184]	54.2	102.3
[185]	49.9	114.2
[186]	45.0	125.3
[187]	38.8	137.4
[188]	23.4	150.2
[189]	40.9	102.2
[190]	36.1	114.4
[191]	30.5	0.0
[192]	18.9	137.6
[193]	27.6	101.5
[194]	21.8	113.2
[195]	15.1	124.8
[196]	14.9	100.2
[197]	8.1	110.9
[198]	280.6	104.7
[199]	282.9	117.1
[200]	286.0	129.8
[201]	290.5	142.4
[202]	290.3	157.6
[203]	291.1	97.6
[204]	294.9	110.1

[205]	299.3	122.4
[206]	304.9	134.2
[207]	310.9	148.4
[208]	325.0	162.0
[209]	305.8	102.3
[210]	310.1	114.2
[211]	315.0	125.3
[212]	321.2	137.4
[213]	336.3	150.2
[214]	319.1	102.2
[215]	323.9	114.4
[216]	329.5	126.0
[217]	341.7	137.6
[218]	332.4	101.5
[219]	338.2	113.3
[220]	344.9	124.8
[221]	345.1	100.2
[222]	351.9	110.9
[223]	259.4	104.7
[224]	257.1	117.1
[225]	254.0	129.8
[226]	249.5	142.4
[227]	249.7	157.6
[228]	248.9	97.6
[229]	245.1	110.1
[230]	240.7	122.4
[231]	235.1	134.2
[232]	229.1	148.4
[233]	215.0	162.0
[234]	234.2	102.3
[235]	229.9	114.2
[236]	225.0	125.3
[237]	218.8	137.4
[238]	203.4	150.2
[239]	220.9	102.2
[240]	216.1	114.4
[241]	210.5	126.0
[242]	198.3	137.6
[243]	207.6	101.5
[244]	201.8	113.2
[245]	195.1	124.8
[246]	194.9	100.2
[247]	188.1	110.9
[248]	100.6	104.7
[249]	102.9	117.1
[250]	106.0	129.8
[251]	110.5	142.4
[252]	110.3	157.6
[253]	111.1	97.6
[254]	114.9	110.1
[255]	119.3	122.4
[256]	124.9	134.2
[257]	130.9	148.4
[258]	145.0	162.0
[259]	125.8	102.3

[260]	130.1	114.2
[261]	135.0	125.3
[262]	141.2	137.4
[263]	156.3	150.2
[264]	139.1	102.2
[265]	143.9	114.4
[266]	149.5	126.0
[267]	161.7	137.6
[268]	152.4	101.5
[269]	158.2	113.3
[270]	164.9	124.8
[271]	165.1	100.2
[272]	171.9	110.9

D.2. CRE source locations

Source positions for the CRE cylindrical apparatus. Defined are the radius, azimuth, and elevation for the 72 source locations used, where 0° azimuth 90° elevation is directly front for the subject. Elevation is measured from 0° being directly above the subject. Azimuth is measured in the clockwise direction. The radius is given in meters.

<u>num</u>	<u>r</u>	<u>az</u>	<u>el</u>
[1]	1.21	180	38
[2]	0.93	180	54
[3]	0.79	180	72
[4]	0.75	180	90
[5]	0.79	180	109
[6]	0.93	180	126
[7]	1.21	-150	38
[8]	0.93	-150	54
[9]	0.79	-150	72
[10]	0.75	-150	90
[11]	0.79	-150	109
[12]	0.93	-150	126
[13]	1.21	-120	38
[14]	0.93	-120	54
[15]	0.79	-120	72
[16]	0.75	-120	90
[17]	0.79	-120	109
[18]	0.93	-120	126
[19]	1.21	-90	38
[20]	0.93	-90	54
[21]	0.79	-90	72
[22]	0.75	-90	90
[23]	0.79	-90	109
[24]	0.93	-90	126
[25]	1.21	-60	38
[26]	0.93	-60	54
[27]	0.79	-60	72
[28]	0.75	-60	90
[29]	0.79	-60	109
[30]	0.93	-60	126
[31]	1.21	-30	38

[32]	0.93	-30	54
[33]	0.79	-30	72
[34]	0.75	-30	90
[35]	0.79	-30	109
[36]	0.93	-30	126
[37]	1.21	0	38
[38]	0.93	0	54
[39]	0.79	0	72
[40]	0.75	0	90
[41]	0.79	0	109
[42]	0.93	0	126
[43]	1.21	30	38
[44]	0.93	30	54
[45]	0.79	30	72
[46]	0.75	30	90
[47]	0.79	30	109
[48]	0.93	30	126
[49]	1.21	60	38
[50]	0.93	60	54
[51]	0.79	60	72
[52]	0.75	60	90
[53]	0.79	60	109
[54]	0.93	60	126
[55]	1.21	90	38
[56]	0.93	90	54
[57]	0.79	90	72
[58]	0.75	90	90
[59]	0.79	90	109
[60]	0.93	90	126
[61]	1.21	120	38
[62]	0.93	120	54
[63]	0.79	120	72
[64]	0.75	120	90
[65]	0.79	120	109
[66]	0.93	120	126
[67]	1.21	150	38
[68]	0.93	150	54
[69]	0.79	150	72
[70]	0.75	150	90
[71]	0.79	150	109
[72]	0.93	150	126