THE PHYSICAL MODELLING OF DRUMS USING DIGITAL WAVEGUIDES

By

Joel Augustus Laird

Department of Electrical and Electronic Engineering

November 2001

A dissertation submitted to the University of Bristol
in accordance with the requirements of the degree
of Doctor of Philosophy in the Faculty of Engineering

Contents

List of Tables

List of Figures

Abstract

In this thesis the physical modelling of percussive drums was approached using digital waveguides. The constituent components of a drum were considered individually before connecting them together to complete the model.

To model the drumskin techniques were created to incorporate smooth curved boundaries, calculate the impedance of a 2D waveguide mesh and include the effect of the bearing edge. The accuracy of the curved boundary model, which utilised 'rimguides', was demonstrated with a good reproduction of the first seven resonant modes of a circular membrane. The impedance was used in a kettledrum model where it correctly controlled the exchange of energy between the drumskin and interior air. Simulations of different bearing edge sizes revealed that a blunt edge takes energy from low frequencies and redistributes it into higher frequencies. These decay faster and so the result is a decrease in sustain.

For the interior air it was necessary to correctly model 3D wave propagation and incorporate diffuse reflections, which occur at rough surfaces. Unlike 3D meshes used in previous studies, the new dodecahedral mesh proposed here was found to exhibit near direction independent dispersion error. The effect of diffusion was adequately simulated with a technique that was shown to be controllable, enabling different types of surface to be modelled.

To complete the drum model a way of connecting different waveguide meshes together was found and a new procedure for modelling a mallet exciter was proposed. The interfacing method enabled a lossless interconnection between two 2D meshes and also 2D and 3D meshes. The procedure used for the mallet exciter incorporated non-linear stiffness and the mallet's contact area. Its behaviour was shown to be almost identical to that of a real mallet.

Finally, a digital waveguide model of a kettledrum was constructed to demonstrate the techniques and the results were promising; the resonant modes were reproduced with good accuracy and their decay was sufficient to give the impression of realism, whilst not exactly matching that found through measurement.

Author's Declaration

I declare that the work in this thesis was carried out in accordance with the Regulations of the University of Bristol. The work is original except where indicated by special reference in the text and no part of the dissertation has been submitted for any other degree. Any views expressed in the dissertation are those of the author and in no way represent those of the University of Bristol. The thesis has not been presented to any other University for examination either in the United Kingdom or overseas.

Joel Augustus Laird

Date: 5 Nov. 2001

Acknowledgements

As the World teeters on the edge of continuous economic growth and War one wonders what the point is in developing the state-of-the-art when it could all be lost in an instant. But then we remember that our aim in life is to be happy and so one can't let the petty desires of a few individuals spoil such a noble pursuit. I, for one, find a little happiness in trying to understand the mechanics of the Universe, and at least I can be confident that my area of research, the modelling of musical instruments, has no aggressive application. Firstly then, I will offer no kind words of acknowledgement to Sky News for their incessant coverage of World Affairs that has kept me up late for the past few weeks.

Doing this research has been great fun but also a bit stressful (ever tried thinking about a single subject for nearly five years every waking hour of the day?). I can only apologise to those of you who have been waiting quite some time for me to finish it. By way of an excuse I'll say that I spent almost one of those years writing the software for an exhibit in the London Science Museum - *go and see it* <-cheeky plug ;-) This was a wonderful experience and hopefully something I can come back to later in life and say, 'I did that!'

The true beginning of these acknowledgements commences from the following colon: first of all I'd like to thank my family because, let's face it, without them I wouldn't even be here. I don't think my brother Reuben (and his wife Kiz!), my uncle Peter, my aunt, (other) uncle and cousins, Jill, Haydn, Kate and Tom had much to do with that but they should know that I think of them very fondly. Mum and Dad (aka. Liz and Paul), thank you very much for all

your support! I believe that I am very fortunate indeed. Many kind regards to my grandparents too, Bob, Maurine, Bill and Joyce. I'm very sorry that two of you are unable to share the finishing of this thesis with me, but I think of you most every day. Naturally I must have many ancestors of whom I know nothing about; I'd particularly like to congratulate those who survived the two great Plagues.

There are some friends that need mentioning here. This motley crew have played a large part in my life in recent years and I only hope that we'll continue to have as much fun, although perhaps the level of decadence will drop as we get older and lose our marbles. Cheers to (in no particular order): Greg (smeg), Matt (beer Matt), Rob (loud), Nic (puts up with Rob), Dan (related to Tom Jones somewhere along the line), Martyn (Goth), Nazy (ambassador from Iran), and the 'other' lot: Pete, Dee, Will, Sam, Kath, Ben... there are others but who's going to read this anyway? Extra special thanks goes to dear Amy for her support and patience. I love you very much! I am grateful for the hot meals from your mum, Anne, and the whisky I shared with your dad, Peter.

Separate from the 'friends' paragraph above is this one dedicated to my band mates - www.banacek.org. Thank you Marc (who's currently finishing his own thesis in digital waveguide modelling and proof-read this one for me) and Katherine, Andrew (hope you come back from your wandering in the Far East soon), Iain and Maria. I'm sure we play the oddest bunch of songs and could really make something of it, if we weren't quite so lazy.

There are many contacts and colleagues that definitely deserve acknowledgement. From Bristol Uni. I'd like to thank Paul Masri, Andy Bateman and Nishan Canagarajah. You showed me how to approach research and write papers, and looked after me whilst I was a postgraduate. Thank you also to Ian Craddock and Damion Murphy for being my examiners at the viva of this thesis. Your suggestions and corrections, whilst initially having the size of something quite mammoth, are now extinct and have certainly added to and

improved my work. Kind regards to Tony, Chris, Shirley, Stavri and Albaina (sp.?), Dimitris, Angelici (my God I have no idea how to spell your names), Ron and Will. A special mention goes to Premier Percussion who sponsored me for a few years of my research. Steve, maybe one day I could write that drum modelling software for you and have a go in your mate's (Eddie Jordan) racing car? I forgot, can you introduce me to Nicko McBrain (from Iron Maiden) too? Finally, I'd like to thank all those at Quicksilva for giving me a job to do: Andy, Paul, Huw, Kathryn, Richard, Terry, Maddy, Mel, Colin, Aidan, Chris, Ben, Gareth, Darren, Vikki, Steve, Gayna and Duncan. I'm enjoying myself very much!

Okay, this is the final set of acknowledgements, honest. To all the animals that have shared part of my life, I love you but it's a pity you don't last for very long. In the cat/dog hall of fame there is Eric (I and II) and Sam. Supporting the budgie contingent: Snoop Budgie-Budge (calcium deficiency, RIP), Billy (left in hot car, RIP), Stinky Beak (amazingly clever but didn't quite make it through a closing door, RIP - www.stinkybeak.com), Floyd (fat, lazy feather duster, vanished without a trace one night, RIP?) - you strange creatures, can't any of you make it past the age of two? One budgie is still going strong after a few months: Timmy! (are you a boy or a girl?). Also, thanks to my mum's killer cats Freddie and Fleur, and A.'s dog Toby (smelly-woofter).

Chapter 1

Introduction

1.1 Physical Modelling

In order to understand the world around us and make predictions we construct models. From the metaphysical description of the universe perceived by our ancient ancestors, to the present day scientific theories of quantum mechanics and relativity the premise is the same: a set of rules are presented, enabling the construction of a model that describes the system under study. Of course, the system needn't be quite so grand as those aforementioned; the affect of climate change on insect populations or failure rate of engine parts, to name but two examples, can both be modelled.

Physical modelling generally refers to a model of a mechanical system, to be used for design and testing before the real appliance is constructed. It can also be used to gain a deeper understanding of physical phenomena. Such a model may be an actual scaled version, or a set of governing formulae that are utilised on paper or a computer. The technological revolution has opened the door for the implementation of complex and accurate models, usually requiring hundreds and thousands of sophisticated calculations.

1.1.1 The Physical Modelling of Musical Instruments

An interesting application of physical modelling can be found in musical instruments. Such models have two uses: an aid to musical instrument design and testing, and real-time synthesis for musical performance. A real-time musical instrument synthesis model has to be fairly simple and therefore suffer in accuracy. It still has advantages over other synthesis techniques in that it can be controlled by meaningful variables, such as the magnitude and position of an applied pressure. Above all, physically modelled instruments are expressive and realistic. This means that they capture the behaviour and nuances of their real counterpart. In order to improve their fidelity a deeper understanding of the underlying physics is required. Bearing this in mind it is feasible that in time, as computer processing power increases and hardware becomes cheaper, real-time musical instrument models will become indistinguishable from reality.

This has far reaching implications for musical recording, as only the control data for the models needs to be stored, such as note and volume. Therefore, high fidelity music could be reproduced from small streams of data which could easily be transmitted over the Internet. Alternatively many hundreds of hours of music could be stored on one CD. Interestingly, such a data format already exists in the form of MIDI (Musical Instrument Digital Interface).

MIDI Control

MIDI was developed in 1983 (Heckroth, 1995) to synchronise and connect keyboards, tone generators, rhythm machines and computers together. The MIDI data for the entire 'Planets Suite' by Gustav Holst is available on the Internet and can be downloaded as a file 246 Kb in size (Trussler), compared to a CD of 650 Mb. Of course, the fidelity of the music depends in part on the MIDI synthesiser used, and is at present far from capturing the sound of a real orchestra. Another aspect that determines the realism of performance is the amount of

expression allowed by the control data. For instance, it can be imagined that a physical model of a violin may require data describing the manner in which it is held, the position, pressure and movement of fingers on the fingerboard and the control of the bow. In addition the instrument design and material properties could be sent as initial control data during model creation.

Unfortunately MIDI is potentially limited in its use as a control format for complex physical models; by definition it can only support 128 different controllers and many of these are already set aside for specific purposes (Glatt). These controllers can be either 'fine' or 'coarse' in terms of the expression they permit. The coarse controllers use numbers in the range: 0 to 127, whereas the fine controllers cover the range: 0 to 16384. Naturally the larger range offered by the fine controllers is the most desirable for physical models. Despite this restriction on the amount of control data supported by MIDI, the Yamaha WX5 MIDI wind controller (IDnet, 2002) is a very expressive device, especially when used in combination with the Yamaha VL-70m Virtual Acoustic tone generator. This tone generator implements physical modelling techniques and can produce highly realistic and expressive wind-instrument sounds when used properly.

Physical Modelling within this Thesis

The objective of this thesis is to model a modern drum, which is a member of the percussion family. In order to do this it is necessary to separate the drum into all its constituent components and find a way to model them individually. These components include the drumskin, shell and interior air, but also the more subtle parts such as the roughness of the shell interior. Even the hardware fittings that are attached to the drum are to be considered. Once modelled all these parts need to be fitted back together to create a single model, capturing the essence of a drum in its entirety.

1.2 The Role of the Master Craftsman

The design of drums, like that of many instruments, has been described as a 'black art'. This is because the choice of size, shape and type of material is made by the master-craftsman based on his (usually inherited) knowledge of what works and what doesn't. His aim is to produce an instrument that sounds 'good' especially under particular playing and environmental conditions.
An instrument modelling technique would be very useful for fast prototyping of ideas and may help in the identification of what a good sound actually is. But rather than replacing the role of the master-craftsman such a modelling technique should work with him. Besides, it is the craftsman who bridges the gap between concepts and their realisations, for he knows in which ways certain materials can be manipulated. For example, there would be no point in designing an instrument with Computer Aided Design (CAD) software if it would actually be impossible to build.

1.3 Thesis Overview

The focus of this thesis has been on the development of an accurate musical instrument modelling technique, rather than develop a real-time physical model that would quickly become obsolete, as suggested by projected increases in processing power. After all, it is easier to reduce the complexity of accurate models to yield real-time implementations, than work the other way round.
Whilst there are already accurate techniques in existence that can be used for musical instrument modelling, many of these are over complicated for the task. The work done for this thesis has been the extension of a recently introduced technique known as 'Digital Waveguide Modelling' (Smith(III), 1987)(Smith(III), 1992). Only with the advent of more powerful computers has this become feasible for modelling anything other than simple systems.
In order to model a percussive drum many new effects and components had to

be incorporated, particularly those that deal with specific acoustic problems. The existing digital waveguide technique is therefore extended, resulting in a modelling method that is useful within complex audio applications.

Audio and visual demonstrations of the results from this thesis are included on the accompanying CD-Rom (see appendix A). These demonstrations are referenced at the relevant points within the text. Also on the CD-Rom there can be found a digital copy of this thesis and the C++ code that was developed. In addition there are some programs to demonstrate wave and pressure flow within a drum, and a full application that creates audio samples from a Tom-Tom drum model with editable parameters.

1.3.1 Background

The background consists of the principles of acoustics, modelling techniques and an overview of digital waveguide modelling.

The first of these introduces the subject of acoustics. It presents this at a level that is necessary for understanding the physical phenomena to be modelled. It explains how sound is generated, transmitted and heard. Some modelling techniques that are already in use are then discussed. The principles behind their implementation are described and their advantages and disadvantages are compared.

The overview of digital waveguide modelling elaborates further on the existing waveguide technique. It presents the development done on digital waveguides prior to the commencement of this work, bringing together research that helped to advance the state of the technique.

1.3.2 The Physical Model

This part contains the majority of the work done for the thesis. It is comprised of sections on drum modelling, a section on instrument excitation and finishes with the analysis, conclusions and future work.

The drum modelling sections identify the areas of the digital waveguide technique that need attention in order to construct a drum model. The extensions to the exisiting methods that have been developed are discussed in detail.
Although some work on instrument exciters had already been done, the excitation section is best suited here rather than in the background part. This is because the exciters are models in their own right and are essential for the analysis of the drum model. In addition, a new implementation of a mallet exciter is described.
For the final analysis the developed digital waveguide technique is used to construct a kettledrum model. This is analysed and compared with an equivalent model that used a finite element technique (Rhaouti, Chaigne, and Joly, 1999). The results from this are discussed and conclusions are formed.
During the course of this work many ideas and possible future directions revealed themselves; the future work section briefly discusses these. The methods that can be extended are mentioned here and ideas are suggested as to how further development could progress.

1.3.3 Appendices and Bibliography

At the start of the appendix is a description of the accompanying CD-Rom. The majority of the mathematics and formulations relevant to this work are also to be found in the appendices section. They were put here so that the bulk of the thesis wasn't broken up by pages of equations, which would have made it difficult to read. In this way the text only shows the mathematical results where necessary.

Part I

Background

Chapter 2

Principles of Acoustics, the Ear and Modelling Techniques

2.1 Sound Generation, Propagation, the Ear and Perception

2.1.1 Sound Generation and Propagation to the Ear

When objects collide kinetic energy is passed between them. This transfer of energy will alter their motions, such as the case of two snooker balls colliding. In addition, suppose the middle snooker ball of a rack were to vibrate then the whole rack would vibrate too, as the energy is transmitted from ball to ball. This transmission is not confined to one medium; the rack vibrates the table, the table the floor and so on. Similarly, the air surrounding the vibrating objects also receives this energy, which then flows from air molecule to air molecule as they collide with each other.

The ear, shown in figure 2.1, enables us to analyse this energy by collecting it with the 'pinna' and guiding it towards the 'eardrum', which is caused to vibrate. From the eardrum the vibrations are propagated along a chain of three small bones (the 'hammer', 'anvil' and 'stirrup' or 'malleus', 'incus' and

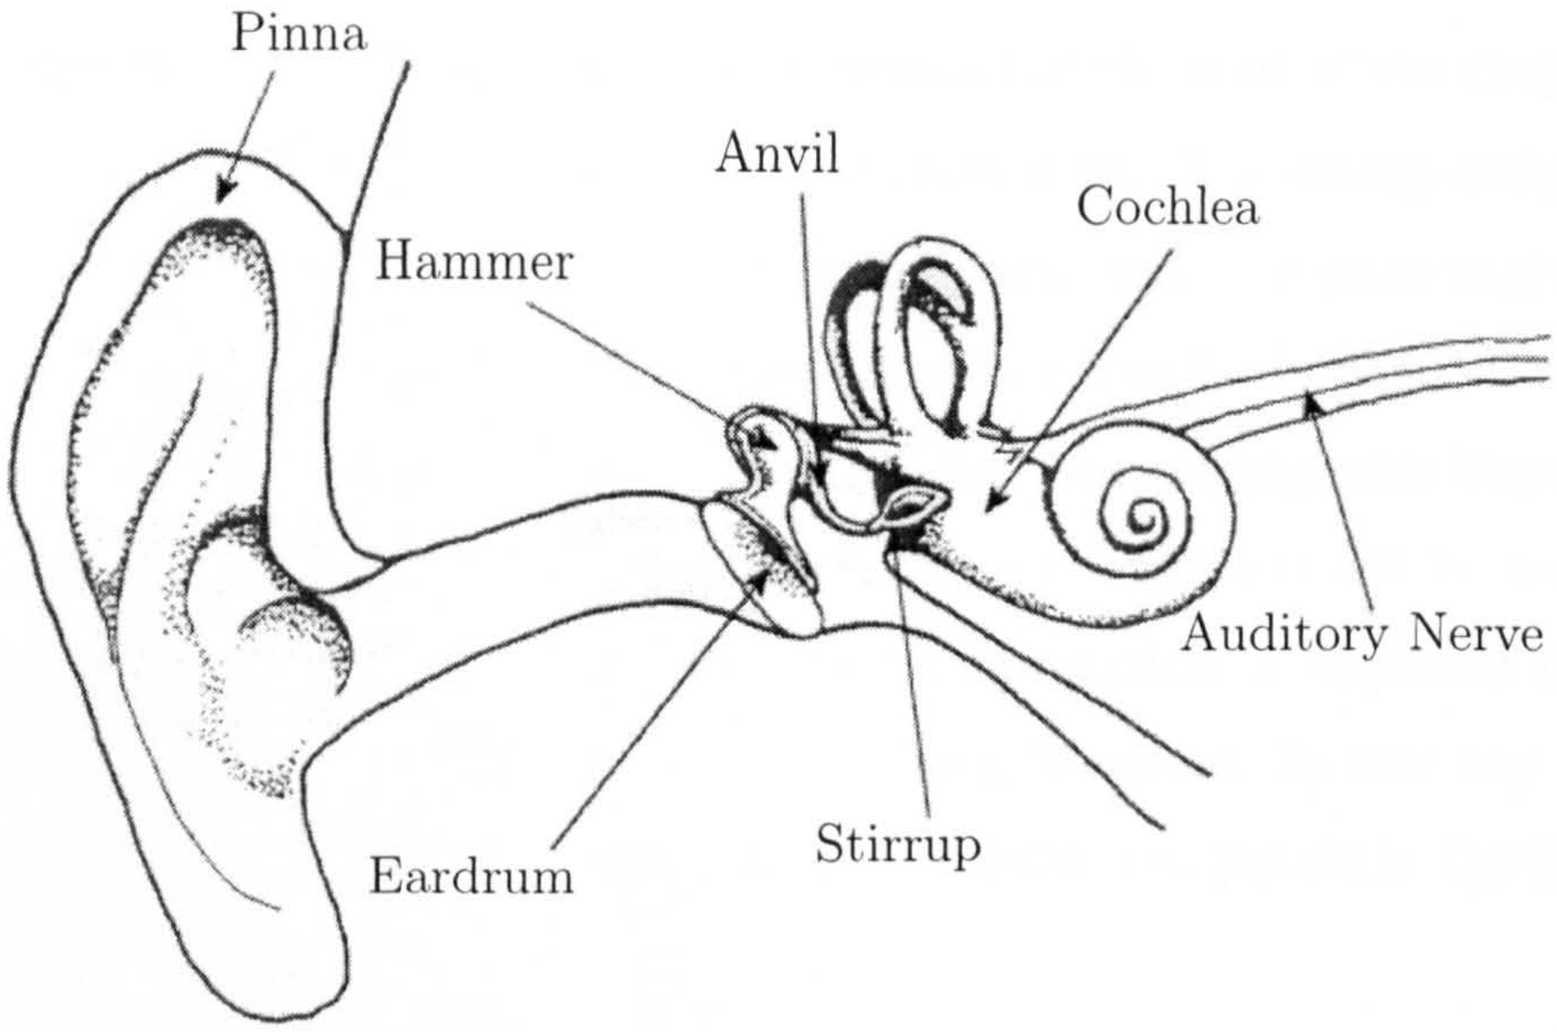

Figure 2.1: *The Human Ear.*

Reproduced from Moore (1990)

'stapes') and then through a small membrane, known as the 'oval window', into the 'cochlea'. The purpose of the three small bones, or 'ossicles', will be addressed later when sound pressure levels are introduced.

It is within the cochlea that a very important function is performed resulting in the conversion from the kinetic energy of the vibrations into electrical energy. This will be discussed later in section 2.1.2, where the limits of auditory perception are considered. From the cochlea the electrical energy is transmitted along the 'auditory nerve' and into the brain where it is processed, resulting in the sensation of 'hearing' sound.

As air envelopes everything on the planet's surface it is responsible for most of our experience of sound, but the rack and table mentioned in the example above were also carrying the sound energy. In fact sound is transmitted through all media to some extent. However, the density and other physical characteristics of the medium affect the speed of sound propagation and its decay of magnitude with time and distance.

In the snooker table vibration example, mentioned at the start of the chapter, it is important to note that there is no net movement. The energy-carrying particles of each medium oscillate about a fixed point, and they do so a certain number of times a second. The rate of oscillation is called the frequency and is measured in *Hertz* (*Hz*), named in honour of the German physicist Heinrich Hertz. As the energy propagates the next particle to receive it will be a step behind the last, or out of 'phase', such that each particle is displaced by a different amount along a line traced to the energy's source. In this way the energy travels not as a single point, but as a collection of points in the form of a 'wave'.

Transverse and Longitudinal Waves

There are different ways in which a sound wave will propagate and this depends on the orientation of the plane it travels along to the plane of vibration. A vibrating string, for example, exhibits transverse vibrations that result from transverse travelling waves. Here the direction of wave travel is perpendicular to the vibration. If this energy were to travel through 3D space such as air it would be in the form of longitudinal travelling waves, where the wave transmission is in the same direction as the vibration (such as compression waves travelling along a 'slinky' spring). Longitudinal waves are characterised by changes in pressure.

Sound Pressure Levels

The range of volume for sound is so great that a special logarithmic scale was conceived to measured it by. It is defined such that every additional power of 10 in the amount of energy carried by a wave is represented by one *Bel*, named after Alexander Graham Bell. One *Bel* is a large step and so it is useful to use a unit that is one tenth of the size, in order to deal with the smaller changes in sound pressure level or intensity. This smaller unit is known as the decibel (*dB*).

The sound pressure level of a noise makes use of this scale, and is defined as (Hall, 1987a)

$$SPL\,(dB) = 20\log_{10}\left(\frac{p}{p_{ref}}\right), \tag{2.1.1}$$

where p represents the pressure amplitude of the wave (Pa),

p_{ref} is the pressure reference standard for air ($2 \times 10^{-5}\,Pa$).

The ossicles within the middle ear act as a lever system to transfer mechanical vibrations from the eardrum, caused by oscillations in sound pressure level, to the inner ear. Without these the oscillations would be too small for the inner ear to detect properly and so any sound would be barely audible.

There are two small muscles within the middle ear know as the 'tensor tympani' and 'stapedius', which are attached to the malleus and stapes respectively. When stimulated by high intensity sounds (greater than approximately $75\,dB\,SPL$ (Howard and Angus, 1996)) these muscles contract; the action of the tensor tympani results in an increase in the tension of the eardrum, and the stapedius pulling on the stapes stiffens the chain of bones. This increases the impedance of the middle ear, reducing the efficiency with which vibrations are transmitted from the eardrum to the inner ear and thus provides protection from loud sounds. The mechanism, known as the 'acoustic reflex', results in approximately 12 to 14 dB of attenuation, but only for frequencies below $1\,KHz$. It takes between 60 and 120 ms (Howard and Angus, 1996) for the muscles to react and so there is a potential for damage to be caused by sudden loud noises.

2.1.2 Perception of Sound

The Choclea and Basilar Membrane

The inside of the choclea has a very complex structure, but for the purposes of this thesis it is sufficient to say that it forms a wound up duct, like a snail's shell, and is partitioned into two tubes by the 'basilar membrane'. The

basilar membrane is the actual sensor organ. One tube of the cochlea has the oval window at its end which is in contact with the stapes, as previously mentioned, and the other tube is sealed off with another elastic membrane, the 'round window'. Both tubes are connected at the far end, or 'apex', by a small hole in the basilar membrane called the 'helicotrema' and filled with an incompressible fluid, the 'perilymph' (Roederer, 1973a)(Jeans, 1937).

The vibrations received at the oval window are converted into oscillations of the perilymph fluid that propagate through the cochlear duct and set the basilar membrane into motion. There are about 30000 sensitive hairs, arranged in 'inner' and 'outer' rows along the basilar membrane, that detect its motion and impart signals to nerve cells (neurons) which they are in contact with (Spoendlin, 1970). Remarkably, the basilar membrane oscillates in regions that correspond to the frequencies of the received sound: the lower the frequency the closer to the apex lies the region of activated hair cells, the higher the frequency the closer it is located to the oval window (Roederer, 1973a). This correlation between frequency and distance along the basilar membrane was found through physiological measurements performed on anesthetized or even dead animals (von Békésy, 1960), and has been demonstrated by damaging selected bits of the basilar membrane and finding that the animal then becomes deaf to notes of certain pitches (Jeans, 1937).

Psychoacoustics

The study of sound perception and its limitations is known as psychoacoustics and can be put to good use when modelling audible sound. For example, frequencies below $20\,Hz$ and above $20\,kHz$ may be ignored or modelled less accurately as they are outside the normal range of the human ear (Wood, 1964)(Moore, 1990). The reasons for this are due to the natural construction of the basilar membrane, where the relationship between a change in frequency and the corresponding displacement of the activated region along the membrane is logarithmic. As a result, roughly two-thirds of its length is sensitive

to a range of frequencies between $20\,Hz - 4000\,Hz$ (Roederer, 1973a) with all other frequencies being squeezed into the remaining one-third. Those above $20\,kHz$ are situated so close to the oval window they are undetectable. It should also be noted that an active region actually spreads a small distance along the membrane, and this places limits on the ear's ability to resolve pitches of very similar frequency.

For two pure tones presented one after the other there is a natural limit below which the difference in frequency is undetectable. This limit, or 'just noticeable difference' (JND), varies from person to person, is a function of musical training, and unfortunately, depends considerably on the method of measurement employed (Roederer, 1973a). However, the JND always lies between two threshold values: below one of which a change in pitch is never detected, and above the other a change is always detected.

When two pure tones of almost identical frequency are played simultaneously the active regions on the basilar membrane are merged into each other. This causes the sensation of 'beating', where the loudness seems to rapidly oscillate. If the difference in frequency is increased slightly then the beating disappears, but it is replaced by a 'rough' and unpleasant tone sensation. Separating the frequencies even further will result in both resonance regions on the basilar membrane being far enough apart to give two distinct pitch signals. However, at this limit the sensation of roughness still persists. When the frequency difference has surpassed a large enough value, known as the 'critical band', the roughness disappears and both pure tones sound 'smooth'. Interestingly, for the human ear the critical band corresponds to an extension on the basilar membrane 'serviced' by roughly 1300 receptor cells (Roederer, 1973a). These groups of cells cover an almost constant distance along the membrane, regardless of whether they are at the high frequency or low frequency end. It therefore follows that combinations of high frequencies are more difficult to resolve than lower ones, because of the logarithmic relationship between frequency and displacement of the activated region.

The ear is also sensitive to a large range of sound pressure levels, giving rise to the sensation of loudness. This sensitivity varies however, such that pure tones of equal sound pressure level, but of different frequency, don't necessarily seem to have equal loudness. Noises of very low sound pressure level may be below the threshold of hearing and therefore undetectable, whereas very loud noises can be painful (Wood, 1964)(Roederer, 1973b). In addition to this sensitivity range there is an interesting effect know as 'masking', where one tone can be undetectable due to the presence of another much louder tone. The minimum sound level that the masked tone must exceed in order to be heard varies with its frequency and depends on the frequency and sound level of the masking tone. For two similar frequencies this masking level can be as close as $15\,dB$ below the level of the masking tone, yet drop to over $60\,dB$ below for larger frequency differences (Roederer, 1973b).

2.2 Modelling Techniques

There are many modelling techniques to choose from that may be used for sound synthesis and analysis. Each has its advantages and disadvantages, depending on the desired application. For example some are very accurate, using variables relevant to the physical system, but take a long time to compute and are complex in their construction. Others are more abstract or simple to construct but are not so accurate. When considering a technique it is obviously necessary to decide on certain criteria, two of which may be: 'how much accuracy is required?' and 'how much computing power is available?'

As mentioned on page 4 in chapter 1.3 an accurate modelling technique is desired for this research. Because of the emphasis on instrument design the technique should utilise physically relevant variables, but it should also be possible to present it in a form that is easily manipulated by a craftsman, who may have no knowledge of physical modelling techniques. It is assumed that computing power is of no consequence, although the technique ought to be

contained within reasonable limits.

A good introduction to sound synthesis methods and comparisons of physical modelling techniques has already been written (Tolonen, Välimäki, and Karjalainen, 1998). The following is a simple description of the techniques that were considered for investigation in this research. In addition a few of their advantages and disadvantages are discussed.

2.2.1 Pure Mathematical and Real Models

Before the advent of powerful computers all modelling was done with either pure mathematics or through the construction of real models to simulate behaviour. An example of a real model could be a network of masses and springs to demonstrate the propagation of waves. Real models such as these can only of course be used for the simplest of simulations.

Modelling a system with pure maths results in a formula that yields an accurate and instantaneous result. That is to say that once the system's initial conditions have been set, its exact state can be found at any point in time without an iterative process. Unfortunately, the more complicated the system the more difficult it is to describe mathematically. Also, a new model needs to be formulated for each new system that is considered. It is often the case that solutions to modelling problems are found intuitively, where the mathematician has a good idea of what he is looking for in the first place. For example, to model a swinging pendulum a sensible trial solution would be in the form of a sinusoidal function.

Mathematical formulae are useful for predicting a system's resonant frequencies, but finding the energy distribution within these frequencies is much more complicated. Any transient behaviour exhibited by the system before it settles down into its resonant frequencies is also very difficult to model in this way.

2.2.2 Finite Element

A finite element model involves breaking the system up into many small elements, each being a description of a simple mechanical system such as a mass and spring. The elements are considered small enough such that their behaviour is well described by the underlying mathematics. The complete model is constructed by sticking the elements together and solving the conditions at their boundaries. Detailed discussions of the theory of finite elements can be found in the works of Petyt (Petyt, 1990) and Macey (Macey, 1987).

Finite element models are very accurate and are often used to solve mechanical and aeronautical systems such as vibrations and stresses of machines and airflow over the wing of an airplane. Such accuracy is probably excessive for modelling musical instruments, where the vibrations are low in amplitude. In addition finite elements are mathematically intensive; the boundary conditions are difficult to solve and creating an interface between models is non-trivial.

2.2.3 Ray-tracing/Image-source

Ray-tracing and image-source methods are borrowed from the discipline of computer graphics but may be used for room acoustics problems too. They both avoid modelling the entire space by only considering the response at the listening point due to at least one excitation point.

In the case of ray-tracing (Kinsler and Frey, 1962)(Kulowski, 1985) different paths are found that lead from the excitation point to the listening point. The lengths of these paths are used to calculate the different times taken by the sound to reach the ear. The response is therefore the superposition of different delayed versions of the original sound.

Similarly, the image-source method (Eyring, 1930)(Allen and Berkley, 1979) also involves a superposition of delayed sounds. However, the lengths of the paths are found in a different way. Each source is reflected about the boundaries of the acoustic space (as shown in figure 2.2) to form a collection of virtual

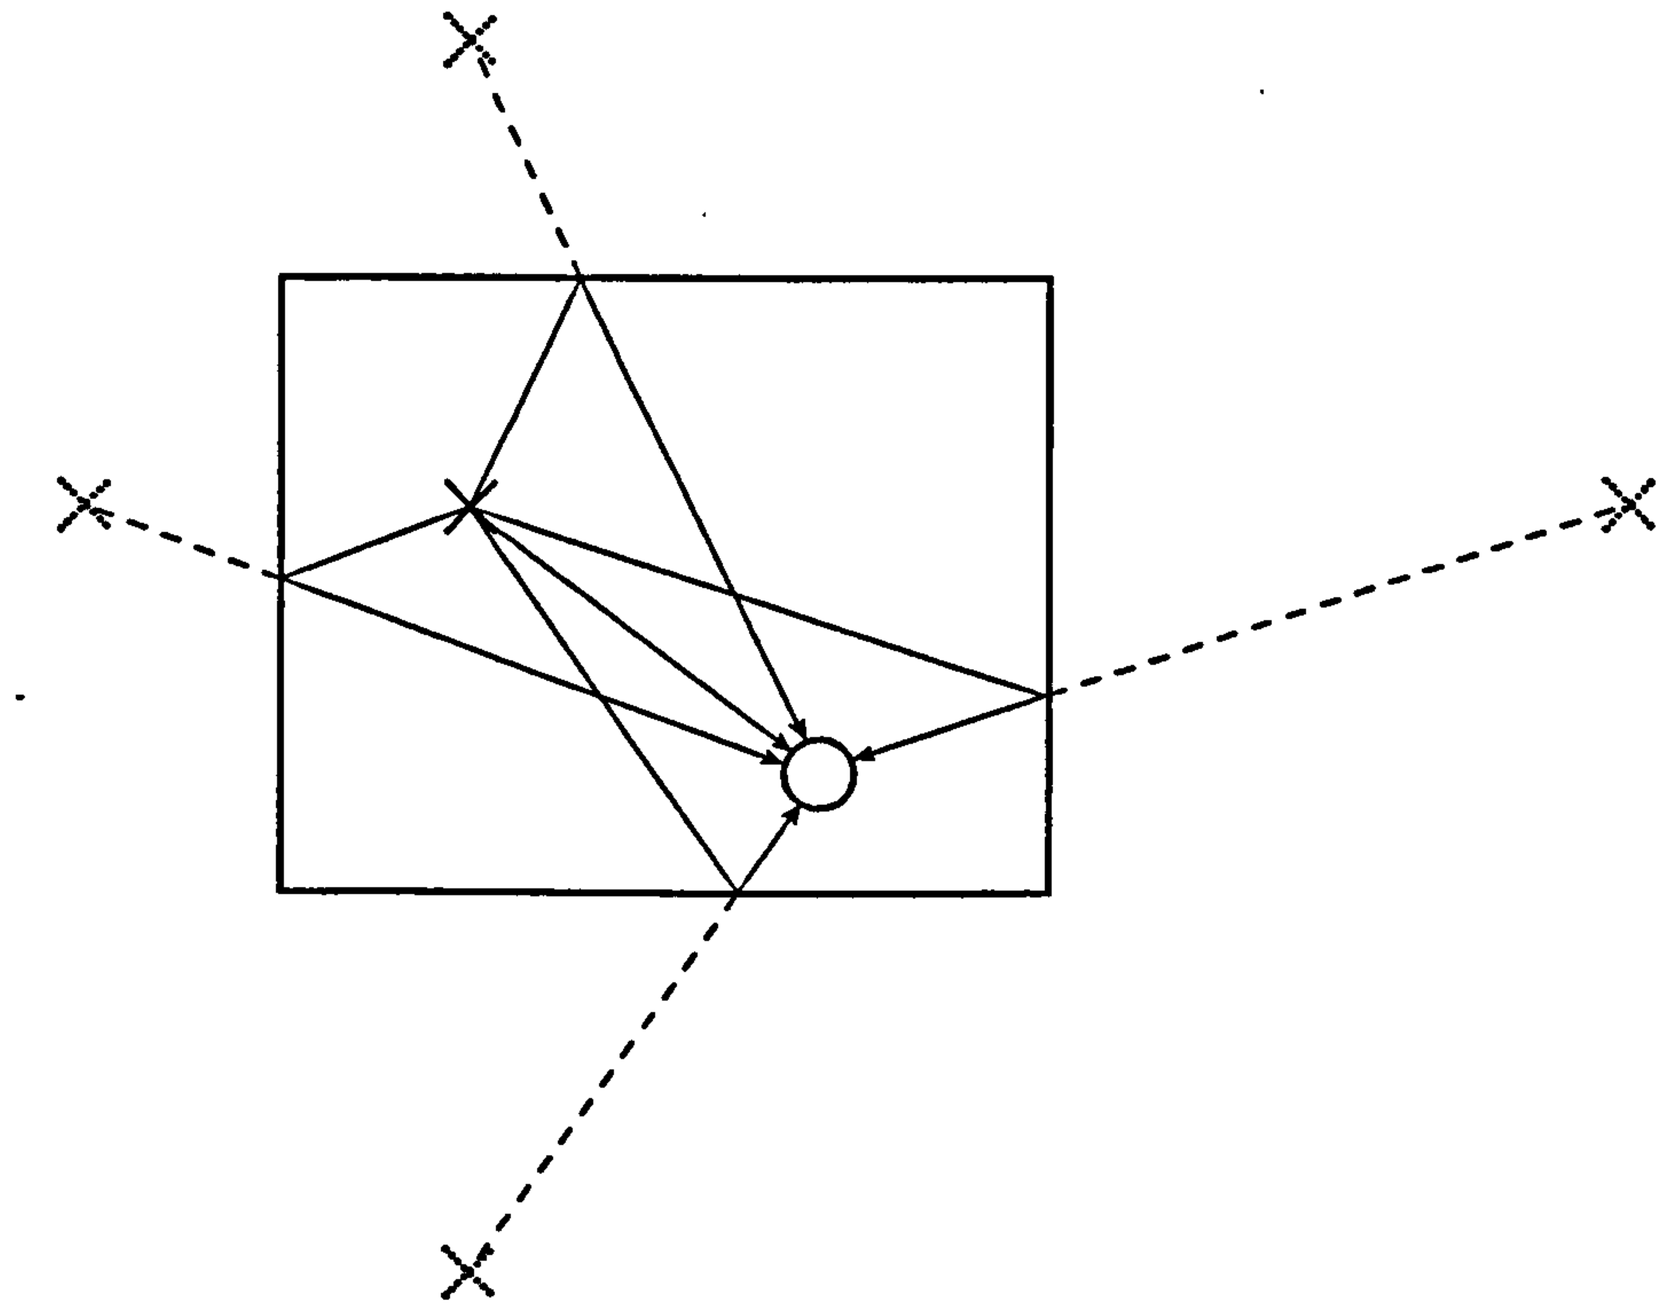

Figure 2.2: *The Image-Source Technique. The crosses represent real and 'image' sources, the circle is a listening point.*

sources or 'images'. A straight line drawn from an image to the listening point represents a possible path to the source as seen from this listening point. Each image is then reflected again and again until the length of the path is so great that the sound can be considered to have decayed to an inaudible level.
These methods both suffer from the problem of the response time having to be explicitly set, and increasing this corresponds to an increase in complexity. For example, in the case of ray-tracing the response time is increased by considering more rays in order to include a greater number of reflections before the sound reaches the listening point. As there are potentially an infinite number of rays the creation of the model can take a long time. For the image-source method the response time is extended by reflecting the images more times about the boundaries and thus it can be imagined that the number of images will increase very rapidly.

Another disadvantage is that important effects such as diffraction (the scattering of waves as they pass an obstacle), diffusion (the scattering of waves reflected at a boundary) and absorption at the boundaries have to be explicitly modelled. These effects that involve scattering pose a particular problem in that the number of rays or images increases dramatically, which in turn will have a noticeable effect on the time taken to create and compute a model. Although potentially useful for modelling the resonance of a drum's air cavity these methods are not applicable to the modelling of vibrating objects such as a drumskin.

2.2.4 Modal Synthesis

Modal synthesis is an abstract approach to modelling that concentrates on the resonant modes of a vibrating system, rather than the materials and their shape. The frequencies and the damping coefficients of the modes of vibration are represented explicitly in the model (Adrien, 1991). In this way the sound of a musical instrument may be accurately reproduced as well as more unusual sounds.

The modal frequencies are generated by a set of damped harmonic oscillators whose behaviour conform to the laws of mechanics. This physical basis gives the models a natural sound and allows a meaningful interaction between musician and instrument. More complicated responses may be incorporated by linking the oscillators together with elastic connections. The basis of model generation and the relationship between real-life vibrating objects and their modal counterparts has previously been described (Djoharian, 1993).

A very useful feature of modal synthesis is that it represents a system spectrally. This means that it is possible to directly control the acoustic characteristics such as frequency, amplitude and damping. Its most endearing quality is the ability to generate models that sound natural but are materially unknown or impossible to construct in reality.

Despite these acoustically attractive features this technique does not explicitly model real vibrating systems without recourse to measurement or numerical simulation. This requirement is necessary if a model is to be used for 'what-if?' analysis, where a designer may want to know the affect on a drum's sound, for example, after changing the shape or material.

2.2.5 Modalys

Modalys is the successor of MOSAIC (Morrison and Adrien, 1993) (the name being changed to avoid conflicts with the Internet navigator 'Mosaic'). Like Modal synthesis, it is another abstract approach to modelling virtual instruments. It too concentrates on the generation of resonant modes, their shape and damping factors. However, it calculates and controls these in a different manner.

A Modalys synthesis model is based on four types of elements (objects, accesses, connections and controllers) which may be assembled in a variety of ways (Eckel, Iovino, and Caussé, 1995). Modalys objects describe vibrating structures that are defined by their geometrical structure and the physical parameters of the material they are made from. The structures may be strings, plates or membranes, for example. From these a modal representation is calculated to yield the frequency, damping factor and mode shape of each resonant mode, which can be controlled independently.

The Modalys objects are assembled using Modalys connections, which describe the interaction types between the objects such as a strike, pluck or bow. The locations of the interaction points are specified using Modalys accesses and these, as well as other time-varying synthesis parameters, are assigned to Modalys controllers.

The use of objects and their connections within Modalys helps to produce sounds that feel natural, linking the sound phenomenon to a production mechanism that could feasibly exist in reality. However, despite this and the expres-

sive and interesting acoustical properties of Modalys it suffers from the same disadvantages as Modal synthesis; it doesn't explicitly model a vibrating system such as a drum. For this the resonances would have to be independently calculated or measured and then imported as modal data.

2.2.6 Tao

Tao (Pearson and Howard, 1995) is another system that is very similar conceptually to MOSAIC. Within Tao objects of various shapes and materials may be interconnected and excited at different points in time and space.
The objects, their connection and excitation are described using a scripting language that consists of two parts: an orchestra language containing the object declarations and a score language describing when, where and how to excite them. The synthesis engine is entirely different to Modalys or MOSAIC, however, as it takes its inspiration from celluar automata modelling techniques and has been developed to exhibit all the wave phenomena observed in the physical world (Pearson and Howard, 1995). This enables pieces of material of specific shape and size to be defined and have damping applied to various regions.
The synthesis engine involves a collection of point masses arranged in a regular grid pattern and connected together with springs. Each mass, or cell, maintains information about its own position, velocity, mass etc. and whether it is free to move or not. In addition all cells are constrained to have one degree of freedom along a single axis which is perpendicular to the grid on which they are laid out (Pearson). In this way it is a numerical implementation of the real models of masses and springs referred to at the start of this section, albeit with the scope for greater complexity. Figure 2.3 depicts an arrangement of masses and springs within a small section of material modelled using Tao.
This technique has been used to produce a variety of high quality timbres and as it models the wave phenomena rather than the modal frequencies it is po-

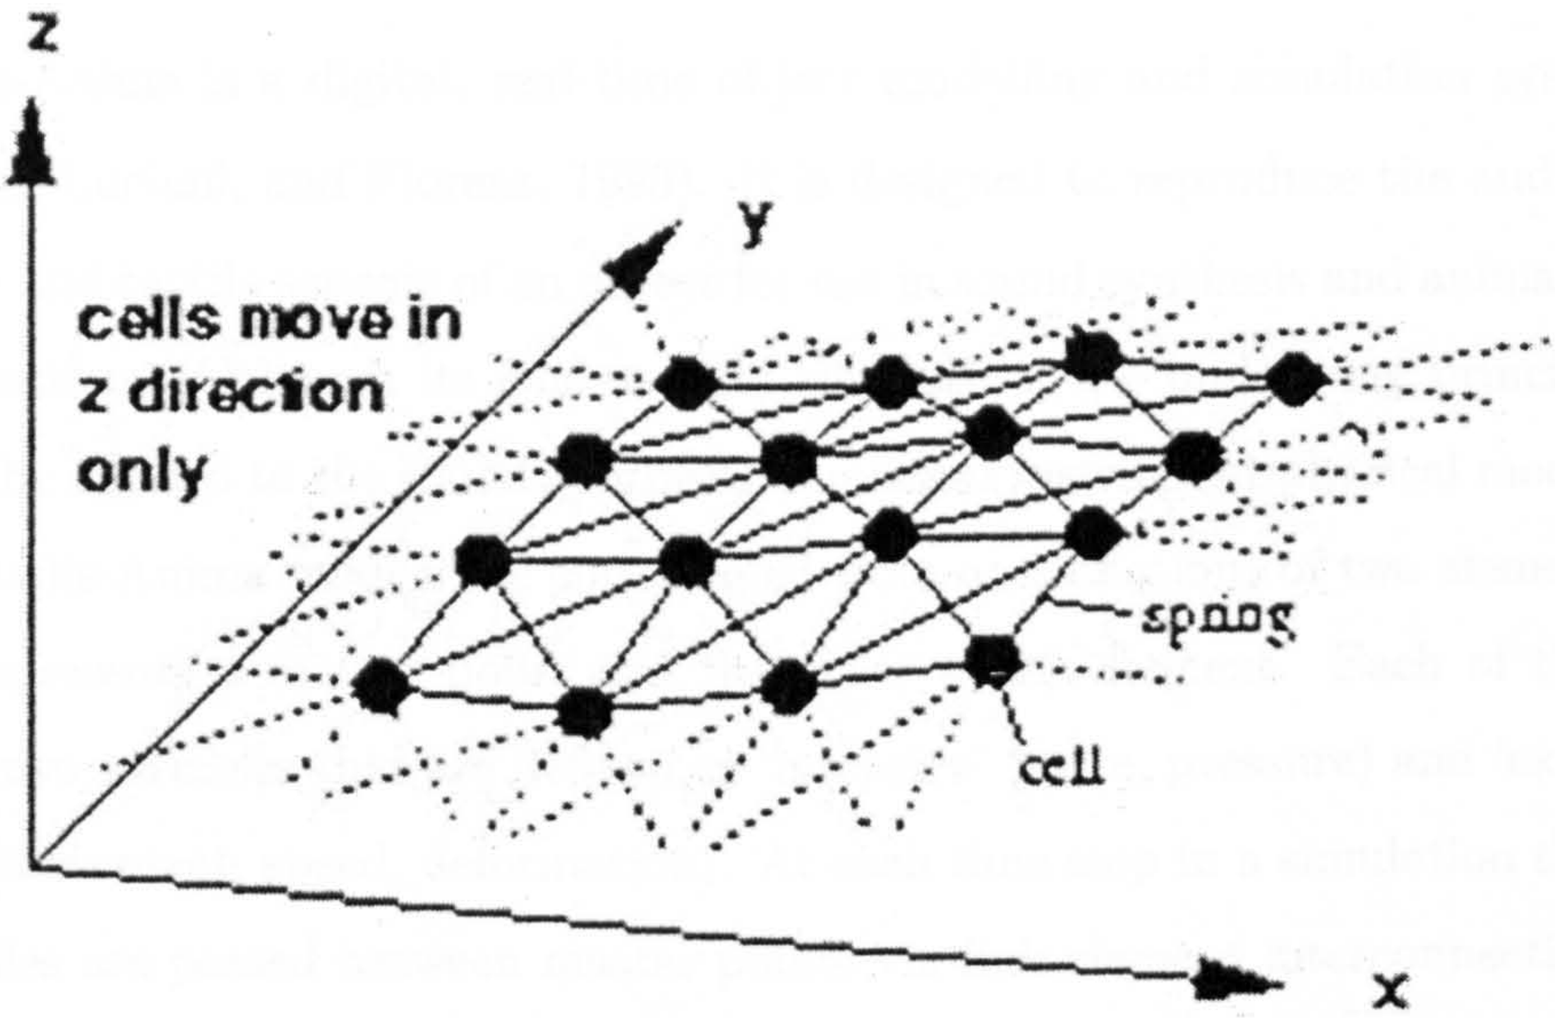

Figure 2.3: *A Small Section of Tao's Acoustic Material showing the Masses and Springs and the Direction in which the Cells are free to move.*
Reproduced from Pearson

tentially useful for physically modelling drums. Of particular relevance to this thesis is the application of the technique to instrument design. For instance, '...with simple string-like instruments consisting of a single row of cells linked together with springs it has been found that damping small regions at either end of the string more highly than the rest of the cells produces a natural string-like decay in the spectrum [of sound] whereas a string with uniform damping exhibits amplitude decay only...' (Pearson and Howard, 1995). This demonstrates the acoustically important absorption effects caused by the connections at the ends of a string to an instrument body.

Unfortunately the models require a large number of calculations for each time step and this leads to long simulation times. Also, the technique is only capable of modelling transverse waves on surfaces or strings rather than longitudinal pressure waves travelling through 3D space. This precludes its use for modelling acoustic spaces such as the interior of a drum.

2.2.7 Cordis-Anima

Cordis-Anima is a digital, real-time object modelling and simulation system (Cadoz, Luciani, and Florens, 1993). It is designed to reproduce the audible, visible and tactile aspects of an object for use in sound synthesis and animation applications. Although its goal is total simulation, the underlying principles could be applied to the sound synthesis of musical instrument physical models. All Cordis-Anima models are constructed from combinations of two elements: one represents a matter point and the other a link element. Each of these holds two variables that are defined as 'intensive' (force, pressure) and 'extensive' (movement, speed, deformation). At each time step in a simulation these variables are passed between matter points via link element interconnections. Computed inside each matter point is an algorithm that defines its behaviour as either a mass, spring or friction element. A material and its shape may therefore be suggested by grouping masses, springs and friction elements together and then arranging these as an interconnected network.

This modelling technique is evidently accurate and also benefits from being intuitive in its manipulation, as the elements are arranged in the shape of the object to be modelled. It is, however, fairly well developed and computationally expensive. Furthermore, even though a model may behave in accordance with actual experience, it would appear to be difficult to mathematically prove its equivalence to a real system.

2.2.8 Digital Waveguide

A new addition to the acoustic modelling family is the digital waveguide. The digital waveguide is an accurate digital form of d'Alembert's travelling wave solution to the wave equation (see section 3.2). 2D and 3D models are made by creating a mesh of waveguides, with each waveguide representing a path of finite length along which a wave may travel. In this way it can be seen that waveguide models are intuitive to construct; the mesh is formed such that it

fits within the desired boundary. In addition, the manner in which waveguides are calculated makes them ideal for implementation on a DSP chip, for they mainly require the copying of data from one location to another. As they are carrying digital signals, various acoustic effects can be included in the models by incorporating techniques from the well established field of digital signal processing.

Unfortunately, the digital waveguide has primarily been researched as a method for real-time synthesis rather than analysis, and so there has been a lack of development in the more complicated aspects of acoustics.

For this investigation it was decided that digital waveguides would be considered for modelling musical instruments. The reason for this was the potential for accuracy and the need for further development of the modelling technique to be useful as an analysis tool.

Chapter 3

Digital Waveguide Modelling - Overview

3.1 Underlying Mathematics

At the core of the digital waveguide modelling technique is the one dimensional wave equation, the derivation of which can be found in appendix C.1. Its general form is

$$\frac{\partial^2 y}{\partial t^2} = c^2 \frac{\partial^2 y}{\partial x^2}, \tag{3.1.1}$$

with

$$c = \sqrt{\frac{T}{\sigma}}. \tag{3.1.2}$$

This equation shows that the acceleration of a point is related to the local change in gradient and the wavespeed, c.

3.1.1 D'Alembert's Solution to the Wave Equation

In 1747 d'Alembert published the travelling wave solution to the 1D wave equation as follows (Lindsay, 1973)

$$y(x,t) = y_r(x - ct) + y_l(x + ct), \tag{3.1.3}$$

where y is the displacement at point x and time t,

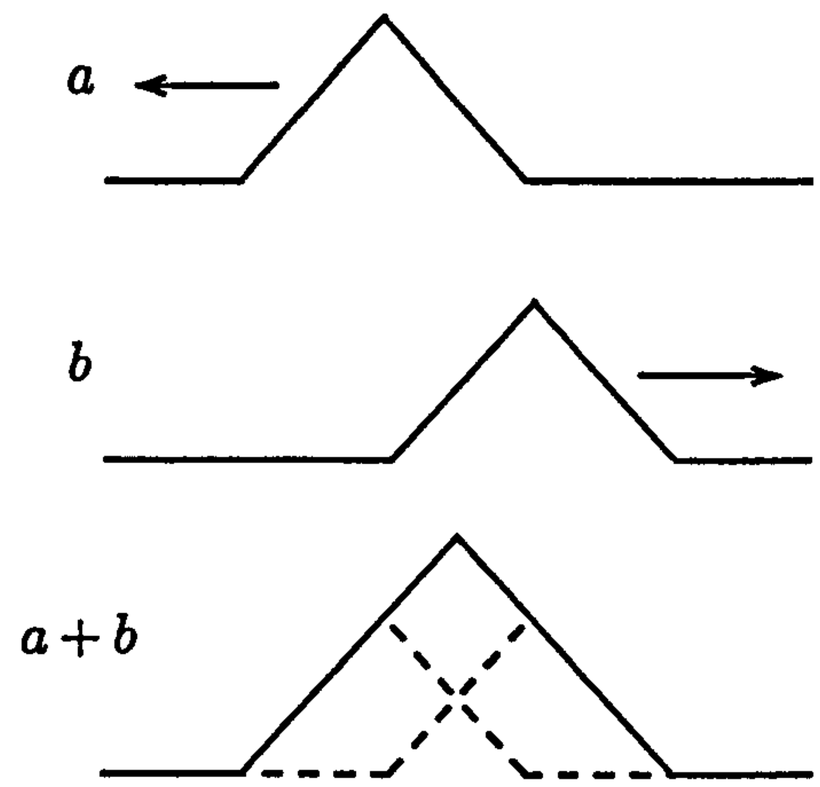

Figure 3.1: *Superposition of Two Waves (Travelling in Opposite Directions) to make a Single Wave.*

y_r is a wave travelling to the the right,

y_l is a wave travelling to the left,

c is the wavespeed.

Essentially this describes an instantaneous wave variable at a single point, by the superposition of two independent travelling wave shapes propagating at speed c. The two waves progress in opposite directions, see figure 3.1, and carry information appropriate to the system under study. This information could be pressure, force or velocity waves.

It is easy to prove that the travelling wave solution satisfies 3.1.1 because the two functions $y_r(x-ct)$ and $y_l(x+ct)$ are twice differentiable

$$\frac{\partial^2 y_r}{\partial t^2} = c^2 \frac{\partial^2 y_r}{\partial x^2}, \tag{3.1.4}$$

$$\frac{\partial^2 y_l}{\partial t^2} = c^2 \frac{\partial^2 y_l}{\partial x^2}, \tag{3.1.5}$$

and therefore

$$\frac{\partial^2 y}{\partial t^2} = \frac{\partial^2 (y_r + y_l)}{\partial t^2}, \tag{3.1.6}$$

$$= c^2 \frac{\partial^2 (y_r + y_l)}{\partial x^2}, \tag{3.1.7}$$

$$\Rightarrow \frac{\partial^2 y}{\partial t^2} = c^2 \frac{\partial^2 y}{\partial x^2}. \tag{3.1.8}$$

This also shows that the wave equation is obeyed regardless of the travelling wave shapes y_r and y_l.

3.2 Digital form of D'Alembert's Solution

The Digital Waveguide was derived from d'Alembert's solution by sampling the travelling wave description (Smith(III), 1987)(Smith(III), 1992). This results in the following formula, graphically described as a delay-line pair in figure 3.2

$$y(t_n, x_m) = y^+(n-m) + y^-(n+m), \tag{3.2.1}$$

where $+$ as a superscript denotes a travelling wave component
propagating to the right,
$-$ as a superscript denotes a travelling wave component
propagating to the left,
n is a sample period number,
m is a position along the delay-line pair.

The distance across each unit of delay is that which a wave will travel in one sample period. It can be expressed as $\Delta x = \frac{c}{f_s}$, where f_s is the sample rate of the model.

At each discrete point the information yielded is accurate, provided the travelling wave shapes are initially bandlimited to less than half the sampling frequency. Information between these points can be obtained through interpolation, and frequency dependent effects (see section 3.5) which occur in real

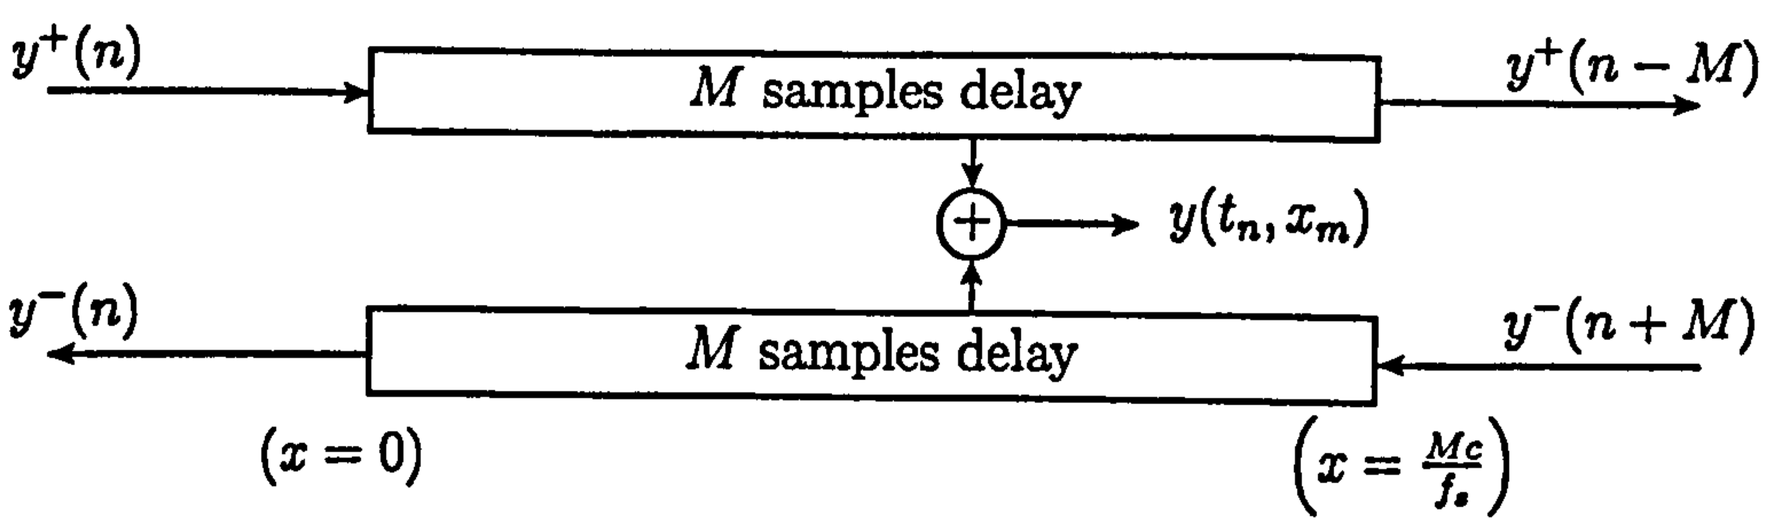

Figure 3.2: *The Digital Waveguide.*

life media may be reproduced by inserting digital filters into the waveguide structure.

3.3 Non-Integer Length Digital Waveguides

It may have occurred to the reader that the waveguide described above is constrained to being a whole number of delays in length. As each delay element corresponds to a spatial distance then this limitation restricts models to specific sizes. One way of alleviating the problem is to increase the model's sample rate, which decreases each delay's spatial distance, effectively enhancing the model's resolution. This enables a greater range of model sizes and allows a broader spectrum of frequencies to travel along the waveguide, but with a heavier computational burden.

An alternative to avoid increases in sample rate is to add fractional delay filters (Laakso, Välimäki, Karjalainen, and Laine, 1996) to the ends of a waveguide. These filters are carefully constructed to give a specific phase response, in order to model a fractional delay. They may be non-recursive in their design, known as Finite Impulse Response (FIR) filters, or recursive Infinite Impulse Response (IIR) filters. The greater their complexity, or order, the more accurate the delay over the entire frequency range. As other filters may be employed to model frequency dependent effects, it is useful to eliminate any frequency dependent gain of a fractional delay filter. This can be done by restricting the filters to an IIR allpass implementation, see appendix F.1 about allpass fractional delay filter design for more information on this. As these filters have a unity gain response over the entire frequency range they ensure that this part of the model is energy conserving. For modelling purposes a first-order allpass filter is usually sufficient to implement a fractional delay (as will become apparent later on, in section 4.1.3).

3.4 Digital Waveguides of Time-Varying Length

In some cases, such as modelling a pitch bend on a guitar string, it is necessary to vary a waveguide's length over time. Performing this variation in length was found to be very important when modelling the effect of a drum's bearing edge, see section 4.2. As previously stated allpass fractional delay filters may be used to make waveguides of non-integer length, and therefore this length can be varied by altering the filters' coefficients. Unfortunately, altering the coefficients of IIR filters such as these will introduce transients into the system, which produces an audible click.

A technique has been proposed (Duyne, Jaffe, Scandalis, and Stilson, 1997) to implement a real-time pitchbend-able model, where the transients caused by coefficient changes are eliminated. The transient elimination method had already been described (Välimäki, Laakso, and Mackenzie, 1995). Here it was found that if a first-order allpass filter's coefficient, a, is kept within the range $-\frac{1}{5} \leq a < \frac{1}{3}$ (which corresponds to a fractional delay, D, in the range: $0.5 < D \leq 1.5$), then the transient caused by a coefficient change is audibly insignificant after six samples. Therefore, by keeping the last six input samples in a buffer, it is possible to eliminate any transients by replaying this buffer through the filter after each coefficient change. This six sample period is known as the filter's 'warm-up' period. In this case the eliminator routine runs at six times the sample rate of the overall system. Alternatively, an eliminator filter (which is a copy of the allpass filter) is run in parallel and its state variables are copied to the allpass filter every sixth sample. The eliminator filter's coefficient is changed six samples before that of the allpass filter used in the system. This restricts the coefficient update to be performed, at most, once every six samples.

3.5 Incorporating the Effect of Internal Viscous Friction and Stiffness

The physical characteristics of the medium through which the sound propagates alters the wave equations, generally adding higher order terms. Suffice to say that d'Alembert's travelling wave functions are no longer solutions. Explicitly analysing these effects is beyond the scope of this thesis; for more information on how these may be incorporated into complex digital waveguide models consult the future work section 8.1 of this thesis. Nevertheless, from analysis of the stiff and lossy wave equations it is possible to draw some information about the resulting behaviour. This may be used to give the impression of stiffness and internal friction.

Stiffness causes the speed of wave travel to become frequency dependent, a phenomenon known as dispersion. Generally, high frequencies propagate faster than low frequencies. This effect may be included by incorporating allpass filters of a non-uniform delay versus frequency.

Another frequency dependent effect is internal friction, which is caused by viscosity within the medium. When this occurs kinetic and potential energy are effectively lost as they are converted into heat. The higher frequencies lose energy at a greater rate than the lower frequencies (Morse and Ingard, 1968). This effect can be modelled with lowpass filters, see section F.2.

3.6 Real-time Synthesis

Real-time synthesis models have been created using the techniques described above. Of particular note is the Karplus-Strong Algorithm (Karplus and Strong, 1983) which preceeded the digital waveguide (it is, in fact, an equivalent but simplified version). Real-time implementations presently use the 1D waveguide to model wave propagation along a string or a tube, with the more complicated effects of body resonance appended to the end of the waveguide as

a bank of filters. Such implementations are used exclusively in Yamaha's VL series of synthesisers. More recently waveguide models have been included in computer soundcards, such as Creative Labs' AWE 64 Gold (Creative, 2001). In order to increase efficiency many components of the waveguide model can be commuted together. As an example, all the filters that are used for modelling losses and dispersion may be lumped into a single filter that is appended to one end of the waveguide. In special cases this can reduce the waveguide to standard lattice/ladder digital filters which have been extensively developed in the past (Gray and Markel, 1975).

3.7 The Digital Waveguide Mesh

The digital waveguide only models wave propagation in one dimension. However, it is possible to extend the technique to two dimensions or more by connecting unit length waveguides together to form a mesh. The 2D mesh has been shown to satisfy the 2D wave equation (Duyne and Smith(III), 1993), which makes it useful for modelling wave propagation along a membrane.

3.7.1 Scattering Junctions

If each waveguide is considered to be a string then their intersection at a point is described by a 'series' junction, where there is no energy loss. Such junctions are known as 'scattering junctions' as they take energy in and redistribute it amongst the waveguides. There are two conditions at each junction that must be satisfied:

Condition 1

The velocities of all the strings (1D waveguides) at the junction must be equal since they are all moving together at that point. This junction has a velocity,

v_J. If each string's velocity is denoted as v_i then this condition may be written

$$v_i = v_J. \tag{3.7.1}$$

Condition 2

There must be no net force at the junction, as all the forces exerted by the strings must balance each other

$$\sum_i^N f_i = 0. \tag{3.7.2}$$

where N is the total number of waveguides attached to the junction.

As the digital form of d'Alembert's solution splits the wave variables into two opposing travelling waves, the velocities and forces from above are equivalent to: $v_i = v_i^+ + v_i^-$ and $f_i = f_i^+ + f_i^-$. Each force is related to the velocity by its impedance (as described in the next section on waveguide impedances) with $f = Rv$. In the case of the travelling waves this relationship becomes

$$f_i^+ = R_i v_i^+, \tag{3.7.3}$$

$$f_i^- = -R_i v_i^-. \tag{3.7.4}$$

Notice that there is an inversion for waves travelling to the left (those with a $-$ superscript). This is because the direction of force is now opposite to the direction of wave travel.

Combining these equations together results in the formulation of the lossless scattering junction (see section D.1 of appendix D)

$$v_J = \frac{2\sum_i^N R_i v_i^+}{\sum_i^N R_i}, \tag{3.7.5}$$

$$v_i^- = v_J - v_i^+. \tag{3.7.6}$$

If the junction is to be clamped then v_J must equal zero. Therefore, the reflected velocity wave is related to the incoming velocity wave by

$$v_i^- = -v_i^+. \tag{3.7.7}$$

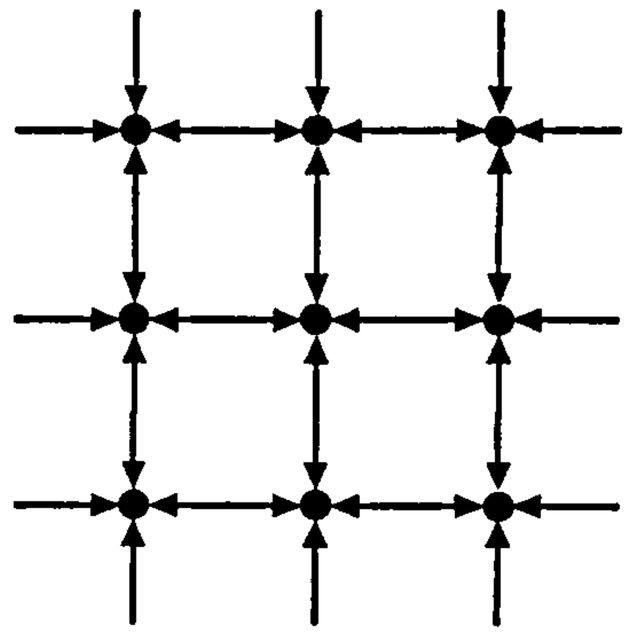

Figure 3.3: *The Structure of a Rectilinear Mesh. The lines represent waveguides and the interconnections represent lossless scattering junctions.*

It should be noted that the impedance of a scattering junction within the mesh is (see section D.2 of appendix D)

$$R_{mesh} = \frac{\sum_i^N R_i}{2}. \tag{3.7.8}$$

3.7.2 Dispersion Error

A common form of the 2D digital waveguide mesh has a rectilinear construction (see figure 3.3). Ideally all waves should travel at the same speed within the mesh structure. However, it can be shown that the mesh is equivalent to a finite difference scheme (see section E.2 and Fontana and Rocchesso (1998)) and similarly it too exhibits dispersion error. This means that the speed of wave propagation depends on its frequency. In the case of the rectilinear mesh the dispersion error is noticeably direction dependent, as shown in figure 3.4. A method for calculating the error will be covered later in section 4.1.2.

The plot in figure 3.4 is orientated in the same plane as the waveguide mesh. Every point represents the wavespeed of a wave travelling in a particular direction with a certain frequency. The frequency and direction is measured as the distance and angle of a line drawn from the centre to a point. Its wavespeed is suggested by the colour at that point, with white representing the maximum wavespeed and black representing the minimum wavespeed.

It can be seen that waves travelling in the diagonal directions of the rectilinear

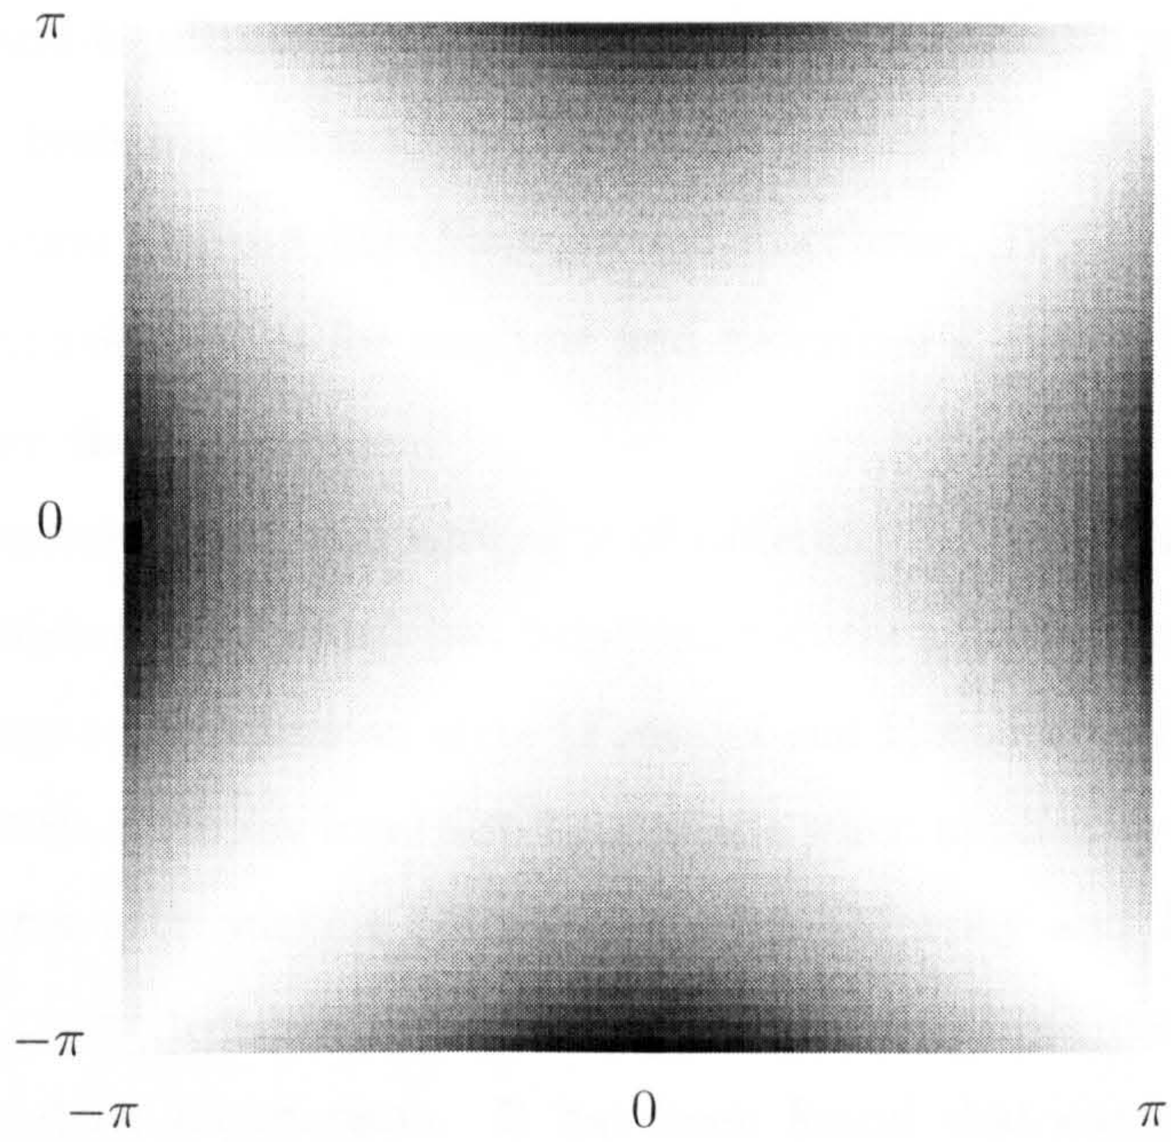

Figure 3.4: *The Dispersion Error of a Rectilinear Waveguide Mesh. The axes are normalised-π frequencies.*

mesh don't suffer any dispersion error. The worst dispersion is experienced by waves travelling in the horizontal and vertical directions. The effect of this is that the actual boundary of the mesh doesn't correspond in shape to a similar boundary applied to a real membrane. The shape of this boundary also depends on the frequency of a travelling wave. This results in the resonant modes of the model being slightly different to those predicted by mathematical analysis.

A method has been proposed to correct this angular dependency by using interpolation techniques (Savioja and Välimäki, 1996). This introduces an extra mathematical burden and so a simpler way to correct the angular dependency is found in section 4.1.2, where a circular membrane is to be modelled.

3.7.3 Applications of Waveguide Meshes

Waveguide meshes have naturally found applications in the area of membrane and room acoustics modelling. The efficiency and accuracy of different mesh

geometries and topologies have been of particular interest to researchers; relevant to this thesis was the investigation of 2D meshes for modelling membrane-based percussion instruments (Fontana and Rocchesso, 1998). This introduces some mathematical tools for analysis and describes a method for simulating an air load on the membrane.

The computational cost and accuracy of different 2D geometries (waveguide and finite difference schemes) has been extensively analysed in terms of sampling efficiency and dispersion error (Fontana and Rocchesso, 1999), and comparisons of their accuracy have also been made when applied to acoustics modelling of rooms with various boundary shapes (Murphy and Howard, 2000). Within this paper frequency warping techniques and high resolution meshes were also used for comparison. It has been found that overall a triangular mesh structure performs the best in terms of its accuracy and bandwidth versus computational efficiency and memory load.

Performance testing and model validation has also been applied to parallel implementations of the tetrahedral waveguide mesh with an aim to real-time synthesis of acoustic spaces (Campos and Howard, 2000). Preliminary listening tests were promising and binaural room responses of a listener's head have confirmed that the model is able to provide appropriate sound localisation cues.

3.8 Waveguide Impedances

The impedance of a medium relates an applied force, or pressure, to the resulting velocity. It is a measure of how easily a medium can be moved or compressed. So far the waveguide models have been homogeneous; that is to say that they model only one type of media. For these cases it has been possible to ignore the impedances of the waveguides because they are all equal. For example equation 3.7.5 above reduces to

$$v_J = \frac{2\sum_i^N v_i^+}{N}, \tag{3.8.1}$$

if all waveguide impedances are the same. As complex models will require different meshes to interact with each other, their impedances will be required. The equations used to calculate the impedances are shown below.

1D Case

The impedance of a 1D digital waveguide is equivalent to the driving point mechanical impedance of a string (Hall, 1987b)

$$R = \frac{T}{c} = \mu c = \sqrt{T\mu}, \tag{3.8.2}$$

where T is the tension applied to the string (N),

c is the wavespeed (ms^{-1}),

μ is the linear density (kgm^{-1}).

The relationship between force and velocity at a point is

$$f = Rv. \tag{3.8.3}$$

2D Case

There is very little about the mechanical impedance of a 2D membrane in the literature. The reason for this is because the applied force has to act on an area instead of an infinitesimally small point, which would cause an infinite acceleration. As the junctions of a waveguide mesh actually represent small tessellated elements, then the force at these junctions is actually acting over a finite area. A method for calculating the value of a 2D waveguide mesh impedance is proposed in section 4.1.5.

3D Case

In the 3D case the impedance relates pressure to velocity, instead of applied force, and so a different symbol, Z, is used

$$p = Zv, \tag{3.8.4}$$

where v is the velocity of a planar front of the medium with applied pressure p.

The impedance of a 3D medium is calculated as (Hall, 1987c)

$$Z = \rho c, \tag{3.8.5}$$

where ρ is the volume density (kgm^{-3}).

3.9 Summary

In this chapter an overview of the existing digital waveguide techniques has been given. A single digital waveguide is a solution to the 1D wave equation. Each unit waveguide has a finite length, restricting the range of model sizes, yet it is possible to append digital filters in order to create waveguides of fractional lengths. With careful control of these filters the length of a waveguide may be altered over time, enabling dynamic effects. The presence of internal friction and stiffness may also be suggested by appending filters.

The simple models used to date are ideal for real-time synthesis of strings and pipes. However, more complicated models may be created by interconnecting waveguides to form a mesh. A 2D waveguide mesh is equivalent to a finite difference scheme of the 2D wave equation and is therefore useful in modelling a membrane, but unfortunately its structure introduces dispersion error. 2D and 3D meshes have been studied and used to model membranes and acoustic spaces.

When interconnecting waveguides and waveguide meshes it is necessary to consider their impedances. The impedances of 1D and 3D media may be found through well-known calculations, but before commencing this thesis there was no method for calculating the impedance of a 2D waveguide mesh.

Using digital waveguides to create a complete drum model will require improvements and additions to the existing techniques. For instance, to model

the circular drumskin it will be necessary to find a way of accurately modelling curved boundaries at the edge of a 2D waveguide mesh. A technique for attaching different waveguide meshes together will also be needed, as it may be envisaged that the drum model will be made up from a number of meshes. There are other aspects of a drum that the current waveguide techniques are not capable of modelling. These will be addressed in the next section.

Part II

The Physical Model

The Physical Model - overview

The focus of this work was to produce a model of a drum using digital waveguide techniques. The state of such techniques, as of commencing the research, has been described in chapter 3. In order to create a drum model new techniques were required and the existing techniques needed adapting and improving.

The process adopted to complete the model involved examining the main components of a drum and then attempting to model them separately. Once done the separate models could be combined together. The main components considered for this research are shown in figure 4.0.

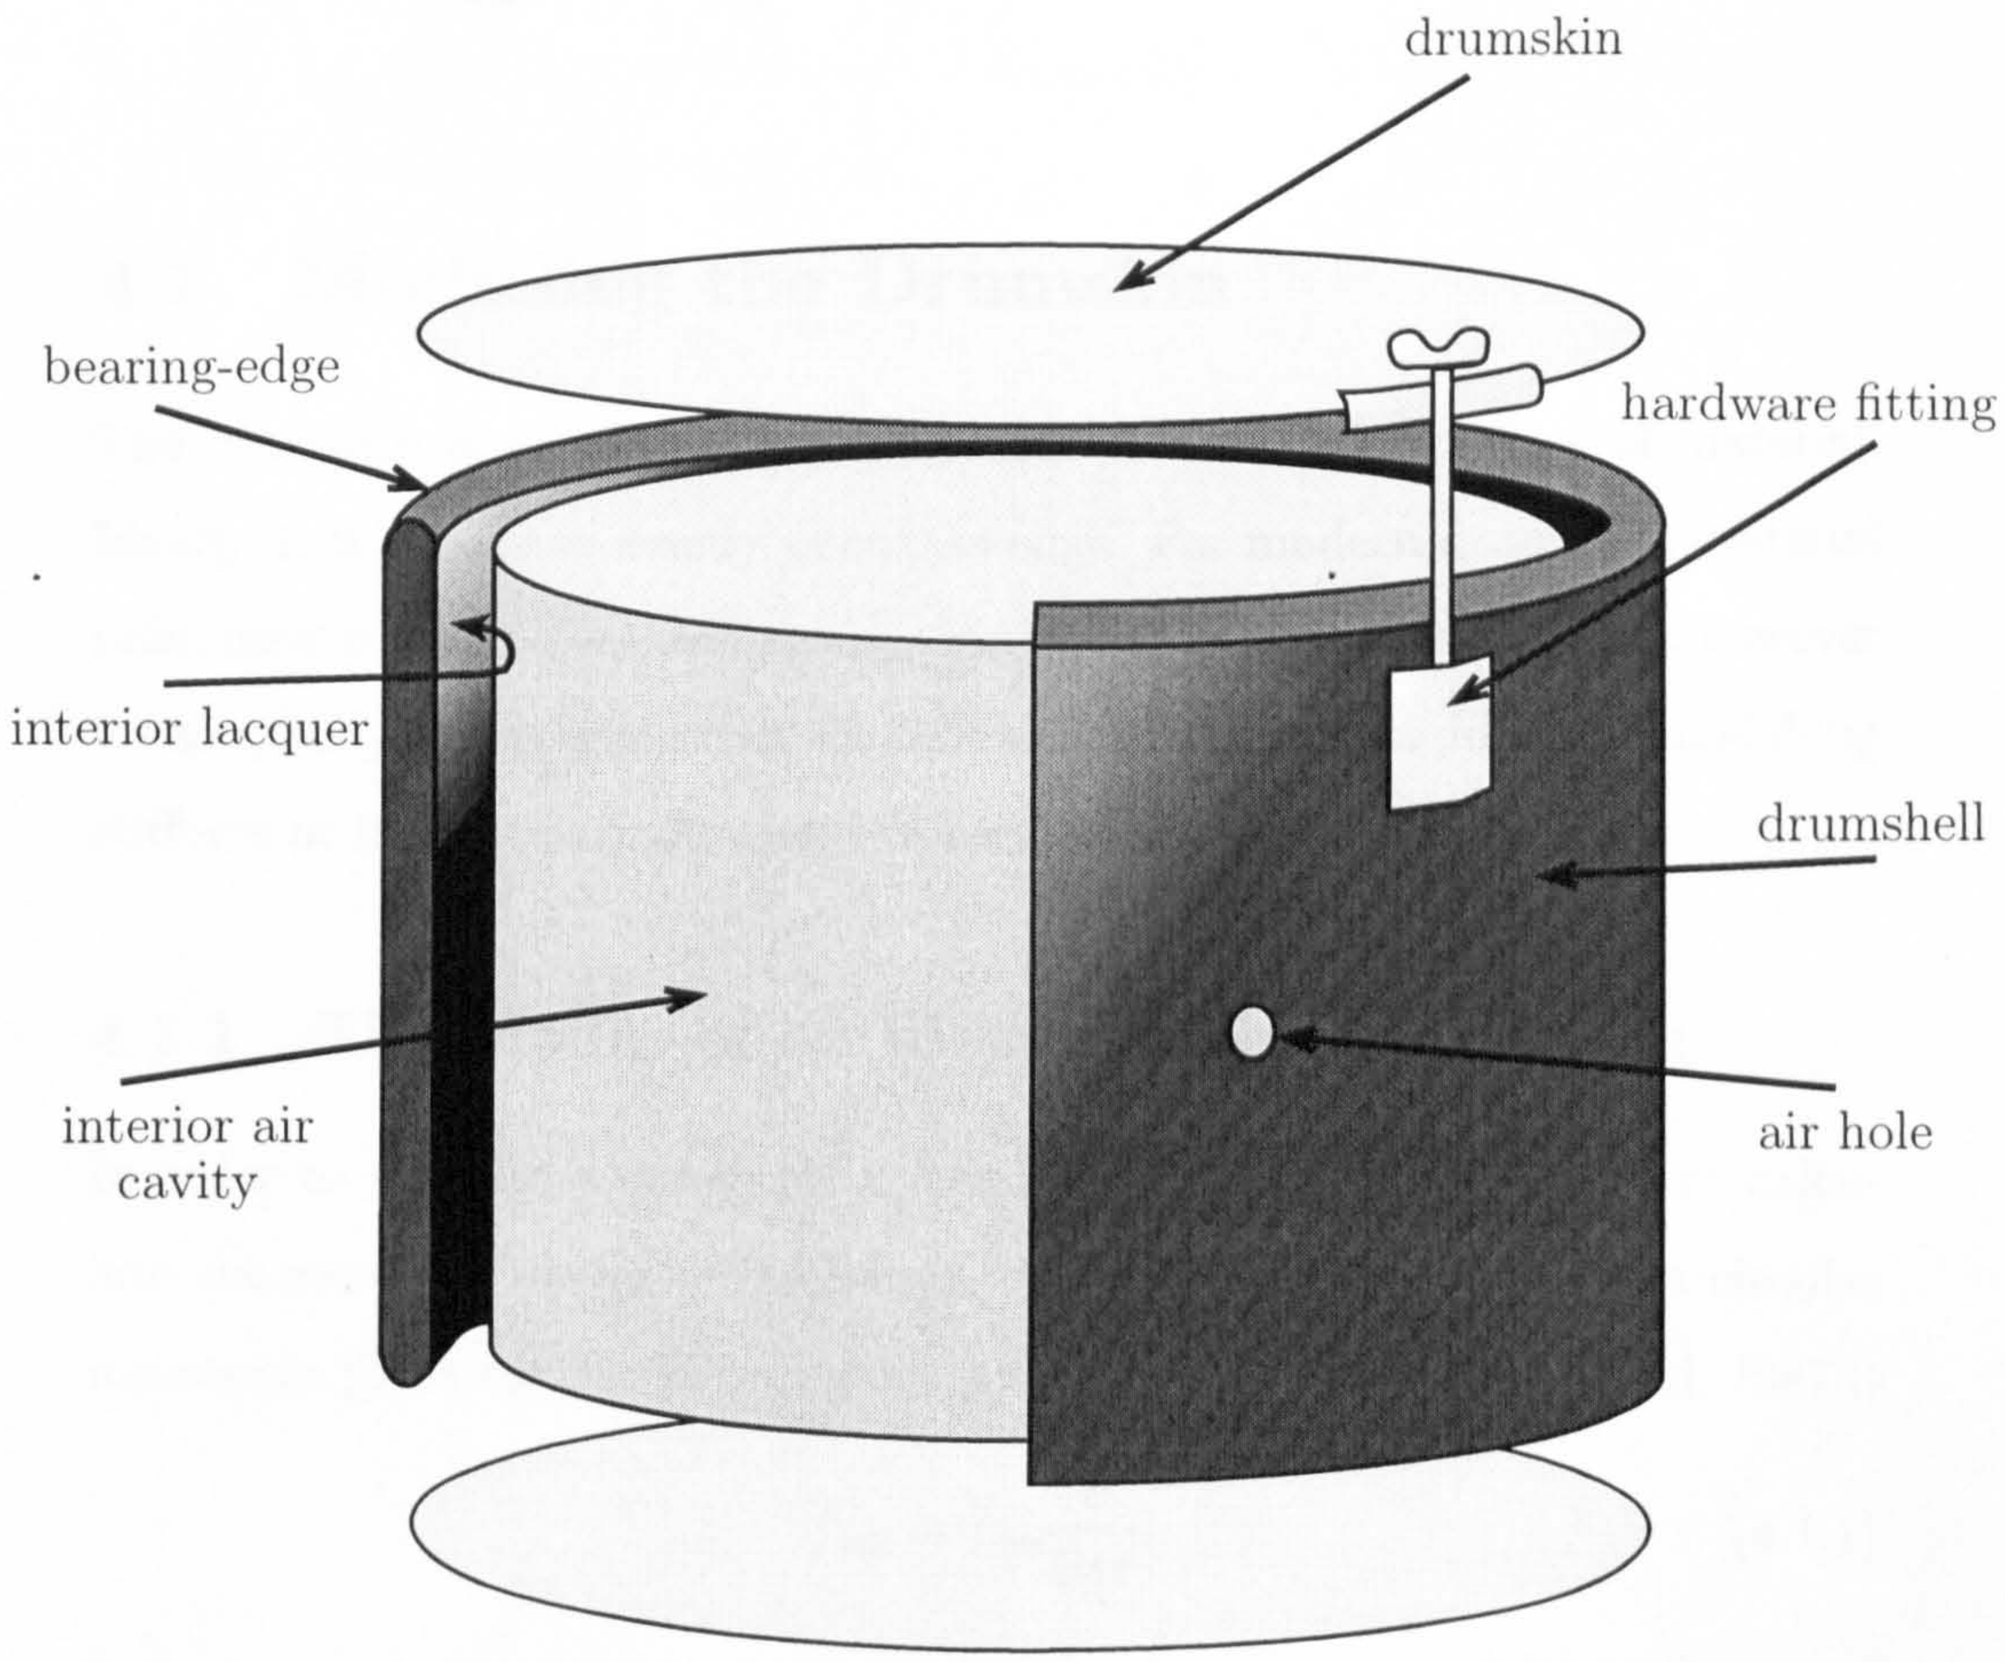

Figure 4.0: *The Main Components of a Drum.*

Chapter 4

Modelling the Drumskin and its Bearing Edge

4.1 Modelling the Drumskin

The drumskin is an almost perfectly flexible and circular piece of material. Ideally, it is tensioned evenly about its edge. For modern drums the material often used is mylar. Naturally this material will exhibit some stiffness, however for simplicity this is ignored in the following investigation. Ideas for modelling stiffness or its effect are discussed in section 8.1.

4.1.1 The Modes of an Ideal Circular Membrane

In order to test the accuracy of a membrane model it is necessary to calculate the expected modes of resonance. A mathematical analysis of a circular membrane produces the following equations to predict the modes (Hall, 1987d)

$$f_{mn} = j_{mn}\frac{c}{2\pi r}, \tag{4.1.1}$$

$$c = \sqrt{\frac{T}{\sigma}}, \tag{4.1.2}$$

where m is the number of nodal lines that cross the membrane diametrically,

n is the number of concentric nodal circles (including the boundary),

f_{mn} is the resonant frequency of the membrane for a particular value of m and n (Hz),

j_{mn} is the n^{th} root of the Bessel function of order m (see appendix B section B.6),

r is the membrane's radius (m),

c is the wave speed $\left(\frac{m}{s}\right)$,

T is the tension $\left(\frac{N}{m}\right)$,

σ is the density $\left(\frac{kg}{m^2}\right)$.

Suitable parameters for a membrane (Fletcher and Rossing, 1991a) are listed in table 4.1. Putting these values into equations 4.1.1 and 4.1.2 results in the modes of resonance shown in table 4.2.

material	Mylar
b (thickness)	$0.0003\,m$
r (radius)	$0.16\,m$
ρ (density)	$\approx 2072\,\frac{kg}{m^3}$
tension	$3500\,\frac{N}{m}$

Table 4.1: *Parameters of a Typical Membrane.*

mn	01	11	21	02	31	12	41
$f_{mn}\,(Hz)$	179.5	285.4	384.1	412.8	475.7	524.1	566.5

Table 4.2: *Resonant Modes of an Ideal Circular Membrane.* $r = 0.16m$, $\sigma \approx 0.6217\,\frac{kg}{m^2}$, $T = 3500\,\frac{N}{m}$.

4.1.2 Modelling the Wave Propagation of the Membrane

The 2D digital waveguide mesh is employed to model the two dimensional transverse wave propagation of a membrane. From looking at the dispersion

error plot of a rectilinear mesh (see figure 3.4 in section 3.7) it can be seen that this technique does not model ideal wave flow correctly. The worst part of the error is the angular dependent dispersion, which causes travelling waves to propagate at different speeds depending on their direction as well as their frequency. This would cause a digital waveguide membrane model to resonate at the wrong frequencies, as though it was a different shape to that intended. Work has recently been done to reduce the amount of angular dispersion error with interpolation techniques (Savioja and Välimäki, 1996), but more accurate models could also be achieved by increasing the sample rate of the model. Doing so increases the maximum possible frequency modelled by the mesh, and so the dispersion error affects frequencies closer to the high-end of the audible range. In addition, as the higher frequency modes of 2D shapes are more dense in the spectrum this error may be imperceptible to the ear.

The 2D Triangular Digital Waveguide Mesh

Raising the sample rate requires the number of nodes in the mesh to be increased, in order to preserve the same spatial size of the modelled object. This is in fact causing the mesh to be more dense, but still the angular dependency of the dispersion error remains the same. Another way to increase the density but keep the sample rate the same and alter the angular dependency is to employ a different mesh structure. Ideally, such a mesh structure would cause the dispersion to be angular independent.

A denser mesh arrangement is found by using a triangular mesh structure, see figure 4.1. Notice that wave propagation is now explicitly modelled in more directions than in a rectilinear mesh (see figure 3.3 in section 3.7). Different 2D mesh structures and their dispersion errors have previously been investigated (Duyne and Smith(III), 1996). Here, in addition to the rectilinear and triangular meshes a hexagonal mesh was considered; this is made of 3-port scattering junctions. However, the triangular mesh can be shown to exhibit near direction independent dispersion error and has been found to support the

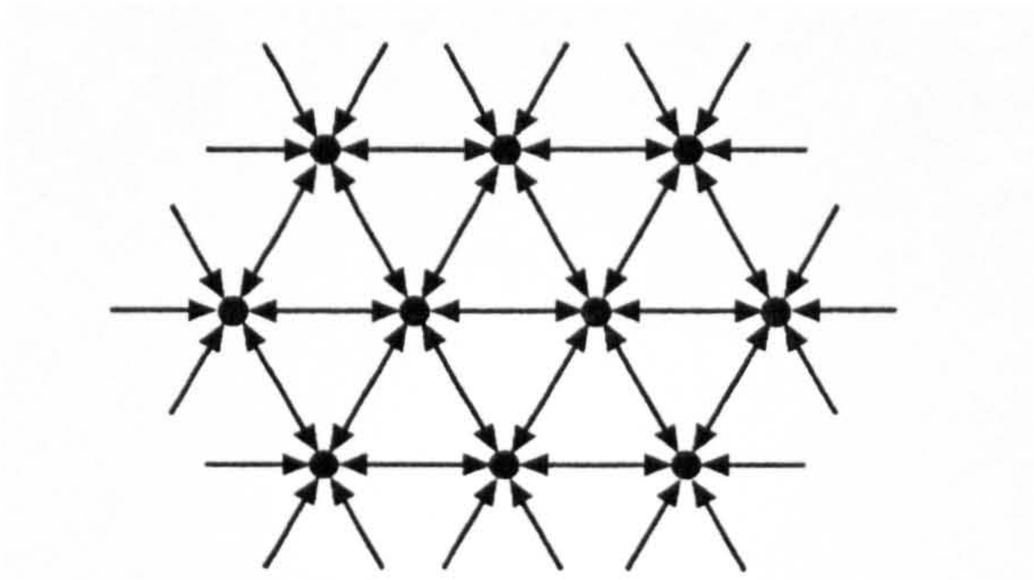

Figure 4.1: *The Structure of a Triangular Mesh. The lines represent waveguides and the interconnections represent lossless scattering junctions.*

widest bandwidth of frequencies (Fontana and Rocchesso, 1999). For these reasons the 2D triangular digital waveguide mesh was chosen to model wave propagation across the drumskin. It is notable that the triangular mesh is the densest 2D arrangement of nodes interconnected by unit length waveguides. This is because the minimum number of equally lengthed lines that can enclose an area is three. The area formed is a triangle, like the triangular elements within the mesh that are bounded by three waveguides each.

The Dispersion Error of a 2D Triangular Digital Waveguide Mesh

The frequency dependent wavespeed (dispersion error) of the triangular mesh is found to be (see appendix E.1 for the derivation)

$$c\left(\underline{\omega}\right) = \frac{\angle G}{|\underline{\omega}|}, \tag{4.1.3}$$

$$\angle G = \arctan \frac{\sqrt{4 - b^2}}{-b}, \tag{4.1.4}$$

$$b = -\frac{2}{3}\left(\cos\omega_x + \cos(\frac{1}{2}\omega_x - \sqrt{\frac{3}{4}}\omega_y) + \cos(-\frac{1}{2}\omega_x - \sqrt{\frac{3}{4}}\omega_y)\right), \tag{4.1.5}$$

where $\underline{\omega}$ is the two-dimensional spatial frequency vector with variables $(\omega_x, \omega_y)^T$,

$c\left(\underline{\omega}\right)$ is the frequency and direction dependent speed of plane wave travel (in space samples per time sample).

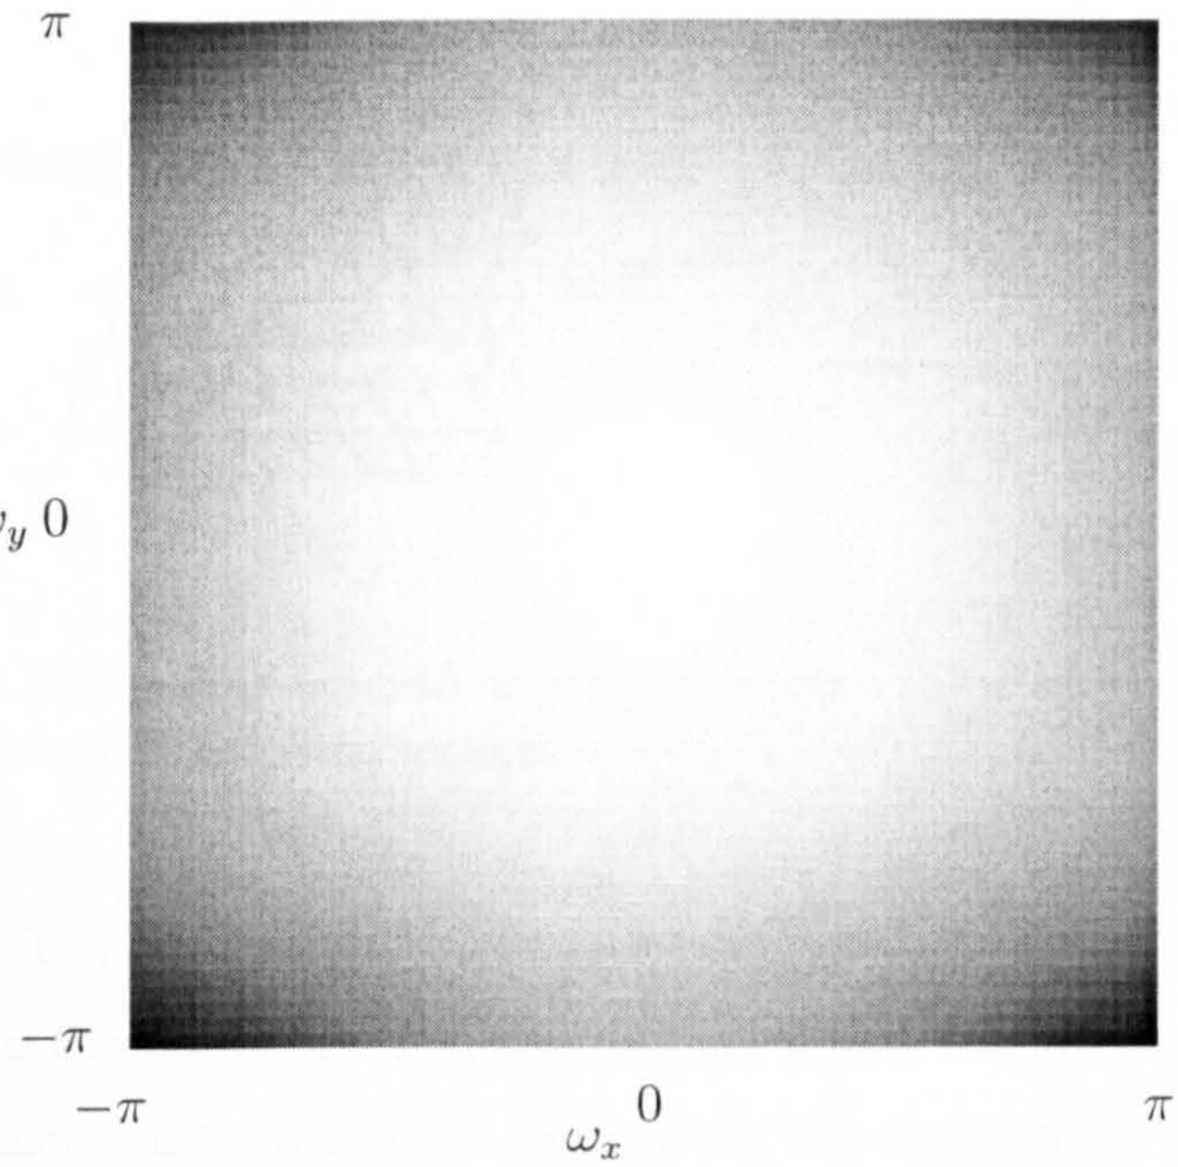

Figure 4.2: *The Dispersion Error of a Triangular Waveguide Mesh.*

The nominal wavespeed within the mesh, where $\omega_x = \omega_y = 0$ has previously been found (Fontana and Rocchesso, 1998)

$$c_0 = \frac{1}{\sqrt{2}}. \tag{4.1.6}$$

Figure 4.2 shows the results of plotting $c\,(\underline{\omega})$ for different values of $\underline{\omega}$. Although the higher frequencies travel slower, as indiciated by the dark regions, the shading is almost circular in shape. This demonstrates that the triangular mesh structure has near direction independent dispersion error.

4.1.3 Modelling the Clamped Boundary

Now that an appropriate mesh structure has been proposed, it is necessary to find some way of adding a circular clamped boundary in order to model the drumskin properly. The following is an investigation that was done in order to advance the state of waveguide modelling techniques to model smooth, curved boundaries.

As previously described a clamped boundary may be modelled by forcing the

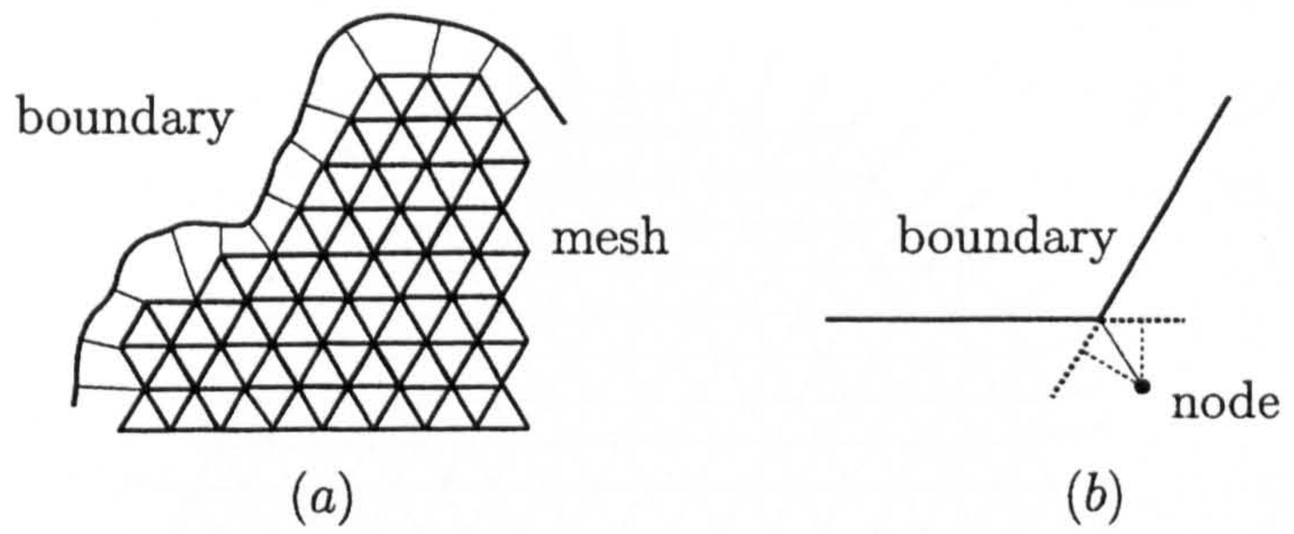

(a) (b)

Figure 4.3: *Rimguides Intersecting Boundaries.* (a) *Perpendicular Intersection,* (b) *Closest Point Intersection.*

nodes' velocities at the edge to zero (see section 3.7.1). However, this will only approximate the boundary shape as not all nodes will lie on this boundary, giving the model a jagged edge. Performing the clamping in this way causes waves travelling up waveguides, that are connected to these edge-most nodes, to be reflected and sent back down again. Alternatively, this effect may be achieved by using special waveguides that perform the reflections at one end, thus disposing of the zero velocity nodes. These special waveguides are looped waveguides that force a sign inversion on the travelling waves. They have been given the name of 'rimguides' (Laird, Masri, and Canagarajah, 1998).

As these rimguides model waves that are reflected back in the direction they have come from, they must lie perpendicular to the gradient at the point where they intersect the boundary, see figure 4.3a. It should be noted that there are instances when a perpendicular intersection of the boundary is not possible, as shown in figure 4.3b. In this case the rimguide is attached to the closest point on the boundary. In fact, attaching to the closest point will also make a perpendicular intersection under normal circumstances, and so this is the preferred method for constructing rimguides.

Another point to notice is that the edge-most nodes may not have as many connections as the nodes within the mesh. This irregularity would result in an impedance discontinuity at the edge, causing inaccurate reflections. To counter the problem the rimguides' impedances must be scaled to remove this

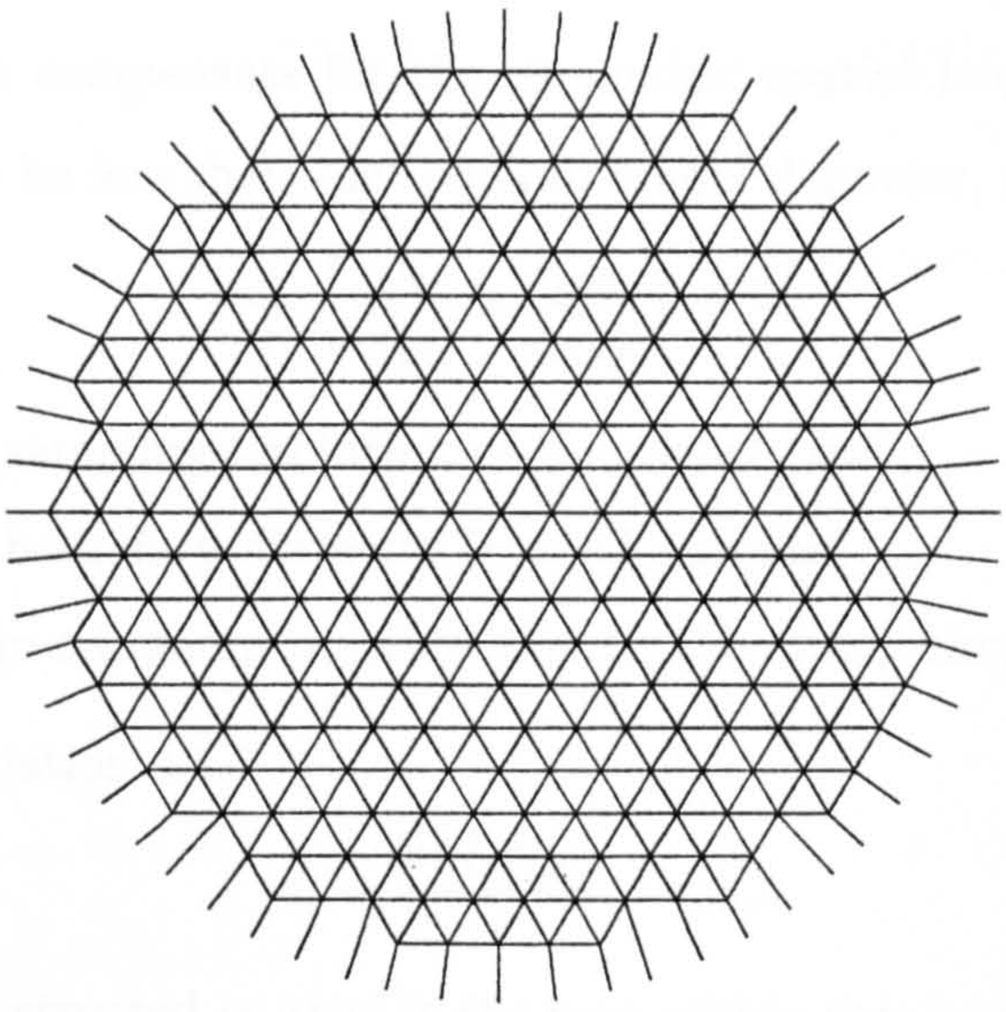

Figure 4.4: *A Circular Membrane modelled by a Waveguide Mesh with Rimguides about its Edge.*

impedance discontinuity. For example, consider an edge-most node with n waveguides attached to it, each of impedance R. If the nodes within the mesh have N waveguides attached, then this edge node should have a rimguide of impedance $R(N-n)$ appended, to preserve the total impedance of R *times* N.

Shaping the Waveguide Mesh

The rimguides used to model a clamped boundary are to be attached to the edge of the digital waveguide mesh. The mesh itself must be shaped such that it fits best within a circular boundary; any nodes outside this boundary are removed from the model. Figure 4.4 depicts a waveguide mesh with rimguides attached to model a circular membrane. The diameter of the boundary is determined by satisfying the following relation: the time it takes for a travelling wave to propagate from one side of the membrane model, through the centre to the opposite side, must equal the time taken in the ideal membrane. As

the rimguides will be an extension of the mesh the wavespeed through them is also $c(\underline{\omega})$. To compensate for the rimguides' spatial length, the boundary diameter has to be less than the required model diameter, resulting in

$$\frac{1}{f_s}\left(\frac{d_b}{c(\underline{\omega})}+\frac{\tau_{min}}{c(\underline{\omega})}\right)=\frac{2r}{c}, \tag{4.1.7}$$

where f_s is the sampling resolution of the model (Hz),

d_b is the boundary diameter (spatial samples),

τ_{min} is the minimum possible loop length of any rimguide (spatial samples).

Rearranging equation 4.1.7

$$d_b=\frac{2rf_sc(\underline{\omega})}{c}-\tau_{min}. \tag{4.1.8}$$

The mesh is constructed so that it fits best within this boundary. Because it is made of triangular shapes its edge will be jagged, resulting in an angular dependent diameter $d_m(\angle\underline{\omega}) \leq d_b$.

Constructing the Rimguides

The aforementioned angular dependent diameter may be substituted in place of the desired boundary diameter in a rearranged form of equation 4.1.7; this gives an expression from which the required rimguide lengths or phase delays can be calculated

$$\tau(\underline{\omega})=\frac{2rf_sc(\underline{\omega})}{c}-d_m(\angle\underline{\omega}), \tag{4.1.9}$$

where $\tau(\underline{\omega})$ is the desired phase delay or loop length of a rimguide situated around the mesh at angle $\angle\underline{\omega}$ from the horizontal.

The desired phase delay of the rimguides is angular dependent and therefore the slight hexagonal shape apparent in figure 4.2 can be compensated for. Ideally the rimguides should match this desired response for all frequencies to make the boundary truly circular. Unfortunately this is not possible, but there are a number of different ways that may be considered in which to construct the rimguides, some yielding better results than others. The following is a description and comparison of four rimguide construction methods.

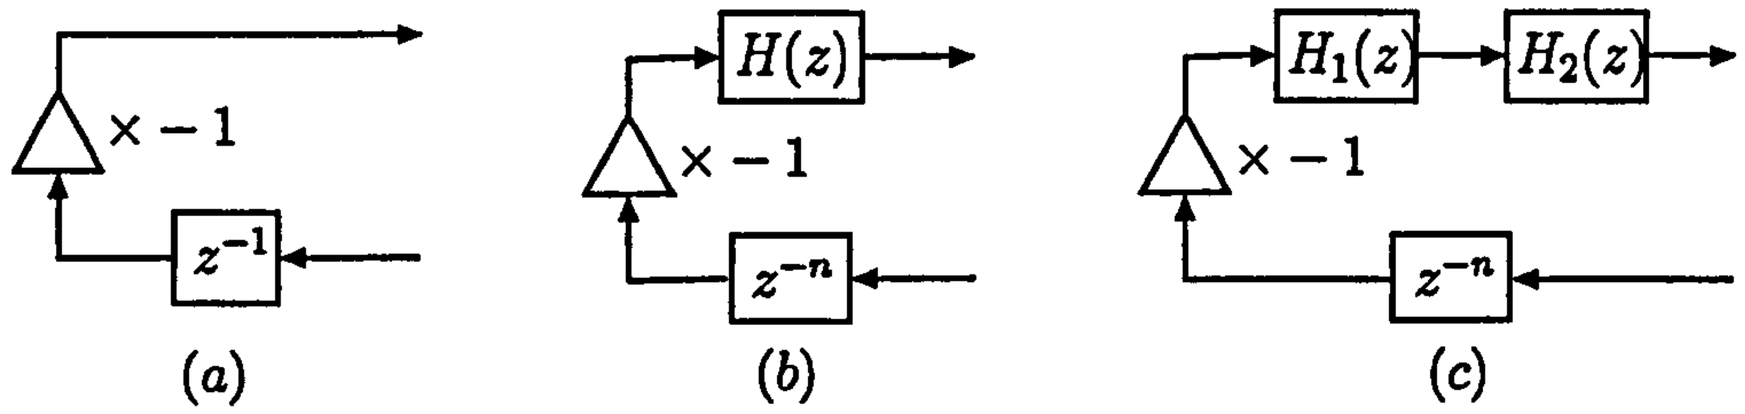

Figure 4.5: *Different Rimguide Structures.*

Method 1

The simplest rimguide consists of just one unit of delay, see figure 4.5*a*, to make $\tau_{min} = 1$ in equation 4.1.8. The mesh is constructed using this equation and equation 4.1.9 with $|\underline{\omega}| = 0$.

As this delay models both the outward and inward bound paths of wave travel, the rimguide extends out from the mesh by half a sample. These rimguides retain the jagged mesh boundary shape of the membrane model, and they certainly don't satisfy equation 4.1.9 as $\tau(\underline{\omega}) = \tau_{min}$. The modelled membrane is therefore not of the right size, nor truly circular, although it may suffice for models of high sample rate.

Method 2

To overcome the limitations of the previous method the rimguides require fractional spatial lengths. The modelled membrane boundary can then be extended out to the desired boundary. As mentioned in section 3.3 a fractional delay, or fractional spatial length, can be modelled with fractional delay filters (Laakso, Välimäki, Karjalainen, and Laine, 1996). The simplest to use is the first order allpass filter (see appendix F.1 for information on the construction and analysis of this filter), which has the transfer function

$$H(z) = \frac{a + z^{-1}}{1 + az^{-1}}. \tag{4.1.10}$$

This filter's phase response causes a frequency dependent fractional delay. For waves close to DC ($0\,Hz$) the coefficient, a, may be approximated with (see

appendix F.1.3 equation F.29)

$$a \approx \frac{1-D}{1+D}, \tag{4.1.11}$$

where D is the desired delay of the filter (samples).

The filter best approximates a fractional delay within the range: $0.5 \leq D \leq 1.5$. The new rimguide has this filter appended to at least one unit of delay, see figure 4.5*b*, making $\tau_{min} = 1.5$. Due to the positioning of nodes and the boundary, it is possible that more than one unit of delay may be required in order to keep the fractional delay within the specified range. Hence, the delay of the filter is calculated with

$$D = \tau(0) - n, \tag{4.1.12}$$

where n, the number of unit delays, is the largest integer

that satisfies $n \leq \tau(0) - 0.5$.

The mesh is constructed using equations 4.1.8, 4.1.9 and 4.1.12, with $|\underline{\omega}| = 0$. This method ensures that a wave close to DC takes the same time to travel diametrically across the membrane model as it would in the real membrane. The accuracy of the model will decrease for higher frequencies. Therefore, this model achieves a circular boundary for low frequency waves.

Method 3

It is possible to tune the rimguide of method 2 so that waves of a particular frequency take the right amount of time to traverse the model, the aim being to improve the rimguide accuracy. For this method the 'focus' frequency is chosen to be the membrane's fundamental mode of resonance. Obviously equations F.29 (see appendix F.1.3) and 4.1.12 can no longer be used as the coefficient, a, needs to be found for $\tau(\theta^c)$ where θ^c is the normalised-pi fundamental frequency, $\frac{2\pi f_{01}}{f_s}$. For the membrane under study the fundamental frequency is $f_{01} = 179.5Hz$ (see table 4.2).

The filter coefficient may be found for a particular frequency and delay using (see appendix F.1.3 equation F.28)

$$a = \frac{B \pm \sqrt{(B^2+1)\sin^2|\underline{\omega}|}}{\sin|\underline{\omega}| - B\cos|\underline{\omega}|}, \tag{4.1.13}$$

$$B = \tan(-D(|\underline{\omega}|) \times |\underline{\omega}|), \tag{4.1.14}$$

where $D(|\underline{\omega}|) = \tau(\underline{\omega}) - n$ is the delay of the filter (samples),

n is the largest integer that satisfies $n \leq \tau(\underline{\omega}) - 0.5$.

There are two possible solutions to equation 4.1.13, however there will be only one coefficient that gives a stable filter, in the range: $-1 < a \leq 1$. When $\tau(\underline{\omega})$ and $|\underline{\omega}|$ are large there are no valid solutions, but as long as the focus frequency is a small fraction of the sample frequency there will always be a solution.

The mesh is constructed using equations 4.1.8, 4.1.9 and 4.1.13, with $|\underline{\omega}| = \frac{2\pi f_{01}}{f_s}$. Once again the minimum possible loop length of the rimguide is $\tau_{min} = 1.5$ spatial samples.

Method 4

As well as tuning the rimguide to the fundamental resonant node, it can also be made to include the second resonant mode. This is achieved by adding another fractional delay filter to the rimguide structure, see figure 4.5c. The two filters have coefficients a_1 and a_2 with corresponding fractional delays $D_1(|\underline{\omega}|)$ and $D_2(|\underline{\omega}|)$. The delay of either filter may be calculated from its coefficient using (see appendix F.1.2 equations F.19 and F.20)

$$D_i(|\underline{\omega}|) = \frac{-\arctan\left(\frac{(a_i^2-1)\sin|\underline{\omega}|}{2a_i+(1+a_i^2)\cos|\underline{\omega}|}\right)}{|\underline{\omega}|}. \tag{4.1.15}$$

The minimum possible loop length of the rimguide is still $\tau_{min} = 1.5$ spatial samples, as one filter may effectively be removed by setting its coefficient to 1. This would leave a rimguide of one filter and one unit of delay.

The two filters combine to give a more complicated phase delay response which

can be controlled to result in specific delays at two different frequencies. It is difficult to find a mathematical solution to predict the necessary coefficients, so a numerical search method may be employed instead. The search method involves choosing trial solutions of a_1 which keep $D_1(|\underline{\omega}|) \leq \tau(\underline{\omega}) - n$ at the frequency of the fundamental resonant mode, f_{01}. As before n, the number of unit delays, is the largest integer that satisfies $n \leq \tau(\underline{\omega}) - 0.5$.

The required delay of the second filter can be found at this frequency

$$D_2(|\underline{\omega}|) = \tau(\underline{\omega}) - D_1(|\underline{\omega}|). \tag{4.1.16}$$

Equations 4.1.13 and 4.1.14 are then used to calculate this filter's coefficient, a_2.

The trial solution can be verified by checking the resultant delay at the second resonant mode, f_{11}

$$D_1(|\underline{\omega}|) + D_2(|\underline{\omega}|) = \tau(\underline{\omega}). \tag{4.1.17}$$

The search continues until equation 4.1.17 holds at both frequencies, depending on whether or not a solution exists. If there is no solution then the first resonant mode is modelled accurately and the second resonant mode is approximated as best as possible.

The mesh is constructed in the same way as the other methods, but initially with $|\underline{\omega}| = \frac{2\pi f_{11}}{f_s}$, the second resonant mode. The reason for this is so that the higher focus frequency will not need less than 1.5 samples delay; due to the dispersion error it travels more slowly through the mesh and therefore requires the least delay in the rimguide.

Results and Comparison of the Four Rimguide Methods

Models of the membrane under study were created using each of the four methods with a sample rate of $f_s = 11025\,Hz$. The models were excited with an off-centre impulse in order to excite all the resonant modes (see section 6.2.1 for more information on implementing an impulse excitation). Results were obtained by recording the wave velocities at the excitation node, and then

Resonant Mode	Theoretical (Hz)	Method 1 (Hz)	Method 2 (Hz)	Method 3 (Hz)	Method 4 (Hz)
1	179.5	185 (+3.0%)	178 (−0.8%)	178 (−0.8%)	178 (−0.8%)
2	285.4	294 (+3.0%)	283 (−0.8%)	283 (−0.8%)	283 (−0.8%)
3	384.1	394 (+2.6%)	379 (−1.3%)	379 (−1.3%)	379 (−1.3%)
4	412.8	423 (+2.5%)	407 (−1.4%)	407 (−1.4%)	407 (−1.4%)
5	475.7	488 (+2.6%)	471 (−1.0%)	471 (−1.0%)	470 (−1.2%)
6	524.1	537 (+2.5%)	518 (−1.2%)	518 (−1.2%)	517 (−1.4%)
7	566.5	581 (+2.6%)	559 (−1.3%)	559 (−1.3%)	559 (−1.3%)

Table 4.3: *Resonances from Membrane Models using Different Rimguide Methods.*

analysing this output in terms of its power spectral density to find the resonant modes. These results are presented in table 4.3.

It can be seen that the first seven resonant modes are modelled with reasonable accuracy for all the methods. The most noticeable difference is the response obtained using the first rimguide method, which consisted of a single unit delay to model the path to the boundary and back. Here the frequencies are higher than intended, due to the modelled membrane diameter being too short in many cases.

The remaining three methods used fractional delays to compensate for the dispersion error of the mesh and to correctly model the distance from the edge of the mesh to the actual boundary. These methods all produce resonances that are slighly lower than intended, yet they are more accurate than method 1. Interestingly, their resonances are identical to each other except for the fifth and sixth resonant modes of method 4, which are slightly lower still.

In a previous study of the four rimguide methods (Laird, Masri, and Canagarajah, 1998) a circular membrane was modelled with a lower density mesh and the impulse was applied at the centre to excite all the modes of frequency f_{0n}. In this case the pattern of results was similar, with the fourth method once again being slightly less accurate than the preceding two. Despite focusing two resonant modes in an attempt to increase accuracy, the fourth method also introduces a greater delay for the higher frequencies. This greater delay

is due to the presence of the extra allpass filter and it effectively lengthens the rimguide. As the rimguides are string models in essence, the longer they become the more the membrane model behaves like a string; lowering the fundamental and altering the pattern of the modes.

The extra filter in the fourth method results in an additional computational burden and also makes the filter coefficients more difficult to calculate. When this is considered along with its apparent reduction in accuracy it may be deduced that the fourth method is undesirable when compared with methods 2 and 3. These two methods are identical in terms of their implementation and appear to yield the same results, yet method 2 has the simplest filter coefficient calculations. Rimguide method 2 therefore, which was designed to approximate the boundary best for waves close to DC, is the most useful of the four methods. It is simple to construct and is also accurate. The CD-Rom (see section A.1.1 of appendix A) has a recording of the sound produced from a circular drumskin model that used this method. The parameters for this model were the same as those used in this section. As the model is lossless the sound produced is continuous and would last forever. This is, of course, a very unnatural noise for a drumskin.

4.1.4 Adding Viscous Friction Losses to the Drumskin Model

The model of the drumskin accurately reproduces the important resonant modes but does not include any losses. The majority of these losses would be due to internal viscous friction and their inclusion in the model is necessary to create a realistic sound.

As mentioned in section 3.5 it is possible to model the effect rather than the cause; the effect being that high frequencies suffer a greater loss than low frequencies. Naturally a lowpass filter can be used for this purpose. The first order IIR lowpass filter (see appendix F.2 for information on the construction

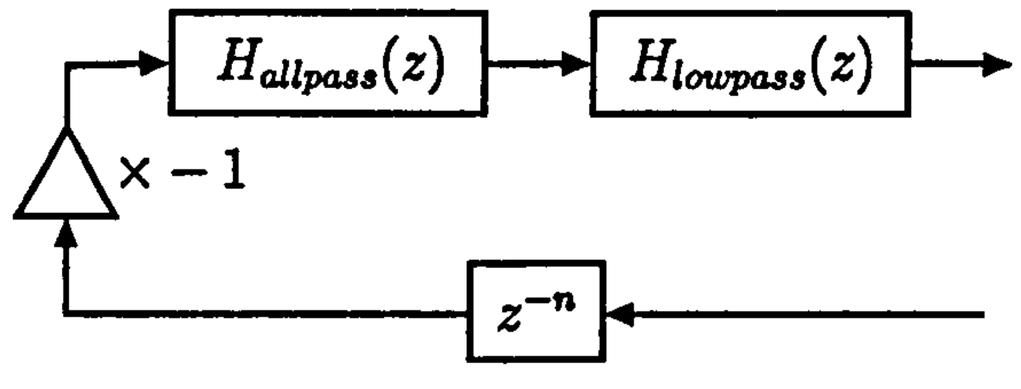

Figure 4.6: *The Lossy Rimguide.*

and analysis of this filter) is simple yet controllable, and has the following transfer function

$$H_{lowpass}(z) = \frac{1+a}{1+az^{-1}}. \tag{4.1.18}$$

Rather than trying to incorporate this filter within the waveguide mesh structure, which would make the calculation of the dispersion error unwieldy, it is better to add it to the rimguides that surround the boundary. Doing so results in the new rimguide structure show in figure 4.6.

Designing the Lowpass Filter

As the effect of a frequency dependent loss is to be modelled it is useful to characterise this in terms of a desired loss at a certain frequency, $L(\omega)$ $(\frac{dB}{s})$. This can be matched against experimental data. The filtering is only to be done at the mesh boundary, so the required loss at this point will be

$$L_b(\omega) = L(\omega)\frac{2r}{c}, \quad (dB) \tag{4.1.19}$$

which can be converted into a gain (see equation 2.1.1 of section 2.1.2)

$$G(\omega) = 10^{\frac{-L_b(\omega)}{20}}. \tag{4.1.20}$$

The lowpasss filter coefficent, a, that results in this gain at frequency ω is found with (see appendix F.2.3 equation F.47)

$$a = \frac{1 - G(\omega)\left(G(\omega)\cos\omega \pm \sqrt{G(\omega)^2(\cos^2\omega - 1) + 2 - 2\cos\omega}\right)}{G(\omega)^2 - 1}. \tag{4.1.21}$$

Note that there are two possible solutions to equation 4.1.21, but there will be only one coefficient that gives a stable filter, in the range: $-1 < a \leq 1$.

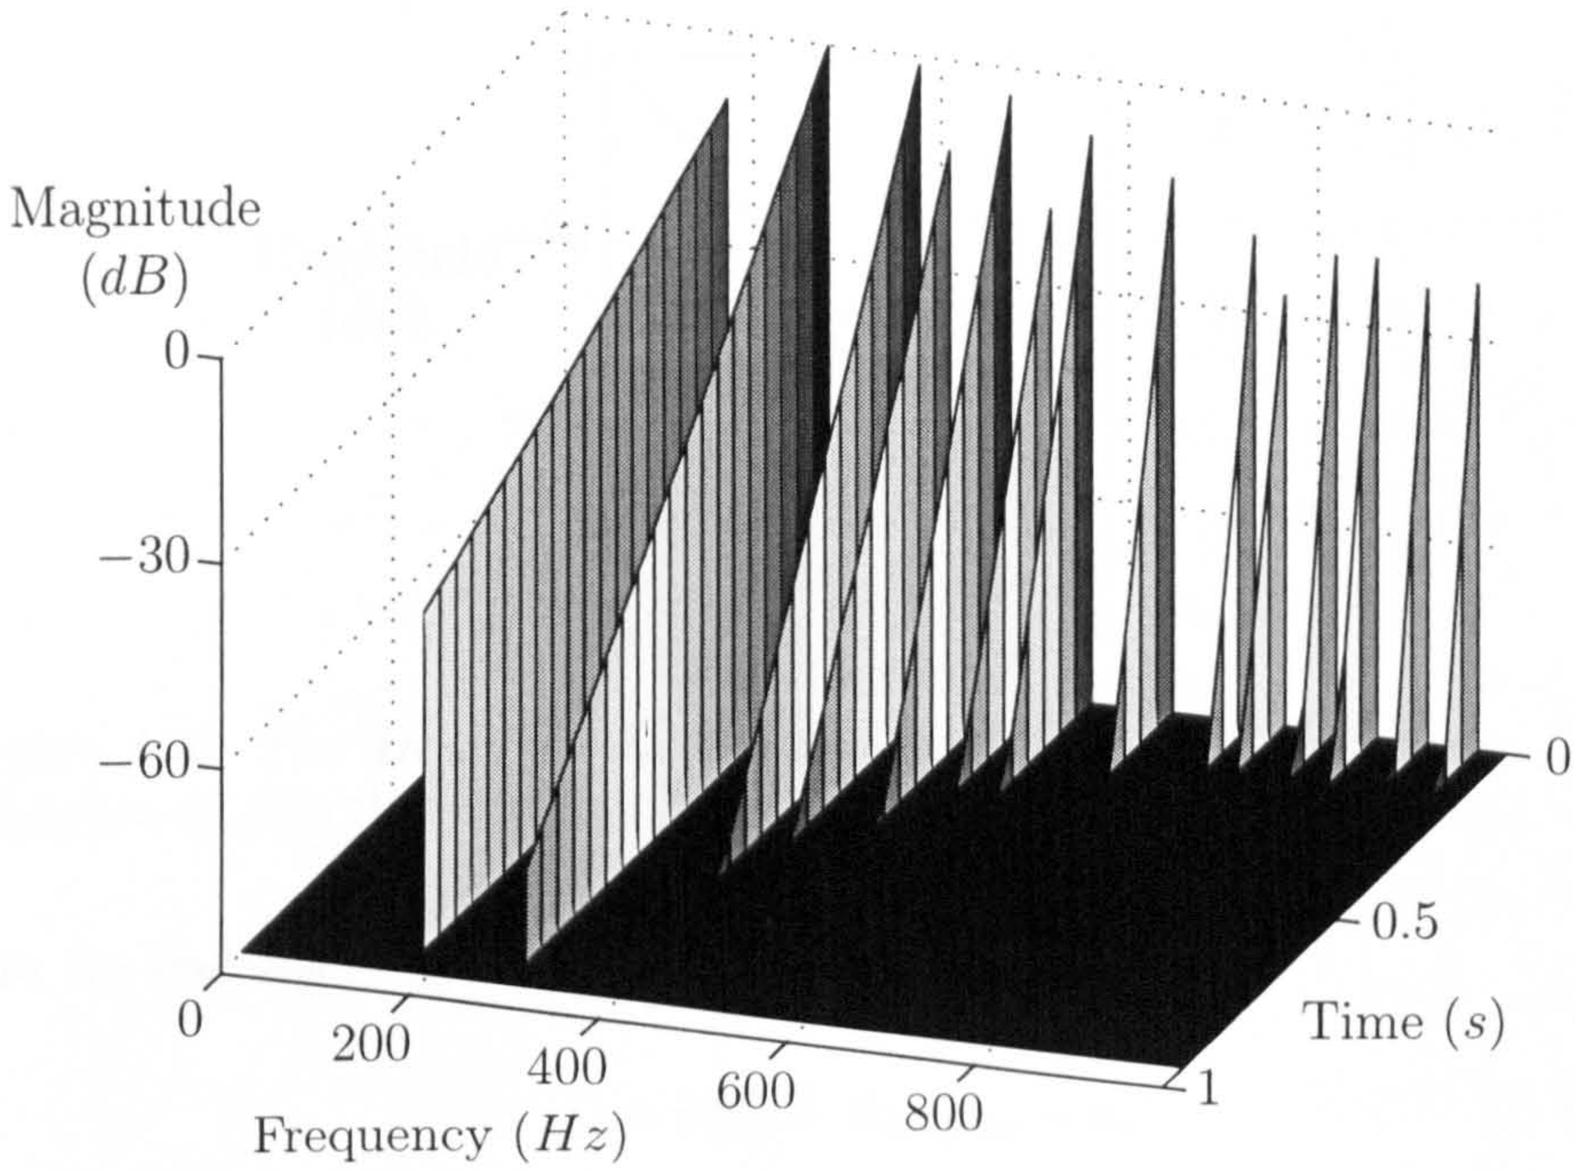

Figure 4.7: *Waterfall Plot of Sound from a Circular Membrane with Viscous Friction Losses.*

Reconstructing the Rimguides

The lowpass filter has an associated delay that will be added to the rimguides' looplength. To account for this the rimguides need to be reconstructed, which involves reworking their equations. The lowpass delay is found to be (see appendix F.2.2 equations F.38 and F.39)

$$D_{lowpass}(\omega) = \frac{\arctan\left(\frac{-a\sin\omega}{1+a\cos\omega}\right)}{\omega}. \tag{4.1.22}$$

As $\omega \mapsto 0$ this becomes

$$D_{lowpass} = \frac{-a}{1+a}. \tag{4.1.23}$$

Assuming that method two is to be used for rimguide construction, the minimum possible loop length is now $\tau_{min} = 1.5 + D_{lowpass}$. This is used in equation 4.1.8 as before to limit the waveguide mesh size. The calculation used to find the allpass filter delay of the rimguide, equation 4.1.12, is changed to incorpo-

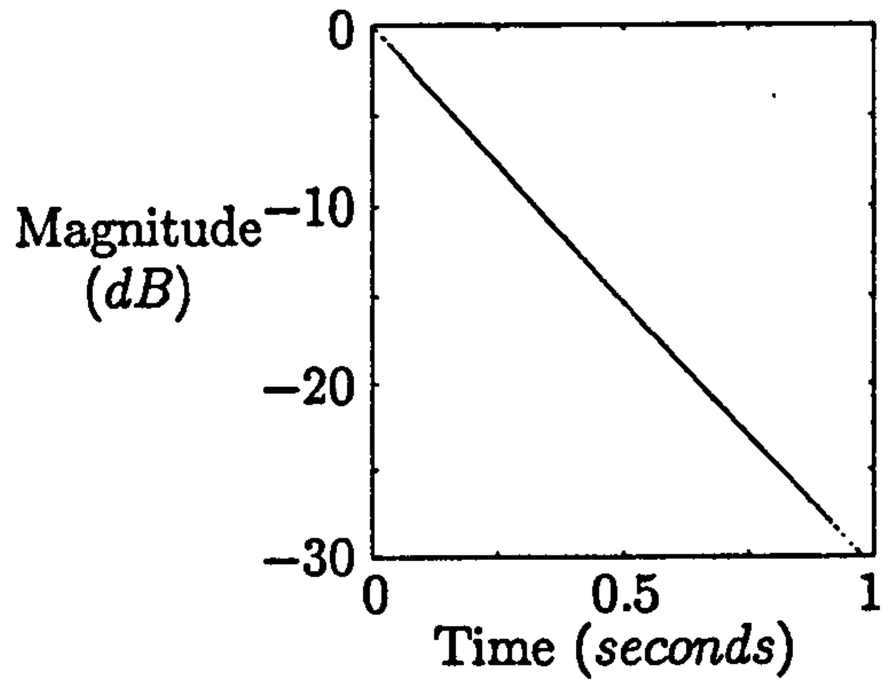

Figure 4.8: *The Decay of the Fundamental Resonant Frequency,* f_{01}*, with* $L(f_{01}) = 30\,dBs^{-1}$.

rate the lowpass delay

$$D = \tau(0) - D_{lowpass} - n, \tag{4.1.24}$$

where n, the number of unit delays, is the largest integer

that satisfies $n \leq \tau(0) - D_{lowpass} - 0.5$.

Results from the Viscous Friction Model

The effect of adding lossy rimguides to the membrane model is shown in the waterfall plot of figure 4.7. Here a circular membrane model was constructed with the same material parameters and dimension as before. This time the sample rate used was $f_s = 22050\,Hz$ in order to better define the resonant modes. The value of the loss used was $-30\,dBs^{-1}$ at the fundamental frequency: $L(f_{01}) = 30\,dBs^{-1}$. The plot was created by performing a Fourier analysis (see appendix B.5) of 2048 points with Hanning window and 50% overlap, plotting the change in magnitude of the resonant frequencies over time. It can be seen that the lossy rimguides simulate the effect of internal viscous friction, where the higher frequencies decay faster than the lower frequencies. The rate of decay of the fundamental frequency was checked by isolating the first mode in the waterfall plot and normalising its magnitude such that it started at $0\,dB$. This is shown in figure 4.8. The decay of the fundamental

frequency can be seen to agree with the desired loss of $-30\,dBs^{-1}$.
A recording of this sound can be found on the CD-Rom (see section A.1.1 of appendix A). Notice that the frequency dependent damping produces a much more natural sound than that from the lossless drumskin model. It is similar to a tight plastic membrane being flicked, such as a mylar drumskin before it is attached to a drum. The impulse excitation injected all possible frequencies that the model could support. Whilst useful for analysing the model's response this method of excitation is unrealistic and causes an initially harsh and distorted sound. Notice also that there is a faint high frequency 'ring' throughout the sample. This has been heard in other models (constructed during this investigation) that incorporate excessive damping and have been excited with an impulse; it is possible that this was due to a coding mistake in the mesh structure.

4.1.5 Impedance of the Membrane

Whilst the impedances of 1D and 3D media are well defined (see section 3.8) there is very little in the literature about the impedance of a 2D membrane, and a method to calculate the impedance of a 2D waveguide mesh did not exist. The following presents a method that was found to calculate this impedance, which is required to find the velocity of the membrane that results from an applied force. Naturally, this force has to be exerted over a finite area or it would cause an infinite acceleration. The impedance is useful when connecting the membrane model to other media, in order to make a drum model for example. It has been found that the impedance can be calculated if the membrane is considered to be a network of interconnecting strings. Afterall, the digital waveguide mesh is made up from a network of single digital waveguides, each representing a string.
Consider a single element of a triangular digital waveguide mesh, as shown in figure 4.9. This element is the fundamental building block of the mesh that

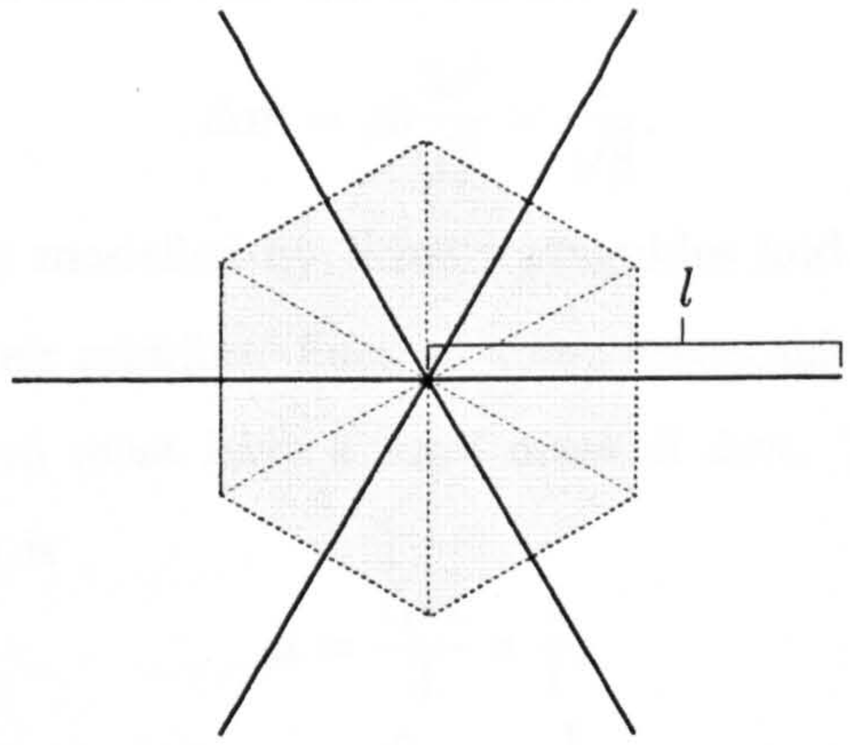

Figure 4.9: *Single Element of a Triangular Waveguide Mesh. The solid lines are waveguides (of length l metres), the shaded part is the fundamental building block of the mesh.*

can be tessellated, as suggested by the arrangement of waveguides. It can be seen that this element is hexagonal in shape. Its mass is

$$\Delta m = \rho b A, \tag{4.1.25}$$

where A is the surface area of the element (m^2),

ρ is its volume density $\left(\frac{kg}{m^3}\right)$,

b is its thickness (m).

As the wave speed through the mesh is $c_0 = \frac{1}{\sqrt{2}}$ (spatial samples per time sample), the actual length of a waveguide must be $l = \frac{\sqrt{2}c}{f_s}\,(m)$. Therefore, the surface area of the element can be found by summing the area of each triangle indicated in the diagram

$$A = \tan\frac{\pi}{6} \times \frac{l}{2} \times \frac{l}{2} \times 6, \tag{4.1.26}$$

$$= \left(\frac{\sqrt{2}c}{2f_s}\right)^2 \tan\frac{\pi}{6} \times 6, \tag{4.1.27}$$

$$= 3\frac{c^2}{f_s^2} \tan\frac{\pi}{6}. \tag{4.1.28}$$

Combining equations 4.1.25 and 4.1.28 results in

$$\Delta m = \rho b \frac{3c^2}{f_s^2} \times \frac{1}{\sqrt{3}}. \tag{4.1.29}$$

Now, this element is modelled by three waveguides laid on top of eachother and connected at their middles. Each of these waveguides represents a string which when combined must have a total mass of Δm. Therefore, the linear density of one string is

$$\mu = \frac{\Delta m}{3} \times \frac{1}{l}, \tag{4.1.30}$$

$$= \rho b \frac{c}{\sqrt{2} f_s} \times \frac{1}{\sqrt{3}}, \tag{4.1.31}$$

$$= \frac{\rho b c}{\sqrt{6} f_s}. \tag{4.1.32}$$

As mentioned in section 3.8 the impedance of a string is found to be $R = \mu c$. This makes the impedance of each unit length waveguide in the waveguide mesh

$$R_{waveguide} = \frac{\rho b c^2}{\sqrt{6} f_s}. \tag{4.1.33}$$

The impedance of an element in the waveguide mesh membrane model can now be found (see section D.2 of appendix D)

$$R_{mesh\,element} = 3 R_{waveguide}, \tag{4.1.34}$$

$$= \sqrt{\frac{3}{2}} \times \frac{\rho b c^2}{f_s}. \tag{4.1.35}$$

When combined with equation 4.1.2 this reduces to

$$R_{mesh\,element} = \sqrt{\frac{3}{2}} \times \frac{T}{f_s}, \tag{4.1.36}$$

where T is the tension applied to the membrane (Nm^{-1}).

Notice that this derivation of the impedance within a 2D waveguide mesh is related to the tension and sampling rate; the density, thickness and wavespeed apparently play no part. However, it is important to note that these quantities are incorporated into the model as they affect the spatial length of the waveguides. It should be remembered that this is not a driving-point impedance, it is the impedance of an element in the waveguide mesh.

4.2 Modelling the Bearing Edge

The bearing edge is the part of the drum that supports the drumskin. It is added on to, or carved out of the edge of the drumshell and is analogous to the bridge of a stringed instrument, such as that found on the violin or guitar. As the drumskin is tensioned it is pulled tightly over the bearing edge, as shown in the cross-sectional diagram in figure 4.10.

Premier Percussion Ltd. have a special drill-bit to cut the bearing edge of their drums. They don't, however, have any information as to how this edge affects the sound produced. There is no existing technique to incorporate this into a digital waveguide implementation, so the following work was done to model and study its effect. Modelling this aspect of the drum is therefore important so that the effect, if any, can be measured. Once modelled it would be possible to test different bearing edge geometries in order to ascertain how the effect may be controlled. The parameters: r, α and β indicated in figure 4.10, are those that can be changed by making adjustments to the machine that produces the drill-bits. The parameters l_1 and l_2 may be controlled by changing the position of the lathe.

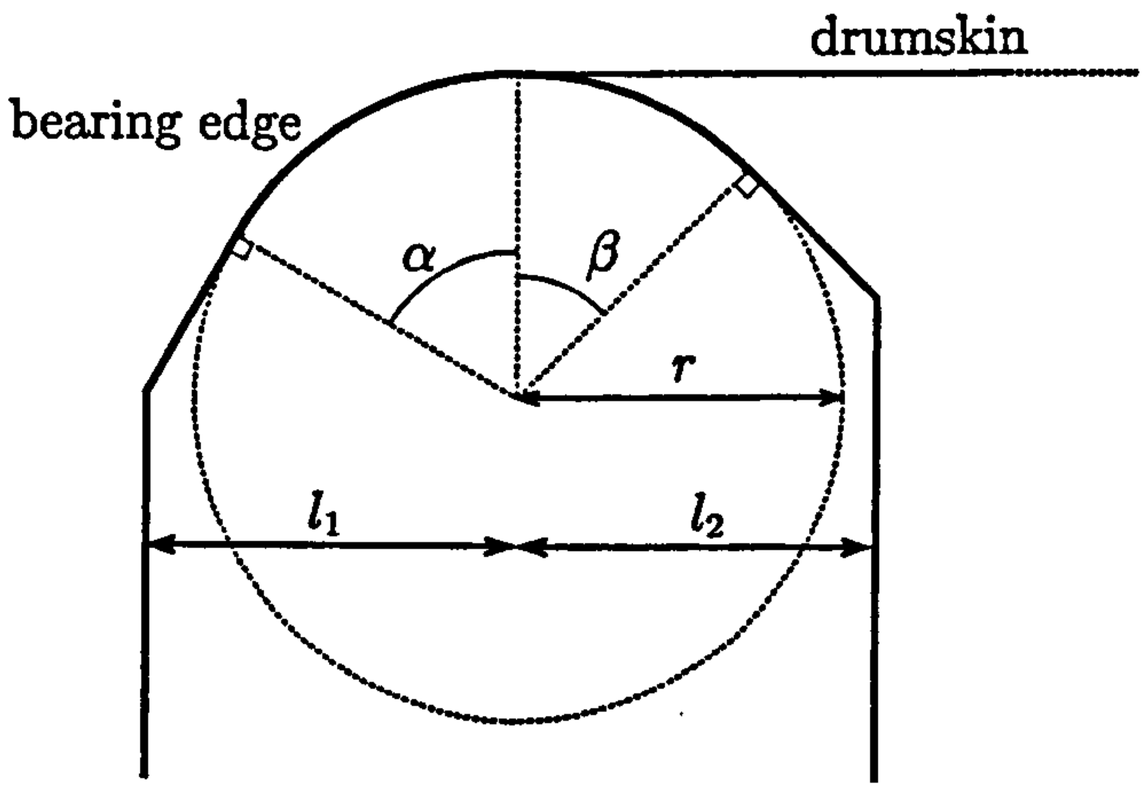

Figure 4.10: *Cross-section of the Bearing Edge.*

It is simpler to consider modelling the 1D case of a string pulled over an instrument bridge. This bridge is the same as the cross-section of the bearing edge. As the string vibrates it can be imagined that its end wraps and unwraps around the edge (compare the position of the point where the string/membrane cross-section meets the bearing edge in figures 4.10 and 4.11); this translates to a shortening and lengthening of the string for vibrations of small amplitude, where the effect of bending can be ignored.

Changing the length of a string model as it vibrates has been used to simulate the effect of tension modulation in the Kantele (Välimäki, Karjalainen, Tolonen, and Erkut, 1999), although the manner in which this was achieved is very different to that proposed here. In this case the lengthening and shortening of the string model was to simulate varying wavespeeds caused by changes in tension rather than a physical change in length. The relationship between vibration and model length was therefore quite different. In addition the fractional changes in length were modelled using FIR filters and these are inherently lossy. For the bearing edge a lossless model was desired such that the losses due to viscous friction could be added and controlled independently using the method described in section 4.1.4.

4.2.1 Mathematical Description of the Bearing Edge

Figure 4.11 indicates the data streams required for the model. A is the moving point where the end of the string, or cross-section of membrane, touches the bearing edge. The stream $s(n)$ holds the wave data that is entering the bearing edge model. Point B is restricted to movement in the vertical plane, and its horizontal coordinate is where wave data is written into stream $s(n)$ and read from stream $s'(n)$. It is the stream $s'(n)$ that needs to be generated in order to implement the bearing edge model. The stream $d(n)$ is the vertical displacement of the string measured from the equilibrium to point B, x samples from the tip of the bearing edge. Note that all distances are measured in spatial

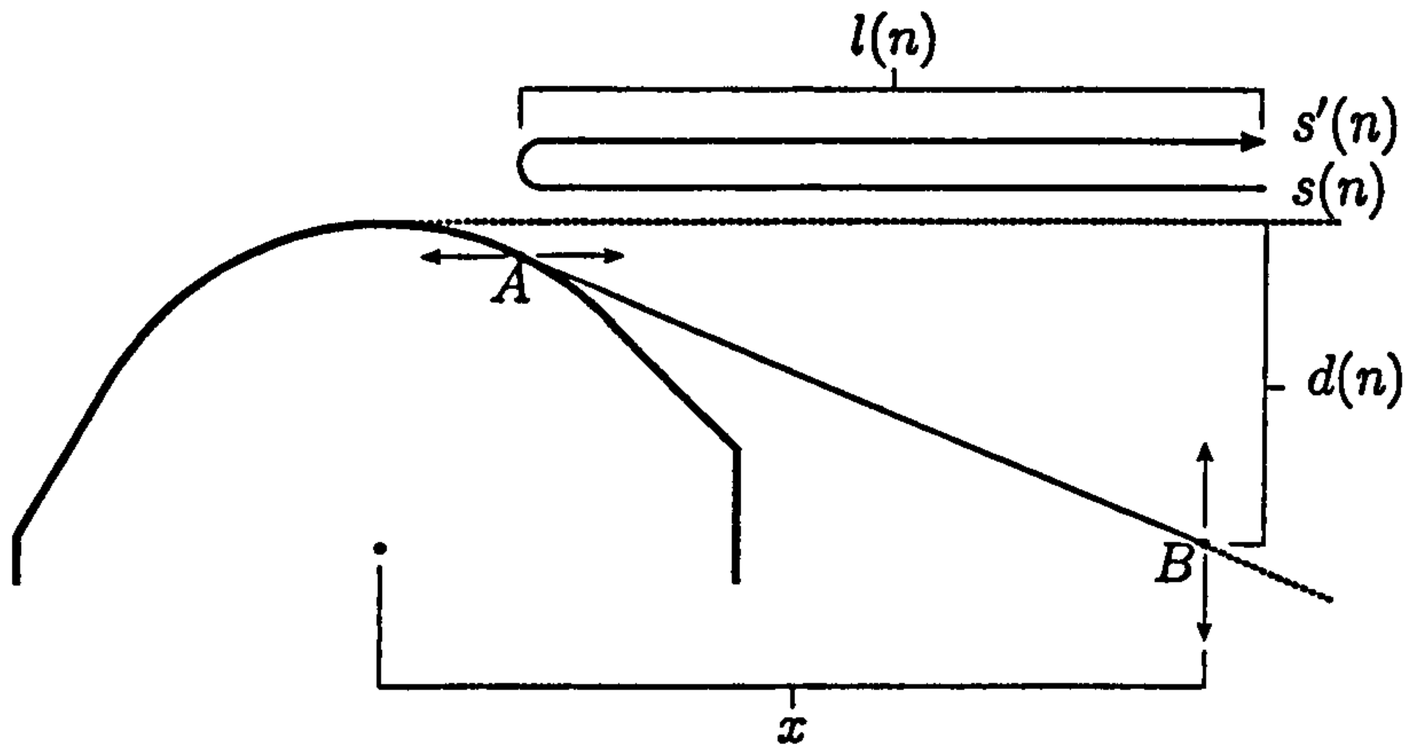

Figure 4.11: *The Data Streams Used to Model the Bearing Edge.*

samples, by multiplying the actual distance by $\frac{f_s}{c}$.

The stream $l(n)$ is the varying horizontal distance from point A to point B. For small amplitude waves, where $\tan^{-1}(\partial d(n)/\partial x) \ll 1$ (as assumed in the formulation of the 1D wave equation, see appendix C.1), $l(n)$ is roughly equal to the actual distance from point A to point B. It is related to $d(n)$ in the following way (see appendix G.1 equations G.22 and G.23)

$$l(n) = \begin{cases} x - l_2, & X > \sin\beta \\ x - rX, & -\sin\alpha \leq X \leq \sin\beta \\ x + l_1, & X < -\sin\alpha \end{cases} \tag{4.2.1}$$

$$X = \frac{rx - (r + d(n))\sqrt{x^2 + d(n)^2 + 2rd(n)}}{x^2 + (r + d(n))^2}; \tag{4.2.2}$$

where r is the radius of the bearing edge (samples),

x is the horizontal distance from the tip of the bearing edge to point B (samples) as shown in figure 4.11,

$l_1 + l_2$ = the thickness of the instrument bridge (samples),

α and β are the absolute angles of the straight lines cut from the bearing edge, as shown in figure 4.10.

The substituted variable, X, is in fact $sin\,\theta$ where θ is the angle between the highest point of the bearing edge and the point where it is touched by the

string/membrane cross-section (see figure G.1 in appendix G.1 for a graphical description). The limits imposed by α and β should be outside the normal mode of operation. These two angles mark the range over which the bearing edge is curved, beyond them the edge is straight until it stops abruptly at either side of the instrument bridge. If the limits are exceeded then it can be imagined that the bearing edge is bypassed and the length of the string/membrane cross-section is determined by the thickness of the bridge alone. This mode of operation is included in equation 4.2.1.

The data at point A, $s(n-l(n))$, must leave point B as $s'(n)$, $l(n)$ samples in the future. One way of performing this would be to take the data $s(n-l(n))$ and insert it into some buffer such that it will arrive $l(n)$ samples later. This would involve insertion at some fractional point along the buffer such that it behaves like a non-integer length delay line running backwards. This process has been called 'deinterpolation' (Välimäki, 1995). Unfortunately, directly manipulating the data stream in this way would corrupt it and cause energy losses as deinterpolation is done with FIR filters. Instead, it would be better to read the appropriate data from stream $s(n)$ using IIR allpass interpolation (see section 3.3 on fractional length delay lines) so that there are no losses. In this way $s'(n)$ is modelled with a rimguide of time-varying length whose input is $s(n)$. The minimum length for this rimguide is that which brings the mesh boundary out to the radius $l_{string} - l_2$, where l_{string} is the length of the string measured to the tip of the bearing edge. Altering the rimguide length in a time-varying way requires the allpass filter to be 'warmed-up' between each sample. The reason for this is to remove transient errors that are caused by changes in the filter coefficient. The method for warming-up a filter is described in section 3.4. The calculation of the read pointer's position in stream $s(n)$, to yield $s'(n)$, is non-trivial and is addressed in the following section on implementational issues.

4.2.2 Implementational Issues of the Bearing Edge Model

Reading the Stream $s(n)$

To aid the indexing of stream $s(n)$ a buffer, $b(n)$, is defined

$$b(n + l(n + j) + j) = l(n + j). \tag{4.2.3}$$

This buffer has the interesting property that the data read at $b(n)$ indicates, in samples, how long ago it was written into the buffer. In this case $l(n)$ is received at point B, after having been sent from point A, $l(n)$ samples in the past. Implementing the buffer $b(n)$ involves deinterpolation into its data stream and this will be dealt with shortly. Inaccuracies in the deinterpolation technique will result in slightly wrong values of $l(n)$ being read from $b(n)$, but will not cause losses in the bearing edge model.
The data leaving point B is found with

$$s'(n) = s(n - 2b(n)). \tag{4.2.4}$$

This is because data that left point A, $b(n)$ samples in the past, must have entered the stream $s(n)$, $2b(n)$ samples in the past. Reading the stream $s(n)$ from this non-integer index is done with the aforementioned IIR allpass interpolation.

Implementing the Buffer $b(n + l(n + j) + j) = l(n + j)$

Figure 4.12 shows example data held in buffer $b(n)$. The value at $b(n+i)$ needs to be calculated and written to the buffer for all integer values of i between $l(n-1)-1$ and $l(n)$. The data is to be deinterpolated into the buffer using first order linear interpolation (see appendix G.2 equation G.28)

$$b(n + i) = \frac{i(l(n) - l(n-1)) + l(n)}{l(n) - l(n-1) + 1}, \tag{4.2.5}$$

$$i \in N,\ l(n-1) - 1 < i \leq l(n). \tag{4.2.6}$$

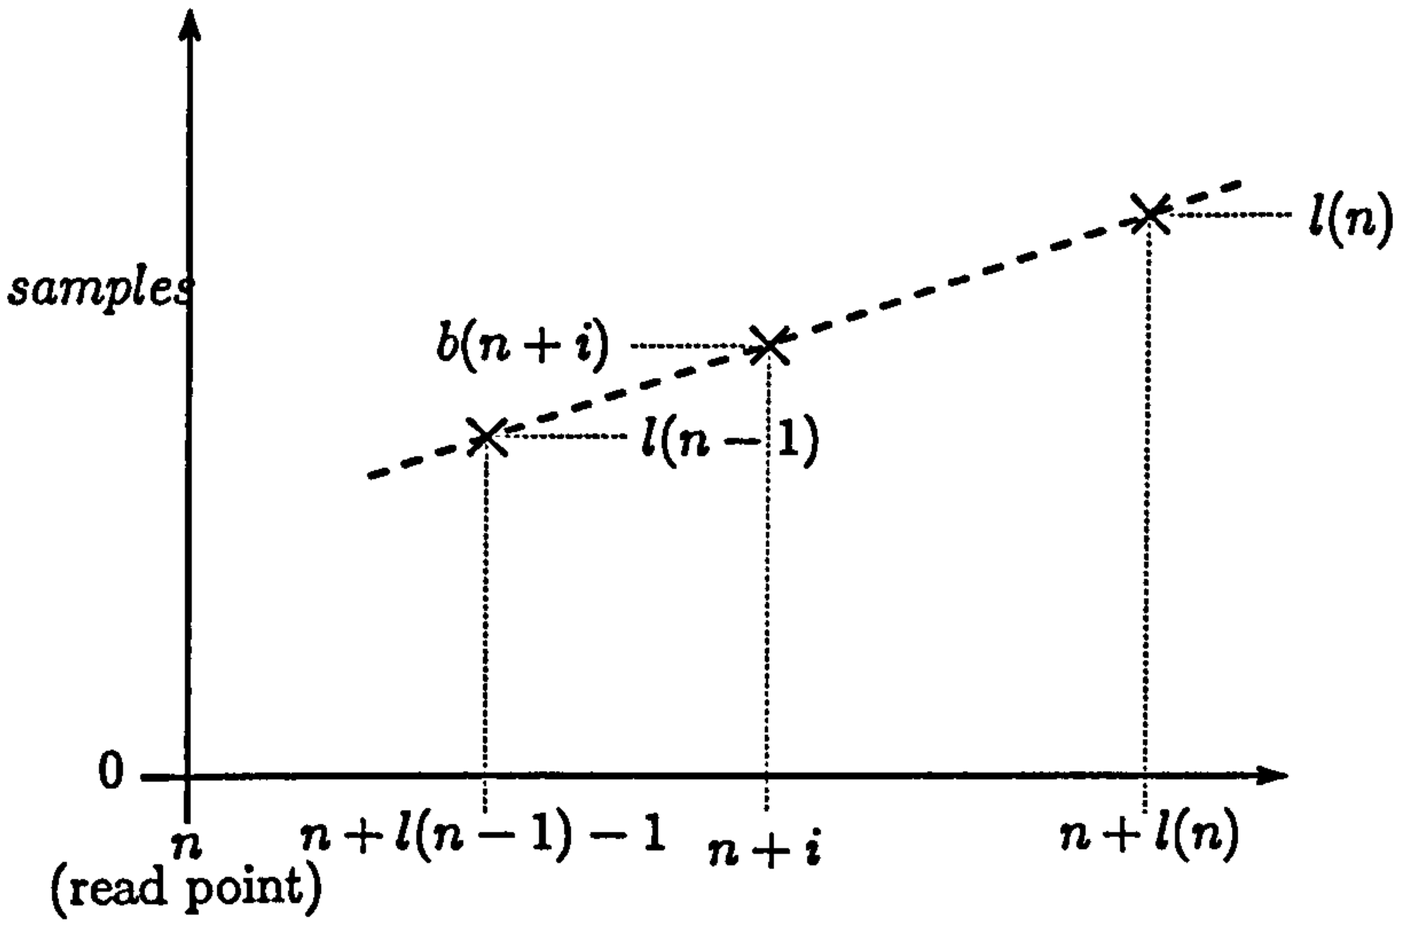

Figure 4.12: *Example Data held in Buffer $b(n)$.*

Writing to the buffer $b(n)$ in this way will approximate equation 4.2.3.

Note that the bearing edge model is only defined for values

$$|l(n) - l(n-1)| \leq 1. \tag{4.2.7}$$

If this relationship were violated then the string would have changed length faster than the speed of wave travel, resulting in unpredictable behaviour.

4.2.3 The Effect of the Bearing Edge

To test the effect of the bearing edge a string model was constructed with the bearing edge model implemented at both ends. The string model's parameters that were used are listed in table 4.4. The model did not incorporate viscous friction.

The first mode of vibration of the string was excited for different values of string displacement and bearing edge radius r. This was done so that the effect of the bearing edge, if any, could be seen in terms of the energy transferred from the fundamental mode of resonance. The string was put into its first mode of vibration by exciting it at its middle with a sinusoidal signal for half a cycle. Its output (the velocity read at the excitation point) was recorded for one

l_{string}	$70\,cm$
linear density	$0.006\,\frac{kg}{m}$
tension	$240\,N$
sample rate	$44100\,Hz$
α	$60°$
β	$45°$
l_1	$r \times 3$
l_2	$r \times 7$

Table 4.4: *Parameters used for the String and Bearing Edge Model.*

second.

The excitation signal was defined as (see appendix G.3)

$$v = \begin{cases} A\sin(\omega t), & \omega t \leq \pi \\ 0, & \omega t > \pi \end{cases} \tag{4.2.8}$$

$$A = \frac{D\omega}{2}, \tag{4.2.9}$$

where v is the velocity (ms^{-1}),

D is the desired initial displacement of the string (m),

ω is the string's fundamental frequency (rad/s).

Figure 4.13 shows a range of frequency plots for the different bearing edge radii and initial string displacements. All the values are normalised such that the greatest magnitude is always $0\,dB$. The plots to the right correspond to models with the greatest bearing edge radii and the lower plots correspond to models with the largest initial displacement.

4.2.4 Conclusions from the Bearing Edge Model

From figure 4.13 it can be seen that the effect of the bearing edge is to excite higher modes of resonance. The greater the bearing edge radius or string displacement the more prominent these frequencies become. However, for a radius of $0.1\,cm$ (as used by Premier Percussion Ltd.) there is little effect on

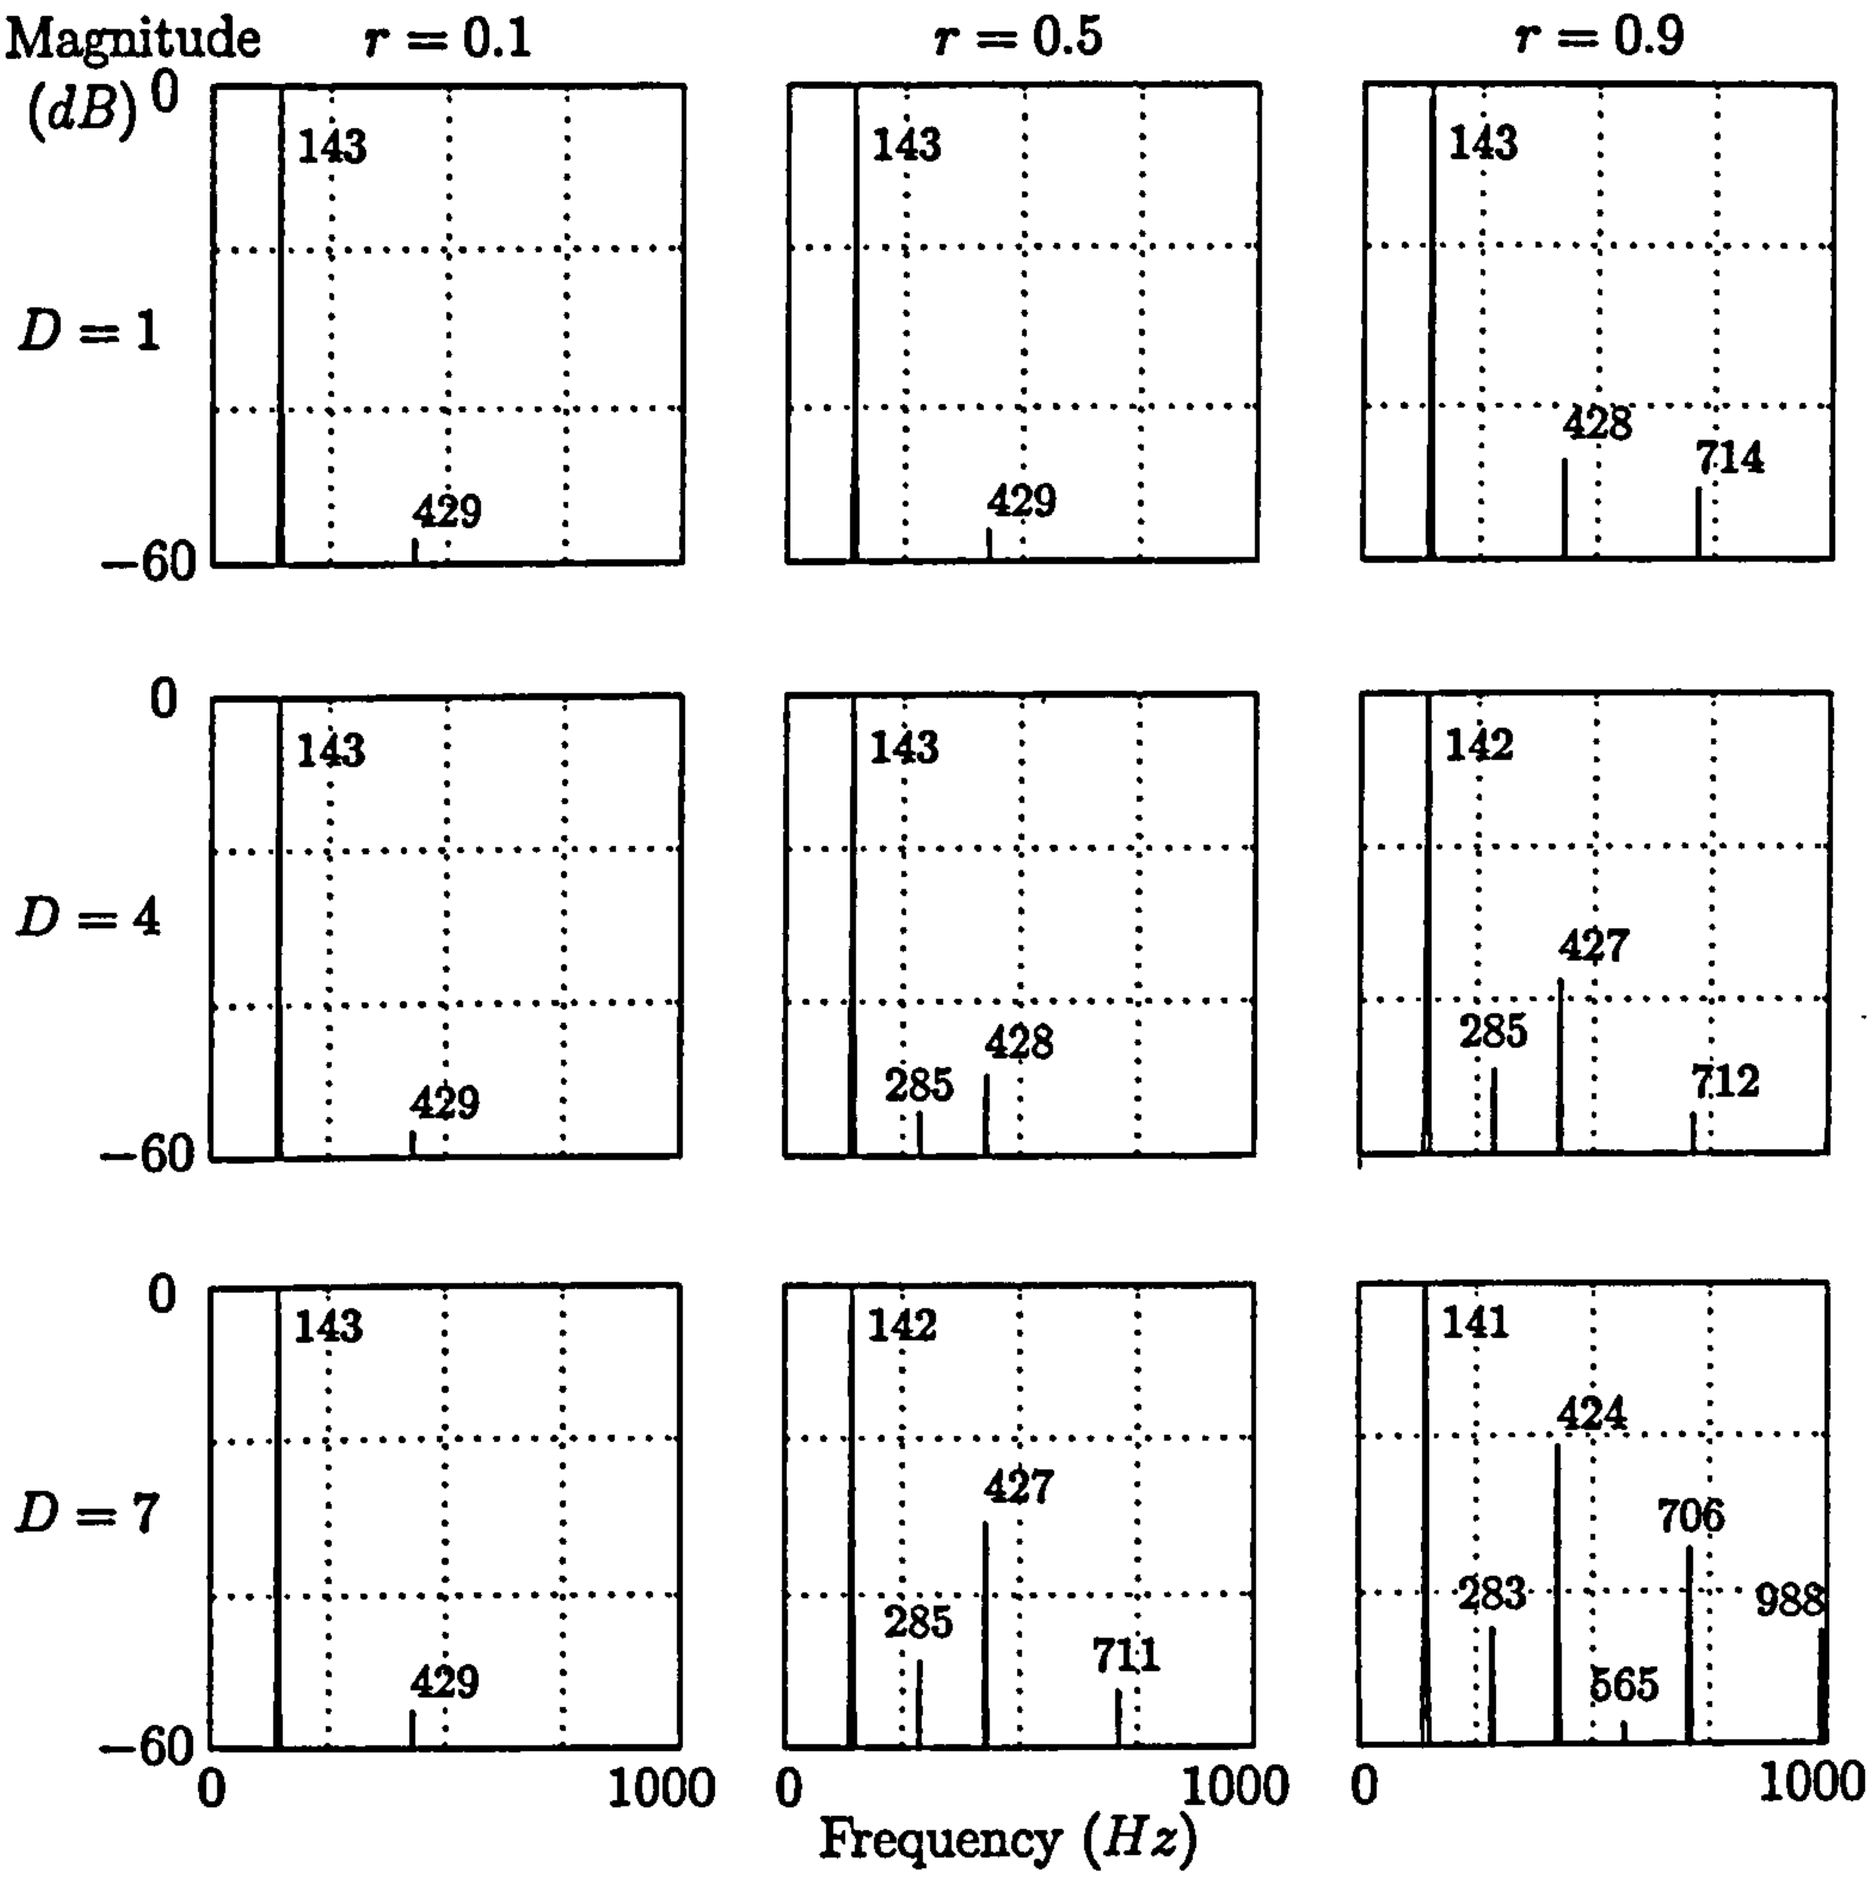

Figure 4.13: *Resonances from String Models with different Initial Displacements, D (cm) and Bearing Edge Radii, r (cm). The numbers close to the resonant modes denote the frequencies (Hz).*

the sound under normal playing conditions; the extra frequencies caused by the bearing edge are close to 60 *dB* below the fundamental frequency. In the examples the higher modes of resonance caused by the bearing edge are predominantly the odd numbered harmonics of the string. This is a consequence of the initial excitation being at the string's middle point. It can be seen that under the more extreme conditions the even numbered harmonics start to appear. These are probably due to small inaccuracies in the interpolation and filter warm-up used in the bearing edge model. Notice also that the frequencies

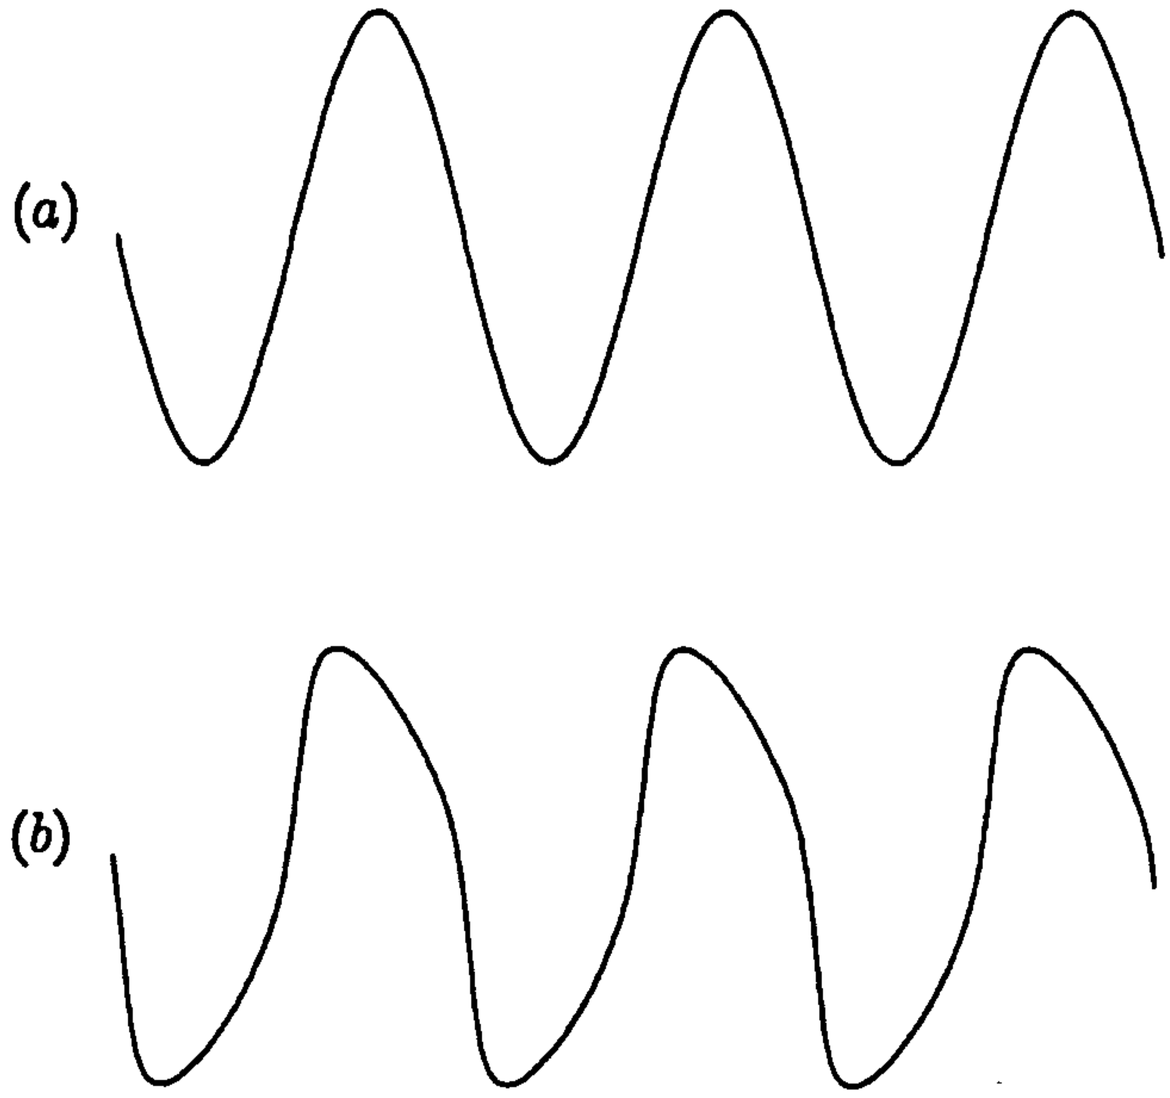

Figure 4.14: *Wave Output of a String Model with Bearing Edge Parameters:* $D = 7\,cm$, $r = 0.9\,cm$. *(a) at* $\approx 0.1\,seconds$, *(b) at* $\approx 0.9\,seconds$.

become slightly lower for these more extreme conditions. This is because the string is spending more time with its ends unwrapped from the bearing edge and so on average it is slightly greater in length.

These results can be heard on the CD-Rom (see section A.1.1 of appendix A). For the smallest initial displacement there is no audible difference between the different bearing edge radii. The sounds are almost a pure 'hum'. As the displacement increases the higher frequencies that enter the sounds become more apparent. In general these sounds have a slight 'buzz'. For the more extreme displacements and bearing edge radii it is possible to hear the sound change over time; the spreading of sound across the resonant frequencies is very noticeable. There is also a very slight change in pitch towards the end of the sample of the most extreme conditions, as indicated in the resonance plots of figure 4.13.

Figure 4.14*a* shows the wave pattern read at about 0.1 seconds into the sample of the most extreme conditions, where $D = 7\,cm$ and $r = 0.9\,cm$. At this

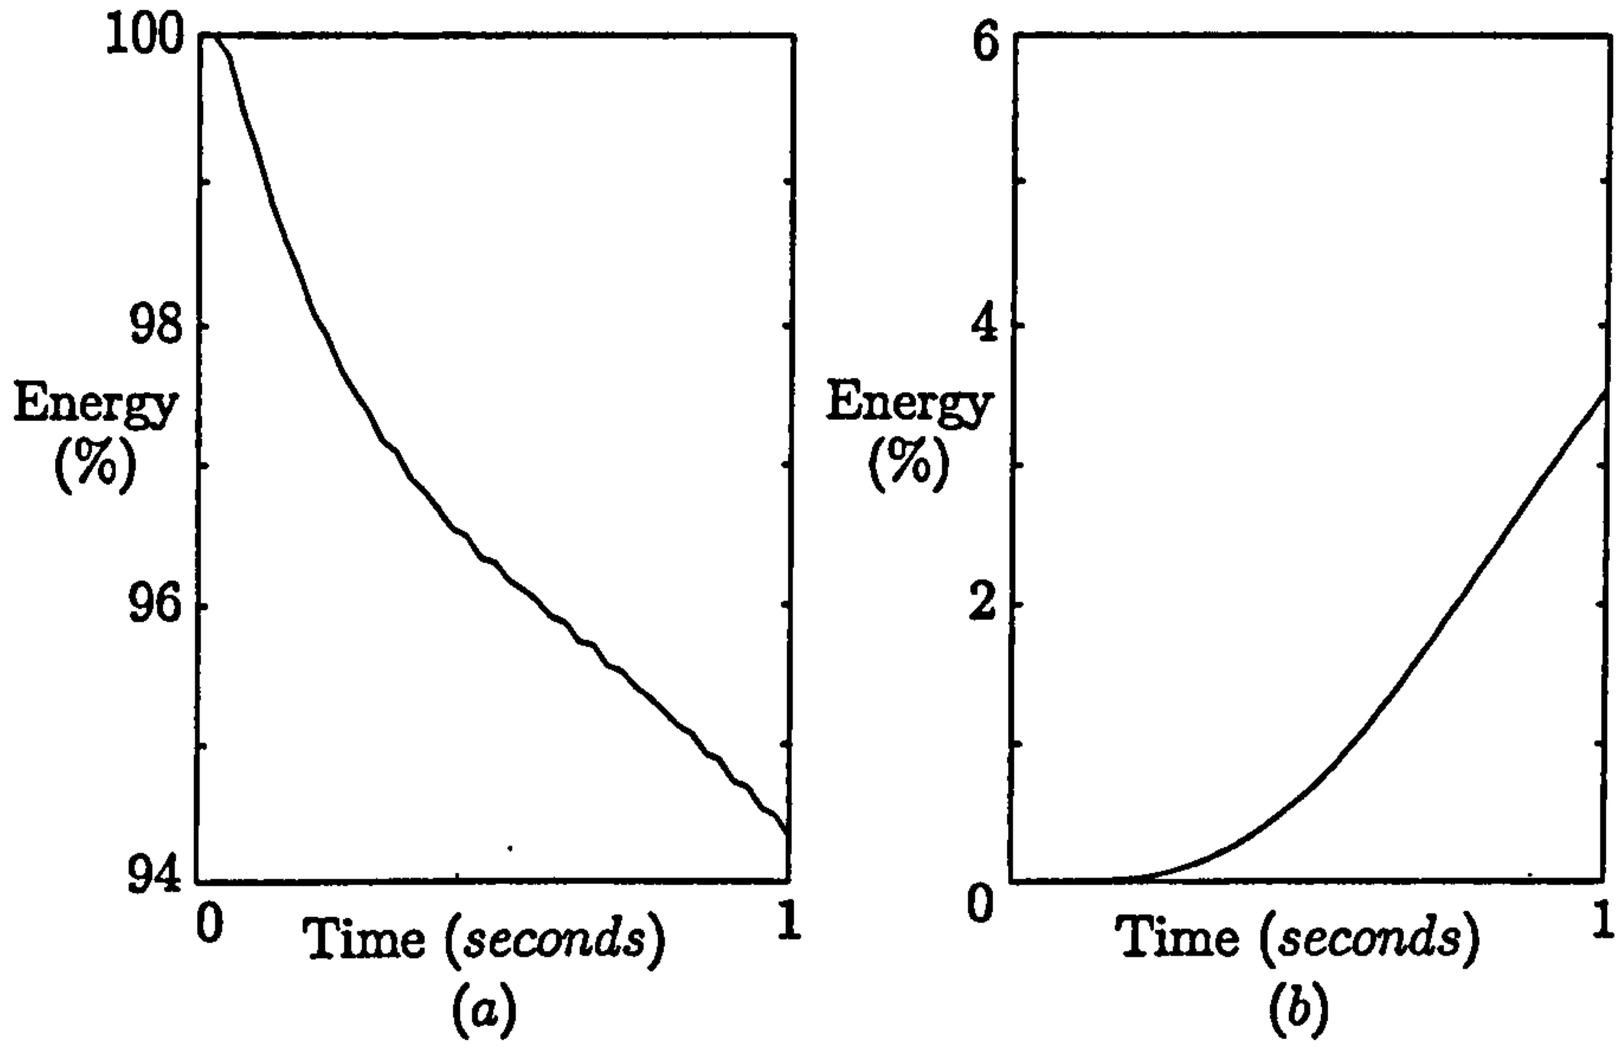

Figure 4.15: *Change in Resonant Mode Energy of a String Model with Bearing Edge Parameters:* $D = 7\,cm$, $r = 0.9\,cm$. (*a*) $f_1\,(141\,Hz)$, (*b*) $f_3\,(424\,Hz)$.

stage the wave is still almost a pure sinusoid. At about 0.9 seconds, see figure 4.14*b*, the wave has started to take the form of a 'saw-tooth,' which is due to the presence of the odd numbered harmonics. It would appear that energy is being transferred from the fundamental into higher modes of resonance. The transfer of energy from the fundamental to the third harmonic is shown in figures 4.15*a* and 4.15*b*. These plots were made in a similar fashion to that of figure 4.8 in section 4.1.4, but in this case rather than plotting the amplitude in decibels and normalising it to $0\,dB$ the amplitude values were squared, to find their energy, and normalised against the initial energy of the fundamental frequency. The values are plotted as percentages of this initial fundamental energy. It can be seen that the fundamental lost more than 5% of its energy over the second of time, with over 3% entering the third mode of resonance. It would appear to be desirable for the bearing edge to have no effect, for stringed instruments such as the guitar have a bridge whose tip is as thin as possible. The reason for this, as suggested by the instrument players themselves, is to increase sustain. The results from the bearing edge model support this hypothesis: a greater bearing edge radius would spread more of the en-

ergy into the higher frequencies, and as higher frequencies suffer greater energy losses through visco-elastic damping (see section 3.5) the instrument would lose energy more quickly resulting in a decrease in sustain.

As an aside it was noticed that the bearing edge model has another application in audio signal processing. If the difference $l(n) - l(n-1)$ is kept at a constant value then the stream $s'(n)$ is a resampled version of $s(n)$. This lossless resampling technique could be done in real-time.

4.3 Summary

In this chapter a model for the drumskin was created. The basis of this model was a triangular 2D waveguide mesh structure, chosen for its near direction independent dispersion error.

A method for modelling a circular boundary was proposed that used new structures given the name of 'rimguides'. These structures enable waves to propagate from the edge-most nodes of a mesh to the desired boundary and back. Increasing orders of complexity were investigated that involved tuning the rimguides to various resonant modes of the membrane. However, it was found that good accuracy was achieved using simple rimguides that correctly modelled the propagation of waves of very low frequency. It was also demonstrated how lowpass filters could be included in the rimguide structure to simulate the effect of viscous friction, which could be controlled in terms of a loss rate at a particular frequency.

The impedance of a 2D waveguide mesh is required to connect the mesh to other media. Before commencing this work no formula had been documented to calculate this impedance and so a way in which this could be done was proposed. The formulation involved considering the mesh to be made up of tessellated elements, each consisting of a collection of strings or 1D waveguides. As the impedance of an individual string is easy to calculate it was then possible to find their combined effect. This resulted in a formula to calculate

the impedance of a single element within the waveguide mesh.

The main vibrating part of an instrument, whether it is a membrane or string, has to be supported by an edge. This edge is called a bearing edge for membranes and a bridge for strings. It is usually carved to form a small curve where it is in contact with the vibrating medium. A technique to incorporate its effect was found and investigated. The results of this investigation showed that the edge causes energy to leak into the higher resonant modes when the curvature radius or excitation is large. Because this results in a reduction in sustain it was argued that it is desirable to have the smallest edge radius possible.

Chapter 5

Modelling the Interior of the Drum

5.1 Modelling the Air Cavity

The air enclosed by the drumshell and skins plays an important role in shaping the sound. Being three dimensional it has a complicated arrangement of resonant modes that interfere with the vibration of the drumskin. The air also has a frequency dependent loading effect on the skin (Fontana and Rocchesso, 1998).

Three dimensional wave propagation is easily modelled using a 3D digital waveguide mesh structure but, like the 2D mesh, the arrangement of waveguides introduces dispersion error.

5.1.1 Tetrahedral and 6-Port Rectilinear 3D Digital Waveguide Meshes

A previous examination of dispersion errors was undertaken for tetrahedral and 6-port rectilinear 3D waveguide meshes (Duyne and Smith(III), 1996). The tetrahedral mesh was created for its multiply-free architecture, to be used in real-time applications.

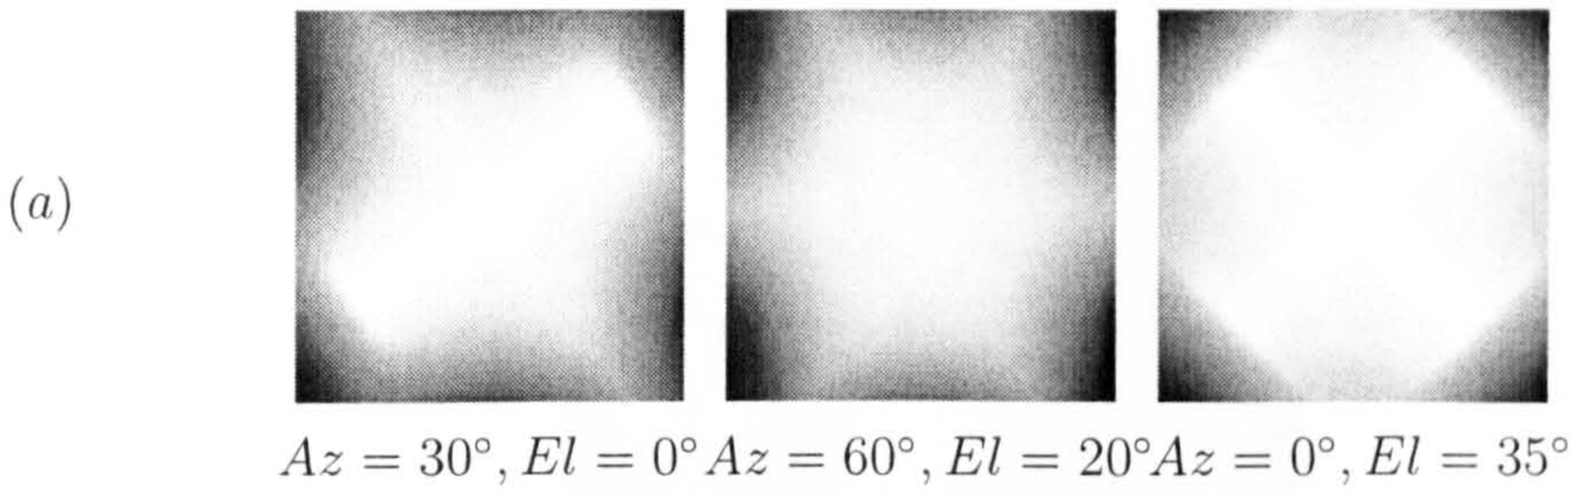

Figure 5.1: *Dispersion Error Plots for 3D Digital Waveguide Meshes.* *(a) Tetrahedral Mesh, (b) 6-Port Recilinear Mesh.*

Figures 5.1a and 5.1b show the dispersion error plots of the tetrahedral and 6-port rectilinear 3D meshes. As the dispersion errors are three dimensional it is only possible to show dispersion slices. These may be taken from any azimuth and elevation and are centred about the point $\omega_x = \omega_y = \omega_z = 0$. The limits of the slices are at the Nyquist frequency, $|\underline{\omega}| = \pi$. The azimuth and elevation angles used in figure 5.1 were chosen to show the worst and best dispersions exhibited by each mesh. It can be seen that both these meshes have widely varying dispersion errors that are dependent on the direction of wave travel. This dependency would cause the modelled shape to exhibit the wrong modes of resonance, invalidating these meshes for use in accurate modelling.

5.1.2 The 3D Dodecahedral Digital Waveguide Mesh

Ideally a mesh structure needs to be found that exhibits near direction independent dispersion error, as was done previously for a 2D membrane. But due to the extra degree of freedom in three dimensions there is a large range

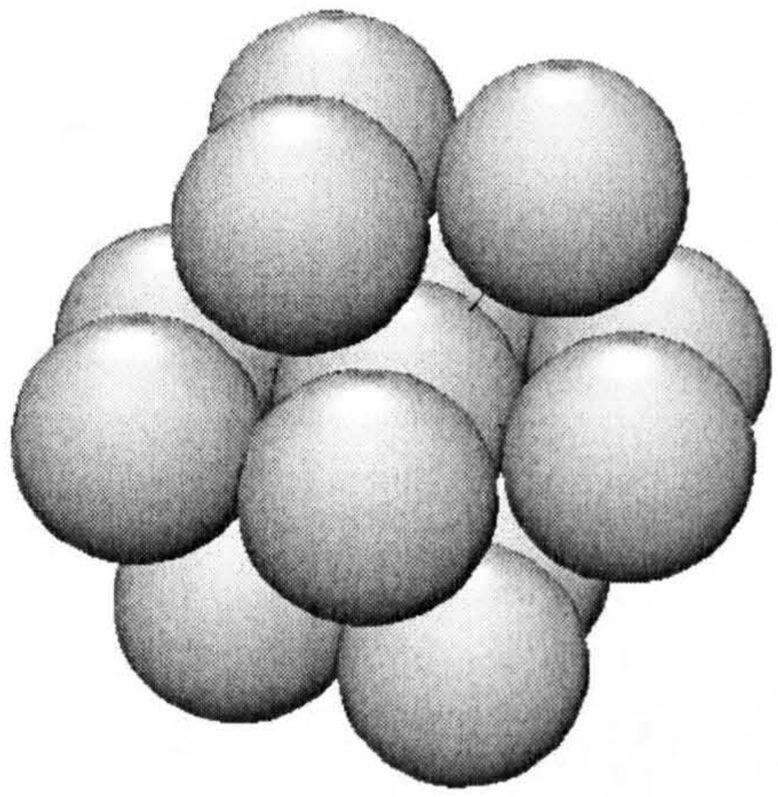

Figure 5.2: *Hexagonal-Close Packing.*

of possible mesh structures. Just as 2D wave propagation was best modelled using the densest arrangement of nodes it was reasoned that the densest 3D arrangement should be used here.

The densest arrangement of nodes within 3D space is the solution to Kepler's cannon ball stacking problem (Plus Online Magazine, 1997). This was recently proved (Hales, 1998) to be face-centred cubic packing or hexagonal-close packing (they have the same density). The structure of hexagonal-close packing is shown in figure 5.2. If each ball is considered to be a node in a mesh, then it defines a 3D digital waveguide mesh of 12 waveguides intersecting each junction node. The coordinates of the nodes for a hexagonal-close packing structure centred about the point $(0, 0, 0)$ are listed in table 5.1.

The fundamental building block of these meshes, with each face intersected perpendicularly by a waveguide, is the dodecahedron. In particular, the regular building block for face-centred cubic packing is the rhombic dodecahedron. Diagrams of this can be found on the World Wide Web (Bulatov, 1999). An example of close packing with tessellated rhombic dodecahedrons can also be found on the World Wide Web (Seven Stones, 2000). These new waveguide mesh structures have therefore both been called 'The 3D Dodecahedral Mesh'. The same 3D mesh structures were discovered in parallel with this investiga-

node no.	coordinate
1	$(-1, 0, 0)$
2	$(-\frac{1}{2}, \frac{\sqrt{3}}{2}, 0)$
3	$(\frac{1}{2}, \frac{\sqrt{3}}{2}, 0)$
4	$(1, 0, 0)$
5	$(\frac{1}{2}, -\frac{\sqrt{3}}{2}, 0)$
6	$(-\frac{1}{2}, -\frac{\sqrt{3}}{2}, 0)$
7	$(0, \frac{1}{\sqrt{3}}, \frac{\sqrt{2}}{\sqrt{3}})$
8	$(-\frac{1}{2}, -\frac{1}{2\sqrt{3}}, \frac{\sqrt{2}}{\sqrt{3}})$
9	$(\frac{1}{2}, -\frac{1}{2\sqrt{3}}, \frac{\sqrt{2}}{\sqrt{3}})$
10	$(0, \frac{1}{\sqrt{3}}, -\frac{\sqrt{2}}{\sqrt{3}})$
11	$(-\frac{1}{2}, \frac{1}{2\sqrt{3}}, -\frac{\sqrt{2}}{\sqrt{3}})$
12	$(\frac{1}{2}, \frac{1}{2\sqrt{3}}, -\frac{\sqrt{2}}{\sqrt{3}})$

Table 5.1: *Coordinates of Nodes within the 3D Dodecahedral Mesh (using hexagonal-close packing centred about the point* $(0, 0, 0)$*).*

tion (Fontana, Rocchesso, and Apollonio, 2000). In this case their construction was not explained in terms of face-centred cubic packing and hexagonal-close packing, and the dodecahedron was not mentioned as the fundamental tessellating building block. The names given to the meshes were the '3D Triangular Waveguide Mesh (3DTWM)' and the 'Interlaced Square Waveguide Mesh.' It was concluded that both geometries yield efficient structures to model 3D resonators and they have very uniform dispersion characteristics.

The Dispersion Error of a 3D Dodecahedral Digital Waveguide Mesh

The frequency dependent wavespeed (dispersion error) of the dodecahedral mesh is found in a similar way to that of the triangular mesh presented earlier (see appendix E.1). This 3D mesh structure results in

$$c(\underline{\omega}) = \frac{\angle G}{|\underline{\omega}|}, \tag{5.1.1}$$

$$\angle G = \arctan \frac{\sqrt{4 - b^2}}{-b}, \tag{5.1.2}$$

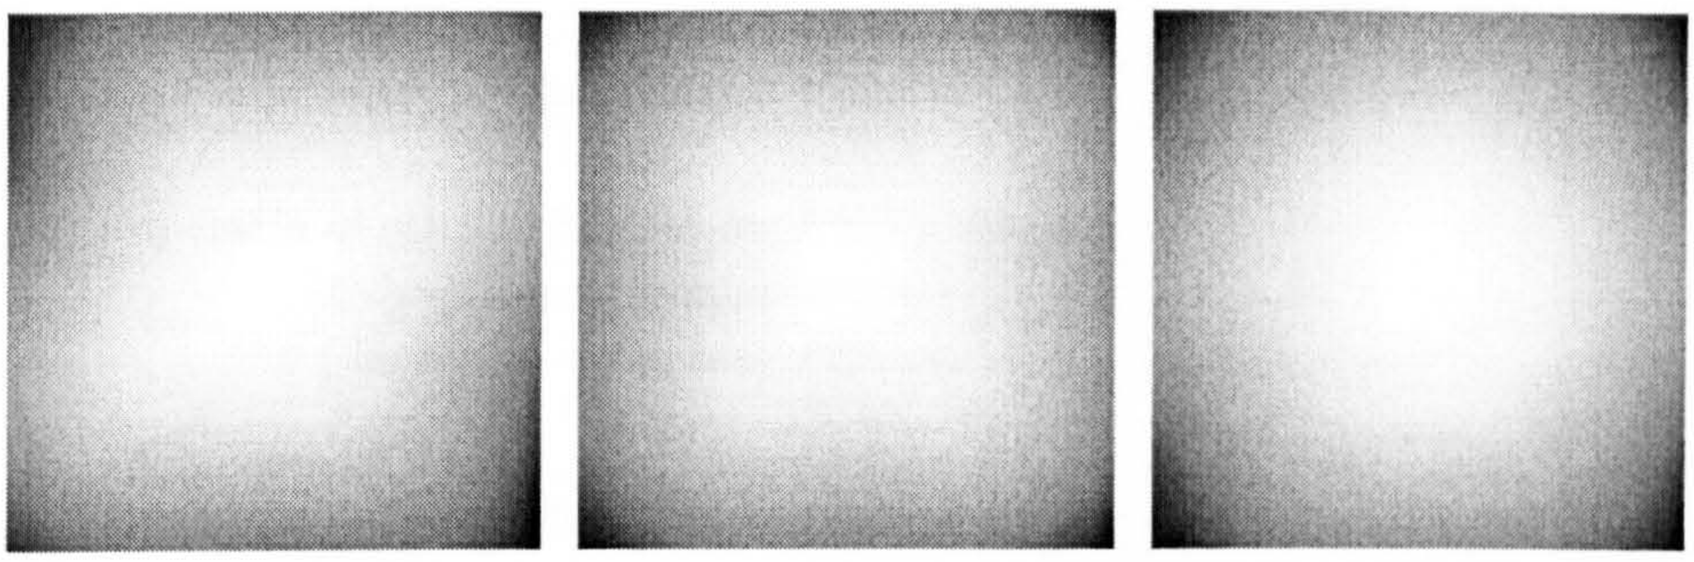

$Az = 30°, El = 0°$ $Az = 0°, 30°$ $Az = 0°, El = 90°$

Figure 5.3: *Dispersion Error Plots for the 3D Dodecahedral Digital Waveguide Mesh.*

$$b \triangleq -\frac{1}{6}\sum_{i=1}^{12} e^{j\underline{P}_i^T \cdot \underline{\omega}}, \tag{5.1.3}$$

where $\underline{P}_i$ are the vector coordinates of the node points,

$\underline{\omega}$ is the three-dimensional spatial frequency vector with variables $(\omega_x, \omega_y, \omega_z)^T$,

$c(\underline{\omega})$ is the frequency and direction dependent speed of plane wave travel (in space samples per time sample).

The nominal wavespeed within the mesh, where $\omega_x = \omega_y = \omega_z = 0$ is (see appendix E.2)

$$c_0 = \frac{1}{\sqrt{3}} \tag{5.1.4}$$

Figure 5.3 shows the dispersion error slices of the dodecahedral mesh from a range of viewing angles, as was done previously in figure 5.1. Once again the viewing angles were chosen to show the worst and best dispersion errors. Notice that near direction independent dispersion error is achieved, with the worst case having a slight diamond shape; this occurs when the viewing angle faces directly along a waveguide.

5.1.3 Modelling Pressure Waves in Air

It is now possible to model, with reasonable accuracy, pressure waves travelling through three dimensional space. If the three dimensional space is a gas then

parameter	value
γ (ratio of specific heat at constant pressure)	1.4
M (average molecular weight)	29.0
ρ (density at room temperature)	1.205 (kgm^{-3})

Table 5.2: *Properties of Air.*

these pressure waves travel with a speed of (Hall, 1987e)

$$c_{gas} = \sqrt{\frac{\gamma R T_0}{M}}, \tag{5.1.5}$$

where γ is the ratio of specific heat at constant pressure to that at constant volume,

R is the universal gas constant $(8314 \frac{J}{kg\,K})$,

T_0 is the temperature (K),

M is the average molecular weight (kg).

In the case of air, whose properties are given in table 5.2, the wavespeed at room temperature $(T_0 = 293\,K)$ is $c_{air} \approx 343 \frac{m}{s}$. Therefore, when modelling air with a dodecahedral mesh each waveguide is equivalent to a length of $\frac{c_{air}}{c_0 f_s} = \frac{343\sqrt{3}}{f_s}\,m$, where f_s is the desired sampling rate.

The values evaluated at node junctions within the mesh represent pressure differences. These differences are relative to the equilibrium pressure of the gas. Hence, a positive value relates to an increase in pressure and a negative value corresponds to a drop in pressure.

Shaping the Mesh Boundary

The air mesh model can be shaped to fit the interior space enclosed by the drumshell. The drumshell boundary is modelled using rimguides, just like the clamped boundary of the membrane described by method 2 in section 4.1.3. The difference here is that there is no sign inversion at the boundary as longitudinal waves are being modelled. Once again the rimguides are attached from the edge-most nodes of the mesh to the closest points on the boundary.

Losses may also be implemented in a similar fashion to that described in section 4.1.4.

Results from a 3D Dodecahedral Mesh

To test the accuracy of the 3D dodecahedral mesh (with hexagonal-close packing) a model of an IEC standard room (Wright, 1995) was constructed. The room's dimensions were $4\,m$ by $6\,m$ by $2.5\,m$, which present a particularly good challenge because modes 9 and 10 occur at exactly the same frequency, and mode 11 lies in close proximity. The sample rate used for the waveguide model was $2000\,Hz$ and a 50 second sample was obtained in order to provide a good resolution of the resonant modes. The results and their percentage errors are shown alongside the theoretical frequencies in table 5.3 below.
It can be seen that despite the low sample rate the higher modes of resonance are quite well reproduced, although the model did not produce the same frequency for modes 9 and 10. The extra complexity of three dimensions seems to have had more affect on the accuracy of the lower modes.

Mode Number	Theoretical (Hz)	Wavguide Model (Hz)	Error (%)
1	28.70	27.42	−4.46
2	43.06	44.30	+2.88
3	51.75	50.56	−2.30
4	57.41	54.44	−5.17
5	68.89	68.58	−0.45
6	71.76	71.64	−0.17
7	74.63	77.68	+4.09
8	81.24	80.58	−0.81
9	86.11	84.58	−1.78
10	86.11	87.54	+1.66
11	86.16	87.62	+1.69
12	89.67	89.58	−0.10

Table 5.3: *Theoretical and Simulated Frequencies of an IEC Standard Room* $(4\,m \times 6\,m \times 2.5\,m)$.
Theoretical modes from Wright (1995).

5.2 Modelling the Interior Lacquer

Premier Percussion Ltd. sand down the interior of the drum shell and then apply a thin layer of lacquer. From observation it has been noticed that this lacquer lifts wood fibres up, making the interior rough to the touch. No doubt further layers of lacquer would begin to make a smooth surface, filling in the gaps between the fibres. Although it is difficult to calculate to what extent the surface roughness is being altered, this is undoubtedly controlling the well-known phenomenon of diffusion.

5.2.1 Diffusion

The boundary method described in section 4.1.3 modelled specular reflections. Specular means 'mirror-like', such that the angle of reflection equals the angle of incidence. Diffuse reflections, which occur at rough boundaries, scatter energy in almost every direction regardless of the angle of incidence (Hall, 1987f). These two types of reflection are shown in figures 5.4a and 5.4b respectively. In order to model a diffuse surface it is necessary to define a coefficient that describes the scattering of the energy. Two such coefficients that have been considered in the past are the diffusion coefficient, d (formed by AES), and the scattering coefficient, s (formed by ISO). The diffusion coefficient is a measure of the spatial uniformity of the sound globally reflected by an object, whereas the scattering coefficient is defined as the ratio between the sound energy

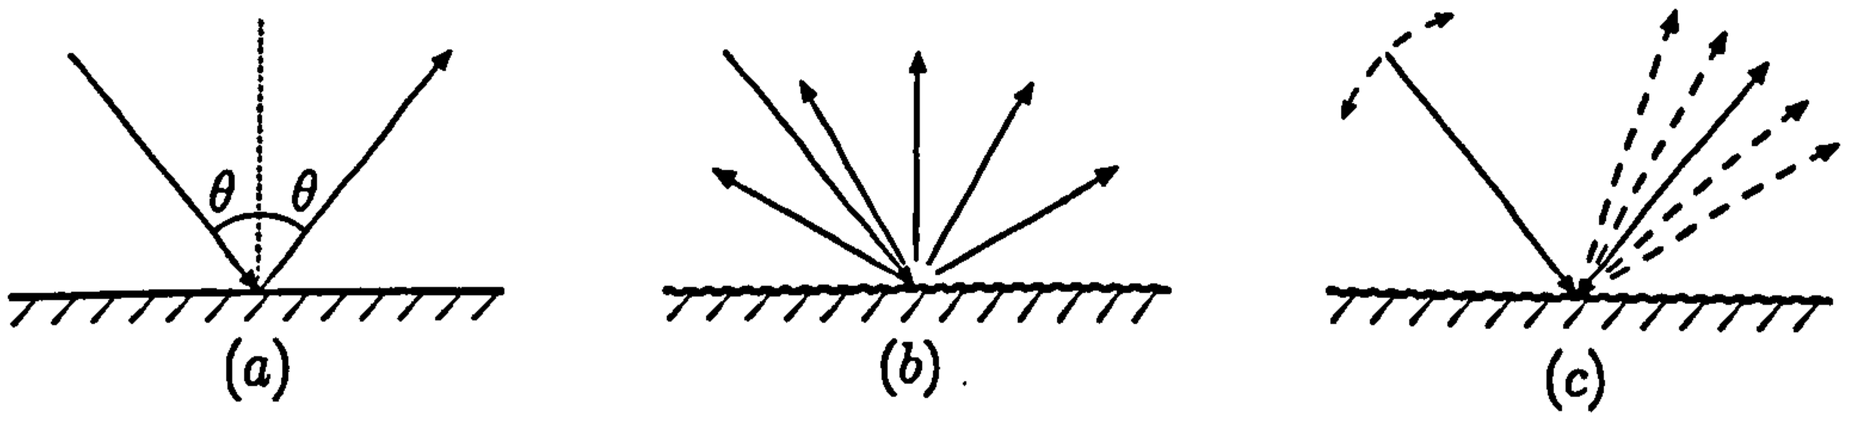

Figure 5.4: *Direction of Wave Fronts reflecting off a Boundary. (a) Specular Reflection, (b) Diffuse Reflection, (c) Varying Incident Wave Angle to model Diffusion.*

reflected diffusely and the total reflected energy. The scattering coefficient is generally a quantity very different from the diffusion coefficient, although work has been done to enable one to be derived from the other (Farina, 2000b). Unfortunately, there is no single standardised method to calculate or measure these coefficients and as a result there are a great variety of diffusion methods for various acoustic models, though as yet there are no methods for use in digital waveguide techniques. A brief discussion and comparision of these models has previously been done (Lam, 1997) and two methods for measuring the scattering coefficient have been studied in detail (Farina, 2000a). It is apparent, however, that diffuse reflections are frequency dependent and that the scattered energy has a particular distribution. For simplicity the frequency dependent aspect will be ignored, but ways in which this could be implemented are discussed in section 8.2.

The model used in this investigation is designed for 2D shapes, once again for simplicity, but a way to extend this to 3D is also considered in section 8.2. A summary of the following work was presented at a conference (Laird, Masri, and Canagarajah, 1999).

The Diffusion Method

Rather than modifying the rimguides at the edge of the digital waveguide mesh, it was found that a diffusion effect could be implemented at the edge-most nodes. This involved pre-altering the angle of incidence of travelling waves in a random fashion over time, such that the reflected waves appeared to have been scattered, as shown in figure 5.4c. Varying the incident angle is not straightforward and care must be taken to ensure that signal strength and power is conserved within the waveguide mesh. In addition, waves at the boundary pose an extra complication due to the irregularity of the mesh structure at this point. Therefore, waves within a mesh are considered first.

5.2.2 Wavefronts within a Mesh

A planar wavefront of amplitude A travelling in the direction θ can be represented by superposing vectors of magnitude a_i, whose directions, $\underline{P_i}$, evenly span a circle

$$A = \left| \sum_{i=0}^{N-1} a_i \underline{P_i} \right|, \tag{5.2.1}$$

$$\theta = \angle \left(\sum_{i=0}^{N-1} a_i \underline{P_i} \right), \tag{5.2.2}$$

where $i = 0, 1, ..., N-1$,

N is the number of vectors used to represent the planar wave.

The vectors and their magnitudes are calculated in the following way

$$\underline{P_i} = \begin{bmatrix} \cos\left(i\frac{2\pi}{N}\right) \\ \sin\left(i\frac{2\pi}{N}\right) \end{bmatrix}, \tag{5.2.3}$$

$$a_i = \frac{2A}{N} \left(\underline{P_i}^T . \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix} \right). \tag{5.2.4}$$

Within a waveguide mesh, A is the amplitude of the signal at a node junction and the vectors, $\underline{P_i}$, represent the directions of the waveguides connected to this node.

Using Circulant Matrices to Rotate Wavefronts

As mentioned above it is important to conserve signal strength and power when altering the direction of wave propagation. These requirements can be met by convolving the incoming signals to a node with a circulant matrix. A circulant matrix has the form

$$M = \begin{bmatrix} x_0 & x_1 & \cdots & x_{N-1} \\ x_{N-1} & x_0 & \cdots & x_{N-2} \\ \cdots & \cdots & \cdots & \cdots \\ x_1 & \cdots & x_{N-1} & x_0 \end{bmatrix}. \tag{5.2.5}$$

These matrices have already been used to model diffusion in a class of artificial reverberators known as feedback delay networks (Rocchesso and Smith(III), 1994). In this investigation, however, they are used in a slightly different way. The stability of any matrix can be ensured by distributing the eigenvalues of its coefficients along the unit circle. For a circulant matrix in particular, calculating the coefficients is easy; the discrete inverse Fourier transform (see appendix B.5) of the eigenvalues computes the first row of the matrix, from which subsequent rows are derived.

Rotation may be achieved by arranging the eigenvalues symmetrically. So, in order to rotate the planar wave through angle φ, the following eigenvalues are used

For $N = 3$,

$$\mathbf{X} = \left[1 \;\; e^{j\varphi} \;\; e^{-j\varphi}\right], \tag{5.2.6}$$

For $N = 4$,

$$\mathbf{X} = \left[1 \;\; e^{j\varphi} \;\; -1 \;\; e^{-j\varphi}\right], \tag{5.2.7}$$

For $N = 5$,

$$\mathbf{X} = \left[1 \;\; e^{j\varphi} \;\; e^{j2\varphi} \;\; e^{-j2\varphi} \;\; e^{-j\varphi}\right], \tag{5.2.8}$$

where N is the number of vectors used to represent the planar wave,

$\mathbf{X}$ is an eigenvector composed of eigenvalues.

Performing the inverse Fourier transform on these eigenvalues gives the coefficients, x, that are used in matrix M of equation 5.2.5. Applying this matrix to the vector magnitudes results in the new convolved vector magnitudes a'_i

$$M. \begin{bmatrix} a_0 \\ a_1 \\ \vdots \\ a_{N-1} \end{bmatrix} = \begin{bmatrix} a'_0 \\ a'_1 \\ \vdots \\ a'_{N-1} \end{bmatrix}. \tag{5.2.9}$$

Replacing these new magnitudes for those in equation 5.2.2 results in a new angle of plane wave propagation, $\theta' = \theta + \varphi$.

5.2.3 Wavefronts at a Mesh Boundary

The circulant matrix will be used to convolve the incoming signals at the edge-most nodes of the waveguide mesh (excluding the signals from the rimguides). Unlike the case within the mesh, the waveguides that connect to an edge-most node are not evenly distributed about a point, as can be seen in figures 4.3 and 4.4 of section 4.1. This causes an error in the rotation, as an even distribution of waveguides, or vectors, was assumed in the rotation method described above.

In the case of the edge-most nodes it is necessary to consider two domains. The first domain is the actual 'real' case of waveguides connected to an edge node. The second domain, to be used in the rotation calculations, is the 'ideal' version of the first which assumes the waveguides are evenly distributed. The difference in vector geometry between these two cases has two effects upon the calculation of wavefront rotation. Firstly, there is a unique but nonlinear mapping between the domains, such that an angle in the ideal case can be translated to an angle in the real case. Secondly, as a result of this nonlinearity, a rotation in the real case is dependent not only upon the amount of rotation in the ideal case, but also the initial incident angle of a wavefront.

Mapping of Wavefront Angles Between Domains

At the edge-most nodes of a waveguide mesh the surrounding waveguides can be equated to the unit length vectors $\underline{P}_{ri}$. Substituting these 'real' vectors for the 'ideal' vectors, $\underline{P}_i$, in equation 5.2.2, results in the direction of an incident wavefront at the boundary

$$\theta_r = \angle \left(\sum_{i=0}^{N-1} a_i \underline{P}_{ri} \right), \tag{5.2.10}$$

where N is the number of waveguides attached to an edge-most node.

The vector magnitudes, a_i, are calculated as before using equation 5.2.4. This maps an ideal wavefront travelling within a mesh to a real wavefront at the

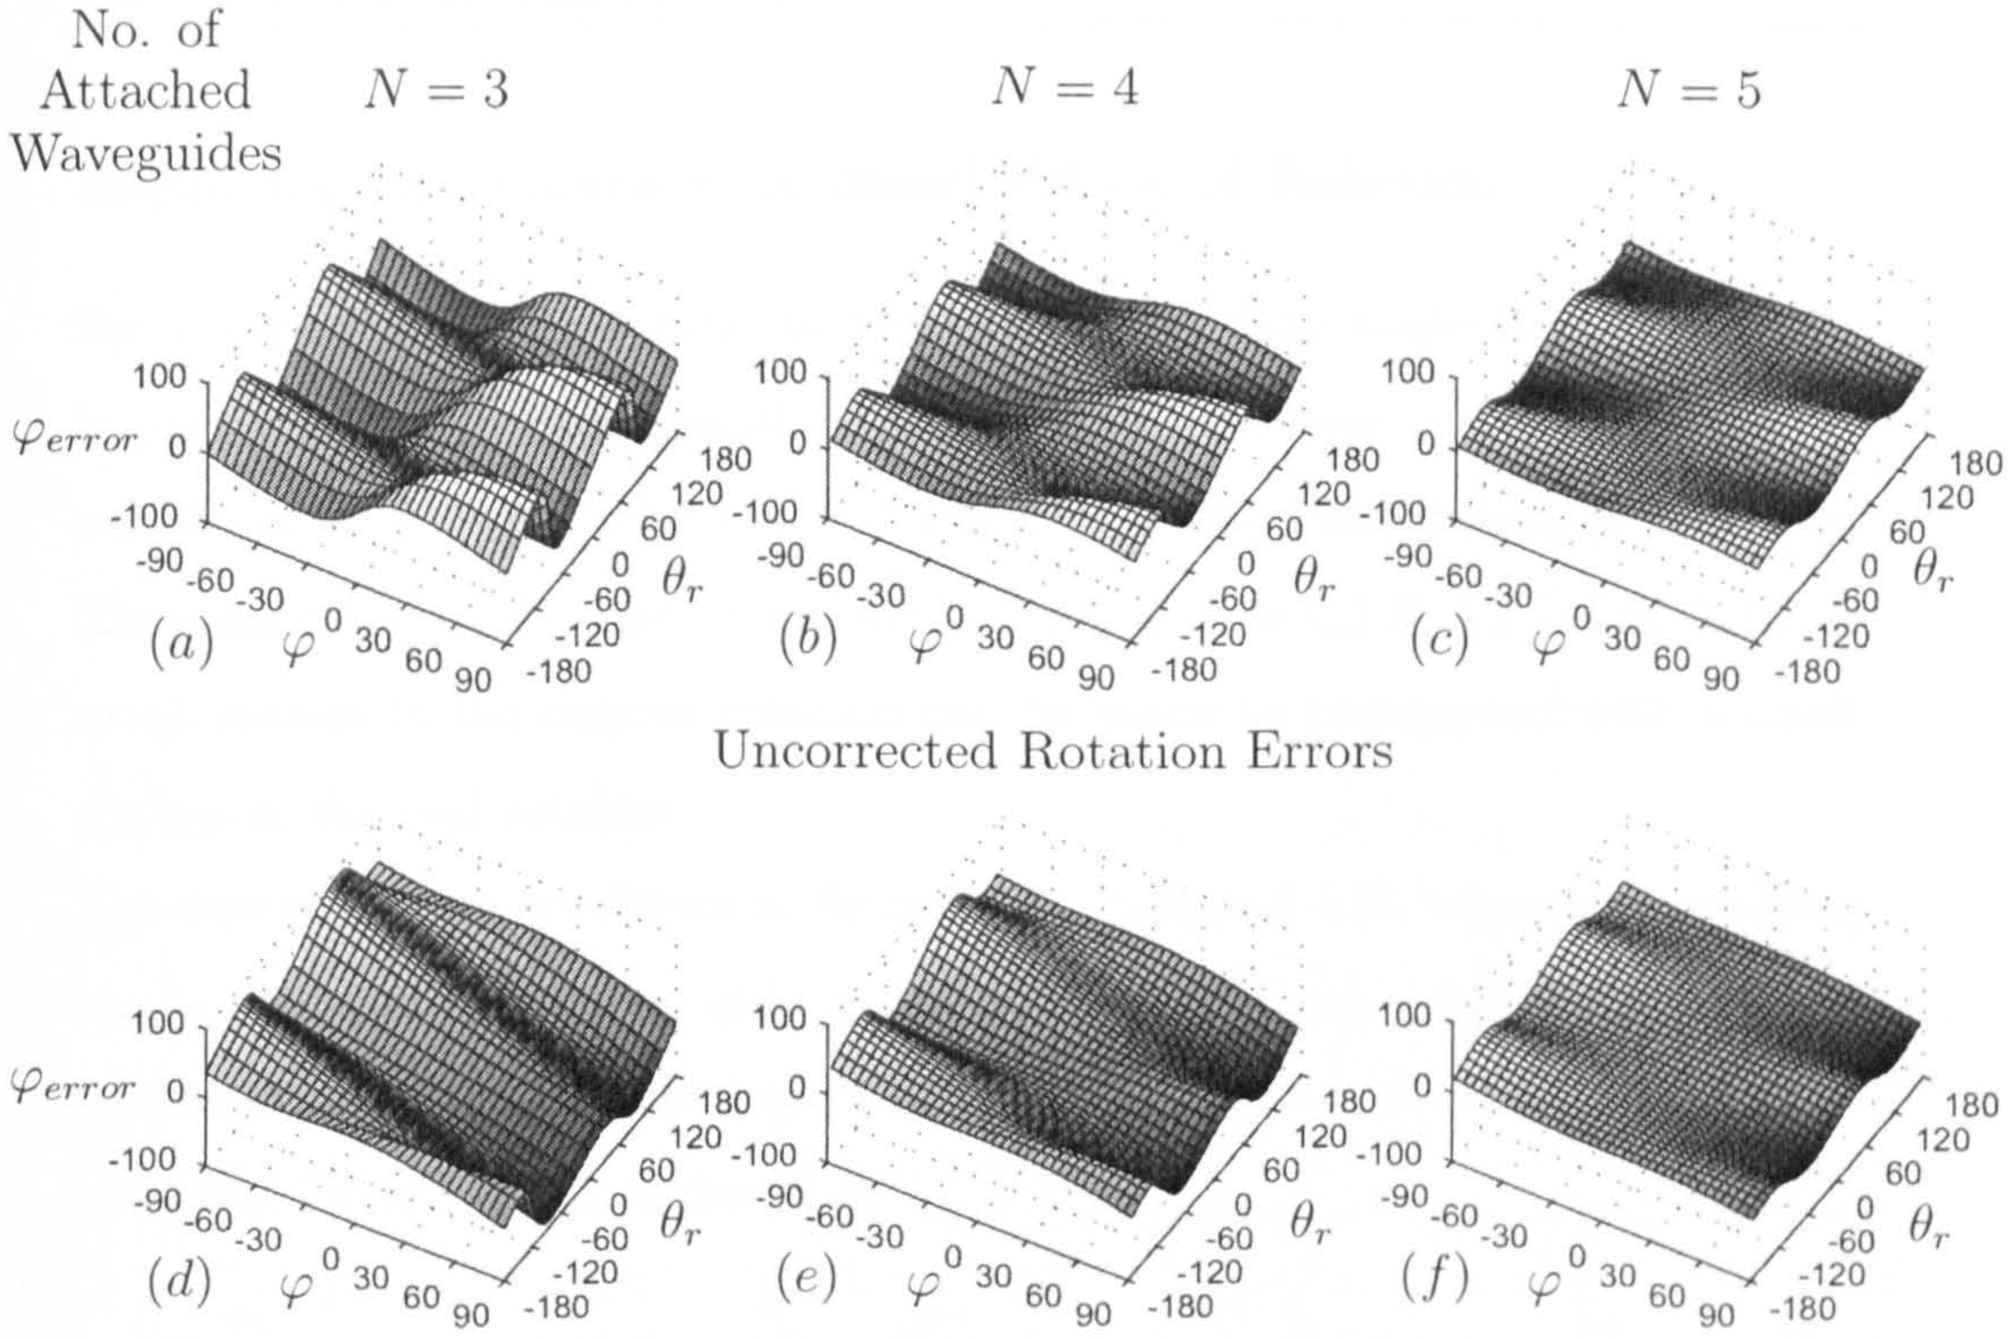

Figure 5.5: *Rotation Errors for Different Perimeter Node Structures (all angles are in degrees).*

mesh boundary.

Finding the Rotation Error

After convolution with matrix M the rotated angle of the real wavefront is found with

$$\theta_r' = \angle \left(\sum_{i=0}^{N-1} a_i' \underline{P}_{ri} \right). \tag{5.2.11}$$

The resulting error in rotation can thus be calcuated as

$$\varphi_{error} = (\theta_r' - \theta_r) - \varphi. \tag{5.2.12}$$

Plotting this error over a range of incident angles and rotations results in an error map. Figures 5.5a, 5.5b and 5.5c show the error maps for $N = 3$, $N = 4$ and $N = 5$ respectively. It can be seen that the greatest errors in rotation occur when the wavefront is described by the least amount of vectors or waveguides ($N = 3$).

Improving the Accuracy for Small Values of Rotation

By applying a constant ratio to the desired rotation angle it is possible to improve the accuracy for small values of rotation, where $\varphi \approx 0$. This ratio is calculated as $g = \frac{d\theta'_r}{d\varphi}$, where $\theta_r = 0$ and $\varphi = 0$. By substituting φ with $\frac{\varphi}{g}$ in the calculation of the matrix eigenvalues (equations 5.2.6, 5.2.7 and 5.2.8) a small change in the desired rotation can be made to correspond with a small change in the real rotation.

The new responses are shown in figures 5.5d, 5.5e and 5.5f, where the resulting errors are now seen to be only a reduction in the achieved rotation angle. Whilst an improvement has been made for small values of φ, the errors have increased for the larger values of φ.

Maximum Rotation Errors

For every angle of an incident wavefront, θ_r, a maximum absolute error is attained somewhere across the whole range of desired rotations, φ. As each plot in figure 5.5 spans a range of incident angles there is a corresponding range of maximum absolute errors. Note also that there is a certain degree of symmetry within the plots, due to the symmetry of the perimeter node structures, such that $\varphi_{error}(\varphi, \theta_r) = -\varphi_{error}(-\varphi, -\theta_r)$.

The different rotation error ranges are shown in table 5.4. Although the corrected rotations have greater maximum errors it should be remembered that these are intended for use with small values of desired rotation, φ. It can be seen that the rotation errors are sometimes quite large. The accuracy, however, is not so important when the waves are to be scattered randomly, as shall be seen when the rotation technique is applied to the modelling of diffusion.

5.2.4 Modelling Diffusion

To rotate the propagation angle of an incident wave the aforementioned circulant matrix is used to convolve the signals arriving at an edge-most node

	$N = 3$	$N = 4$	$N = 5$
uncorrected	$30° \rightarrow 60°$	$16° \rightarrow 31°$	$7° \rightarrow 13°$
corrected	$34° \rightarrow 89°$	$25° \rightarrow 62°$	$15° \rightarrow 30°$

Table 5.4: *Maximum Absolute Rotation Error Ranges for Different Perimeter Node Structures (all angles are approximate and in degrees).*

(excluding those arriving from the rimguides). The signals leaving this node are left unaltered. By randomly varying the amount of rotation at each sample step the reflected waves are scattered into different directions, as though they had been diffused by a rough surface. The amount of diffusion may be controlled by limiting the range of the rotation angles. The probability distribution used to generate the random rotation angles can be tailored to mimic the desired distribution of the scattered waves.

Results from the Diffusion Model

The same lossless drumskin model as used in section 4.1 was created and the diffusion model was appended to the boundary of the mesh. To help demonstrate the effect of the diffusion the angle range was ramped up from 0° to ±90° over a period of ten seconds. This is of course impossible to recreate in reality, but shows just how versatile a computer model can be. A uniform probability distribution was used to generate the random rotation angles and the angle correction method from section 5.2.3 was used. The model was excited with an off-centre impulse and the velocities over its entire area were summed to simulate an external listening point. The results from this simulation are shown as a spectrogram plot in figure 5.6.

Whilst the reflections are specular the energy in each resonant mode remains constant. It is interesting to see that as the diffusion starts to take effect the energy enters frequency modes that weren't initially excited. As diffusion increases further, significant (random) amplitude modulations are caused at each mode. This results in a widening in frequency of the modes and the presence

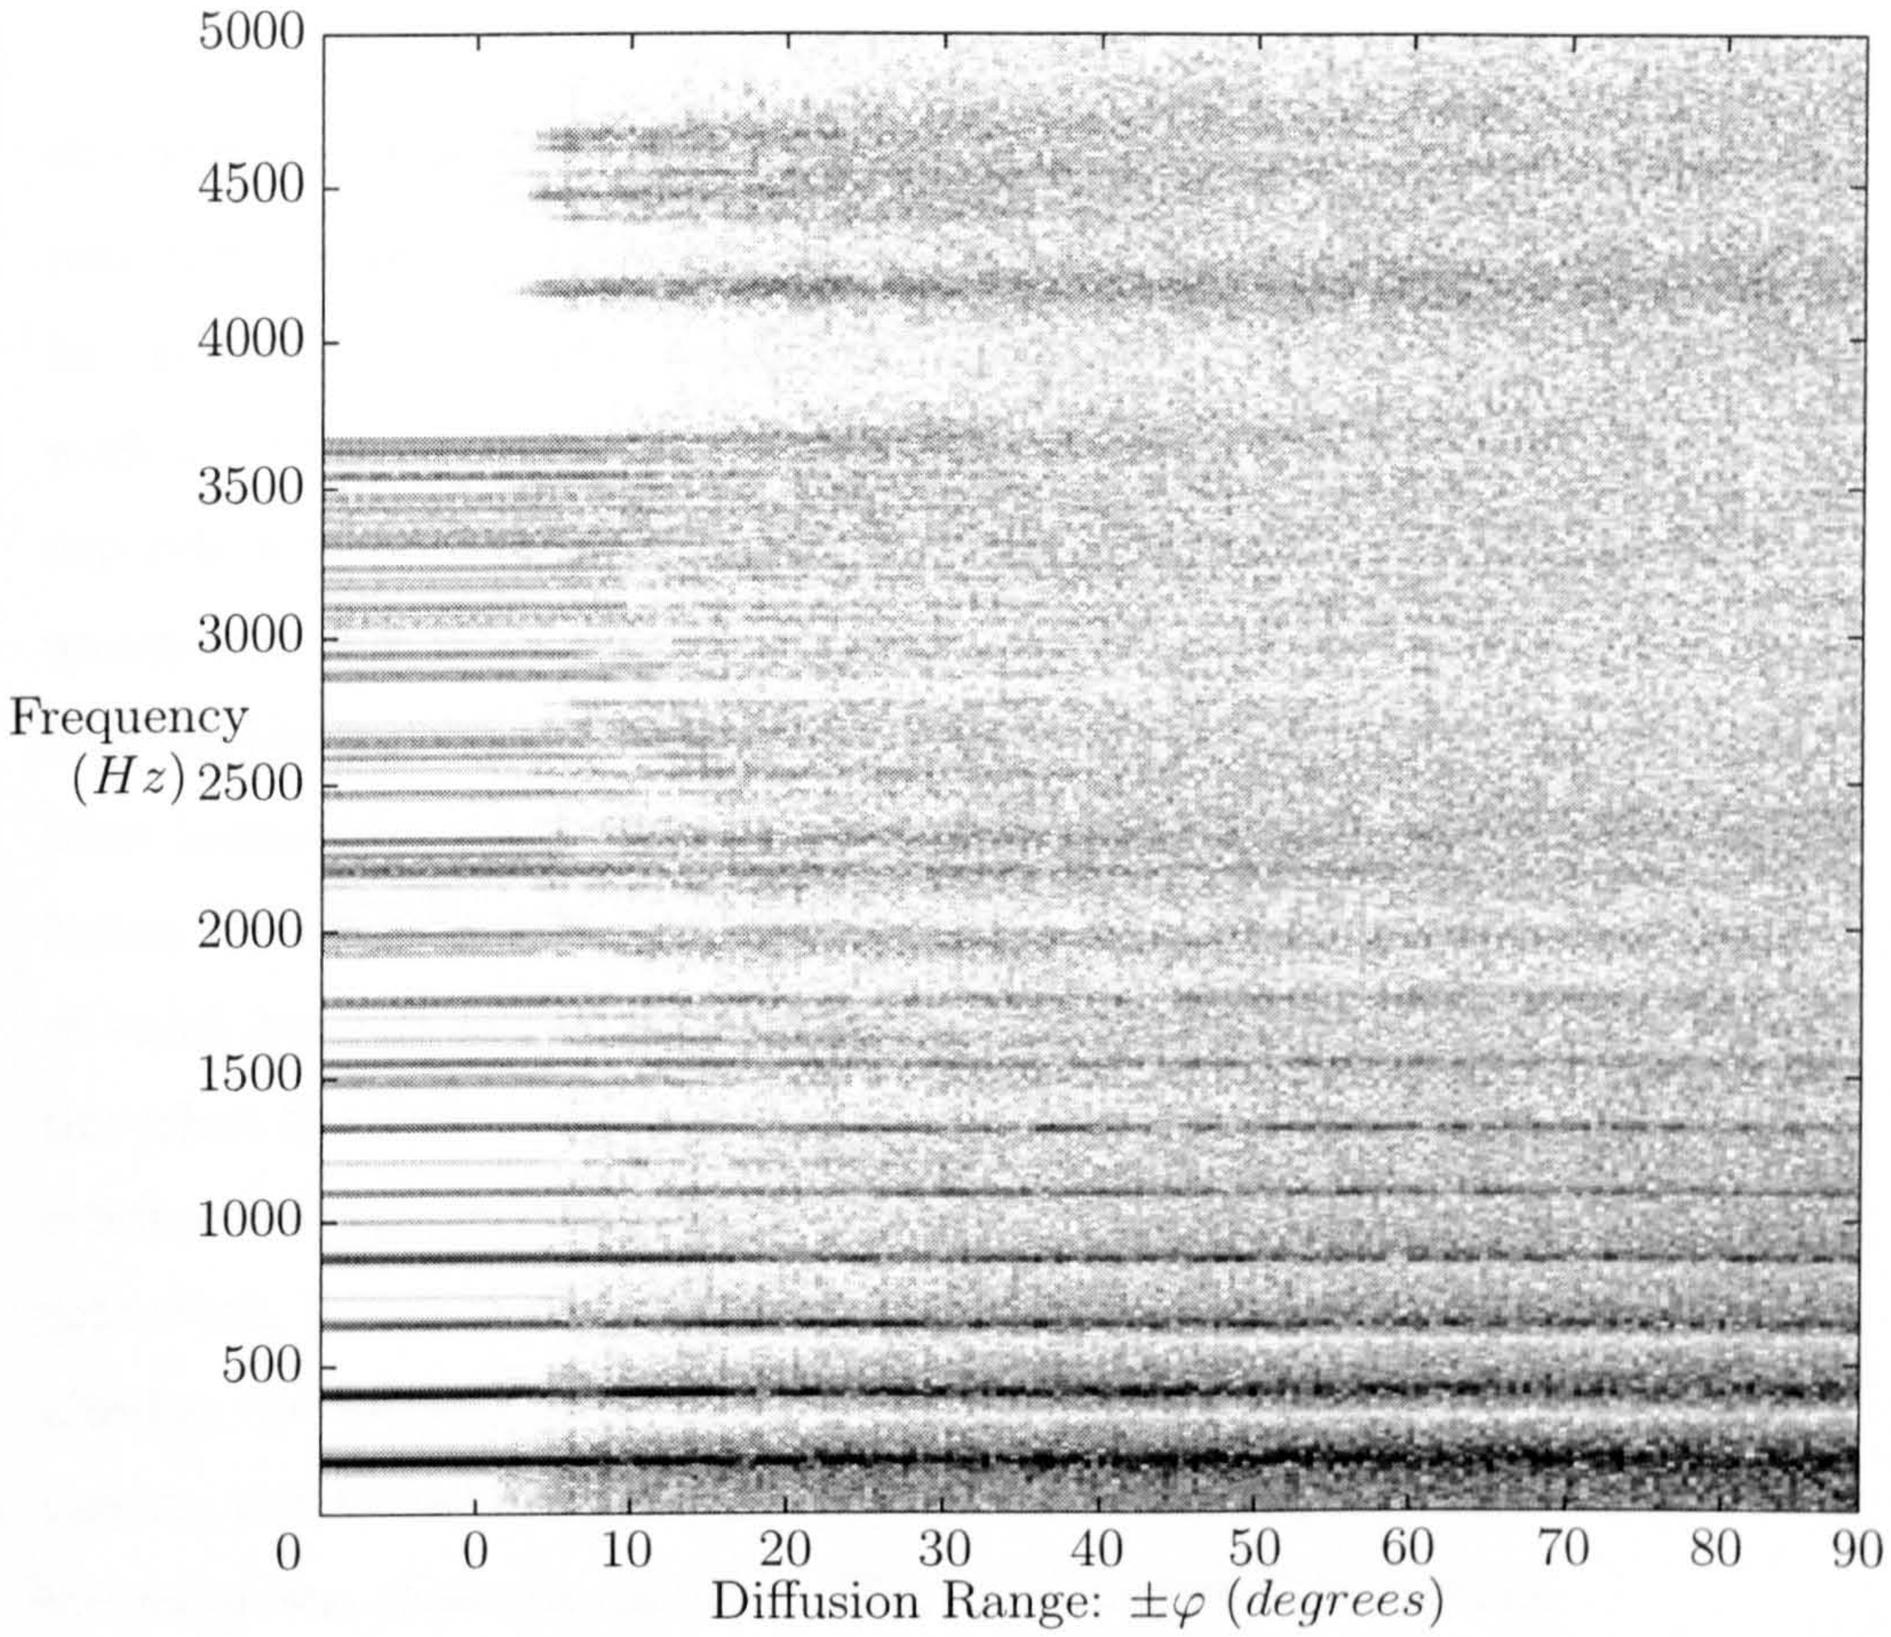

Figure 5.6: *Spectrogram of Sound from a Lossless Circular Membrane with increasing Diffusion at its Boundary.*

of additional noise. By the end of the simulation noise is the dominant sound. A simulation of time-varying diffusion can be heard on the accompanying CD-Rom (see section A.1.1 of appendix A). The model utilised was the same as that used for the results in this section. Here the diffusion angle range is increased with time until it reaches its maximum and then it is gradually reduced back to the specular reflection case. The final sound is similar to the initial sound, demonstrating that the technique is lossless and stable.

5.3 Summary

The modelling of wave propagation in 3D can be achieved using 3D digital waveguide meshes. Previous research in this area had utilised simple mesh

structures or those with multiply-free architectures, often for the purpose of real-time simulation. These meshes exhibit poor dispersion error and are therefore not sufficient for the accurate modelling of 3D wave propagation. In this work a new mesh structure was discovered that exhibited near direction independent dispersion error. This mesh has the most dense arrangement of waveguides that is possible in three dimensions. Its fundamental tessellating building block is the dodecahedron, hence it is called the 'dodecahedral' mesh. Most boundaries of space are rough which causes a phenomenon known as diffusion. Its effect is to scatter reflected waves into many directions, the range of which depends on the surface roughness and frequency of the waves. Whilst this effect has been incorporated into other modelling techniques there was no existing method for simulating diffusion at the boundary of a digital waveguide mesh. A technique was found to achieve this in 2D meshes which involved altering the incident angle of waves as they propagate to the boundary. By varying this incident angle randomly over time, within a specified range, the technique was shown to simulate diffusion in a controlled manner.

Through mathematical analysis of the rotation technique fundamental to the diffusion model it was found that there was a difference between the desired rotation and the actual rotation. This error generally increased for the larger values of rotation but was kept small when the boundary nodes were well described with a greater number of connecting waveguides.

Chapter 6

Completing the Drum Model

6.1 Interfacing the Different Media

The air cavity and drumskin models described in the previous chapters constitute the main components required to build a drum model. However, a method for interfacing these different media is required. Previous waveguide investigations had only considered meshes in isolation and so there was no existing interfacing method before commencing this work. As both models use meshes of different density and structure the interfacing is not a straightforward task. In order to pass contributions from a low-density mesh to a higher one it can be imagined that some form of interpolation will be required at the low-density surface. When passing information the other way a deinterpolation method will be needed. The simplest interpolation/deinterpolation methods to implement are those that are zero order; in which case the wavefronts are non-continuous. Improvements to this interfacing method are discussed in the future work section 8.5.

6.1.1 Interfacing the Edges of two 2D Waveguide Meshes

To help explain the proposed technique consider the interface between the edge of two 2D waveguide meshes that are supporting transverse waves. Each

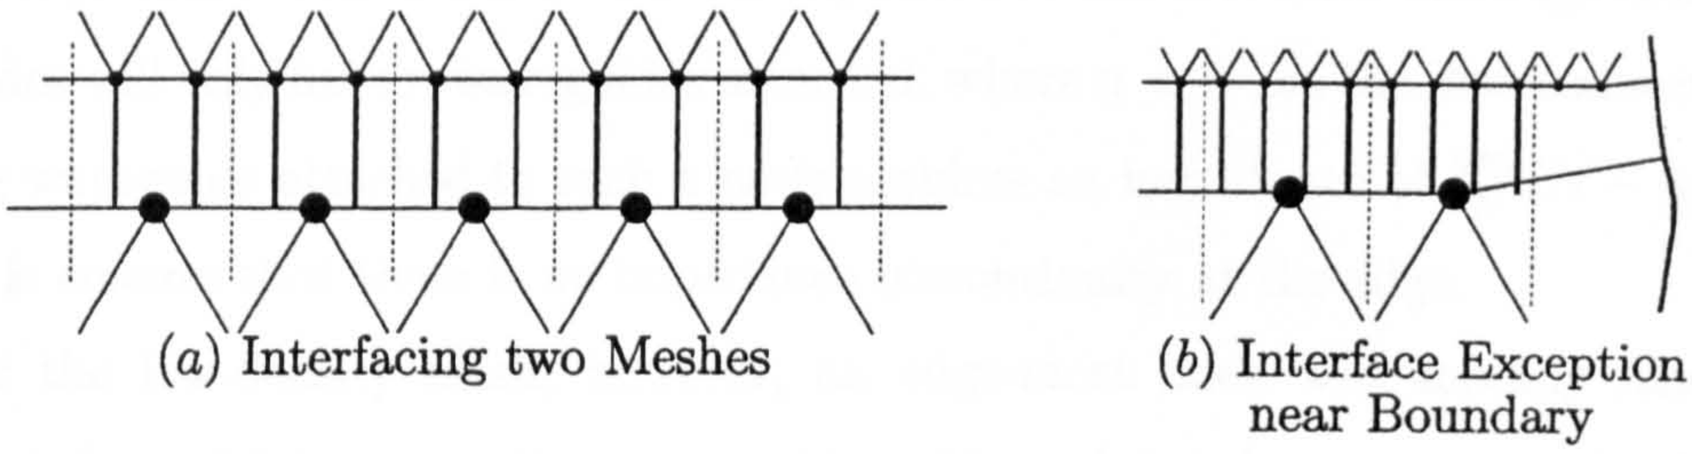

(a) Interfacing two Meshes

(b) Interface Exception near Boundary

Figure 6.1: *Interface Connections between two 2D Triangular Waveguide Meshes. The circles are node junctions and the lines represent waveguides.*

node at the interface boundary requires a least one connection to carry the travelling wave contributions between the different media. These connections are non-integer length waveguides (see section 3.3), the arrangement of which can be seen in figure 6.1a. The minimum length of these waveguides must be 1.5 spatial samples, as each path is made of a unit of delay and an allpass filter. Notice that the interconnections are all aligned to the same direction and don't appear to attach directly to the node junctions of the lower density waveguide mesh. This is done to emphasize the required length of the interconnecting waveguides; they model waveflow perpendicularly across the boundary and hence take the shortest distance. In the model, however, these waveguides are connected to the nearest nodes. This assumes that the wavefront measured at a node is constant across the element centred over this node (bounded by the vertical dashed lines), in accordance with piecewise constant interpolation.
It is important to note that if the two meshes are of very different densities then some interconnections could be outside of this element. An example of this is shown in figure 6.1b. It has not been possible to include these connections without causing instabilities, and so they are simply ignored in the model.

The Interface Connection at the Low-Density Mesh

Assuming that the low-density mesh represents a homogeneous material of impedance R_1, then its waveguides all have an impedance of $\frac{2R_1}{N}$ (see equation D.16 in section D.2), where N is the number of waveguides connected to each

node junction within the mesh. As explained in section 4.1.3 the edge-most nodes will only have n waveguides attached, where $n < N$, so the interconnecting waveguide attached to such a node requires an impedance of $\frac{2R_1}{N}(N-n)$. This ensures that there is no impedance discontinuity at the edge.

For the low-density mesh, however, an edge-most node will actually have a number of interconnecting waveguides, M, arriving from many nodes in the high-density mesh. Each of these must be scaled such that their total impedance equals that mentioned above. Therefore, the interconnecting waveguides have an impedance of

$$R_{connection} = \frac{2R_1}{MN}(N-n), \tag{6.1.1}$$

where R_1 is the impedance of the low-density mesh,

N is the number of waveguides connecting to a node junction within the mesh,

n is the number of waveguides connected to the edge-most node being considered,

M is the number of interconnecting waveguides arriving at this node from the high-density mesh.

The Interface Connection at the High-Density Mesh

To complete the interface the interconnecting waveguides are attached to their corresponding node junctions in the high-density mesh. The impedance of these waveguides remains the same as that defined above. The impedance discontinuity caused by doing this is, in fact, the appropriate impedance mismatch between the two meshes; an interconnection is an extension of the low-density mesh.

An Example Simulation

To demonstrate the interface model two 2D triangular waveguide meshes were created. They both had the same impedance yet one was twice the density of

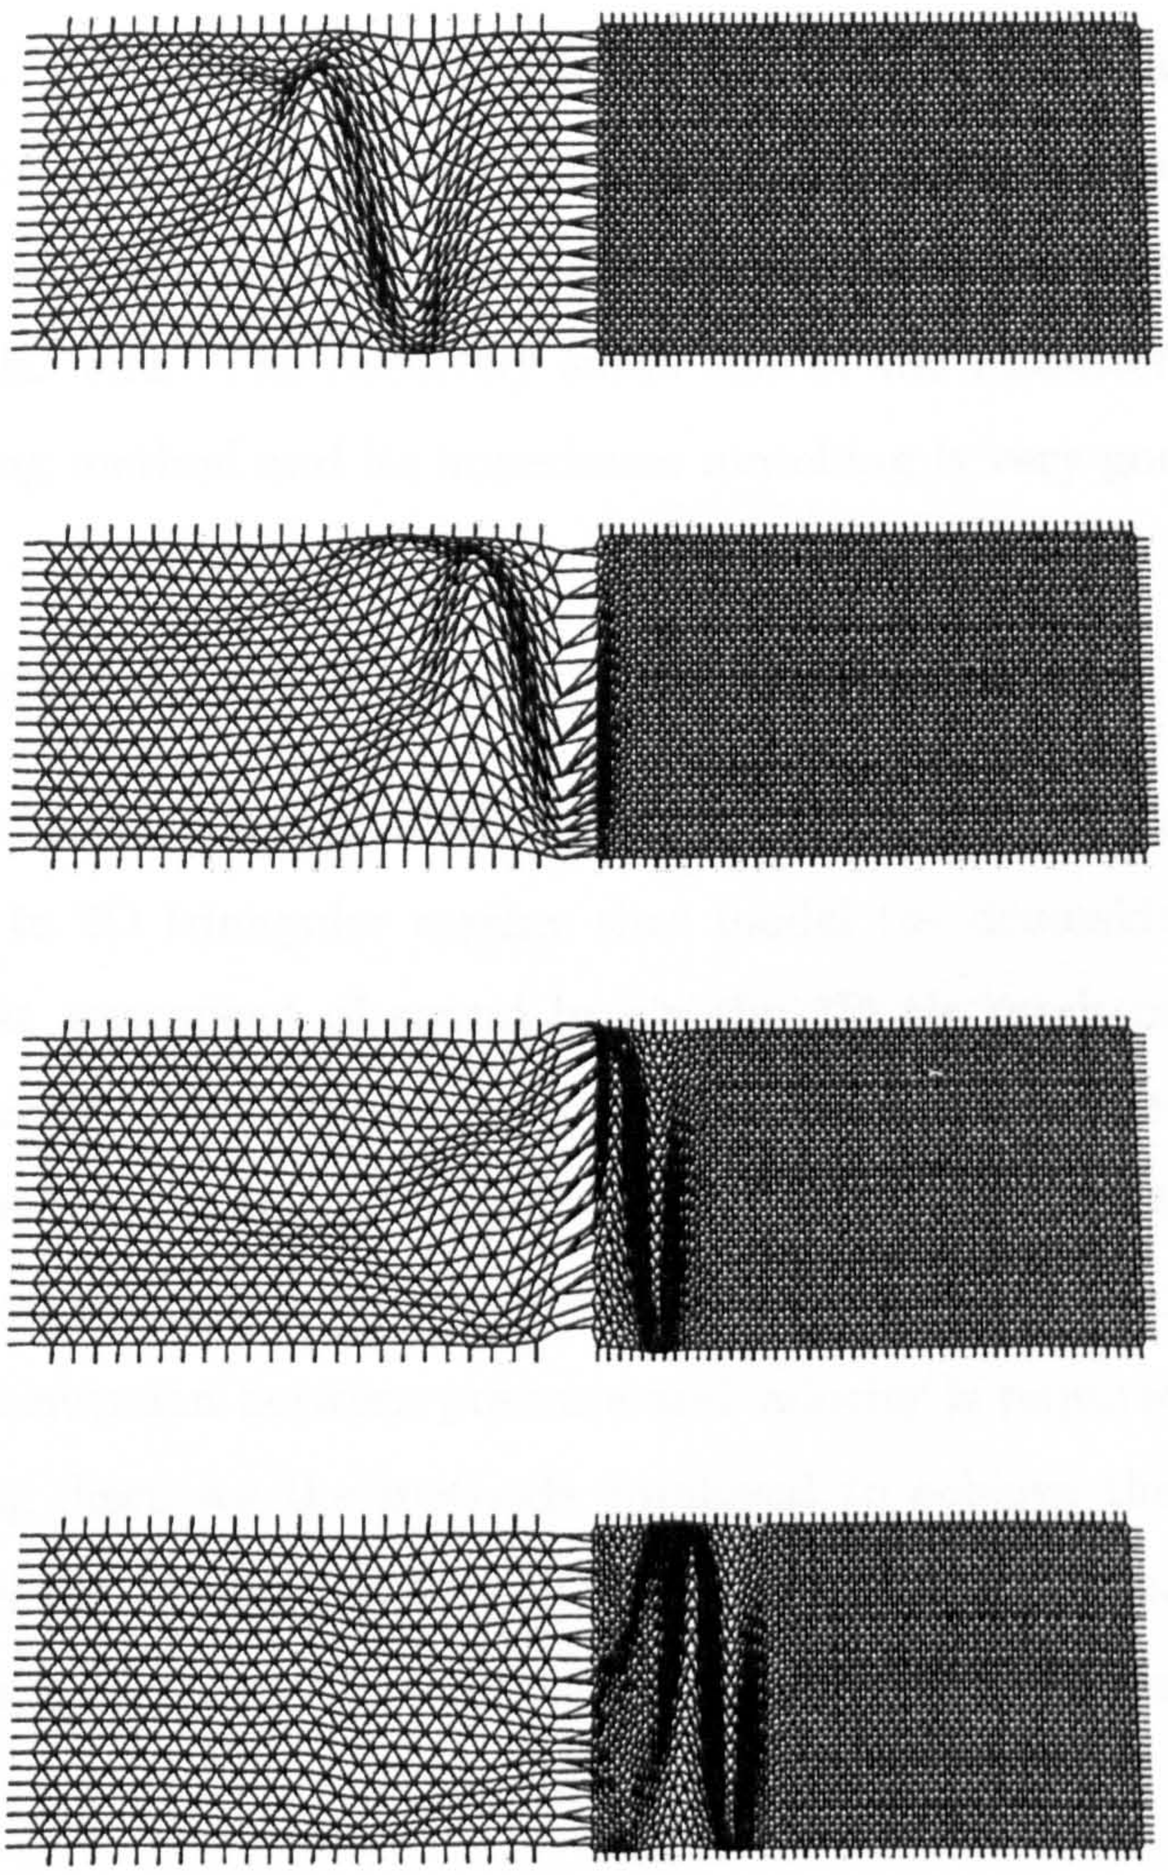

Figure 6.2: *Progression of a Wavefront across a Digital Waveguide Mesh Interface. The time intervals between each frame are identical.*

the other. The mesh boundaries were rectangular in shape and three sides of each mesh were clamped with rimguides. The remaining sides were attached together using the interfacing technique. A wavefront was generated by exciting the nodes along one boundary with one cycle of a sinusoidal function. Figure 6.2 shows the progression of the wavefront across the interface.

An animation of this can be found on the accompanying CD-Rom (see section

A.1.3 of appendix A). Notice how the wavefront travels more slowly in the denser mesh. It is possible to discern a very small reflection at the interface. This reflection returns to the start, where it is reflected again, and then travels back to the interface entering the high density mesh just as the initial wave reaches the far end. The relatively small size of the reflection suggests that the interfacing method and its impedance matching is very good.

6.1.2 Interfacing 2D and 3D Meshes

In the case of a drum model a 3D dodecahedral mesh, modelling the air, will be attached to 2D triangular meshes that model the drumskins. Due to the relatively fast wavespeed of sound in air the 3D air mesh model will have the lowest density. The interfacing connections are to be created in the same manner as before, but special consideration needs to be given to the fact that longitudinal waves have to be converted to transverse waves and vice-versa. In addition, a conversion between pressure and velocity is required.
The following discusses the methods employed to achieve these conversions and to model the 2D-3D interface. The interface is used in the simulation of a kettledrum, in chapter 7.1, where the interior air is connected to a circular membrane.

Converting between Longitudinal and Transverse Waves

In order to convert longitudinal waves to transverse waves and back it is first necessary to define the directions in which the transverse waves act. As an example, consider a tom-tom drum model with positive displacement relating to 'up'. The tom-tom has two membranes, one on the top and the other on the bottom of a cylindrical air column.
At the top membrane an increased (positive) value of pressure within the air column will cause an upwards (positive) deflection. Because a positive value in the air model results in a positive value in the membrane model there is no

need to alter the sign of the travelling wave from air to membrane. Conversely, a positive displacement in the top membrane will decrease the pressure below in the air column. In this case the returning path of the waveguide, which connects the membrane model to the air column model, must have a sign inversion attached to it.

The interconnections between the air column and the bottom membrane behave in the opposite manner. Here, a positive pressure within the air column forces the membrane downwards, causing a negative displacement. Therefore, the path of the waveguide that connects the air model to this membrane model has a sign inversion associated with it. The return path obviously leaves the sign of the travelling wave unaffected.

Converting between Pressure and Velocity

Before data in an air model can be passed to a membrane model a conversion from pressure to velocity is required. This can be done by dividing the pressure by the air's impedance, Z (see equation 3.8.4 in section 3.8). Conversely, the conversion from velocity back to pressure can be done by multiplying by Z.

Impedance of a 2D-3D Interconnection

The impedance at the air end of the interconnecting waveguides is calculated as before using equation 6.1.1 with $R_1 = Z$. However, at the membrane end the 3D model is being converted to a 2D model, and so the values of impedance are expected to relate velocity to force instead of pressure.

Each connection from the air can be thought of as a small pipe with cross-sectional area $\frac{A}{M}$, where A is the area of an element centred over an air mesh node (m^2). Multiplying the air's impedance by this cross-sectional area results in the impedance experienced at one end of the pipe, which is connected to an element of the membrane mesh. This impedance will relate velocity to force. The impedances of an interconnection, as seen by both meshes, are therefore

calculated in the following way

$$Z_{air\ connection} = \frac{2Z}{MN}(N-n), \tag{6.1.2}$$

$$R_{membrane\ connection} = \frac{ZA}{M}. \tag{6.1.3}$$

6.2 Instrument Excitation

The exciter provides the interface between player and instrument. It is the manipulation of the exciter that enables the player to generate and control the instrument's sound. Consequently, the method of excitation is the most important factor in defining the tonal qualities and evolution of this sound. After all, it is the excitation that transfers all the energy to the instrument and it does this in a time-varying manner.

As an example consider a piano, violin and guitar. All three instruments share the same topology; there are strings connected to a sound-box, albeit with different dimensions. However the methods of excitation are fundamentally different, which results in the characteristics we expect to hear from each instrument. The piano strings are struck with a felt mallet, the violin strings are generally bowed and the strings of a guitar may be plucked or picked.

The interaction between player, exciter and instrument is also dynamic with one affecting the other. This is an important point to consider, as the current state of the instrument will influence the exciting mechanism and, in turn, this will result in slight changes in tonal quality. Including such feedback interactions in a model is not only physically correct but also necessary to capture certain aspects of realism, the absence of which produces an unnatural sound. An example of this may be heard in sampled drum synthesisers. These synthesisers replay real recorded drum samples and are therefore very realistic when used for isolated drum hits. But, if a drum roll is performed then the sound is stuttered and unnatural, because each hit triggers the same sample regardless of the drum's state.

6.2.1 Excitation Modelling

Impulse Exciter

This is the simplest form of excitation to model. Ideally, it lasts for an infinitesimally short period of time (at most one sample period during modelling), and therefore does not respond to feedback from the instrument. The impulse exciter is modelled by adding a single value, for one sample period, to a junction's velocity within the digital waveguide mesh. Obviously, such an excitation is impossible to achieve in reality and its usefulness may not be immediately apparent.

It is the study of digital filters that highlights the importance of impulses and their resulting impulse responses, for these can be used to characterise the filter's behaviour; its frequency dependent gain and phase response can be found in this way. Essentially, the ideal impulse is a spike of energy made up from a superposition of all frequencies. The impulse response shows how all these frequencies are affected by the filter.

A digital waveguide model of an instrument is essentially a very large and complex filter, and thus its behaviour can be similarly studied through excitation with an impulse. In this case, the impulse injects energy into all possible frequencies up to the Nyquist frequency. Fourier analysis (see section B.5) may then be performed on the model's output in order to find all its resonant frequencies.

Mallet Exciter

A mallet is the general term used to describe hammers and drumsticks. These exciters are used to play pianos and percussive instruments. The mallet is thrown against the instrument, which it interacts with for a short while and then bounces off. The interaction can be viewed as a non-linear mass and spring system (Chaigne and Askenfelt, 1994), see figure 6.3. Here the mass models the dynamic mass of the mallet and the spring models the compression

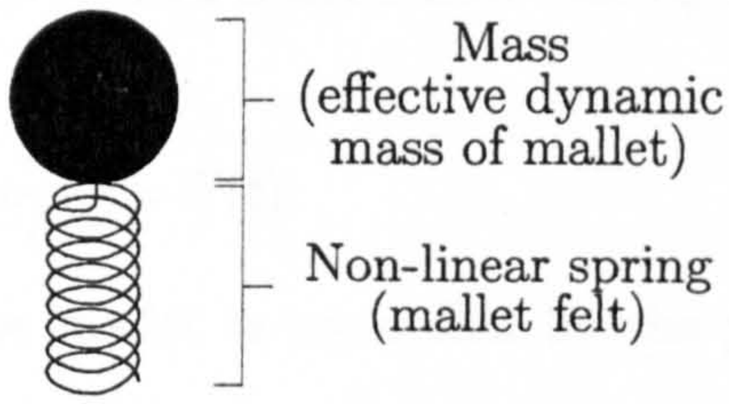

Figure 6.3: *The Mallet Model*

of the part that touches the drumskin (in this case the mallet felt). The free end of the spring connects to the model of the drumskin.

Governing Equations

When the mallet pushes against the membrane of a drum part of it compresses like a spring. The relationship between force and compression is generally non-linear, but well described by the equation(Chaigne and Askenfelt, 1994)

$$F(t) = K[(\delta - u(t) + W(t))^{+}]^{\alpha}, \tag{6.2.1}$$

where K is the coefficient of mallet stiffness $(Nm^{-\alpha})$,

δ is the initial position of the mallet's centre of gravity (m),

$u(t)$ is the time-varying position of the centre of gravity (m)

(both measured perpendicularly from the membrane),

$W(t)$ is the mean displacement of the membrane's area in contact

with the mallet (m),

the symbol + means 'positive part of',

α is the stiffness nonlinear component.

The parameters K and α can be found through measurement of force-deformation curves, following the same method as used for piano hammers (Chaigne and Askenfelt, 1994).

It can be seen that the force, $F(t)$, increases with an increase in compression. This force acts on both the membrane and the mallet, pushing them apart.

The motion of the mallet obeys Newton's Second Law (see equation C.10 in appendix C.1)

$$m\frac{\partial^2 u}{\partial t^2} = F(t). \tag{6.2.2}$$

Because of the way in which the mallet is controlled by the percussionist, the effective dynamic mass, m, is not equal to the static mass. The effective dynamic mass (Chaigne and Askenfelt, 1994)(Hall, 1992) and its initial velocity can be found by measurement, using the same procedure as that used by Chaigne and Doutaut for xylophones (Chaigne and Doutaut, 1997).

Modelling the Mallet

In order to excite a digital waveguide model with a mallet, a model has to be created that is governed by equation 6.2.1. A travelling wave model of a mallet had already been created (Duyne, Pierce, and Smith(III), 1994), but it is difficult to incorporate the nonlinear component α of equation 6.2.1. The mallet's contact area was also not included in this model.

For this investigation a simple but effective finite difference method was formed to implement the mallet. Initially, its velocity is set to v_0 and the mesh nodes that it will interact with are found, such that the mallet's contact area can be included. The number of these mesh nodes is stored in $N_{contact\ nodes}$. The force, $F(t)$, starts at 0. At each time interval the following steps are performed.

Applying Interaction Force to Mesh Nodes

The current value of $\frac{F(t)}{N_{contact\ nodes}}$ (the force applied to each node) is included in the velocity equations of the mesh nodes that the mallet interacts with, as shown below (see equation D.18 in appendix D.3)

$$V_{j\ contact\ nodes}(t) = \frac{2\sum_n v_n^+ R_n + \frac{F(t)}{N_{contact\ nodes}}}{\sum_n R_n}. \tag{6.2.3}$$

The total force applied by the mallet has now entered the digital waveguide model.

Calculating $u\left(t+\frac{1}{f_s}\right)$

The mallet's acceleration, $a(t) = \frac{\partial^2 u}{\partial t^2}$, is calculated using equation 6.2.2. This acceleration is assumed to be constant over the sample step, so the next velocity of the mallet will be

$$v\left(t+\frac{1}{f_s}\right) = v(t) + \frac{a(t)}{f_s}. \tag{6.2.4}$$

The new position of the mallet can then be found through first order integration

$$u\left(t+\frac{1}{f_s}\right) = u(t) + \frac{v(t) + v\left(t+\frac{1}{f_s}\right)}{2f_s}. \tag{6.2.5}$$

Calculating $W\left(t+\frac{1}{f_s}\right)$

The average velocity of the mesh nodes that the mallet interacts with is calculated as

$$V_{avg.\,contact\,nodes}(t) = \frac{\sum_j V_{j\,contact\,nodes}(t)}{N_{contact\,nodes}}, \tag{6.2.6}$$

so that the updated mean displacement, $W\left(t+\frac{1}{f_s}\right)$, can be found

$$W\left(t+\frac{1}{f_s}\right) = W(t) + \frac{V_{avg.\,contact\,nodes}(t)}{f_s}. \tag{6.2.7}$$

Iterating the Mallet Model

The waveguide model is now iterated for one sample period $\left(\frac{1}{f_s}\right)$. During this time the travelling waves move to their next positions and the new values of displacement, velocity, force and pressure are calculated within the model. Once this has been done equation 6.2.1 is used to find the new value of $F(t)$ and the above processes are repeated.

The accuracy of this new mallet modelling technique will be evaluated in the next chapter where it is used to excite a waveguide model of a kettledrum. Its acceleration curve is compared with that from measurement and a finite element simulation.
It is evidently straightforward to change or refine this mallet physical model.

All that is required is to substitute a new equation in place of equation 6.2.1 that lies at the heart of the model, defining its behaviour. This gives it an advantage over the travelling wave model that was previously mentioned, which would have to be completely revised if such a change were to be made.

6.3 Summary

The ability to attach digital waveguide meshes of different density and structure together is essential when creating composite waveguide models. There had been no documented attempt of achieving this and so a new technique was proposed in this chapter. When attaching two meshes together the wave variable over an element of the least dense mesh was considered to be constant. This enabled a simple interpolation/deinterpolation at the interface, which was modelled using fractional length waveguides. Careful calculation of the interfacing waveguides' impedance was required to ensure stability, although this could result in mesh nodes close to an edge remaining unattached. The technique was demonstrated by connecting two 2D waveguide meshes together and its extension to 2D-3D interfaces was described. Such an interface is used to connect the drumskin to the interior air in the next chapter, where a kettledrum model is constructed.

Without a method of exciting a model there would be no way of evaluating its response with measurement from reality. In the case of a drum the mallet interaction needs to be modelled. A drum mallet model that used a travelling wave description had previously been researched, but this did not incorporate an accurate mathematical description of the interaction. A new model was therefore created that used a finite difference implementation with this mathematical description at its core. This mallet model is easy to attach to a digital waveguide mesh and it also incorporates the size of the contact area.

Chapter 7

Model Analysis and Conclusions

7.1 Analysis of a Kettledrum Model

The new waveguide techniques described in chapters 4, 5 and 6 are sufficient to construct a simple drum model. In order to evaluate these techniques a comparison is made with the results from a previous study of a kettledrum (Rhaouti, Chaigne, and Joly, 1999). Here, a finite element model was constructed and an analysis was done on its approximation to measurements read from the real instrument.

7.1.1 Model Creation

The parameters of the kettledrum model are presented in table 7.1. All measurements are taken from the centre of the membrane with the z-coordinate corresponding to the vertical axis.

A digital waveguide model of this kettledrum was created using the techniques described in the previous chapters. These cover all aspects of the model's construction, however there is as yet no technique to model the air surrounding the drum. This is necessary to record the output at the external listening point defined in table 7.1. Whilst this could be done by interfacing the drum to a large external 3D mesh, a simpler method was utilised. This involved summing

Mallet	
$v_0 = 1.4\,ms^{-1}$	$\delta = 0.025\,m$
$m = 0.028\,kg$	$K = 1.6 \times 10^8\,Nm^{-\alpha}$
$\alpha = 2.54$	Excitation Point:$(x, y) = (0.21\,m, 0\,m)$
Membrane	
$r = 0.3175\,m$	$\sigma = 0.262\,kgm^{-2}$
$T = 3325\,Nm^{-1}$	
Kettle	
$H = 0.5\,m$	Boundary Shape:$z = -\frac{H}{r}\sqrt{r^2 - x^2 - y^2}$
Sampling	
$f_s = 24\,kHz$	Listening Point:$(x, y, z) = (-0.31\,m, 0\,m, 0.1\,m)$

Table 7.1: *Parameters for a Kettledrum.*
Measured by Rhaouti, Chaigne, and Joly (1999).

the contribution at the listening point from each node of the membrane mesh. The signals from the nodes were delayed using fractional length delay lines (see section 3.3), each measured according to the distance between node and listening point. This technique does not include the interaction between the external air and membrane surface.

7.1.2 Model Analysis

The waveguide model was excited with the mallet and its output recorded for 9000 samples. Its resonant frequencies up to $700\,Hz$ were obtained through Fourier analysis (see appendix B.5) and compared in table 7.2 with those of the real kettledrum and finite element model.

It can be seen that the higher frequencies correspond well with those recorded from the real instrument and actually surpass, in some cases, the accuracy of those predicted by the finite element model. However the lower frequencies, especially the first two resonant modes, do not share this level of accuracy. A possible reason for this may be due to the interfacing problem demonstrated

Mode (mn)	f_{mn} (measured) (Hz)	f_{mn} (finite element simulation) (Hz)	f_{mn} (digital waveguide simulation) (Hz)
01	136	139	155
11	147	147	168
21	221	221	227
02	248	245	256
31	288	288	296
12	315	315	315
41	357	355	357
22	395	395	...
03	403	408	400
51	424	419	419
61	480	483	480
71	552	544	549

Table 7.2: *Comparison between Measured and Simulated Frequencies of a Kettledrum.*
Measured and finite element frequencies obtained by Rhaouti, Chaigne, and Joly (1999).

in figure 6.1b of section 6.1.1. For this waveguide model of a kettledrum the problem occurs close to the membrane boundary, where the outer membrane nodes are not connected to nodes of the interior air cavity. This will affect the load experienced by this part of the membrane, altering the resonant frequencies. Also the exterior air wasn't included in the waveguide model and this would apply an air load to the top surface of the skin. It has been calculated and demonstrated that the loading caused by an unconfined sea of air lowers the modes of resonance, with the modes of lowest frequency being affected the most (Fletcher and Rossing, 1991b). Indeed the membrane model on its own with no air-loading has been demonstrated to be very accurate at these low frequency modes (see section 4.1), so the errors are likely to have arisen from the membrane-air interface.

Figures 7.1 and 7.2 show the acceleration of the mallet during the kettledrum's excitation. The former compares the measured acceleration with the simulated acceleration from the finite element model, the latter is the mallet acceleration

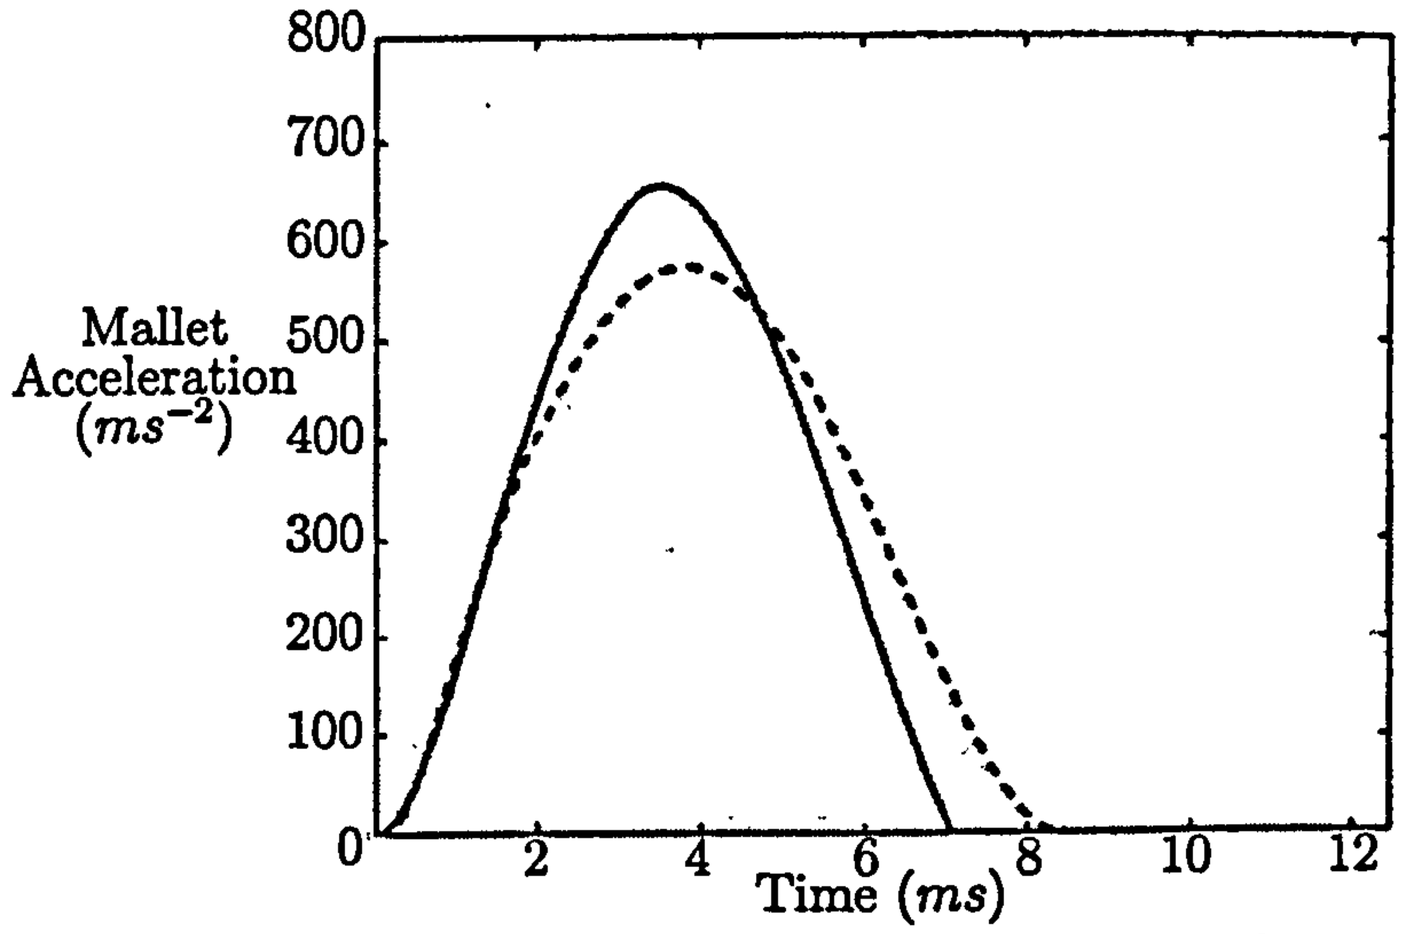

Figure 7.1: *Acceleration of the Mallet. Comparison of measured (solid line) and finite element simulation (dashed line).*
Reproduced from Rhaouti, Chaigne, and Joly (1999)

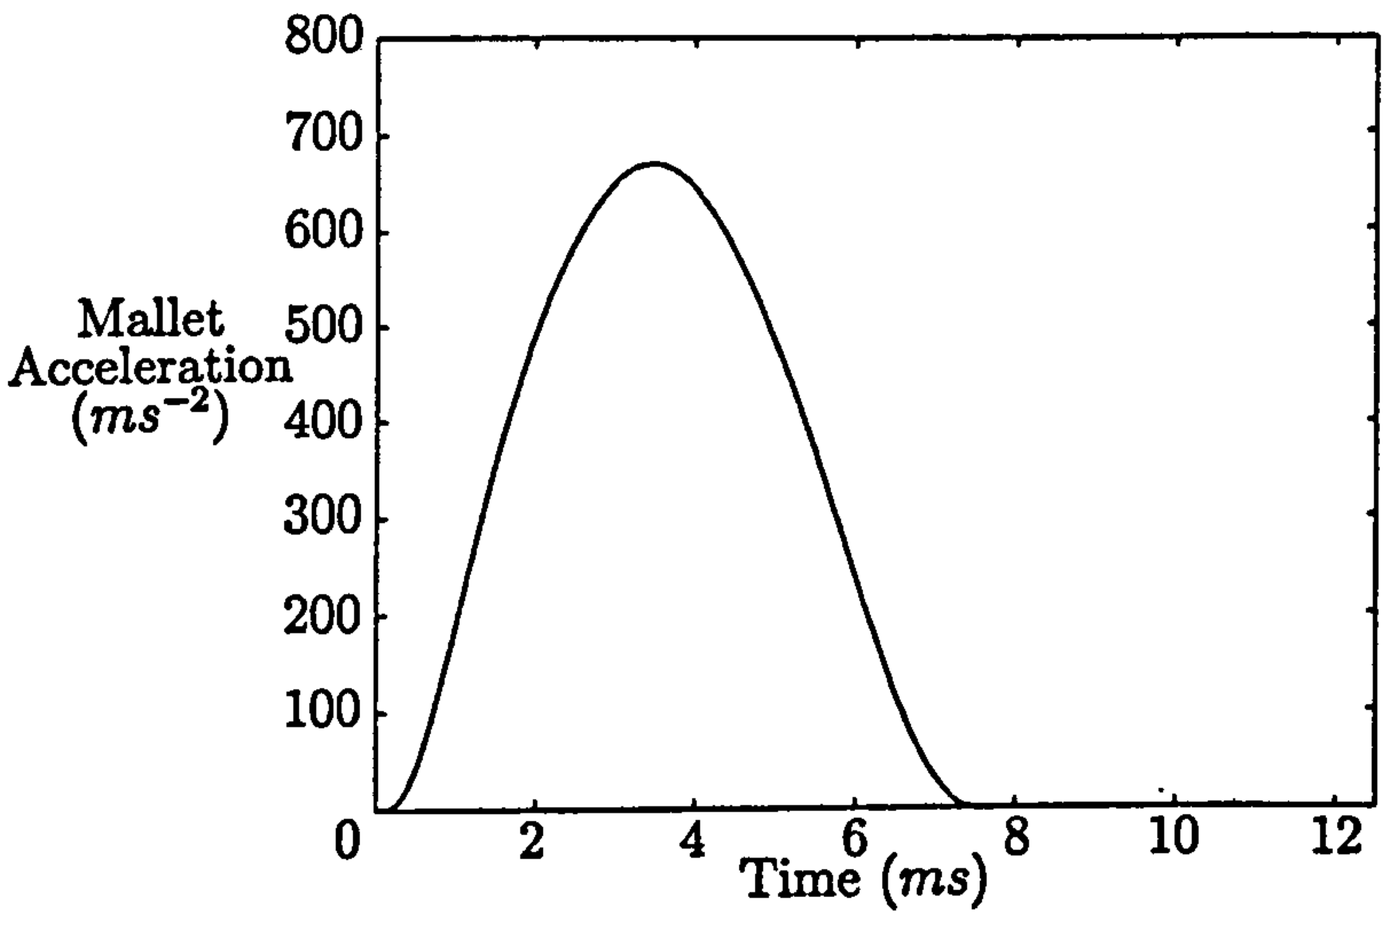

Figure 7.2: *Acceleration of the Mallet. Recorded from the digital waveguide kettledrum model.*

read from the waveguide model. The radius of the mallet's contact area was estimated to be $r_{mallet} = 0.015\,m$. This estimate results in an acceleration curve from the waveguide model that is remarkably similar in amplitude and shape to the real measured curve. This is a good justification of the mallet model and its interaction with the membrane because of the complexity of the dynamics involved; the acceleration of the mallet depends on the compression

of the non-linear felt, which in turn affects and is affect by the vibration of the membrane. If either part behaved incorrectly then this acceleration curve would not match the measured curve so closely.

To complete the kettledrum model it is necessary to add losses such that the decay times of the resonant frequencies match those of the real drum. As described in section 4.1.4 losses may be simply incorporated into waveguide models by adding lowpass filters and gains to the mesh boundary. Unfortunately, as this technique does not make use of measured material parameters the loss has to be adjusted through trial-and-error. It was found that a reasonable approximation could be made by adding first order IIR lowpass filters to the boundary of the interior air. A filter coefficient was used that would cause an attenuation of $-71\,dBs^{-1}$ for a wave of fundamental frequency (f_{01}) travelling diametrically in the air just below the membrane. This coefficient was used in each of the lowpass filters. In addition to this a pure gain was applied at the boundary of the membrane mesh. The value of this gain was chosen such that waves close to DC in frequency, travelling diametrically across the membrane, would suffer an attenuation of $-40\,dBs^{-1}$. Clearly this is not an ideal method for modelling losses, but, despite this shortcoming it should be noted that the finite element model incorporated a viscous friction constant that also had to be found through trial-and-error (Rhaouti, Chaigne, and Joly, 1999).

The models were run for three seconds (72000 samples) in order to let the sound decay to $-60\,dB$ of its original magnitude. The results are plotted in figure 7.3 (output from real kettledrum and finite element model) and figure 7.4 (output from digital waveguide model). Each figure has three plots to convey different information about the sound. The leftmost plot was obtained through Fourier analysis of the entire sample, indicating the resonant modes. The 3D plots in the middle were created through Fourier analysis of 9000 points with Hanning window and 50% overlap, in order to give a time-frequency response. The plots on the right show the signal amplitude response over time.

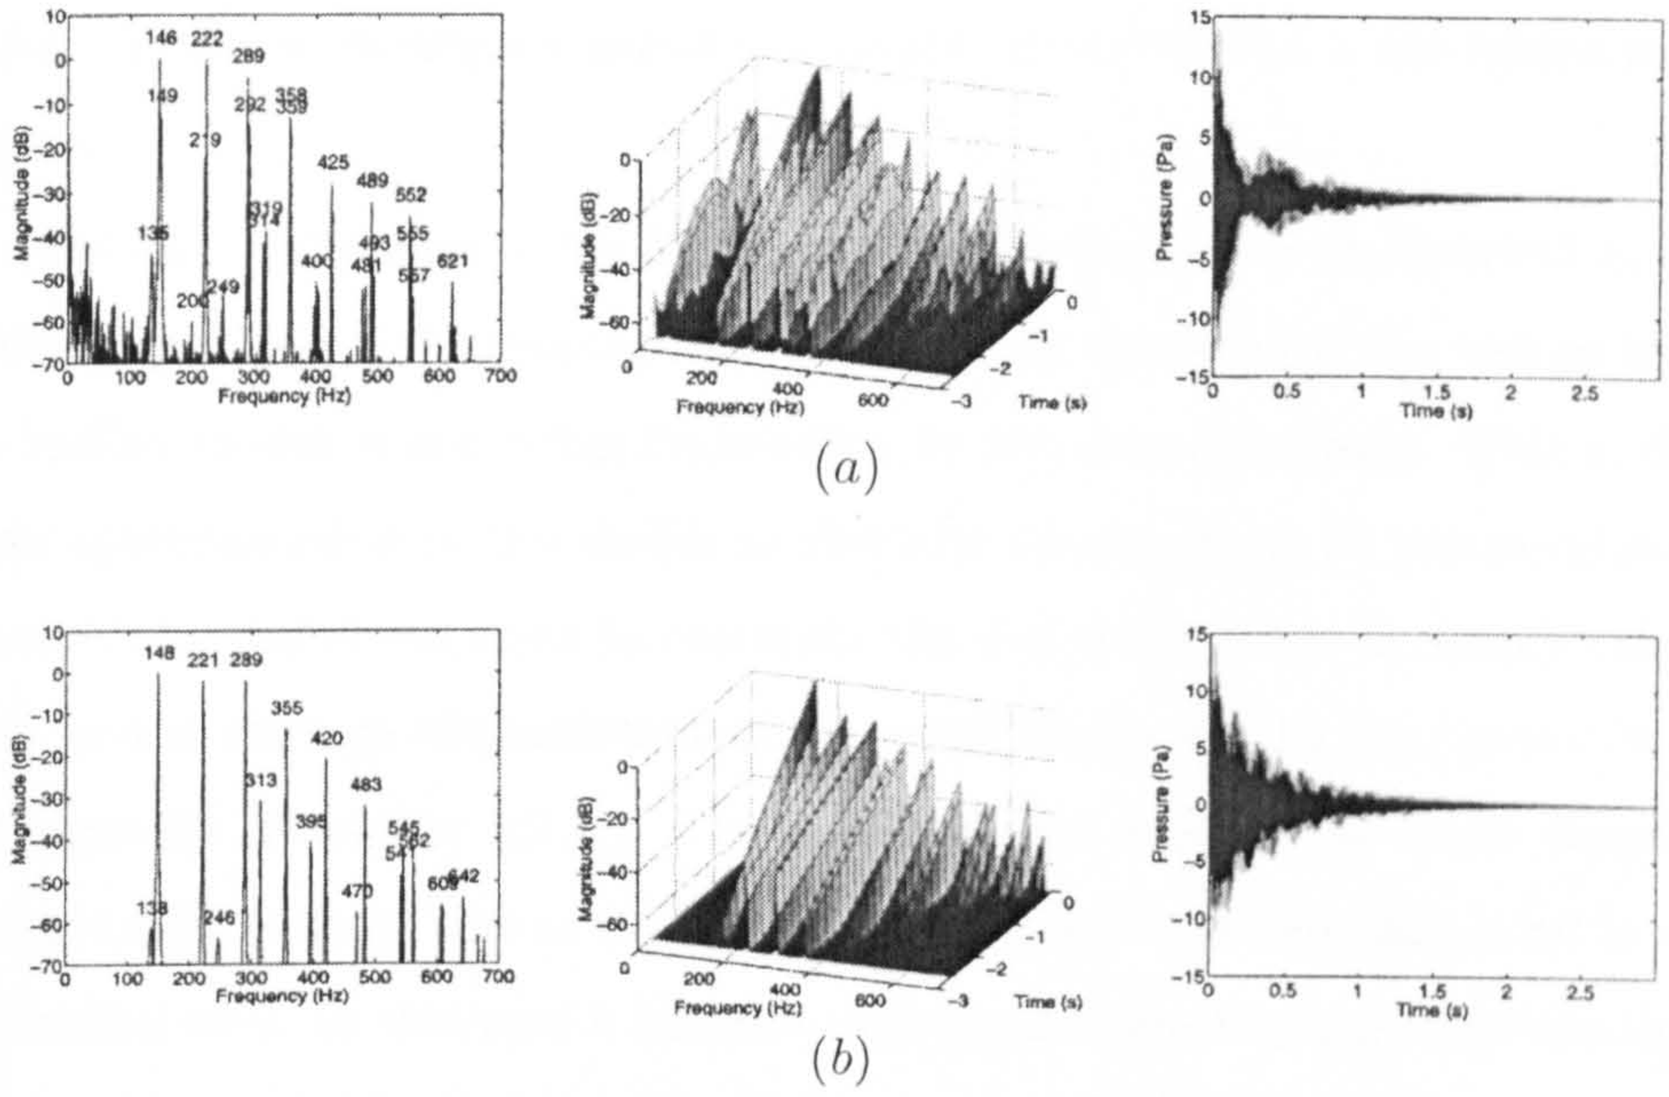

Figure 7.3: *Power Spectra, 3D Plots and Waveforms of Kettledrum Sound. (a) Real Kettledrum, (b) Finite Element Model.*
Reproduced from Rhaouti, Chaigne, and Joly (1999)

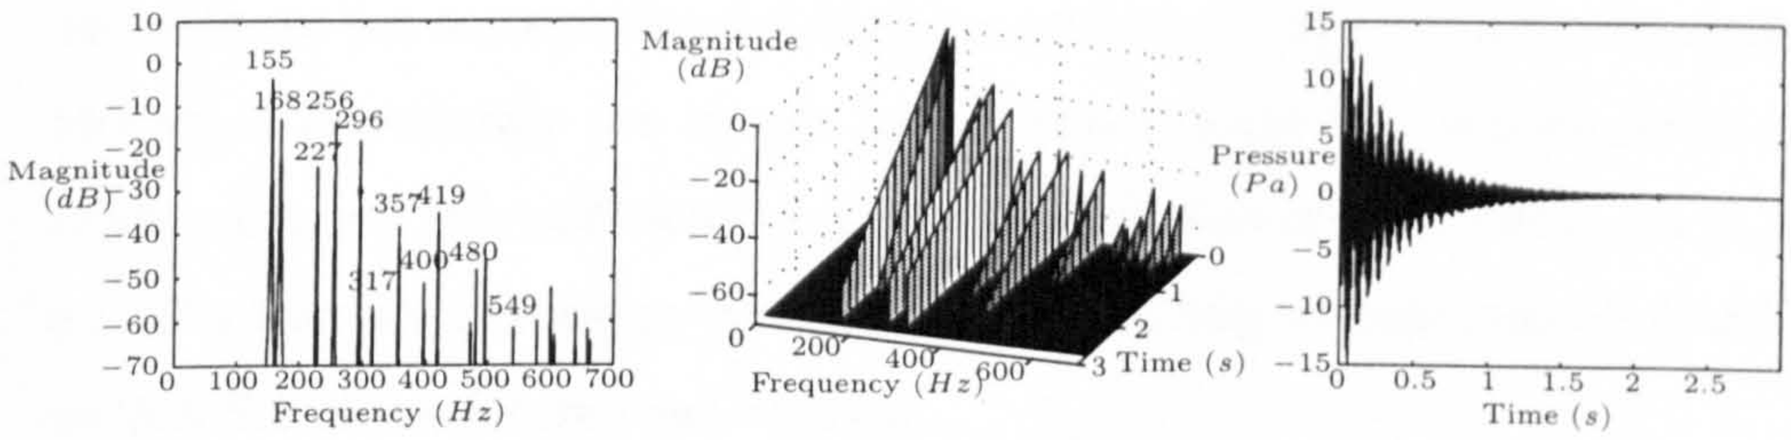

Figure 7.4: *Power Spectra, 3D Plots and Waveforms of Kettledrum Sound (recorded from the digital waveguide model).*

As previously noted the frequency content of the waveguide model conforms well to that of the finite element model and the real instrument, as shown in the leftmost plots. The extra low frequencies in the real instrument's plot were due to noise in the electronic channel during recording.

Doubled peaks can also be seen for the real instrument which result in a different shape of the signal amplitude response. It was found that these were due to imperfect tuning, where the tension of the membrane is not uniform (Rhaouti, Chaigne, and Joly, 1999). A technique for incorporating non-uniform tension

into a waveguide membrane model is proposed in section 8.4 of the future work chapter.

An important difference of the digital waveguide model can be observed in the initial magnitude of the resonant modes. It would appear that the energy from the mallet model is not being transferred to the drum properly. This is due to an approximation in the mallet model; the displacement of the membrane beneath the mallet, required to calculate the excitation force, is simply calculated as the average displacement of the mesh nodes within the contact area (see page 99 of section 6.2.1). The mallet model used in the finite element analysis differed from this in that the mean displacement was calculated in an alternative way. In this case a function was applied to the displacements that put greater emphasis on those close to the centre of the mallet. This was done to take the mallet's curved shape into account.

It can be seen from the time-frequency plots that some of the frequencies' decay times are too short in the digital waveguide model. The loss modelling technique is evidently too simple to accurately copy the response of a real instrument, yet it is sufficient to give the impression of the instrument to the ear. Further information on more accurate modelling of losses can be found in section 8.1 of the future work chapter.

The bearing edge and diffusion models developed earlier in this thesis were not included in the waveguide kettledrum model. This was because they were not considered for the finite element model of the previous investigation, from which results for comparison were taken. Their affect on the output of the model would have unnecessarily complicated the analysis. Nevertheless, the bearing edge would have been easy to incorporate as no modification would be required to its implementation. The diffusion model, on the other hand, would need to be extended to 3D before attaching it to the interior of the kettledrum. Considering the results from their analyses earlier in the thesis, it is possible that they would have caused a greater spread of energy across the resonant frequencies and increased the noise in the system. In addition there would be

a slight decrease in sustain.

The sound of this digital waveguide kettledrum model can be heard on the accompanying CD-Rom (see section A.1.1 of appendix A). Also included is the same model excited at the centre and another excited with a short drum roll. In addition there is the sound from a kettledrum model of the same dimensions but with less tension in the drumskin. In this case its fundamental is audibly lower.

The quality of the sounds is very good and they do indeed give the impression of a kettledrum being struck. In particular, the drum roll sound has a noticeable low frequency 'punch' for each strike of the mallet which sounds very realistic when played loud through a good amplifier. The samples do, however, have a slight 'metallic' feel to them and this is most likely due to the over-simplistic energy loss model employed. An improvement to this part of the model would no doubt result in extremely high fidelity sounds. Nevertheless, these sounds are still musically useful as drum samples. If they could be produced in real-time then the level of control and expression provided by the physical model would be invaluable to a performer.

The kettledrum model was created by specifying the shape and material properties. No other information was fed to the computer to help define the sound and so in this way such a modelling technique could be very useful to a Master Craftsman. In its present form the technique would enable such a craftsman to hear the pitch and relationship between the resonant frequencies for a drum of any shape or size with a variety of drum skins attached. With accurate modelling of the viscous friction losses the craftsman would be able to make a decision on what materials work well before entering the workshop. In fact the entire design and testing could be done on a computer, allowing greater freedom for imagination and creativity. Only the fine tuning and finishing touches would need to be applied to the actual physical product.

7.2 Conclusions

In this thesis the physical modelling of percussive drums was investigated using digital waveguides. The result of this research was a set techniques that could be used to construct a drum model, the accuracy of which was demonstrated with the analysis of a kettledrum. This advanced state of digital waveguide modelling required many extensions and additions to be made to the existing digital waveguide techniques. The culmination of the work was a permanent exhibit, named the 'Virtual Drum', in the Wellcome Wing of the London Science Museum. Its purpose was to demonstrate the cutting-edge in the physical modelling of musical instruments.

7.2.1 Modelling the Drumskin

In chapter 4 a 2D digital waveguide mesh was used to construct a model of a drumskin. In order to achieve a circular boundary techniques were created that extended the mesh to the appropriate shape and size. These techniques involved the attachment of rimguides to the edge-most nodes, which modelled wave propagation to the boundary and back. The techniques that were investigated employed increasing orders of complexity, however it was found that a simple fractional delay was sufficient, which corrected the propagation of waves close to DC in frequency. This technique was shown to reproduce the first seven resonant modes with good accuracy, see method 2 in table 4.3 on page 52.

A method was also proposed for calculating the impedance values of a 2D waveguide mesh. This was done by considering a small element of a mesh which represented its fundamental building block. The accuracy of this calculation is confirmed in the results from the kettledrum model, shown in table 7.2; the interaction between interior air and drumskin is evidently being reproduced correctly.

A technique for modelling the curvature of the bearing edge was also created

and its effect was investigated when attached to a string, like the bridge of a guitar or violin. The technique involved modelling the perceived shortening and lengthening of the string as it wrapped and unwrapped itself around the edge. It was shown in figures 4.13 and 4.15 on pages 67 and 69 respectively that the resulting effect is to leak energy into higher modes of resonance with an increase in bearing edge radius or string excitation. It was argued that this effect reduces sustain and is therefore undesirable. This conforms with experience of playing the guitar, where the bridge is made as sharp as possible to help increase sustain.

7.2.2 Modelling the Interior of a Drum

The modelling of the interior air was considered in chapter 5. Here a new waveguide mesh structure was proposed to model the wave propagation in 3D correctly. This structure, the dodecahedral mesh, was shown to have near direction independent dispersion error in figure 5.3 on page 76. Its response was found to be quite accurate when compared to the theoretical resonances of a room (see table 5.3 on page 78).

Previously all models were restricted to boundaries that exhibited specular reflection. A diffuse reflection modelling technique was created to incorporate into waveguide models the effect of rough boundaries. A method was found to alter the angle of an incident wave to the boundary, and this was performed randomly over time for different angle ranges in order to simulate a diffusion effect. Although it was only investigated for 2D it was shown, in figure 5.6 on page 87, that this technique is controllable and does indeed model the effect of diffusion. With further investigation this will become invaluable for simulating the acoustics of a room.

7.2.3 Interfacing Meshes and the Excition of Models

An area of waveguide modelling that was previously unexplored was the interfacing of different meshes. Naturally, as each part of a model will use waveguide

meshes of different structure and density, it is neccessary to be able to attach them together. Chapter 6 proposed a technique to achieve such an interface, which utilised simple interpolation/deinterpolation to create interconnections. These interconnections are lossless and were demonstrated in figure 6.2 on page 92 to accurately model wave propagation across an interface between two 2D meshes. The technique was advanced to model a 2D-3D interface and was put to good use in the model of a kettledrum, where the interior air needed to be attached to the drumskin.

Finally, a method was described for modelling a mathematical description of a drum mallet. This model is easy to attach to a waveguide model for the purpose of causing an excitation. Its behaviour was found to be very accurate when compared to measurements taken from reality, as shown in figure 7.2.

7.2.4 The Digital Waveguide Modelling Technique

Overall the digital waveguide techniques proposed in this thesis are quite accurate and already useful for musical instrument design. With further development their accuracy could be increased to match more complicated and established modelling techniques such as the finite element. Digital waveguide models enjoy many advantages over such techniques giving them the edge when chosen as a modelling tool. Firstly their execution requires simple computations, the majority of which involve copying data from one location to another. The remaining operations are digital filtering and the calculation of the scattering junction velocities. Another advantage is that they are intuitive and simple to construct. Essentially, they involve filling a boundary of the desired shape with a regular mesh. This can be done with little or no knowledge of the underlying physics and involves hardly any recourse to mathematics. Therefore an instrument designer using this technique could focus more on choosing materials and shape, as though they were building an instrument in reality.

Chapter 8

Future Work

8.1 Modelling Real Materials

All materials exhibit some stiffness and internal viscosity. These properties result in dispersion (frequency dependent wavespeed) and frequency dependent damping. However, for this investigation the waveguide model of the drumskin assumed that the membrane was perfectly flexible. In addition, the damping of the kettledrum model was tailored to fit the response from an actual recording. Whilst the effect of stiffness could have also been included with the addition of dispersive allpass filters, it would be preferable to model the properties of a material explicitly. By doing this the frequency dependent effects would be inherent within the model and there would be no need for the time consuming trial-and-error response matching required to tune such filters.

8.1.1 Modelling Internal Viscous Friction

It has been shown that the damping factors are invariant to a material's stiffness and shape (Djoharian, 1999). Therefore, this may be modelled separately from the other properties of a material. As a result, a material's visco-elasticity is characterised purely in terms of its storage and loss moduli. A model of a series-parallel assembly of springs and dashpots can be constructed to match

this property, and this has already been implemented in Cordis-Anima simulations and modal synthesis (see sections 2.2.4 and 2.2.7). Once modelled it should be possible to incorporate this into digital waveguide structures.

8.1.2 Modelling Bending Stiffness

In order to incorporate stiffness higher order mathematical terms have to be included within the model, whereas the digital waveguide models used in this work only solve the second order wave equation. It has been shown that stiffness in 1D media can be modelled with digital waveguide techniques by connecting two waveguides together in a special way. This inter-connection forms a Digital Waveguide Network (DWN) (Bilbao, 2001), where one waveguide carries velocity waves and the other carries curvature waves. Results from this new waveguide structure are promising, although extending it to 2D or 3D is quite complicated. However, it is possible to model varying stiffness in different directions, such as that caused by the direction of the grain in wood.

Recently, work was done to combine the theory of DWNs with the viscoelastic damping method mentioned above (Aird and Laird, 2001)(Aird, 2002), and work will be forthcoming to develop this further. Once done it will be possible to make digital waveguide models of real materials.

8.2 Extensions to the Diffusion Technique

The method proposed in section 5.2 for modelling diffusion was only applied to the boundary of a 2D digital waveguide mesh. This is quite restrictive, especially as diffusion is most likely to occur at the boundary of 3D spaces. The method did not include any frequency dependent effects either.

8.2.1 Extending to 3D

Extending the method to 3D must involve a convolution of all the incoming signals to a boundary node, but these cannot all be done at once as at least two different rotations are required to pre-alter a wave's angle of incidence: one in the horizontal plane and one in the vertical plane. In fact, the available planes of rotation will depend on the 3D mesh structure that is used and the boundary shape.

A good approach would be to find the minimum number of planes that intersect all waveguides surrounding a node at the boundary. The angle of incidence could then be changed by performing a combination of rotations in each and every plane (using the same 2D rotation method as before). It is likely that the resulting angle of incidence will be even less defined than that for the 2D method. Once again the error should be found so that the diffusion can be performed to a reasonable approximation.

8.2.2 Frequency Dependent Diffusion

In order to include frequency dependent effects in the diffusion method it will be necessary to alter the circulant matrices used in the rotation convolutions. At present the coefficients within the matrices are real numbers, acting as pure gains. If these were to be replaced with filters then the re-distribution of the signals could be done in a frequency dependent way. Great care would have to be taken to ensure that all the filters are complimentary (combinations of lowpass and highpass) in order to preserve signal strength and power.

Another factor that has an important affect on the reverberation of acoustic spaces is frequency dependent absorption at the boundaries. This has already been modelled with good accuracy by matching digital filters to air absorption characteristics and appending them to the edge of a digital waveguide mesh (Huopaniemi, Savioja, and Karjalainen, 1997). The implementation is identical to that proposed earlier in this thesis for simulating viscous friction losses

in a membrane. The absorption is also angular dependent and so it could be possible to extend the diffusion technique to incorporate this.

8.3 Modelling the Extra Drum Components

8.3.1 The Shell

The kettledrum model constructed for this work did not include any interaction with the shell. The boundaries to the drumskin and the interior air were simply modelled with reflections, which assume the shell is very hard like steel. As the shell is made from a stiff material that would also have a certain degree of internal damping, it should first be constructed using the material modelling techniques discussed in section 8.1 above. Once complete it could then be interfaced to the interior air and drumskin using the method presented in section 6.1.

It may be the case, however, that the resonance of the shell has little effect on the vibration of the system as a whole (this could be investigated using the physical model). If so, then only the shell's absorption would need to be modelled at the air and membrane boundaries.

8.3.2 Air Holes

Many drums have at least one air hole punched through the shell. The reason for this is probably to help the pressure inside the drum equalise with that on the outside more quickly. Its effect could be investigated by including it in the drum model.

Adding this to the model would involve removing the reflecting rimguides of the air mesh boundary at the location of the air hole. In their place waveguide connections of the appropriate impedance should be made to the exterior air. This impedance would be that of air in a cylinder, where the cross-sectional area of the cylinder is the area occupied by each waveguide connection.

8.3.3 Hardware Fittings

Hardware fittings are the extra components of a western drum that are there by necessity rather than design. They are required to pull the drumskin tight over the shell and to attach the drum to a stand. When taking the shell's vibrations into consideration the number and position of hardware fittings may have some effect. As these are solid fittings with very little room for movement they will hold the shell still at the point they are attached. Therefore, their best position is most likely to be at the shell's anti-nodes of resonance where there is no vibration anyway. Including their effect would involve forcing to zero the junction velocities of the relevant nodes within the shell model.

8.4 Modelling Non-Uniform Tension of the Membrane

The membrane model used in this work assumed that the tension was uniform across its surface. In reality it is impossible to achieve this as there are a finite number of tuning-lugs distributed around the rim. Between them, therefore, the rim will relax a little and so each lug will be surrounded by a region of higher tension. In addition to this the lugs have to be adjusted separately, which makes it difficult to apply the same tension at every point.

The wavespeed through the membrane is related to the tension, and so a non-uniform tension will cause different wavespeeds in different places. Perhaps one way of modelling this would be to initially draw a warped waveguide mesh with the correct boundary shape. The warping would involve positioning mesh nodes at varying distances from each other, so that the waveguides are of different lengths. In this way nodes that are closer together will cause waves to travel slower, corresponding to regions of low tension. The next step would be to make all these waveguides the same length, creating a regular mesh structure. It can be imagined that this mesh will now have a different

boundary shape to that intended, and it is this new shape that will model the non-uniform tension.

When making observations from this membrane model, or interfacing it to other meshes, it is necessary to view it in its warped form. Thus there will be a translation that maps points in space to points on the mesh.

8.5 Improvements to the Interfacing Method

The waveguide mesh interfacing method, as described in section 6.1, was limited to a zero order interpolation across the wavefront surface of the low-density mesh. This meant that the pressure or force was considered to be constant across an element centred on a mesh node. The wavefront was therefore non-continuous in that it jumped in value at each element boundary. This is clearly not the case in reality where a wavefront surface is curved and smooth.

In order to model a continuous wavefront it is necessary to interpolate between the mesh nodes. These interpolated values could then be sent to the high-density mesh across the interface. The waves reflected from the high-density mesh would need to be deinterpolated back to the nodes in the low-density mesh. Methods for interpolation and deinterpolation within a waveguide mesh have been used to help reduce the dispersion error in a previous study (Savioja and Välimäki, 1996). This would be useful for the mesh interface proposed here, but it is stressed that special care would need to be taken to ensure that energy is conserved within the system.

There is still the issue of the interfacing problem that occurs at the boundary, as demonstrated in figure 6.1b of section 6.1.1. Here it can be seen that not all the nodes in the high-density mesh connect to the low-density mesh. This was suggested as one possible reason for the mis-tuning of the lower modes in the kettledrum's simulation results (see page 102 of section 7.1). It is possible that the interpolating/deinterpolating method may also be useful in overcoming this problem and is definitely worth considering for further research.

8.6 Evaluating Models with Real Measurements

As digital waveguide techniques increase in complexity it will become necessary to evaluate the models with real measurements. For instance, when the drumshell and its stiffness is incorporated into the model its effect will have to be verified. Only by changing the material and comparing the result to reality will it be possible to make future predictions of its effect with confidence.

An important part to verify with measurement is the diffusion model, proposed in section 5.2 with suggested improvements given in this chapter. Taking measurements of the effect will involve recording how the frequency of a sound, and the material's surface off which it reflects, affects the angles into which the sound is scattered. It will also be important to relate some parameter of the surface's roughness to the strength and distribution of the diffusion, such that predictions can be made when using new materials.

Techniques already exist for measuring material properties because many materials have been investigated in the past with the intention of modelling their behaviour. Whilst this will be useful when simulating the individual components of a system, evaluating their combined response requires careful experimentation. Invaluable equipment when working with sound is a sound-proof room, high quality microphones and transducers. As the ear is the best tool for extracting information from sound it is necessary to have high fidelity speakers for audio reproduction.

Part III

Appendices and Bibliography

Appendix A

CD-ROM Guide

The CD-Rom included with this thesis contains a multimedia application, PDF and EPS electronic versions of the thesis, and the C++ waveguide model code that was written during the course of the investigation. The electronic documents are in the folder 'Thesis Documents' and the source code is in the folder 'Waveguide Model Code'. The source code was written for Borland C++ v5.02

A.1 The Multimedia Application

The multimedia application, for use on machines running Microsoft's Windows® operating system, is started by 'double clicking' on the executable file 'start.exe.' On doing so a screen should appear with three options: 'sounds', 'applications' and 'movies' (see figure A.1). The multimedia application is explored by clicking the titles on the screen. Most screens contain a 'back' title to navigate to the previous screen and a 'home' title to return to the start.

A.1.1 Sounds

The sounds section is divided further into four of the main investigations of the thesis. These are the membrane model, the bearing edge, diffusion and finally

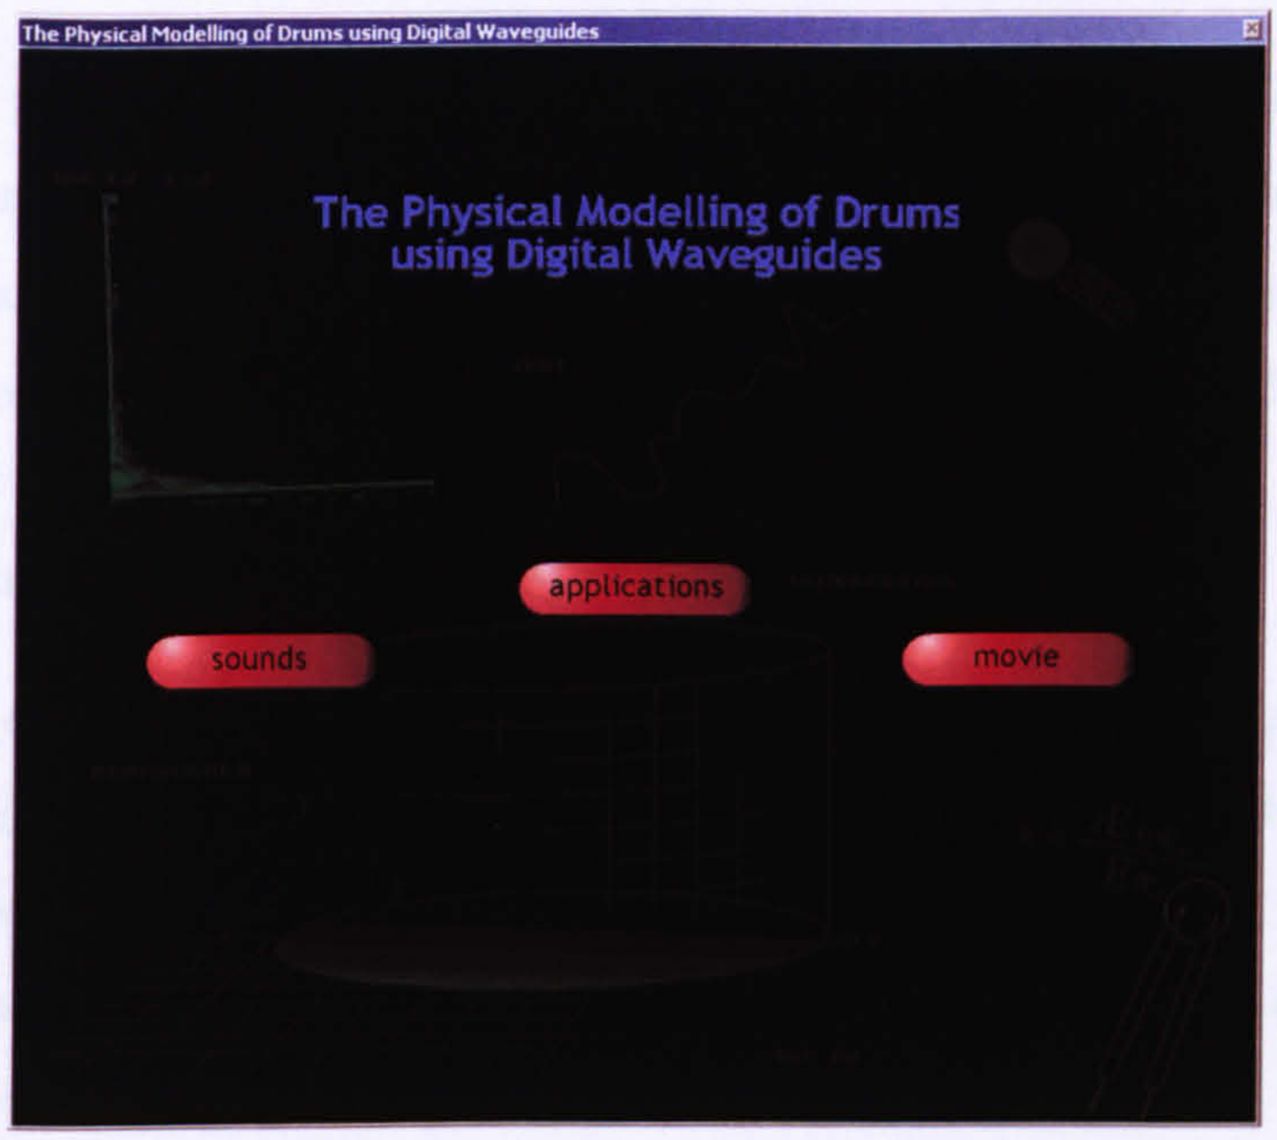

Figure A.1: *Initial Screen of the Multimedia Application.*

the complete kettledrum model. By clicking on each of these a new screen is presented that contains the relevant sounds analysed in the thesis.

Membrane Samples

The membrane samples screen has two sounds available. One is of a lossless membrane and the other has losses to simulate the effect of internal viscous friction. The models that created these sounds are described in section 4.1.

Bearing-edge Effect

This screen has nine sounds to demonstrate the effect of the bearing edge. Each sound was created with a different initial displacement and/or bearing edge radius. The sounds correspond to those that were analysed in section 4.2 of this thesis.

Diffusion Example

The sound presented here is from a membrane excited with an off-centre impulse with a diffusion model appended to the boundary. During the evolution of the sound the diffusion angle range is ramped up to a maximum value and then reduced back to zero. Section 5.2 describes the diffusion model used.

Kettledrum Samples

Section 7.1 of this thesis analysed an implementation of a kettledrum with digital waveguides. The model created was used to generate the sounds presented on this screen. In addition to the sound used in the analysis there are sounds from the same model but excited at the centre, with a drum roll and also with a membrane of lower tension.

A.1.2 Applications

There are three applications to choose from: an interactive membrane, a pressure flow simulation and a drum simulator.

Interactive Membrane

The interactive membrane application (see figure A.2) builds a digital waveguide mesh of a circular membrane using the techniques described in this thesis. The membrane can be excited at any point whilst simultaneously applying a constant pressure elsewhere. The membrane's properties can be changed and it can be viewed from any direction and distance. Full instructions are included with the application.

Pressure Flow

This application graphically demonstrates the flow of pressure within a Tom-Tom drum and plots the waveform recorded at an external listening point.

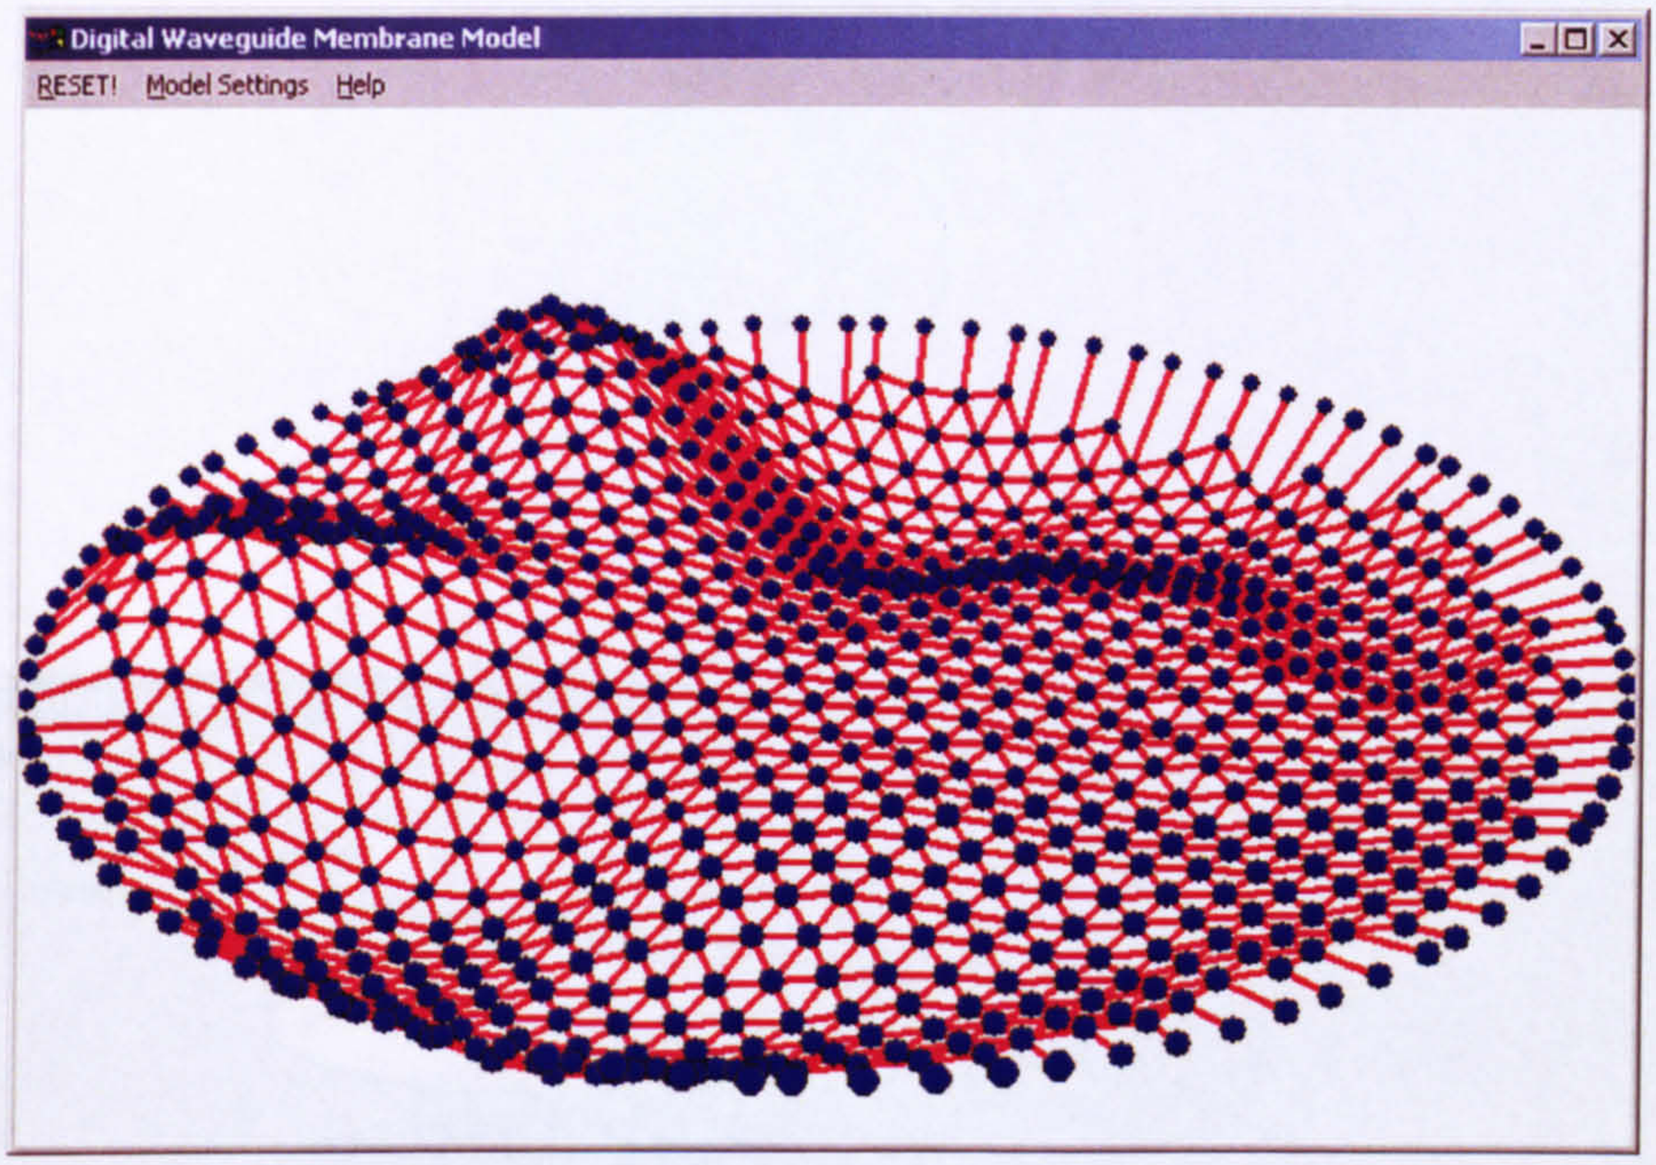

Figure A.2: *The Interactive Membrane Application.*

Once started the top membrane is automatically excited at its centre. Only the contrast of the screen can be controlled.

Drum Simulator

The drum simulator (see figure A.3) brings together most of the elements of this thesis to construct a waveguide model of a Tom-Tom drum. Every parameter can be altered and the sound generated by the model is recorded to a WAV file. The parameter settings can also be saved. The application includes a complete help file.

A.1.3 Movie

The movie clip shows an animation of a plane wave as it travels across two membranes of different density but equal impedance. The membranes are attached together using the interfacing technique described in section 6.1.

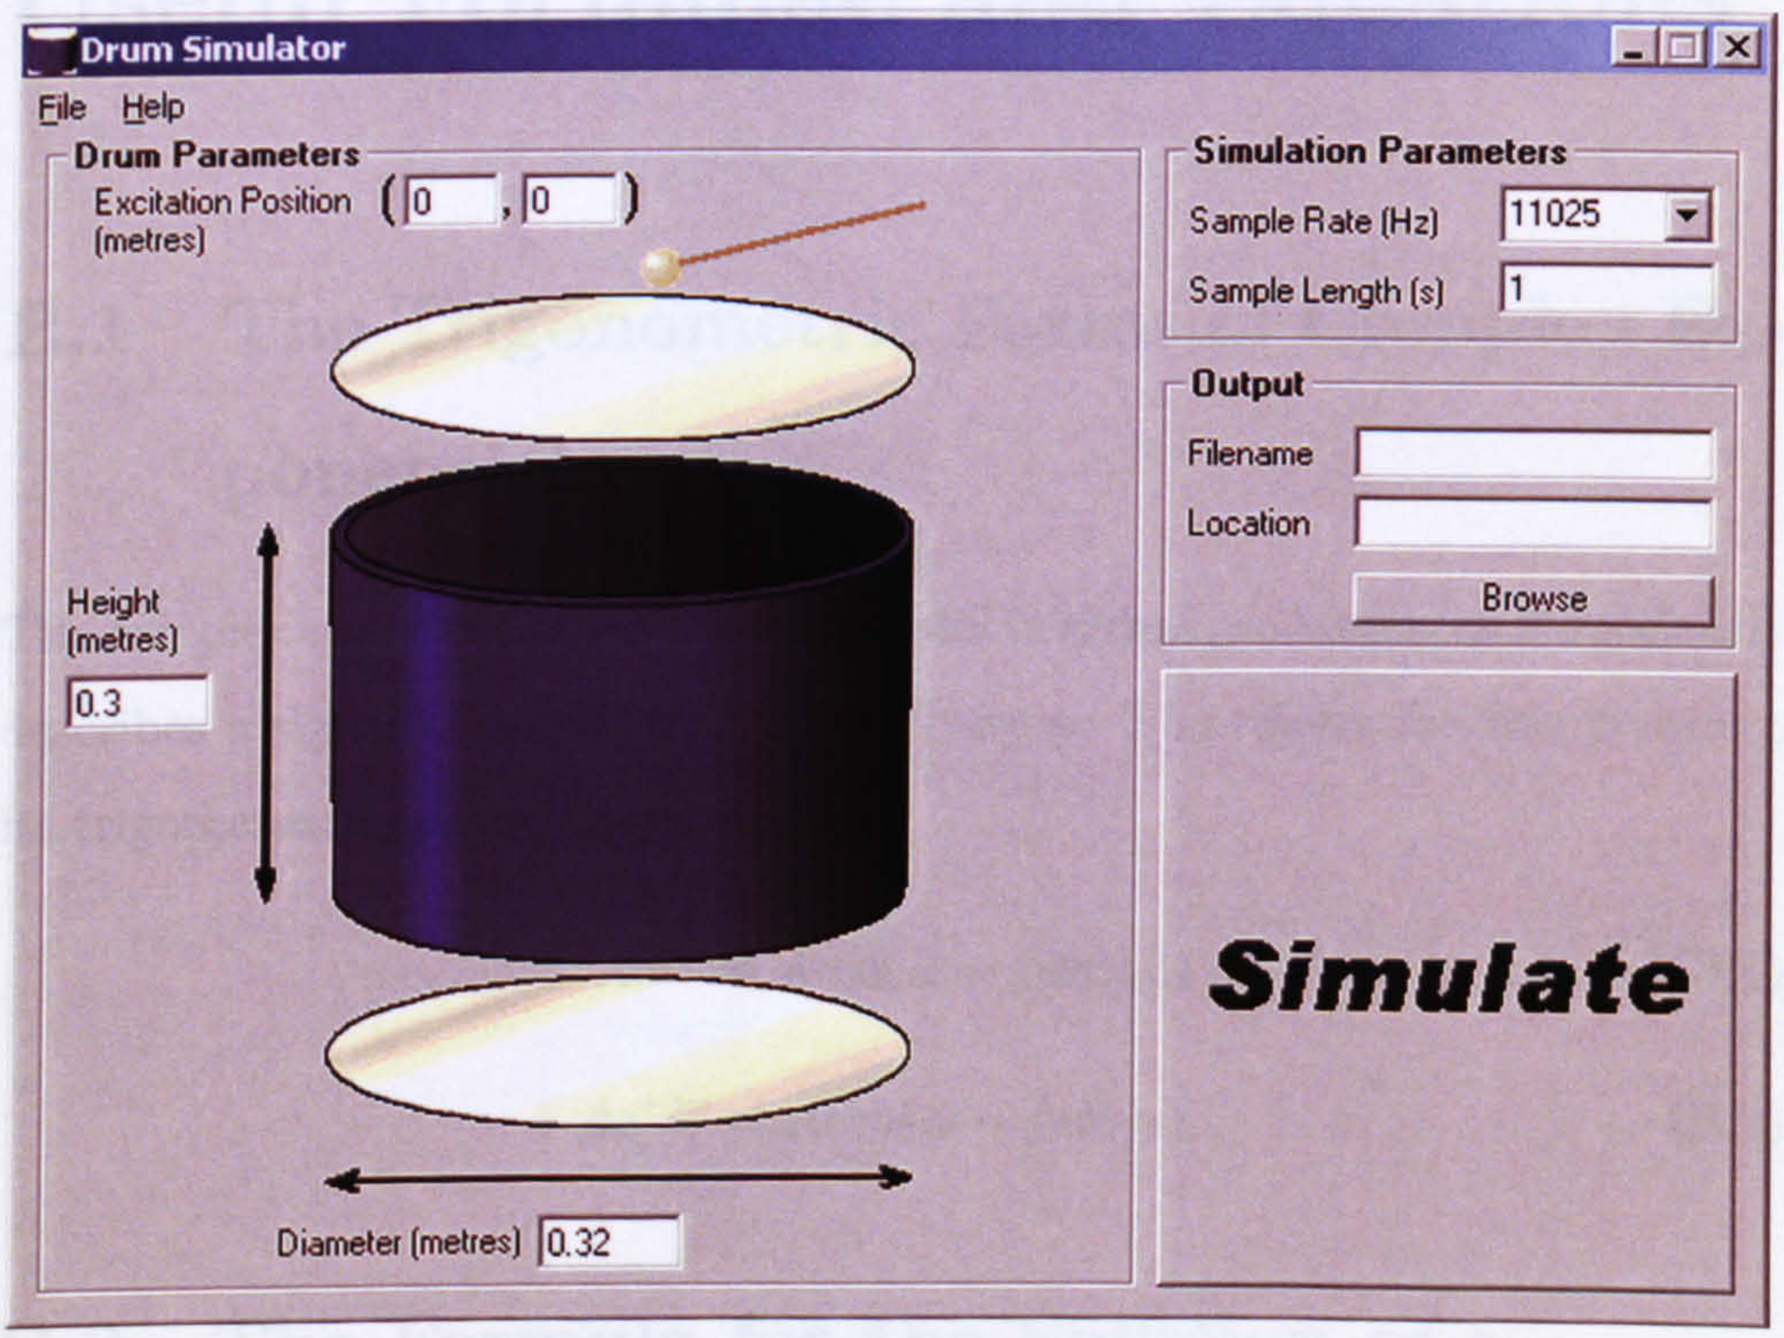

Figure A.3: *The Drum Simulator.*

Appendix B

Useful Formulae and Equations

B.1 The Trigonometric Forms of Complex Exponentials

The complex exponential $Ae^{j\omega}$ is very useful in signal processing as it succinctly describes a signal of amplitude A and phase ω. The reason for this is due to its trigonometric forms shown below

$$Ae^{j\omega} = A(\cos\omega + j\sin\omega), \tag{B.1}$$

$$Ae^{-j\omega} = A(\cos\omega - j\sin\omega). \tag{B.2}$$

B.2 The Formula for the Solution of a Quadratic

Consider the quadratic expression

$$ax^2 + bx + c = 0. \tag{B.1}$$

Its solution can be found using the formula

$$y = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}. \tag{B.2}$$

B.3 The General Gain Expression

The gain (or transfer function) of any system can be found using

$$G = \frac{\sum_k G_k \Delta_k}{\Delta}, \tag{B.1}$$

where G_k is the gain of the k^{th} forward path,

$\Delta = 1$−(sum of all the different loop gains)

+(sum of gain products of all combinations of two non-touching loops)

−(sum of gain products of all combinations of three non-touching loops)

+etc.,

$\Delta_k =$ the value of Δ for that part of the graph not touching the k^{th} forward path,

B.4 Trigonometric Identities

The following is a list of useful trigonometric identities. Despite being trivial they are important for solving many equations.

$$\cos^2\theta + \sin^2\theta = 1. \tag{B.1}$$

$$\tan(\theta_A - \theta_B) = \frac{\tan\theta_A - \tan\theta_B}{1 + \tan\theta_A \tan\theta_B}. \tag{B.2}$$

B.5 The Discrete Fourier Transforms

The Discrete Fourier Transform (DFT) and the Inverse Discrete Fourier Transform (IDFT) are often used in signal processing. One such use is to convert amplitude-time data into frequency/phase-time data and back. In section 5.2.2 the IDFT is used to compute the coefficients of a circulant matrix from its eigenvalues.

The DFT is calculated with

$$X(k) = \sum_{n=1}^{N} x(n) exp\left(\frac{-j2\pi(k-1)(n-1)}{N}\right), \tag{B.1}$$

where $1 \leq k \leq N$.

The IDFT is calculated with

$$x(n) = \frac{1}{N} \sum_{k=1}^{N} X(k) exp\left(\frac{j2\pi(k-1)(n-1)}{N}\right), \tag{B.2}$$

where $1 \leq n \leq N$.

B.6 Bessel Functions

The derivations of and solutions to Bessel functions are beyond the scope of this thesis; information on these can be found in the literature (see Abramowitz and Stegun (1972)).

Bessel functions are defined as solutions to the differential equation

$$x^2 \frac{\partial^2 y}{\partial x^2} + x \frac{\partial y}{\partial x} + \left(x^2 - n^2\right) y = 0. \tag{B.1}$$

One class of solution is called the 'Bessel function of the first kind', $J_n(x)$, and is used within this thesis for finding the resonances of a circular membrane. This solution may be expressed by the formula

$$J_n(x) = \sum_{k=0}^{\infty} \frac{(-1)^k (\frac{x}{2})^{2k+n}}{k!\Gamma(k+n+1)}, \tag{B.2}$$

where $\Gamma(z)$ is the Gamma function

$$\Gamma(z) = \int_0^{\infty} t^{z-1} e^{-t} dt. \tag{B.3}$$

Appendix C

Derivation of the Wave Equations

C.1 The 1D Wave Equation

Figure C.1 depicts part of a string that is in motion, held at tension T_0. The string is assumed to be perfectly flexible such that it does not exhibit any resistance to bending. Consider the two arbitrary points at horizontal distance x and $x+\Delta x$. If the slope was constant then the tensions at these points would be equal in magnitude and opposite in direction. They would therefore cancel eachother to make zero net force and acceleration. However, if the string had

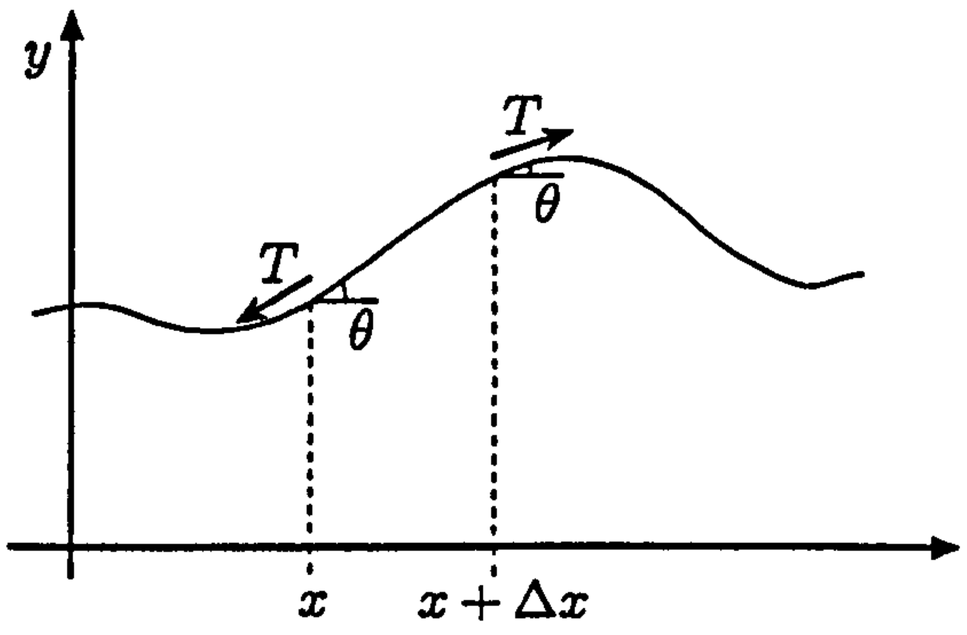

Figure C.1: *Part of a String in Motion.*

curvature then the tension would act in slightly different directions, marked as T, which combine to make a restoring force. As the tension is parallel to the slope of the string, it is possible to calculate the horizontal and vertical components of this force.

The net horizontal force is

$$F_x = T(x+\Delta x)\cos\theta(x+\Delta x) - T(x)\cos\theta(x), \tag{C.1}$$

and the vertical force

$$F_y = T(x+\Delta x)\sin\theta(x+\Delta x) - T(x)\sin\theta(x). \tag{C.2}$$

For simplicity it is desirable to restrict the motion to the vertical direction. This can be done by making the tension constant throughout the string, imposing the following condition along its length

$$T\cos\theta = T_0, \tag{C.3}$$

to make $F_x = 0$.

Inserting equation C.3 into C.2

$$F_y = T_0\tan\theta(x+\Delta x) - T_0\tan\theta(x), \tag{C.4}$$

$$= T_0\left(\frac{\partial y(x+\Delta x)}{\partial x} - \frac{\partial y(x)}{\partial x}\right). \tag{C.5}$$

Solutions to the above two derivatives can be approximated using the Taylor Series Expansion formula

$$f(x) \approx f(x_0) + (x-x_0)f_x(x_0) + \frac{x-x_0}{2!}f_{xx}(x_0) + \frac{x-x_0}{3!}f_{xxx}(x_0) + \ldots, \tag{C.6}$$

where the x subscript denotes a partial derivative.

Setting $f(x) = \frac{\partial y(x)}{\partial x}$ results in

$$f(x_0+\Delta x) \approx \frac{\partial y(x_0)}{\partial x} + \Delta x\frac{\partial^2 y(x_0)}{\partial x^2} + \ldots, \tag{C.7}$$

$$f(x_0) \approx \frac{\partial y(x_0)}{\partial x}. \tag{C.8}$$

All terms of third order or higher can be discarded as Δx is intended to be very small. Also, if small-amplitude waves are assumed, $\theta \ll 1$, then $cos\theta \simeq 1$ and $T \simeq T_0$ at any point along the string. So, inserting equations C.7 and C.8 into C.5 gives

$$F_y = T\Delta x \frac{\partial^2 y}{\partial x^2}. \tag{C.9}$$

The force F_y will accelerate the section of string according to Newton's Second Law of motion

$$F = ma. \tag{C.10}$$

As the mass of the section is $\Delta m = \mu \Delta x$, where μ is the linear density, the relationship between force and acceleration can be written as

$$\mu \Delta x \frac{\partial^2 y}{\partial t^2} = T\Delta x \frac{\partial^2 y}{\partial x^2}. \tag{C.11}$$

Rearranging and cancelling Δx

$$\frac{\partial^2 y}{\partial t^2} = \frac{T}{\mu} \frac{\partial^2 y}{\partial x^2}. \tag{C.12}$$

Notice that $\frac{T}{\mu}$ has the dimensions

$$\frac{[N][m]}{[kg]} = \frac{[kg][m]^2}{[kg][s]^2}, \tag{C.13}$$

$$= \frac{[m]^2}{[s]^2}. \tag{C.14}$$

This is obviously the square of some velocity, that can be defined as

$$c^2 = \frac{T}{\mu}. \tag{C.15}$$

Inserting this into equation C.12 results in the general formulation of the second order one dimensional wave equation

$$\frac{\partial^2 y}{\partial t^2} = c^2 \frac{\partial^2 y}{\partial x^2}. \tag{C.16}$$

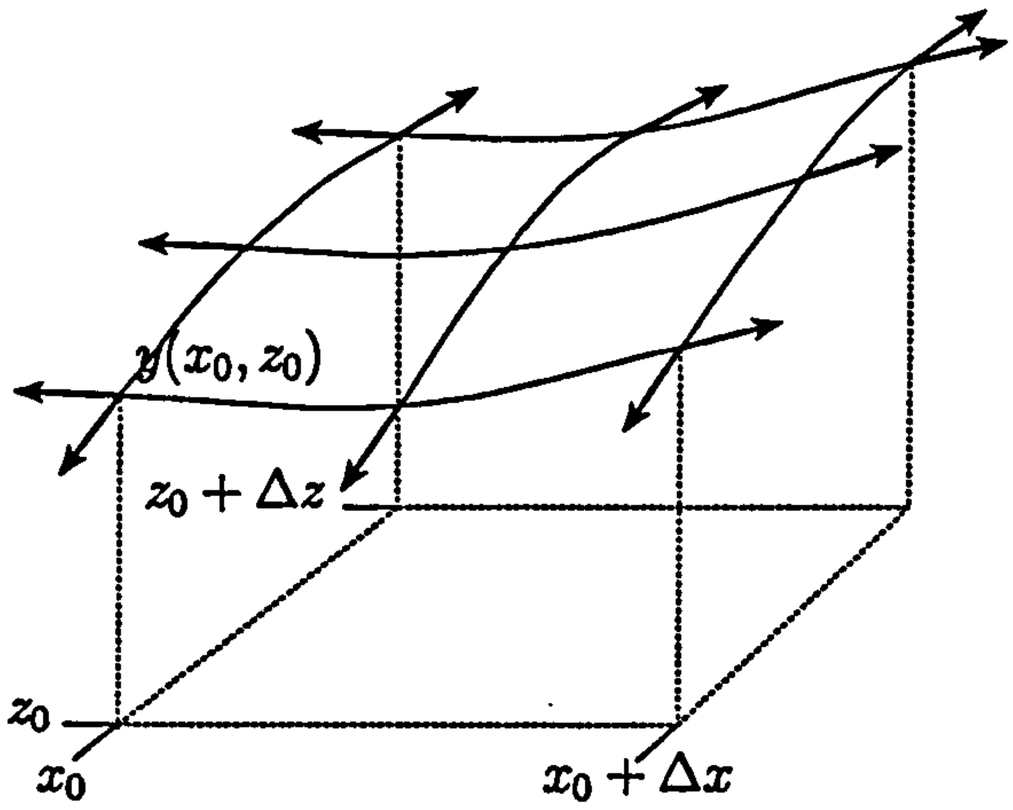

Figure C.2: *Part of a Membrane in Motion.*

C.2 The 2D Wave Equation

The derivation of the two dimensional wave equation follows a similar process to the one dimensional case, as discussed in the previous section. Figure C.2 depicts part of a membrane that is in motion. Because the restoring force is applied at the membrane boundary rather than at specific points, the measurement of tension, T_0, is in force per unit length (units $\frac{N}{m}$).

As in the case of the string, it is useful to restrict the motion to the vertical direction so that T anywhere within the boundary is the same as T_0. Once again this allows the substitution of $T\sin\theta$ with $T_0\tan\theta = T_0\frac{\partial y}{\partial x}$ (see equations C.4 and C.5 in appendix C.1).

The total vertical force along the boundary needs to be found, and this can be done by combining each side's contribution. At an arbitrary point x the combined vertical force component of the sides z_0 and $z_0 + \Delta z$, which is effectively that for a string held in the y-z plane, is defined as

$$F(x) = Tdx\left(\frac{\partial y(x, z_0 + \Delta z)}{\partial z} - \frac{\partial y(x, z_0)}{\partial z}\right). \tag{C.17}$$

This is an adapted version of equation C.5 from the last section. Note the

term Tdx is now used to express the tension, as T has units force per unit length. The total force along these sides is the integral of this expression from x_0 to $x_0 + \Delta x$.

As before, the two derivatives can be approximated using the Taylor Series Expansion. Now we have two variables, so the expansion formula is

$$\begin{aligned} f(x,z) \approx \; & f(x_0,z_0) + (x-x_0)f_x(x_0,z_0) + (z-z_0)f_z(x_0,z_0) + \\ & \tfrac{(x-x_0)(z-z_0)}{2!}\left[f_{xx}(x_0,y_0) + 2f_{xz}(x_0,y_0) + f_{yy}(x_0,y_0)\right] + \ldots \end{aligned} \tag{C.18}$$

Therefore equation C.17 becomes (ignoring the leftmost derivative for now by setting $z = z_0$)

$$F(x,z_0) \approx -Tdx\left(y_z(x_0,z_0) + (x-x_0)y_{zx}(x_0,z_0)\right), \tag{C.19}$$

where the subscripts z and x denote partial derivatives. Once again, the terms of third order or higher have been discarded.

The integral from x_0 to $x_0 + \Delta x$ is

$$-T\left[xy_z(x_0,z_0) + \left(\frac{x^2}{2} - x_0x\right)y_{zx}(x_0,z_0)\right]_{x=x_0}^{x=x_0+\Delta x}, \tag{C.20}$$

$$\begin{aligned} = -T\left((x_0+\Delta x)y_z + \left(\tfrac{(x_0+\Delta x)^2}{2} - x_0(x_0+\Delta x)\right)y_{zx}\right) + \\ T\left(x_0y_z + \left(\tfrac{x_0^2}{2} - x_0^2\right)y_{zx}\right), \end{aligned} \tag{C.21}$$

$$\begin{aligned} = -T\left((x_0+\Delta x)y_z + \left(\tfrac{x_0^2+\Delta x^2+2x_0\Delta x}{2} - \tfrac{2x_0^2+2x_0\Delta x}{2}\right)y_{zx}\right) + \\ T\left(x_0y_z - \left(\tfrac{x_0^2}{2}\right)y_{zx}\right), \end{aligned} \tag{C.22}$$

$$= -T\left(\Delta xy_z + \frac{\Delta x^2}{2}y_{zx}\right). \tag{C.23}$$

If the rightmost derivative is now ignored, by setting $z = z_0 + \Delta z$, the Taylor Series Expansion of equation C.17 is very similar to C.19, except that it is positive and has an extra term

$$F(x,z_0+\Delta z) \approx Tdx\left(y_z(x_0,z_0) + (x-x_0)y_{zx}(x_0,z_0) + \Delta zy_{zz}\right). \tag{C.24}$$

Naturally, the similar terms of equations C.19 and C.24 will cancel eachother out, resulting in the total force contributed by the two sides z_0 and $z_0 + \Delta z$

$$T[\Delta z x y_{zz}]_{x=x_0}^{x=x_0+\Delta x}, \tag{C.25}$$

$$= T((x_0 + \Delta x)\Delta z y_{zz} - x_0 \Delta z y_{zz}), \tag{C.26}$$

$$= T\Delta x \Delta z y_{zz}, \tag{C.27}$$

$$F(x) = T\Delta x \Delta z \frac{\partial^2 y}{\partial z^2}. \tag{C.28}$$

By following the same process, the combined force of the sides x_0 and $x_0 + \Delta x$ can also be found

$$F(z) = T\Delta x \Delta z \frac{\partial^2 y}{\partial x^2}. \tag{C.29}$$

Adding these two results together produces the total force along the boundary

$$\begin{aligned} F &= T\Delta x \Delta z \left(\tfrac{\partial^2 y}{\partial z^2} + \tfrac{\partial^2 y}{\partial x^2}\right), \\ &= T\Delta x \Delta z \nabla^2 y. \end{aligned} \tag{C.30}$$

The mass of the section of membrane is $\Delta m = \sigma \Delta x \Delta z$, where σ is the area mass density. By substituting this and the result of C.30 into Newton's Second Law, the equation of motion for the membrane is found to be

$$\sigma \Delta x \Delta z \frac{\partial^2 y}{\partial t^2} = T\Delta x \Delta z \nabla^2 y, \tag{C.31}$$

$$\frac{\partial^2 y}{\partial t^2} = \frac{T}{\sigma} \nabla^2 y. \tag{C.32}$$

Notice the similarity between this and equation C.12 in the previous section. Just as with $\frac{T}{\mu}$ before, a velocity can be defined

$$c^2 = \frac{T}{\sigma}, \tag{C.33}$$

remembering here that T is measured in force per unit length and σ is mass per unit area.

Inserting this into equation C.32 results in the general formulation of the second order two dimensional wave equation

$$\frac{\partial^2 y}{\partial t^2} = c^2 \nabla^2 y. \tag{C.34}$$

C.3 The 3D Wave Equation

Wave propagation in three dimensions differs from the previous two cases in that the travelling component is a change in pressure, not displacement. In addition the pressure change disperses into all directions and does so as a longitudinal wave. Nevertheless, a similar process is happening as with wave travel in one and two dimensions.

This is best visualised as a cube which has a force and resulting displacement applied to one face. This displacement will increase the pressure within the cube such that the other faces are forced against the neighbouring space. Just as the tension in a 2D membrane applies a restoring force to any displacement, the neighbouring space applies a restoring force to the displaced faces of the cube. It follows that the three dimensional wave equation is a natural progression of terms that becomes (Hall, 1987g)

$$\frac{\partial^2 p}{\partial t^2} = c^2 \nabla^2 p, \tag{C.35}$$

where ∇^2 is the combination of the 3D variables: $\frac{\partial^2 p}{\partial x^2} + \frac{\partial^2 p}{\partial y^2} + \frac{\partial^2 p}{\partial z^2}$.

Appendix D

The Lossless, Series Scattering Junction

D.1 Formulation of the Scattering Junction Equation

A series scattering junction (see section 3.7.1) connects waveguides together and is defined by the following conditions

$$v_i = v_J, \tag{D.1}$$

$$\sum_i^N f_i = 0, \tag{D.2}$$

where v_i is the velocity of waveguide i,

v_J is the velocity of the junction,

f_i is the force exerted by waveguide i,

N is the number of waveguides attached to the junction.

This means that all the velocities are equal at the junction and that there is no net force.

The waveguide variables may be split into opposing travelling waves and re-

lated by their impedances, R_i

$$v_i = v_i^+ + v_i^-, \tag{D.3}$$

$$f_i = f_i^+ + f_i^-, \tag{D.4}$$

$$f^+ = R_i v_i^+, \tag{D.5}$$

$$f^- = -R_i v_i^-. \tag{D.6}$$

Combining equations D.2 and D.4 gives

$$\sum_i^N \left(f_i^+ + f_i^-\right) = 0. \tag{D.7}$$

If equations D.5 and D.6 are inserted into the above and relationship D.3 is applied then the following can be deduced

$$\sum_i^N R_i \left(v_i^+ - v_i^-\right) = 0, \tag{D.8}$$

$$\sum_i^N R_i \left(2v_i^+ - v_i\right) = 0, \tag{D.9}$$

$$2\sum_i^N R_i v_i^+ = \sum_i^N R_i v_i. \tag{D.10}$$

As each velocity is equal to the junction velocity this may be rewritten in the standard form of the lossless scattering junction

$$v_J = \frac{2\sum_i^N R_i v_i^+}{\sum_i^N R_i}. \tag{D.11}$$

D.2 The Impedance of the Scattering Junction

It is possible to write the scattering junction equation from the previous section in terms of v_J and v_i^- instead. In this case it works out to be almost identical except for the change in superscript from + to −

$$v_J = \frac{2\sum_i^N R_i v_i^-}{\sum_i^N R_i}. \tag{D.12}$$

By inserting either equation D.5 or D.6 the scattering equations can also be written in terms of the total travelling forces (in this case the total outgoing force)

$$v_J = \frac{-2\sum_i^N f_i^-}{\sum_i^N R_i}. \tag{D.13}$$

It follows, therefore, that the total incoming force at the junction is equal in magnitude yet opposite in sign to the total outgoing force. The forces are analogous to the 'action' and 'reaction' forces of Newton's Third Law. In addition, they are both related to the junction velocity by multiplication with the constant

$$\frac{2}{\sum_i^N R_i}. \tag{D.14}$$

Alternatively the junction velocity may be related to either force by multiplication with the inverse of the above. The constant that relates velocity to force in this way is known as the impedance, so the impedance of a mesh junction is

$$R_{mesh} = \frac{\sum_i^N R_i}{2}. \tag{D.15}$$

For homogeneous waveguide meshes, where the waveguides all have the same impedance, this reduces to

$$R_{mesh} = \frac{N}{2} R_{waveguide}. \tag{D.16}$$

This means that the total impedance experienced at the junction is half the total impedance of all the waveguides it connects together. This makes sense, because a junction that connects two waveguides of equal impedance to form a straight line is in fact a point on one long waveguide. The impedance at that point should be the impedance of either waveguide, not twice the value. For an even number of waveguides it is as though the junction is comprised of half as many waveguides which intersect each other at their middles, rather than many waveguides arranged around the junction in a symmetrical pattern.

D.3 Applying a Force to the Scattering Junction

In order to connect excitation mechanisms to a waveguide model it is necessary to be able to apply a force to a scattering junction. By considering Newton's Third Law it can be stated that the force of the junction must be opposite but equal to the applied force; in this way all the forces cancel each other to produce no net force. If all forces are defined as acting in the same direction equation D.2 can be rewritten to give

$$\sum_{i}^{N} f_i = -F, \tag{D.17}$$

where F is the applied force at the junction.

After following through the formulation of the lossless scattering junction equation its modified form, that incorporates the external force, F, is found to be

$$v_J = \frac{2\sum_i^N R_i v_i^+ + F}{\sum_i^N R_i}. \tag{D.18}$$

Appendix E

Calculation of Dispersion Error in Digital Waveguide Meshes

E.1 The Triangular Digital Waveguide Mesh

Consider the segment of a triangular mesh centered about the point $A(0,0)$, see figure E.1 below. It has surrounding junctions: $B(-1,0)$, $C(-\cos\frac{\pi}{3}, \sin\frac{\pi}{3})$, $D(\cos\frac{\pi}{3}, \sin\frac{\pi}{3})$, $E(1,0)$, $F(\cos\frac{\pi}{3}, -\sin\frac{\pi}{3})$, $G(-\cos\frac{\pi}{3}, -\sin\frac{\pi}{3})$.

The equations relating these lossless 6-port scattering junctions are

$$v_A = \frac{1}{3}\sum_{\Gamma} v_A^{\Gamma+}, \tag{E.1}$$

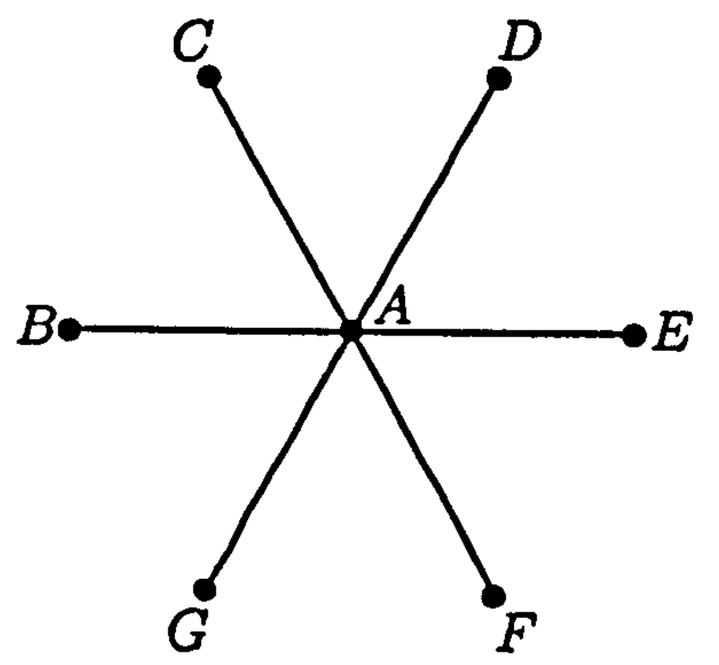

Figure E.1: *Junctions in a Triangular Digital Waveguide Mesh. The lines represent waveguides.*

$$v_A^{\Gamma-} = v_A - v_A^{\Gamma+}, \tag{E.2}$$

where $\Gamma \in \{B, C, D, E, F, G\}$,

v_A represents the junction velocity at point A,

$v_A^{\Gamma+}$ and $v_A^{\Gamma-}$ represent the input and output signals to junction A from junction Γ.

As the junctions are connected by bi-directional units of delay (waveguides) the input to junction A from the direction Γ is equal to the output from Γ delayed by one sample

$$v_A^{\Gamma+} = z^{-1} v_\Gamma^{A-}. \tag{E.3}$$

Using equations E.3 and E.2

$$v_A^{\Gamma+} = z^{-1}\left(v_\Gamma - v_\Gamma^{A+}\right), \tag{E.4}$$

$$= z^{-1}\left(v_\Gamma - z^{-1}\left(v_A - v_A^{\Gamma+}\right)\right), \tag{E.5}$$

$$= z^{-1}\left(v_\Gamma - z^{-1}v_A + z^{-1}v_A^{\Gamma+}\right), \tag{E.6}$$

$$= z^{-1}v_\Gamma - z^{-2}v_A + z^{-2}v_A^{\Gamma+}, \tag{E.7}$$

$$v_A^{\Gamma+}\left(1 - z^{-2}\right) = z^{-1}\left(v_\Gamma - z^{-1}v_A\right), \tag{E.8}$$

$$v_A^{\Gamma+} = \frac{z^{-1}}{1 - z^{-2}}\left(v_\Gamma - z^{-1}v_A\right). \tag{E.9}$$

Inserting equation E.9 into E.1 results in an expression for v_A entirely in terms of the six surrounding junction velocities

$$v_A = \frac{1}{3}\sum_\Gamma \left(\frac{z^{-1}}{1 - z^{-2}}\left(v_\Gamma - z^{-1}v_A\right)\right), \tag{E.10}$$

$$= \frac{1}{3}\frac{z^{-1}}{1 - z^{-2}}\sum_\Gamma \left(v_\Gamma - z^{-1}v_A\right), \tag{E.11}$$

$$= \frac{1}{3}\frac{z^{-1}}{1 - z^{-2}}\left(\sum_\Gamma v_\Gamma - 6z^{-1}v_A\right), \tag{E.12}$$

$$= \frac{1}{3}\frac{z^{-1}}{1 - z^{-2}}\sum_\Gamma v_\Gamma - \frac{2z^{-2}}{1 - z^{-2}}v_A, \tag{E.13}$$

$$v_A\left(1+\frac{2z^{-2}}{1-z^{-2}}\right)=\frac{1}{3}\frac{z^{-1}}{1-z^{-2}}\sum_{\Gamma}v_\Gamma, \tag{E.14}$$

$$v_A\frac{1+z^{-2}}{1-z^{-2}}=\frac{1}{3}\frac{z^{-1}}{1-z^{-2}}\sum_{\Gamma}v_\Gamma, \tag{E.15}$$

$$v_A=\frac{1}{3}\frac{z^{-1}}{1+z^{-2}}\sum_{\Gamma}v_\Gamma. \tag{E.16}$$

To find the dispersion error the spatial positions of the junctions are replaced with their corresponding linear phase terms (Duyne and Smith(III), 1996)

$$v_\Gamma \leftrightarrow v\left(\underline{\omega}\right)e^{j\underline{P}_\Gamma^T\cdot\underline{\omega}}, \tag{E.17}$$

where $\underline{\omega}$ is the two-dimensional spatial frequency vector,

$\underline{P}_\Gamma^T$ is the transposed vector from point A to its surrounding points:

$\underline{P}_1=(-1,0)^T$, $\underline{P}_2=(-\cos\frac{\pi}{3},\sin\frac{\pi}{3})^T$, $\underline{P}_3=(\cos\frac{\pi}{3},\sin\frac{\pi}{3})^T$,

$\underline{P}_4=(1,0)^T$, $\underline{P}_5=(\cos\frac{\pi}{3},-\sin\frac{\pi}{3})^T$, $\underline{P}_6=(-\cos\frac{\pi}{3},-\sin\frac{\pi}{3})^T$.

Equation E.16 is therefore rewritten

$$v\left(\underline{\omega}\right)=\frac{1}{3}\frac{z^{-1}}{1+z^{-2}}v\left(\underline{\omega}\right)\sum_{i=1}^{6}e^{j\underline{P}_i^T\cdot\underline{\omega}}. \tag{E.18}$$

This can be rearranged to give the quadratic expression in z^{-1}

$$1+z^{-2}=\frac{1}{3}z^{-1}\sum_{i=1}^{6}e^{j\underline{P}_i^T\cdot\underline{\omega}}, \tag{E.19}$$

$$z^{-2}-\frac{1}{3}z^{-1}\sum_{i=1}^{6}e^{j\underline{P}_i^T\cdot\underline{\omega}}+1=0, \tag{E.20}$$

$$z^{-2}+bz^{-1}+1=0, \tag{E.21}$$

where $b\triangleq-\frac{1}{3}\sum_{i=1}^{6}e^{j\underline{P}_i^T\cdot\underline{\omega}}$.

Because of the relationships: $\underline{P}_1=-\underline{P}_4$, $\underline{P}_2=-\underline{P}_5$, $\underline{P}_3=-\underline{P}_6$ and the trigonometric form of complex exponentials (see section B.1), the summation may be rewritten as

$$\sum_{i=1}^{6}e^{j\underline{P}_i^T\cdot\underline{\omega}}=2\sum_{i=4}^{6}\cos\left(\underline{P}_i^T\cdot\underline{\omega}\right). \tag{E.22}$$

If $\underline{\omega} = (\omega_x, \omega_y)^T$ then b may be calculated as

$$b = -\frac{2}{3}\left(\cos\omega_x + \cos(0.5\omega_x - \sqrt{0.75}\omega_y) + \cos(-0.5\omega_x - \sqrt{0.75}\omega_y)\right). \quad \text{(E.23)}$$

Equation E.21 is solved using the formula for the solution of a quadratic (see appendix B.2) to find the spectral amplification factor of the spatial spectrum after one time sample

$$G(\underline{\omega}) = -\frac{b}{2} \pm \frac{j\sqrt{4-b^2}}{2}. \quad \text{(E.24)}$$

The real gain is

$$|G| = \sqrt{\frac{b^2}{4} + \frac{4-b^2}{4}}, \quad \text{(E.25)}$$

$$= 1, \quad \text{(E.26)}$$

which means that plane waves propagate with no loss.

The spatial phase shift (taking the positive root) of a plane wave in one time sample is

$$\angle G = \arctan\frac{\sqrt{4-b^2}}{-b}. \quad \text{(E.27)}$$

The frequency and direction dependent speed of plane wave travel, in space samples per time sample, is therefore

$$c(\underline{\omega}) = \frac{\angle G}{|\underline{\omega}|}. \quad \text{(E.28)}$$

E.2 Calculation of the DC Wavespeed of The Dodecahedral Digital Waveguide Mesh

The DC wavespeed of a two dimensional triangular digital waveguide mesh has been calculated as $c = \frac{1}{\sqrt{2}}$ spatial samples per time sample (Fontana and Rocchesso, 1998). The DC wavespeed of the three dimensional dodecahedral mesh may be calculated in a similar fashion.

As well as the orthogonal axes x, y and z one may define the dependent axes l, m, n, o and p for the dodecahedral mesh. The vectors of these axes are

$$\underline{l} = \left(\tfrac{1}{2}, \tfrac{\sqrt{3}}{2}, 0\right)^T, \underline{m} = \left(\tfrac{1}{2}, -\tfrac{\sqrt{3}}{2}, 0\right)^T, \underline{n} = \left(0, \tfrac{1}{\sqrt{3}}, \tfrac{\sqrt{2}}{\sqrt{3}}\right)^T,$$
$$\underline{o} = \left(-\tfrac{1}{2}, -\tfrac{1}{2\sqrt{3}}, \tfrac{\sqrt{2}}{\sqrt{3}}\right)^T, \underline{p} = \left(\tfrac{1}{2}, -\tfrac{1}{2\sqrt{3}}, \tfrac{\sqrt{2}}{\sqrt{3}}\right)^T.$$

The partial derivatives along these axes are therefore related to the partial derivatives along the orthogonal axes through the equations

$$\frac{\partial v}{\partial l} = \frac{1}{2}\frac{\partial v}{\partial x} + \frac{\sqrt{3}}{2}\frac{\partial v}{\partial y}, \tag{E.29}$$

$$\frac{\partial v}{\partial m} = \frac{1}{2}\frac{\partial v}{\partial x} - \frac{\sqrt{3}}{2}\frac{\partial v}{\partial y}, \tag{E.30}$$

$$\frac{\partial v}{\partial n} = \frac{1}{\sqrt{3}}\frac{\partial v}{\partial y} + \frac{\sqrt{2}}{\sqrt{3}}\frac{\partial v}{\partial z}, \tag{E.31}$$

$$\frac{\partial v}{\partial o} = -\frac{1}{2}\frac{\partial v}{\partial x} - \frac{1}{\sqrt{12}}\frac{\partial v}{\partial y} + \frac{\sqrt{2}}{\sqrt{3}}\frac{\partial v}{\partial z}, \tag{E.32}$$

$$\frac{\partial v}{\partial p} = \frac{1}{2}\frac{\partial v}{\partial x} - \frac{1}{\sqrt{12}}\frac{\partial v}{\partial y} + \frac{\sqrt{2}}{\sqrt{3}}\frac{\partial v}{\partial z}. \tag{E.33}$$

A combination of these partial derivatives is required that will equate to $\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} + \frac{\partial^2 v}{\partial z^2}$. As a trial solution, consider $\frac{\partial^2 v}{\partial l^2} + \frac{\partial^2 v}{\partial m^2} + \frac{\partial^2 v}{\partial n^2} + \frac{\partial^2 v}{\partial o^2} + \frac{\partial^2 v}{\partial p^2}$

$$\begin{aligned}\frac{\partial^2 v}{\partial l^2} + \frac{\partial^2 v}{\partial m^2} + \frac{\partial^2 v}{\partial n^2} + \frac{\partial^2 v}{\partial o^2} + \frac{\partial^2 v}{\partial p^2} = {} & \tfrac{1}{4}\tfrac{\partial^2 v}{\partial x^2} + \tfrac{3}{4}\tfrac{\partial^2 v}{\partial y^2} + \tfrac{\sqrt{3}}{2}\tfrac{\partial^2 v}{\partial x \partial y} + \\ & \tfrac{1}{4}\tfrac{\partial^2 v}{\partial x^2} + \tfrac{3}{4}\tfrac{\partial^2 v}{\partial y^2} - \tfrac{\sqrt{3}}{2}\tfrac{\partial^2 v}{\partial x \partial y} + \\ & \tfrac{1}{3}\tfrac{\partial^2 v}{\partial y^2} + \tfrac{2}{3}\tfrac{\partial^2 v}{\partial z^2} + \tfrac{2\sqrt{2}}{3}\tfrac{\partial^2 v}{\partial y \partial z} + \\ & \tfrac{1}{4}\tfrac{\partial^2 v}{\partial x^2} + \tfrac{1}{12}\tfrac{\partial^2 v}{\partial y^2} + \tfrac{2}{3}\tfrac{\partial^2 v}{\partial z^2} + \\ & \tfrac{1}{\sqrt{12}}\tfrac{\partial^2 v}{\partial x \partial y} - \tfrac{\sqrt{2}}{\sqrt{3}}\tfrac{\partial^2 v}{\partial x \partial z} - \tfrac{\sqrt{2}}{3}\tfrac{\partial^2 v}{\partial y \partial z} + \\ & \tfrac{1}{4}\tfrac{\partial^2 v}{\partial x^2} + \tfrac{1}{12}\tfrac{\partial^2 v}{\partial y^2} + \tfrac{2}{3}\tfrac{\partial^2 v}{\partial z^2} - \\ & \tfrac{1}{\sqrt{12}}\tfrac{\partial^2 v}{\partial x \partial y} + \tfrac{\sqrt{2}}{\sqrt{3}}\tfrac{\partial^2 v}{\partial x \partial z} - \tfrac{\sqrt{2}}{3}\tfrac{\partial^2 v}{\partial y \partial z}, \end{aligned} \tag{E.34}$$

$$= \frac{\partial^2 v}{\partial x^2} + 2\frac{\partial^2 v}{\partial y^2} + 2\frac{\partial^2 v}{\partial z^2}. \tag{E.35}$$

Therefore

$$\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} + \frac{\partial^2 v}{\partial z^2} = \frac{1}{2}\left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial l^2} + \frac{\partial^2 v}{\partial m^2} + \frac{\partial^2 v}{\partial n^2} + \frac{\partial^2 v}{\partial o^2} + \frac{\partial^2 v}{\partial p^2}\right). \tag{E.36}$$

Combining this with the three dimensional wave equation (see appendix C.3 equation C.35) results in the following

$$\frac{\partial^2 v}{\partial t^2} = c^2 \frac{1}{2}\left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial l^2} + \frac{\partial^2 v}{\partial m^2} + \frac{\partial^2 v}{\partial n^2} + \frac{\partial^2 v}{\partial o^2} + \frac{\partial^2 v}{\partial p^2}\right). \tag{E.37}$$

Now, from equation E.16 in the previous section, it can be seen that the signal at a junction A may be expressed entirely in terms of the surrounding junctions Γ. For the dodecahedral mesh the expression becomes

$$\left(1+z^{-2}\right) V_A = \frac{1}{6}\sum_{\Gamma} z^{-1} V_\Gamma. \tag{E.38}$$

The inverse Z-transform of this, after multiplying by z, is

$$V_A(n+1) + V_A(n-1) = \frac{1}{6}\sum_{\Gamma} V_\Gamma(n). \tag{E.39}$$

The left member is obviously the second order finite time difference, and the right member is the second order finite space difference along the axes x, l, m, n, o and p.

Comparing equation E.37 with equation E.39 proves the equivalence of the mesh algorithm with a finite difference scheme, whose nominal or DC wavespeed is found to be

$$\frac{c^2}{2} = \frac{1}{6}, \tag{E.40}$$

$$c = \frac{1}{\sqrt{3}}\,(\textit{spatial samples per time sample}). \tag{E.41}$$

It is interesting to note that the DC wavespeed of the three dimensional tetrahedral waveguide mesh is also $c = \frac{1}{\sqrt{3}}$ (Duyne and Smith(III), 1996). Considering that the two dimensional triangular waveguide mesh has a DC wavespeed of $c = \frac{1}{\sqrt{2}}$ (Fontana and Rocchesso, 1998) could it not be possible that the wavespeed in any waveguide mesh is

$$c = \frac{1}{\sqrt{N}}, \tag{E.42}$$

where N is the number of spatial dimensions that the mesh exists within?

The same relationship (rewritten in standard units) has been stated for multi-dimensional rectilinear meshes (Savioja, Karjalainen, and Takala, 1996). Here it is in the form

$$f = \frac{c\sqrt{N}}{dx}, \tag{E.43}$$

where c is the wavespeed (ms^{-1}),

dx is the spatial distance of a waveguide (m).

Appendix F

Filter Design and Analysis

Digital filters play a prominent role in digital waveguide modelling. Their frequency dependent effects on the gain and phase of a signal are particularly useful when modelling phenomena such as viscous friction and dispersion (see section 3.5). It is therefore important to be able to construct, analyse and control them. The following concentrates on two simple but effective filters: the first order allpass filter and the first order lowpass filter.

F.1 The First Order Allpass Filter

Allpass filters have the important quality of unity gain for all frequencies; in other words they are totally lossless. This allows the filter designer to concentrate purely on tailoring the phase response.

For digital waveguide modelling the phase may be used to cause a delay that is not an integer number of time samples (Laakso, Välimäki, Karjalainen, and Laine, 1996). This enables the construction of non-integer length waveguides and the modelling of dispersion. The first order implementation is the simplest.

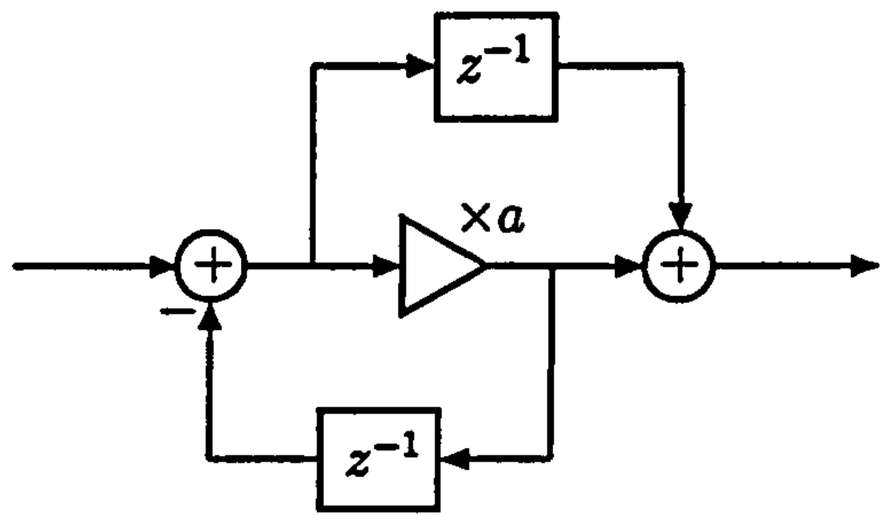

Figure F.1: *The Structure of a First Order Allpass Filter.*

F.1.1 Transfer Function

Figure F.1 depicts a first order allpass filter. It is made up from two forward paths, see figures F.2*a* and F.2*b*, and one loop path as shown in figure F.2*c*. The general gain expression from section B.3 finds the variables: $G_1 = a$, $G_2 = z^{-1}$, $\Delta_1 = \Delta_2 = 1$, $\Delta = 1 + az^{-1}$ resulting in the filter's transfer function

$$H(z) = \frac{a + z^{-1}}{1 + az^{-1}}. \tag{F.1}$$

F.1.2 Response

The effect of this filter on an input signal can be found by replacing z with the arbitrary signal $e^{j\omega}$

$$H\left(e^{j\omega}\right) = \frac{a + e^{-j\omega}}{1 + ae^{-j\omega}}, \tag{F.2}$$

and using the trigonometric form of complex exponentials (see appendix B.1)

$$H\left(e^{j\omega}\right) = \frac{a + \cos\omega - j\sin\omega}{1 + a(\cos\omega - j\sin\omega)}. \tag{F.3}$$

Gain

The gain response can be found by calculating the absolute values of the complex quantities

$$\left|H\left(e^{j\omega}\right)\right| = \frac{\sqrt{(a + \cos\omega)^2 + (-\sin\omega)^2}}{\sqrt{(1 + a\cos\omega)^2 + (-a\sin\omega)^2}}, \tag{F.4}$$

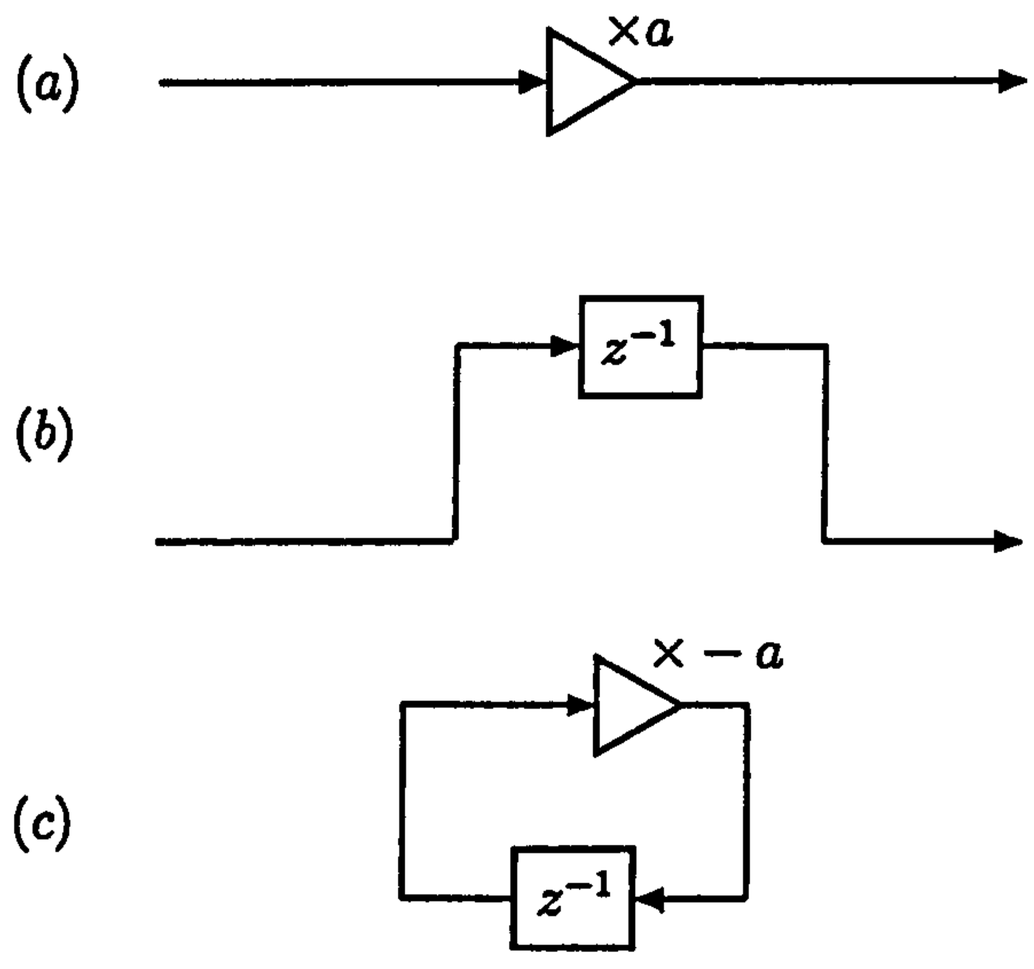

Figure F.2: *The Paths of the First Order Allpass Filter.* (a) *&* (b) *Forward Paths,* (c) *Loop Path.*

$$= \sqrt{\frac{a^2 + 2a\cos\omega + \cos^2\omega + \sin^2\omega}{1 + 2a\cos\omega + a^2\cos^2\omega + a^2\sin^2\omega}}, \tag{F.5}$$

$$= \sqrt{\frac{a^2 + 2a\cos\omega + 1}{a^2 + 2a\cos\omega + 1}}, \tag{F.6}$$

$$= 1. \tag{F.7}$$

Hence the filter has a unity gain for all frequencies.

Phase

The phase response is found by calculating the angle difference between the two complex quantities

$$\phi = \arctan\left(\frac{-\sin\omega}{a + \cos\omega}\right) - \arctan\left(\frac{-a\sin\omega}{1 + a\cos\omega}\right). \tag{F.8}$$

Using trigonometric identity B.2 in appendix B.4

$$\tan\phi = \frac{\tan\phi_A - \tan\phi_B}{1 + \tan\phi_A \tan\phi_B}, \tag{F.9}$$

where $\phi_A = \arctan\left(\frac{-\sin\omega}{a+\cos\omega}\right)$,

$\phi_B = \arctan\left(\frac{-a\sin\omega}{1+a\cos\omega}\right)$.

Inserting the variables ϕ_A and ϕ_B into equation F.9

$$\tan\phi = \frac{u-v}{1+uv}, \tag{F.10}$$

where $u = \frac{-\sin\omega}{a+\cos\omega}$,

$v = \frac{-a\sin\omega}{1+a\cos\omega}$.

As there are a lot of mathematical terms it is best to consider the denominator and numerator of equation F.10 separately. Firstly, solving $u-v$

$$u - v = \frac{-\sin\omega}{a+\cos\omega} + \frac{a\sin\omega}{1+a\cos\omega}, \tag{F.11}$$

$$= \frac{-(1+a\cos\omega)\sin\omega + (a+\cos\omega)a\sin\omega}{(a+\cos\omega)(1+a\cos\omega)}, \tag{F.12}$$

$$= \frac{a^2\sin\omega + a\cos\omega\sin\omega - \sin\omega - a\cos\omega\sin\omega}{a+\cos\omega+a^2\cos\omega+a\cos^2\omega}, \tag{F.13}$$

$$= \frac{(a^2-1)\sin\omega}{a+(1+a\cos\omega+a^2)\cos\omega}. \tag{F.14}$$

Next, solving $1+uv$

$$1 + uv = 1 + \frac{-\sin\omega}{a+\cos\omega}\cdot\frac{-a\sin\omega}{1+a\cos\omega}, \tag{F.15}$$

$$= 1 + \frac{a\sin^2\omega}{a+(1+a\cos\omega+a^2)\cos\omega}. \tag{F.16}$$

Inserting equations F.14 and F.16 into equation F.10

$$\tan\phi = \frac{(a^2-1)\sin\omega}{a+(1+a\cos\omega+a^2)\cos\omega}\cdot\frac{a+(1+a\cos\omega+a^2)\cos\omega}{a+(1+a\cos\omega+a^2)\cos\omega+a\sin^2\omega}, \tag{F.17}$$

$$= \frac{(a^2-1)\sin\omega}{a+\cos\omega+a^2\cos\omega+a\cos^2\omega+a\sin^2\omega}, \tag{F.18}$$

$$\phi = \arctan\left(\frac{(a^2-1)\sin\omega}{2a+(1+a^2)\cos\omega}\right). \tag{F.19}$$

Phase Delay

The phase delay for this filter can be found with

$$D(\omega) = \frac{-\phi}{\omega}, \tag{F.20}$$

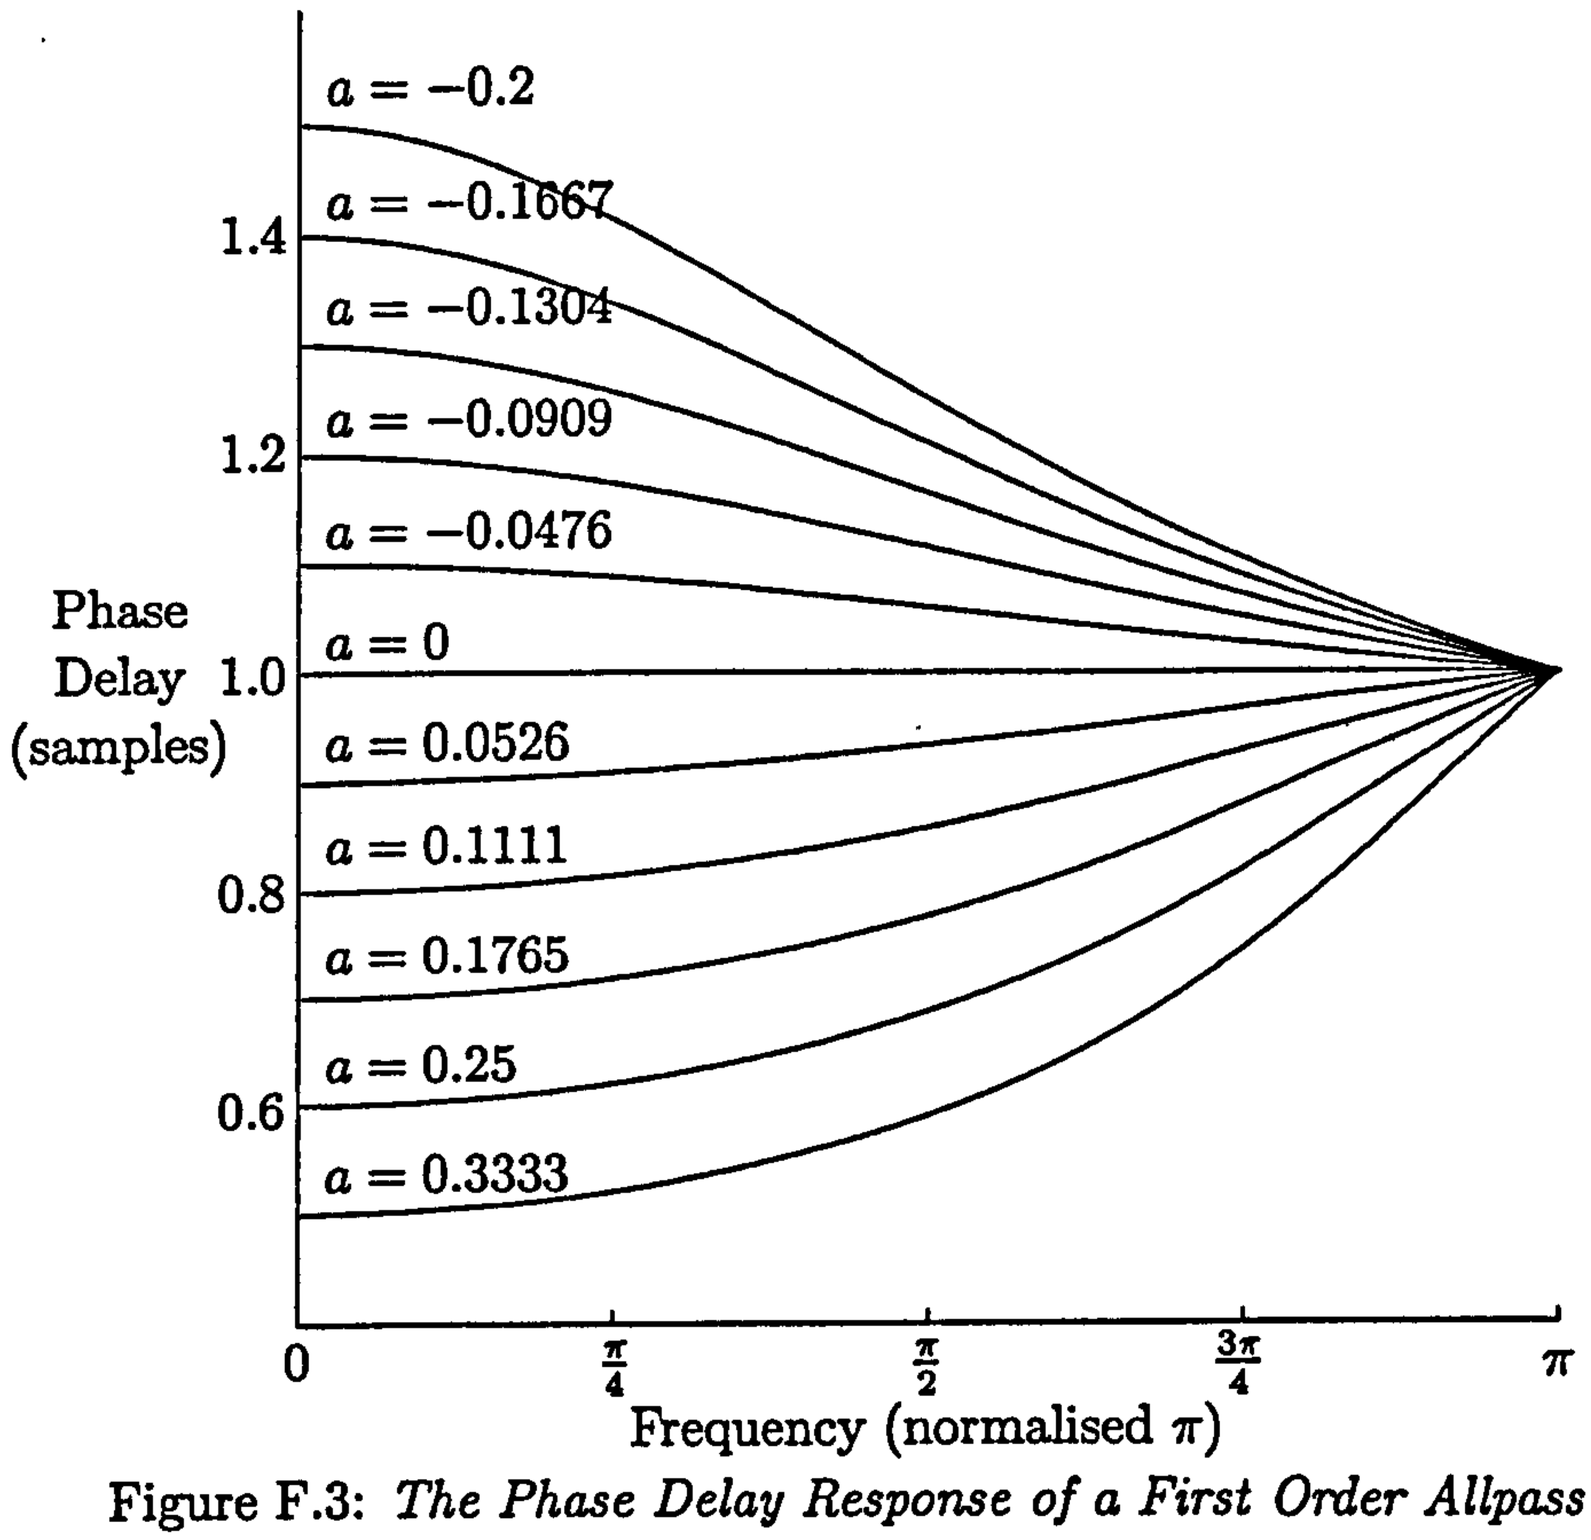

Figure F.3: *The Phase Delay Response of a First Order Allpass Filter.*

where $D(\omega)$ is the phase delay (samples).

Figure F.3 shows the phase delay of the first order allpass filter at different coefficient values. Notice that a fractional delay is approximated well for low frequencies, but the highest frequencies all have a delay close to one sample in length.

F.1.3 Control

In many cases it is desirable to control the filter such that it has a specific phase delay at a chosen frequency. This can be done by combining equation F.20 with equation F.19 and rearranging to give an expression for the filter coefficient, a, in terms of the phase delay, $D(\omega)$, and frequency, ω. The combined equations

give

$$B = \frac{(a^2-1)\sin\omega}{2a+(1+a^2)\cos\omega}, \tag{F.21}$$

where $B = \tan(-D(\omega)\omega)$.

Rearranging

$$2Ba + B\cos\omega + Ba^2\cos\omega = a^2\sin\omega - \sin\omega, \tag{F.22}$$

$$(B\cos\omega - \sin\omega)a^2 + 2Ba + B\cos\omega + \sin\omega = 0. \tag{F.23}$$

Using the formula for the solution of a quadratic (see appendix B.2) and rearranging further

$$a = \frac{-2B \pm \sqrt{4B^2 - 4(B\cos\omega - \sin\omega)(B\cos\omega + \sin\omega)}}{2(B\cos\omega - \sin\omega)}, \tag{F.24}$$

$$= \frac{2B \pm \sqrt{4B^2 - 4(B^2\cos^2\omega - \sin^2\omega)}}{2(\sin\omega - B\cos\omega)}, \tag{F.25}$$

$$= \frac{2B \pm \sqrt{4(B^2(1-\cos^2\omega) + \sin^2\omega)}}{2(\sin\omega - B\cos\omega)}, \tag{F.26}$$

$$= \frac{2B \pm 2\sqrt{B^2\sin^2\omega + \sin^2\omega}}{2(\sin\omega - B\cos\omega)}. \tag{F.27}$$

$$\Rightarrow a = \frac{B \pm \sqrt{(B^2+1)\sin^2\omega}}{\sin\omega - B\cos\omega}. \tag{F.28}$$

Alternatively, it has been shown that a particular analytic solution for the coefficients of a lowpass filter (Thiran, 1971) can be applied to allpass filter design to give maximally flat group delay response at zero frequency (Laakso, Välimäki, Karjalainen, and Laine, 1992). This leads to the following approximation of the filter coefficient, a, required to cause a phase delay, D, at low frequencies (Välimäki, 1995)(Jaffe and Smith(III), 1983)

$$a \approx \frac{1-D}{1+D}. \tag{F.29}$$

F.2 The First Order IIR Lowpass Filter

Unlike an allpass filter a lowpass filter has a frequency dependent gain. This gain is close to unity for low frequencies and tends towards zero for the highest frequencies; the result is that low frequencies pass through the filter with little or no alteration. These filters are useful in digital waveguide modelling where a frequency dependent gain mimics the effect of internal viscous friction (see section 3.5).

There are two possible types of implementation for a lowpass filter: finite impulse response (FIR) and infinite impulse response (IIR). FIR filters have no feedback loops and are therefore easier to construct and design. However, IIR filters yield a more useful response for less implementational cost. The filter described here is a first order IIR lowpass filter.

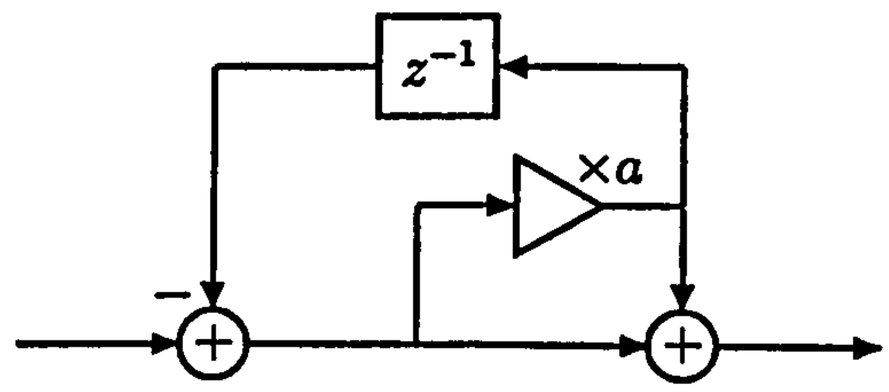

Figure F.4: *The Structure of a First Order IIR Lowpass Filter.*

F.2.1 Transfer Function

Figure F.4 depicts a first order IIR lowpass filter. It is made up from two forward paths, see figures F.5*a* and F.5*b*, and one loop path as shown in figure F.5*c*.

The general gain expression from section B.3 finds the variables: $G_1 = 1$, $G_2 = a$, $\Delta_1 = \Delta_2 = 1$, $\Delta = 1 + az^{-1}$ resulting in the filter's transfer function

$$H(z) = \frac{1+a}{1+az^{-1}}. \tag{F.30}$$

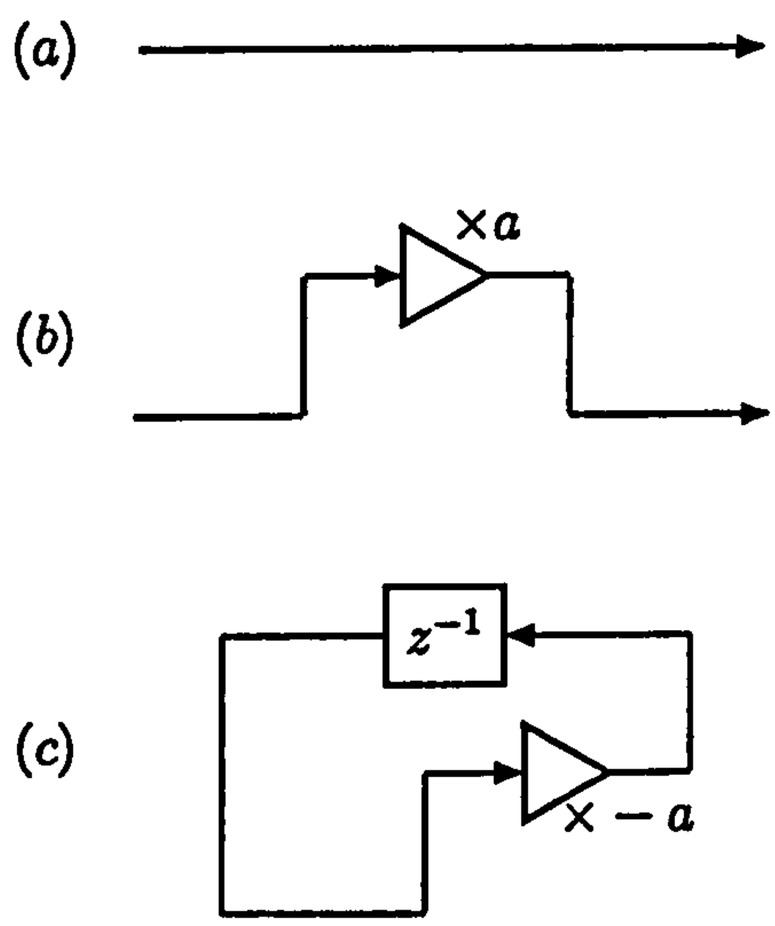

Figure F.5: *The Paths of the First Order IIR Lowpass Filter.* (a) *&* (b) *Forward Paths,* (c) *Loop Path.*

F.2.2 Response

The effect of this filter on an input signal can be found in a similar fashion to that described for an allpass filter. Replacing z with the arbitrary signal $e^{j\omega}$ yields

$$H\left(e^{j\omega}\right) = \frac{1+a}{1+ae^{-j\omega}}, \tag{F.31}$$

and using the trigonometric form of complex exponentials (see appendix B.1)

$$H\left(e^{j\omega}\right) = \frac{1+a}{1+a(\cos\omega - j\sin\omega)}. \tag{F.32}$$

Gain

Calculating the absolute values of the complex quantities results in the gain response

$$\left|H\left(e^{j\omega}\right)\right| = \frac{1+a}{\sqrt{(1+a\cos\omega)^2+(-a\sin\omega)^2}}, \tag{F.33}$$

$$= \frac{1+a}{\sqrt{1+2a\cos\omega+a^2\cos^2\omega+a^2\sin^2\omega}}, \tag{F.34}$$

$$G = \frac{1+a}{\sqrt{1+2a\cos\omega+a^2}}. \tag{F.35}$$

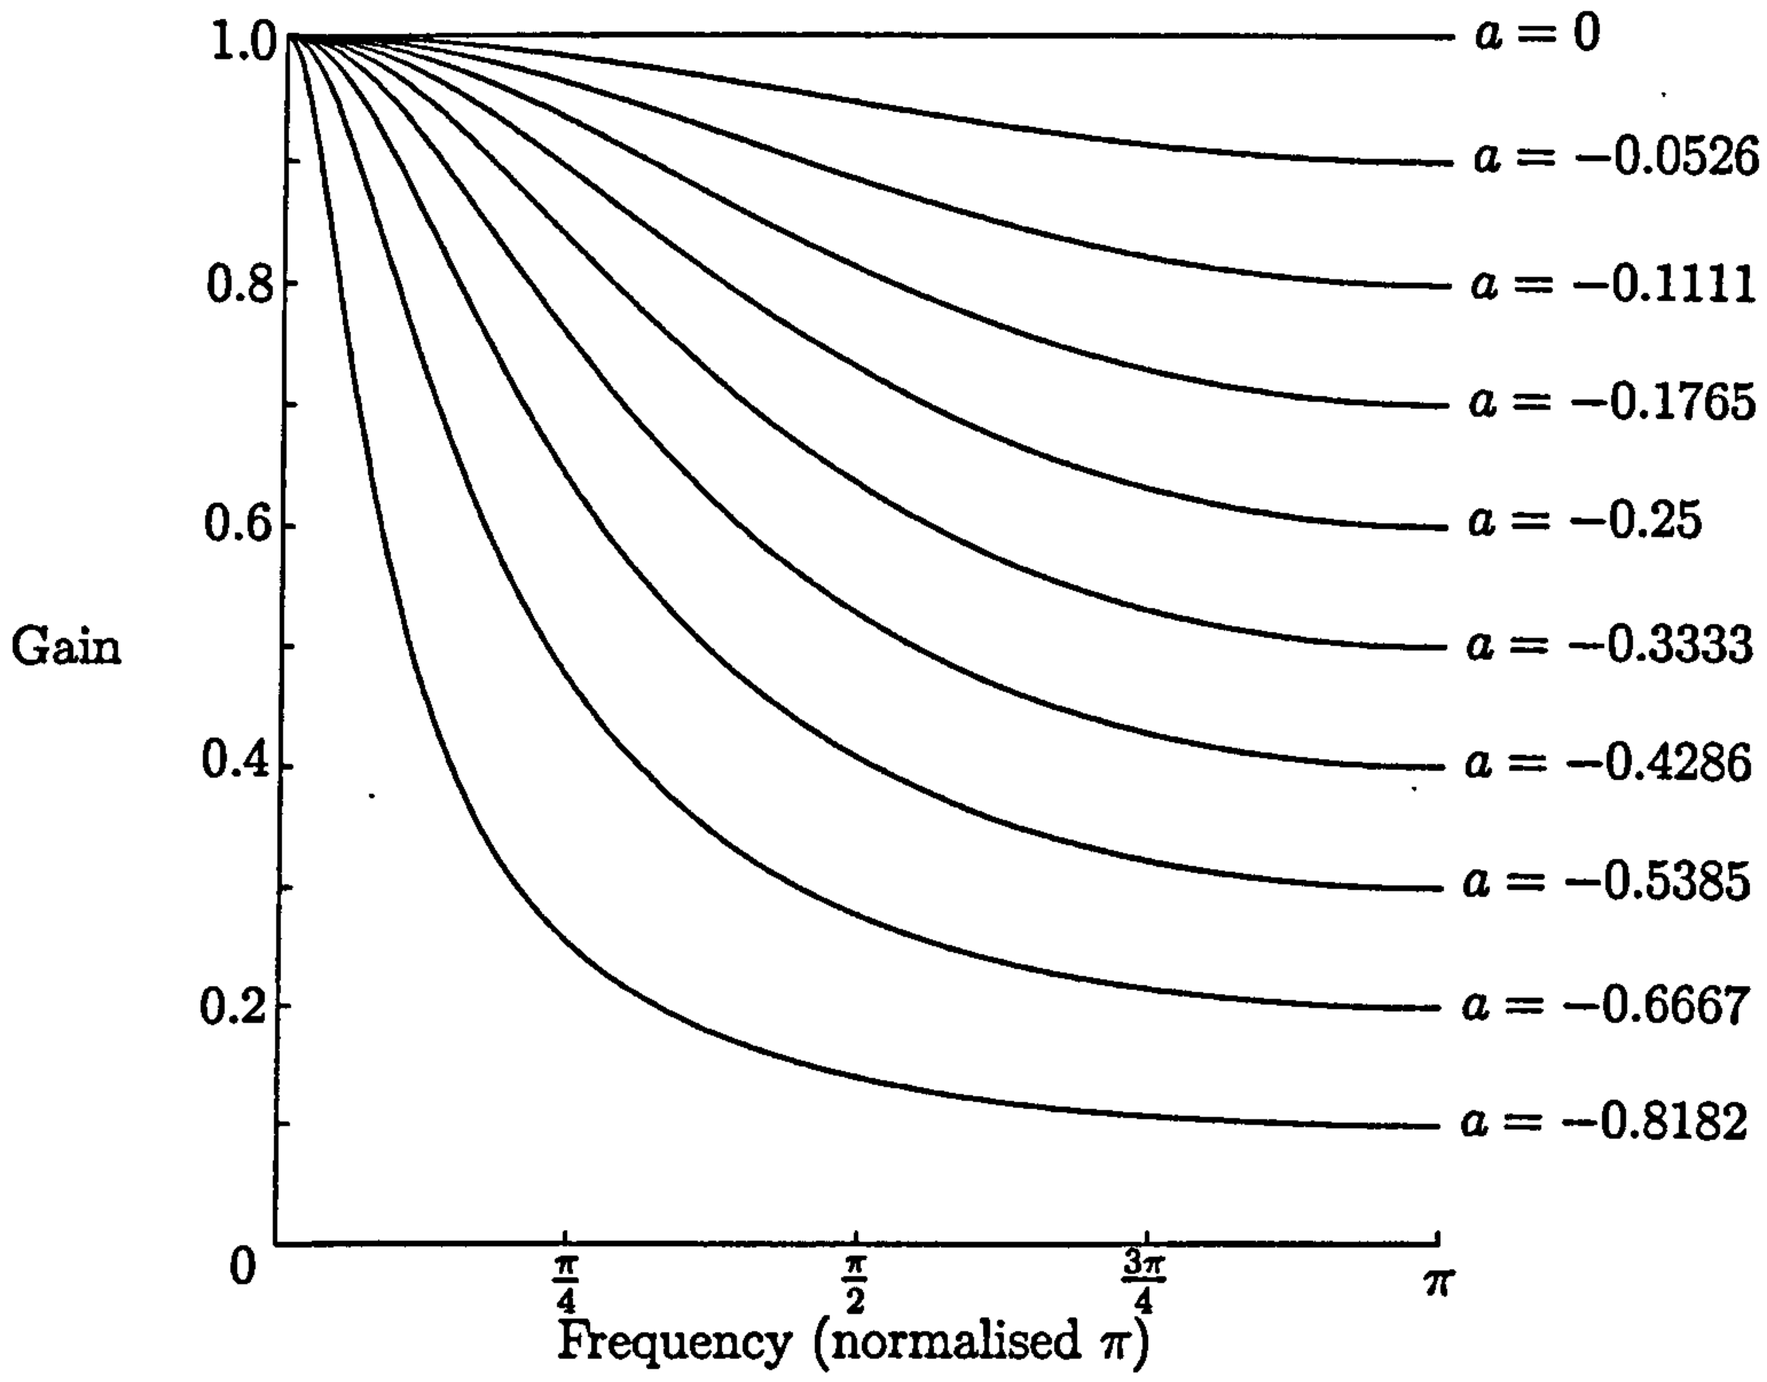

Figure F.6: *The Gain Response of a First Order IIR Lowpass Filter.*

Figure F.6 shows the gain of the first order IIR lowpass filter at different coefficient values. Notice that for waves close to DC the gain is unity.

Phase

As before the phase response is found by calculating the angle difference between the phases of the numerator and denominator of the filter's complex response (equation F.32)

$$\phi = \arctan\left(\frac{0}{1+a}\right) - \arctan\left(\frac{-a\sin\omega}{1+a\cos\omega}\right), \tag{F.36}$$

$$= 0 - \arctan\left(\frac{-a\sin\omega}{1+a\cos\omega}\right), \tag{F.37}$$

$$\phi = -\arctan\left(\frac{-a\sin\omega}{1+a\cos\omega}\right). \tag{F.38}$$

Phase Delay

The phase delay for this filter is found with

$$D(\omega) = \frac{-\phi}{\omega}, \tag{F.39}$$

where $D(\omega)$ is the phase delay (samples).

F.2.3 Control

It is useful to be able to control the filter such that it has a particular gain at a specific frequency. This can be done by rearranging equation F.35 to give an expression for the filter coefficient, a, in terms of the desired gain, G, and frequency, ω.

$$G^2\left(1 + 2a\cos\omega + a^2\right) = 1 + 2a + a^2, \tag{F.40}$$

$$a^2G^2 + 2aG^2\cos\omega + G^2 - a^2 - 2a - 1 = 0, \tag{F.41}$$

$$(G^2 - 1)a^2 + 2(G^2\cos\omega - 1)a + G^2 - 1 = 0. \tag{F.42}$$

Using the formula for the solution of a quadratic (see appendix B.2) and rearranging further

$$a = \frac{2(1 - G^2\cos\omega) \pm \sqrt{(2G^2\cos\omega - 2)^2 - 4(G^2 - 1)(G^2 - 1)}}{2(G^2 - 1)}, \tag{F.43}$$

$$= \frac{2(1 - G^2\cos\omega) \pm \sqrt{4G^4\cos^2\omega - 8G^2\cos\omega + 4 - 4(G^4 - 2G^2 + 1)}}{2(G^2 - 1)}, \tag{F.44}$$

$$= \frac{2(1 - G^2\cos\omega) \pm \sqrt{4(G^4(\cos^2\omega - 1) + G^2(2 - 2\cos\omega))}}{2(G^2 - 1)}, \tag{F.45}$$

$$= \frac{2(1 - G^2\cos\omega) \pm 2\sqrt{G^2(G^2(\cos^2\omega - 1) + 2 - 2\cos\omega)}}{2(G^2 - 1)}, \tag{F.46}$$

$$\Rightarrow a = \frac{1 - G\left(G\cos\omega \pm \sqrt{G^2(\cos^2\omega - 1) + 2 - 2\cos\omega}\right)}{G^2 - 1}. \tag{F.47}$$

Appendix G

Formulation of Equations used to model the Bearing Edge

G.1 Formulation of the Relationship between $d(n)$ and $l(n)$

Figure G.1 shows a string attached to a bearing edge at a particular point in time. The known points, lengths and angles are: $B(0, d)$, r, x, l_1, l_2, α and β.

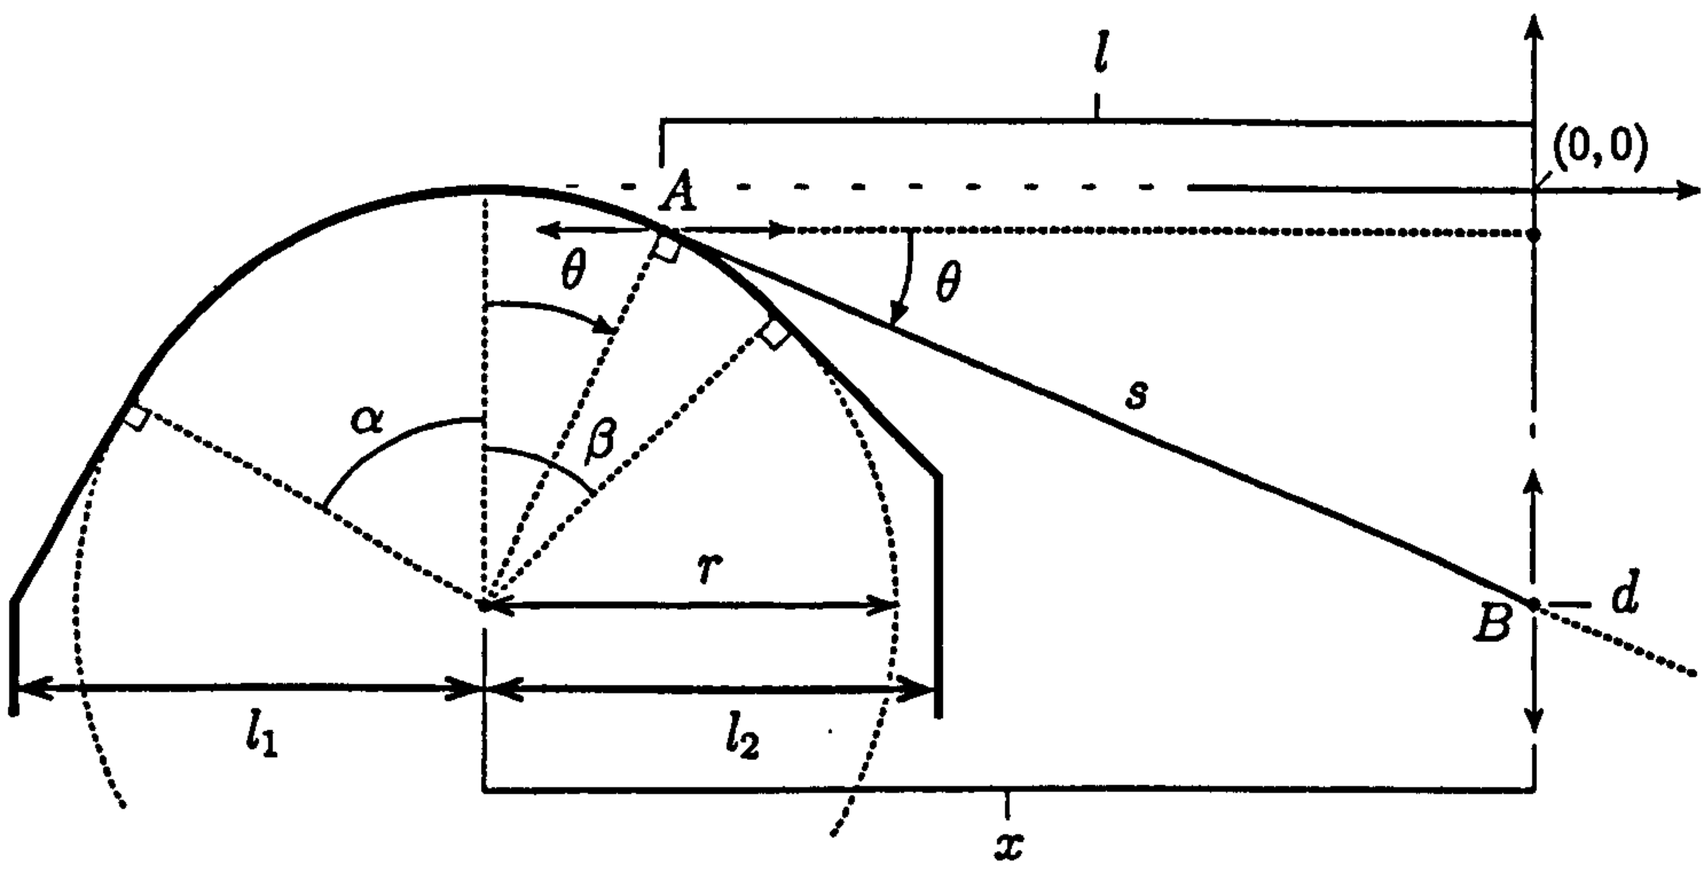

Figure G.1: *Useful Points and Lengths on a Cross-section of the Bearing Edge.*

Appendix G. Formulation of Equations used to model the Bearing Edge

The streams, $d(n)$ and $l(n)$, have been converted into instantaneous values, d and l, for simplicity. A solution is required that will give l for any value of d. The following points and lengths can be deduced

$$l = x - r\sin\theta, \tag{G.1}$$

$$A_y = r\cos\theta - r, \tag{G.2}$$

$$s = \frac{x - r\sin\theta}{\cos\theta}. \tag{G.3}$$

The value of d, which is already known, can also be found

$$d = A_y - s\sin\theta, \tag{G.4}$$

$$= r\cos\theta - r - \frac{x\sin\theta - r\sin^2\theta}{\cos\theta}, \tag{G.5}$$

$$= \frac{r\cos^2\theta - r\cos\theta - x\sin\theta + rsin^2\theta}{\cos\theta}, \tag{G.6}$$

$$= \frac{r - r\cos\theta - x\sin\theta}{\cos\theta}, \tag{G.7}$$

$$\Rightarrow d = \frac{r - x\sin\theta}{\cos\theta} - r. \tag{G.8}$$

Rearranging

$$r + d = \frac{r - x\sin\theta}{\cos\theta}, \tag{G.9}$$

$$(r + d)\cos\theta = r - x\sin\theta. \tag{G.10}$$

Squaring and then rearranging

$$(r + d)^2\cos^2\theta = r^2 + x^2\sin^2\theta - 2rx\sin\theta, \tag{G.11}$$

$$(r + d)^2(1 - \sin^2\theta) = r^2 + x^2\sin^2\theta - 2rx\sin\theta. \tag{G.12}$$

Putting this in the form of a quadratic in $\sin\theta$

$$(x^2 + (r + d)^2)\sin^2\theta - 2rx\sin\theta + r^2 - (r + d)^2 = 0. \tag{G.13}$$

This can be solved using the formula for the solution of quadratic (see appendix B.2)

$$y = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}, \tag{G.14}$$

where $y = \sin\theta$,

$$a = x^2 + (r+d)^2,$$

$$b = -2rx,$$

$$c = r^2 - (r+d)^2 = -d^2 - 2rd.$$

$$\therefore \sin\theta = \frac{2rx \pm \sqrt{4r^2x^2 - 4(x^2+(r+d)^2)(-d^2-2rd)}}{2(x^2+(r+d)^2)}, \tag{G.15}$$

$$= \frac{2rx \pm \sqrt{4(r^2x^2 - (x^2+z^2)(r^2-z^2))}}{2(x^2+z^2)}, \tag{G.16}$$

where $z = r + d$.

$$= \frac{2rx \pm 2\sqrt{r^2x^2 - (r^2x^2 - x^2z^2 + r^2z^2 - z^4)}}{2(x^2+z^2)}, \tag{G.17}$$

$$= \frac{2rx \pm 2\sqrt{x^2z^2 - r^2z^2 + z^4}}{2(x^2+z^2)}, \tag{G.18}$$

$$= \frac{2rx \pm 2\sqrt{z^2(x^2 - r^2 + z^2)}}{2(x^2+z^2)}, \tag{G.19}$$

$$= \frac{2rx \pm 2z\sqrt{x^2 - r^2 + z^2}}{2(x^2+z^2)}, \tag{G.20}$$

$$= \frac{rx \pm z\sqrt{x^2 - r^2 + z^2}}{x^2+z^2}. \tag{G.21}$$

The correct formulation was found to be (by testing with real numbers for both signs of the square root)

$$\sin\theta = \frac{rx - (r+d)\sqrt{x^2+d^2+2rd}}{x^2+(r+d)^2}. \tag{G.22}$$

The length, l, can now be found

$$l = \begin{cases} x - l_2, & \sin\theta > \sin\beta \\ x - r\sin\theta, & -\sin\alpha \le \sin\theta \le \sin\beta \\ x + l_1, & \sin\theta < -\sin\alpha \end{cases} \tag{G.23}$$

where $l_1 + l_2$ is the thickness of the drumshell,

α and β are absolute values.

This formula finds the horizontal distance to the contact point on the bearing edge, for a given displacement of the membrane measured a short distance away. Notice the limits imposed by the angles, α and β, of the straight lines that are cut from the bearing edge. These, combined with the drumshell's thickness, define the most extreme points of the edge that the membrane can be in contact with.

G.2 Finding $b(n+i)$ through Linear Interpolation

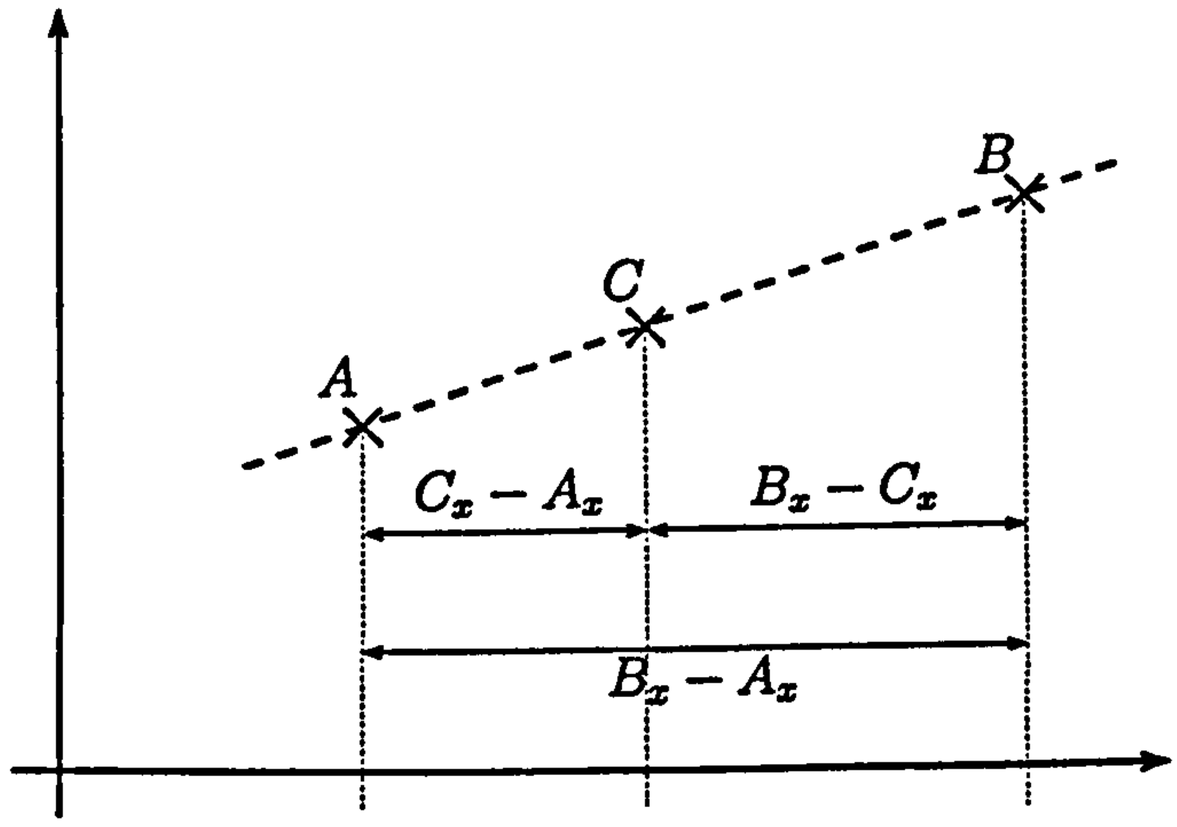

Figure G.2: *Intersecting the Line AB at Point C_x.*

Figure G.2 shows three points on a straight line, AB. The heights of these points represent data within the buffer $b(n+i)$ for different values of i. Both points A and B are known. The x-coordinate of point C is also known, but its y-coordinate needs to be calculated. This can be done using first order linear interpolation

$$C_y = A_y \frac{B_x - C_x}{B_x - A_x} + B_y \frac{C_x - A_x}{B_x - A_x}. \tag{G.24}$$

Expanding

$$C_y = \frac{A_y B_x - A_y C_x + B_y C_x - B_y A_x}{B_x - A_x}. \tag{G.25}$$

Using the values relevant to the bearing edge problem (see figure 4.12 of section 4.2.2) results in the points: $A(L'-1, L')$, $B(L, L)$ and $C(i, b(n+i))$, where $L' = l(n-1)$ and $L = l(n)$ for simplicity.

Inserting these into equation G.25

$$b(n+i) = \frac{L'L - L'i + Li - L(L'-1)}{L-(L'-1)}, \tag{G.26}$$

$$= \frac{Li - L'i + L'L - L'L + L}{L - L' + 1}, \tag{G.27}$$

$$\Rightarrow b(n+i) = \frac{i(L-L') + L}{L - L' + 1}. \tag{G.28}$$

G.3 Calculating the Amplitude A

Section 4.2.3 gives the string excitation as

$$v = \begin{cases} A\sin(\omega t), & 0 < \omega t \leq \pi \\ 0, & \pi < \omega t \end{cases} \quad (ms^{-1}) \tag{G.29}$$

where v is the velocity added to the excitation point (ms^{-1}),

ω is the frequency of the string's 1^{st} mode of resonance $(radians/second)$.

A value of A is required that will result in a given value of initial displacement, $D(metres)$.

Integrating equation G.29 over the range: $0 < t \leq \frac{\pi}{\omega}$

$$D = \int_0^{\frac{\pi}{\omega}} A\sin(\omega t)dt = \left[-\frac{A}{\omega}\cos(\omega t)\right]_{t=0}^{t=\frac{\pi}{\omega}}, \tag{G.30}$$

$$= \left(-\frac{A}{\omega}\cos\pi\right) - \left(-\frac{A}{\omega}\cos(0)\right), \tag{G.31}$$

$$= \frac{A}{\omega}(\cos(0) - \cos\pi), \tag{G.32}$$

$$\Rightarrow D = \frac{2A}{\omega}. \tag{G.33}$$

Rearranging

$$A = \frac{D\omega}{2}. \tag{G.34}$$

References

Abramowitz, M. and Stegun, I. A. (Eds.) (1972). Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables (9 ed.)., Chapter 9 (Bessel Functions of Integer Order), pp. 355–389. National Bureau of Standards Applied Mathematics Series 55. New York: Dover.

Adrien, J. M. (1991). The missing link: Modal synthesis. In G. de Poli, A. Picalli, and C. Roads (Eds.), Representations of Musical Signals, Cambridge, Massachusetts. MIT Press.

Aird, M. and Laird, J. (2001, June). Towards material modelling in physical models using digital waveguides. In Music Without Walls? Music Without Instruments? International Conference, DeMontfort University, Leicester, UK. Published on CDRom.

Aird, M.-L. (2002). Musical Instrument Modelling Using Digital Waveguides. Ph. D. thesis, University of Bath.

Allen, J. B. and Berkley, D. A. (1979). Image methods for efficiently simulating small-room acoustics. Acoust. Soc. Am. 65(4), 943–950.

Bilbao, S. (2001, March). Wave and Scattering Methods for the Numerical Integration of Partial Differential Equations. Ph. D. thesis, Stanford University. Published on Web: `http://ccrma-www.stanford.edu/~bilbao/`.

Bulatov, V. (1999). Rhombic dodecahedron. `http://www.physics.orst.edu/~bulatov/polyhedra/dual/ud12.html`.

Cadoz, C., Luciani, A., and Florens, J. L. (1993, Spring). CORDIS-ANIMA: A modeling and simulation system for sound and image synthesis - the general formalism. Computer Music Journal 17(1), 19–29.

Campos, G. and Howard, D. (2000, December 7-9). A parallel 3D digital waveguide mesh model with tetrahedral topology for room acoustic simulation. In COST G-6 Conference on Digital Audio Effects (DAFX-00), Verona, Italy.

Chaigne, A. and Askenfelt, A. (1994). Numerical simulations of piano strings. part I. a physical model for a struck string using finite difference methods. Journal of Acoust. Soc. Am., **95**, 1112–1118.

Chaigne, A. and Doutaut, V. (1997). Numerical simulations of xylophones. I. time-domain modelling of the vibrating bars. Journal of Acoust. Soc. Am., **101**, 539–557.

Creative (1998-2001). Showcase: Sound blaster AWE64 gold. `http://www.soundblaster.com/products/sbawe64gold/`. Copyright: Creative Technology Ltd.

Djoharian, P. (1993, Spring). Generating models for modal synthesis. Computer Music Journal 17(1), 57–65.

Djoharian, P. (1999, December). Material design in physical modeling sound synthesis. In 2nd COST G-6 Workshop on Digital Audio Effects DAFx99, Trondheim, pp. W99 1–4.

Duyne, S. A. V., Jaffe, D. A., Scandalis, G. P., and Stilson, T. S. (1997, September). A lossless, click-free, pitchbend-able delay line loop interpolation scheme. In ICMC Proceedings, pp. 252–255.

Duyne, S. A. V. and Smith(III), J. O. (1993). Physical modeling with the 2-d digital waveguide mesh. In ICMC Proceedings, pp. 40–47.

Duyne, S. A. V. and Smith(III), J. O. (1996, August). 3D tetrahedral digital waveguide mesh with musical applications. In ICMC Proceedings, Hong

Kong, pp. 9–16.

Duyne, S. V., Pierce, J. R., and Smith(III), J. O. (1994, September). Travelling wave implementation of a lossles mode-coupling filter and the wave digital hammer. In ICMC Proceedings, pp. 411–418.

Eckel, G., Iovino, F., and Caussé, R. (1995). Sound synthesis by physical modelling with modalys. In Proceedings of the International Symposium of Music Acoustics, Le Normant, Dourdan.

Eyring, C. F. (1930, Jan). Reverberation time in 'dead' rooms. Journal of Acoust. Soc. Am. 1(2), 217–241.

Farina, A. (2000a, April). Measurement of the surface scattering coefficient: Comparison of the Mommertz/Vorlander approach with the new Wave Field Synthesis method. In International Symposium on Surface Diffusion in Room Acoustics, Liverpool. Published on Web: `http://pcangelo.eng.unipr.it/Public/Papers/138-Sem00.PDF`.

Farina, A. (2000b, December). A new method for measuring the scattering coefficient and the diffusion coefficient of panels. Acustica 86(6), 928–942. Published on Web: `http://pcangelo.eng.unipr.it/Public/Papers/150-ACTA2000.PDF`.

Fletcher, N. H. and Rossing, T. D. (1991a). The Physics of Musical Instruments, Chapter 18, pp. 515–517. Springer-Verlag New York Inc.

Fletcher, N. H. and Rossing, T. D. (1991b). The Physics of Musical Instruments, Chapter 3,18, pp. 70,502–503. Springer-Verlag New York Inc.

Fontana, F. and Rocchesso, D. (1998). Physical modeling of membranes for percussion instruments. Acustica, 84, 529–542.

Fontana, F. and Rocchesso, D. (1999). Signal-theoretic characterization of waveguide mesh geometries for models of two-dimensional wave propagation in elastic media. Submitted to the IEEE Trans. on Speech and

Audio.

Fontana, F., Rocchesso, D., and Apollonio, E. (2000, December 7-9). Using the waveguide mesh in modelling 3D resonators. In COST G-6 Conference on Digital Audio Effects (DAFX-00), Verona, Italy.

Glatt, J. (unknown year). MIDI specification: Controller numbers. `http://www.borg.com/~jglatt/tech/midispec.htm`.

Gray, A. H. and Markel, J. D. (1975, June). A normalized digital filter structure. IEEE Trans. on Acoust., Speech, and Signal Proc., **ASSP-23**, 268–277.

Hales, T. C. (1998, November). Abstract math/9811071-9811078. `http://arxiv.org/abs/math.MG/9811071 - http://arxiv.org/abs/math.MG/9811078`. Proof of the solution to Kepler's cannon ball stacking problem.

Hall, D. E. (1987a). Basic Acoustics, Chapter 1 (Acoustics: A Preliminary Acquaintance), pp. 8. New York: John Wiley & Sons, Inc.

Hall, D. E. (1987b). Basic Acoustics, Chapter 7 (Transverse Waves on a String), pp. 115. New York: John Wiley & Sons, Inc.

Hall, D. E. (1987c). Basic Acoustics, Chapter 11 (Sources of Radiation), pp. 161. New York: John Wiley & Sons, Inc.

Hall, D. E. (1987d). Basic Acoustics, Chapter 8 (Waves: Further Examples), pp. 127. New York: John Wiley & Sons, Inc.

Hall, D. E. (1987e). Basic Acoustics, Chapter 9 (Sound Waves: Basic Properties), pp. 139. New York: John Wiley & Sons, Inc.

Hall, D. E. (1987f). Basic Acoustics, Chapter 12 (Room Acoustics), pp. 186. New York: John Wiley & Sons, Inc.

Hall, D. E. (1987g). Basic Acoustics, Chapter 9 (Sound Waves: Basic Properties), pp. 136–138. New York: John Wiley & Sons, Inc.

Hall, D. E. (1992, January). Piano string excitation. part VI: Non-linear modelling. Journal of Acoust. Soc. Am., **92**, 95–105.

Heckroth, J. (1995). Tutorial on MIDI and music synthesis. `http://www.harmony-central.com/MIDI/Doc/tutorial.html`.

Howard, D. M. and Angus, J. (1996). Acoustics and Psychoacoustics. Focal Press. ISBN 0-240-51428-9.

Huopaniemi, J., Savioja, L., and Karjalainen, M. (1997, October 19-22). Modeling of reflections and air absorption in acoustic spaces - a digital filter design approach. In IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, Mohonk Mountain House, New Paltz, New York.

IDnet (2000-2002). Yamaha music products: Midi controllers. `http://www.yamaha-music.co.uk/PRODUCTS/MUSIC_PRODUCTION/MIDI_CONTROLLERS/WX5.ASP?sectionid=65`. Copyright: Yamaha-Kemble Music (UK) Ltd.

Jaffe, D. and Smith(III), J. O. (1983). Extensions of the karplus-strong plucked string algorithm. Computer Music Journal 7(2), 56–69.

Jeans, J. (1937). Science and Music, Chapter 7 Hearing, pp. 246–248. London: Cambridge University Press.

Karplus, K. and Strong, A. (1983, Summer). Digital synthesis of plucked-string and drum timbres. Computer Music Journal 7(2), 43–55.

Kinsler, L. E. and Frey, A. R. (1962). Fundamentals of Acoustics (2nd ed.). New York: Wiley.

Kulowski, A. (1985). Algorithmic representation of the ray tracing technique. Applied Acoustics 18(6), 449–469.

Laakso, T. I., Välimäki, V., Karjalainen, M., and Laine, U. K. (1992, October 14-18). Real-time implementation techniques for a continuously variable digital delay in modeling musical instruments. In ICMC Proceedings, San Jose, California, pp. 140–141.

Laakso, T. I., Välimäki, V., Karjalainen, M., and Laine, U. K. (1996, January). Splitting the unit delay. IEEE Signal Processing Magazine 13(1), 30–60.

Laird, J., Masri, P., and Canagarajah, C. N. (1998, November 18). Efficient and accurate synthesis of circular membranes using digital waveguides. In IEE Colloquium, Audio and Music Technology: The Challenge of Creative DSP, pp. 12/1–12/6.

Laird, J., Masri, P., and Canagarajah, C. N. (1999, October). Modelling diffusion at the boundary of a digital waveguide mesh. In ICMC Proceedings, Beijing, pp. 492–495.

Lam, Y. W. (1997, February). On the modelling of diffuse reflections in rooms acoustics prediction. In BEPAC & EPSRC Conference on Sustainable Building, England, pp. 106–113.

Lindsay, R. B. (Ed.) (1973). Acoustics: Historical and Philosophical Development, Chapter J. l. d'Alembert: Investigation of the curve formed by a vibrating string, 1747, pp. 119–123. Stroudsburg: Dowden, Hutchinson & Ross.

Macey, P. C. (1987). Acoustic and Structure Interaction Problems Using Finite and Boundary Elements. Ph. D. thesis, Nottingham University, Nottingham, UK.

Moore, F. R. (1990). Elements of Computer Music, Chapter 1 Introduction, pp. 19. Prentice Hall, Englewood Cliffs, New Jersey.

Morrison, J. and Adrien, J. M. (1993). MOSAIC: A framework for modal synthesis. Computer Music Journal 17(1), 45–56.

Morse, P. M. and Ingard, K. U. (1968). Theoretical Acoustics. New York: McGraw-Hill.

Murphy, D. T. and Howard, D. M. (2000, December 7-9). 2-D digital waveguide mesh topologies in room acoustics modelling. In COST G-6

Conference on Digital Audio Effects (DAFX-00), Verona, Italy.

Pearson, M. The tao physical modelling synthesis program: New features added at ACAT, ANU. `http://www.anu.edu.au/ITA/ACAT/mdp/acmapaper.html`.

Pearson, M. and Howard, D. M. (1995, September). A musicians approach to physical modelling. In ICMC Proceedings, pp. 578–580.

Petyt, M. (1990). Introduction to Finite Element Vibration Analysis. Cambridge, UK: Cambridge University Press.

Plus Online Magazine (1997). Mathematical mysteries: Kepler's conjecture. `http://pass.maths.org.uk/issue3/xfile/index.html`.

Rhaouti, L., Chaigne, A., and Joly, P. (1999, June). Time-domain modeling and numerical simulation of a kettledrum. Journal of Acoust. Soc. Am. 105(6), 3545–3562.

Rocchesso, D. and Smith(III), J. O. (1994, September). Circulant feedback delay networks for sound synthesis and processing. In ICMC Proceedings, Denmark, pp. 378–381.

Roederer, J. G. (1973a). Introduction to the Physics and Psychophysics of Music, Chapter 2 (Sound Vibrations, Pure Tones, and the Perception of Pitch), pp. 18–32. London: The English Universities Press Ltd.

Roederer, J. G. (1973b). Introduction to the Physics and Psychophysics of Music, Chapter 3 (Sound Waves, Acoustical Energy, and the Perception of Loudness), pp. 75–81. London: The English Universities Press Ltd.

Savioja, L., Karjalainen, M., and Takala, T. (1996, September 24-27). DSP formulation of a finite difference method for room acoustics simulation. In IEEE Nordic Signal Processing Symposium, Finland, pp. 455–458. NORSIG'96.

Savioja, L. and Välimäki, V. (1996, September 24-27). The bilinearly deinterpolated waveguide mesh. In IEEE Nordic Signal Processing

Symposium, Finland, pp. 443–446. NORSIG'96.

Seven Stones (2000). Hexagonal close packing. `http://www.7stones.com/Homepage/Publisher/rrRotate.html`.

Smith(III), J. O. (1987). Waveguide filter tutorial. In ICMC Proceedings, pp. 9–16.

Smith(III), J. O. (1992, Winter). Physical modeling using digital waveguides. Computer Music Journal 16(4), 74–87.

Spoendlin, H. (1970). Structural Basis of Peripheral Frequency Analysis. Frequency Analysis and Periodicity Detection in Hearing. Leiden: A. W. Suithoff.

Thiran, J. P. (1971, November). Recursive digital filters with maximally flat group delay. IEEE Trans. Circuit Theory 18(6), 659–664.

Tolonen, T., Välimäki, V., and Karjalainen, M. (1998, March). Evaluation of modern sound synthesis methods. Technical Report 48, Helsinki University of Technology, Otaniemi. Published on Web: `http://www.acoustics.hut.fi/~ttolonen/sound_synth_report.html`.

Trussler, B. (unknown year). Gustav holst: The planets suite. `http://www.aquarianage.org/lore/holst.html`.

Välimäki, V. (1995, December 18). Discrete-Time Modeling of Acoustic Tubes Using Fractional Delay Filters. Ph. D. thesis, Helsinki University of Technology. Published on Web: `http://www.acoustics.hut.fi/~vpv/publications/vesa_phd.html`.

Välimäki, V., Karjalainen, M., Tolonen, T., and Erkut, C. (1999, October). Nonlinear modeling and synthesis of the kantele - a traditional finnish string instrument. In ICMC Proceedings, Beijing, pp. 220–223.

Välimäki, V., Laakso, T. I., and Mackenzie, J. (1995, September). Elimination of transients in time-varying allpass fractional delay filters with

application to digital waveguide modeling. In ICMC Proceedings, pp. 327–334.

von Békésy, G. (1960). Experiments in Hearing. New York: McGraw Hill Book Co.

Wood, A. B. (1964). A Textbook of Sound (third ed.)., Chapter Reception, Transformation, and Measurement of Sound Energy, pp. 385–393. London: G. Bell and Sons Ltd.

Wright, J. R. (1995, October). An exact model of acoustic radiation in enclosed spaces. Journal of Audio Eng. Soc 43(10), 813–820. Presented at the 96th Convention of the Audio Engineering Society, Amsterdam, The Netherlands, 1994 February 26-March 1; revised 1995 May 20.

THESIS CONTAINS CD ROM