Acoustic Inspection of Timber

Simon N. Woods
B.E. (Hons)

A thesis submitted in partial fulfilment
of the requirements for the degree of
Master of Engineering
in
Electrical and Computer Engineering
at the
University of Canterbury,
Christchurch, New Zealand.

Dec 2006

ABSTRACT

The ability to determine wood quality using non-destructive tests has enormous potential for the forestry industry in both research and commercial applications. This thesis describes some of the theory of acoustic waves in wood and how wood stiffness can be estimated by measuring the velocity of acoustic waves. Attention is paid to both resonance and stress wave timer technologies and the benefits and problems with both. A detailed description is given of the design of a new tool (Treetap 5.0) to aid in future, acoustic based, timber inspection research.

ACKNOWLEDGEMENTS

The list of people who I need to thank for their support and assistance while completing this thesis could go on forever. To all those I am unable to mention here personally: thank you all for your part in helping me finish this chapter of study. You assistance never went unnoticed or unappreciated.

Particular thanks must go to WQI for their financial support of the acoustic work done in Canterbury, which I am grateful to have been apart of.

To John Walker in the forestry department: the opportunity to partake in fieldwork in Australia was invaluable to my understanding of the forestry industry as a whole. Thanks also to Nigel Pink (also in the forestry department) who was always willing to help acquire whatever wood samples were needed for experimentation.

My partner, Denise. From the start of this project you have been a constant source of encouragement and support. Thank you. I'm looking forward to starting on our new adventures in Dunedin in the New Year.

Michael Hayes, my supervisor. Mike has been a tower of technical knowledge and assistance throughout this project. I literally cannot thank you enough for all your help, time, effort and patience. It's been both a privilege and a pleasure working with you. Couldn't have done it without you. Cheers.

CONTENTS

ABSTRACT **iii**

ACKNOWLEDGEMENTS **v**

CHAPTER 1 INTRODUCTION **1**
1.1 Outline 3

CHAPTER 2 ACOUSTICS IN SOLIDS **5**
2.1 Introduction 5
2.2 Stress waves 5
2.2.1 Stress and strain 6
2.2.2 Hooke's law 6
2.2.2.1 Reduced notation 7
2.2.3 Velocity notation 7
2.3 Isotropic solids 8
2.4 Orthotropic solids 9
2.4.1 Hooke's law for orthotropic solids 9
2.4.2 Wave propagation in an orthotropic solid 11
2.4.2.1 Christoffel's equations 12
2.4.2.2 Axial propagation solution 13
2.4.2.3 Off axis propagation solution 13
2.5 Dispersion 14
2.5.1 Model conversion 15
2.5.2 Transmission line model 15
2.5.2.1 General solutions 17
2.5.2.2 Transfer function 17
2.5.2.3 Attenuation and velocity 18
2.5.3 Lossless model 18
2.5.4 More complex models 19
2.5.5 Mechanical notation 21
2.5.6 Using the model to simulate materials 22

CHAPTER 3 ACOUSTIC MEASUREMENT OF WOOD **25**
3.1 Introduction 25
3.2 Wood as an acoustic medium 26

3.2.1 Orthogonality 26
3.2.2 Homogeneity 27
3.3 Wood stiffness 30
3.3.1 Industry test 31
3.3.2 Stiffness using acoustics 31
3.3.3 Stress wave timers 32
3.3.3.1 Using two receivers 33
3.3.4 Resonance method 35
3.4 Inducing acoustic pulses 36
3.5 Receiving acoustic pulses 37
3.5.0.1 Force to voltage conversion 37
3.5.0.2 Filtering and amplification 38
3.5.0.3 Analog to digital conversion 38
3.6 Accelerometer 38
3.6.1 Accelerometer model 39
3.6.1.1 Reciprocal transducer 40
3.6.1.2 Three port model 40
3.6.1.3 Admittance characterisation 43
3.6.1.4 Estimating component values 45
3.6.2 Transfer function 46
3.6.3 Redesigning the accelerometer 47
3.7 Dispersion 48
3.7.1 Stress wave timer measurements 51
3.7.2 Resonance measurements 53
3.7.3 Measurement comparison 55
3.8 Problems 55
3.8.1 Shear waves 55
3.8.2 Multiple paths 56
3.8.3 Defects 57
3.9 Summary 58
3.9.1 Resonance testers 58
3.9.1.1 Low frequency measurements 59
3.9.2 Stress wave timers 59
3.9.2.1 Dispersion for the stress wave timer 60

CHAPTER 4 SYSTEM HARDWARE 61
4.1 Introduction 61
4.1.1 System overview 61
4.2 Probes 62
4.2.1 Accelerometer 62
4.2.2 Front end 63
4.2.2.1 Signal detector 64
4.2.2.2 Wakeup circuit 65
4.2.3 Memory 65

4.2.4 Software - time 66
4.2.5 Software - measurement loop 66
4.2.5.1 Initialisation 66
4.2.5.2 Loop functions 67
4.2.5.3 Loop timing 67
4.2.5.4 Finishing 69
4.2.6 Power 70
4.2.6.1 Future work 70
4.2.7 Temperature sensor 71
4.2.8 Housing 71
4.3 Handunit 71
4.3.1 Microcontroller 72
4.3.2 Display 72
4.3.3 Pushbuttons 73
4.3.4 Memory 73
4.3.4.1 Future work 73
4.3.5 USB 73
4.3.6 Power 74
4.3.6.1 Charge circuits 74
4.3.7 Housing 76
4.4 RF communications 77
4.4.1 RF hardware 77
4.4.1.1 RF band 77
4.4.1.2 RF channels 78
4.4.1.3 RF addresses 78
4.4.2 RF protocol 78
4.4.2.1 Hardware protocol 79
4.4.2.2 Commands and acknowledgments 79
4.4.2.3 Timestamp generation 80
4.4.3 Probe enumeration 80
4.4.3.1 Alternative scheme 83
4.4.3.2 Data download 83
4.5 Synchronisation 85
4.5.1 Local time and timestamps 85
4.5.1.1 Clock sources 85
4.5.1.2 Relating local time on two units 87
4.5.1.3 Timestamps 87
4.5.1.4 Timestamp jitter 88
4.5.1.5 Estimating parameters 89
4.5.2 Clock stability 89
4.5.2.1 Period of sampling and period of stability 90
4.5.2.2 Optimal length of timestamp sampling 91
4.5.2.3 Minimum length of timestamp sampling 91
4.6 Summary 93

4.6.1 Positive features 93
4.6.2 Suggested future work 93
4.6.2.1 Accelerometer bandwidth 94
4.6.2.2 Wireless communications 95
4.7 Conclusions 97
4.7.1 Suggestions for future experiments 97
4.7.2 Treetap 5.0 98

REFERENCES **99**

Chapter 1

INTRODUCTION

Wood has been a commodity among humans for centuries. Throughout our history it has been employed as fuel, tools, weapons, shelter, and transportation. It has been used to furnish homes, produce music, make paper, entertain children, and many other applications.

Over the years we have learnt a huge amount about wood. We have learned how to shape, bend, and process it to make a huge variety of products, found ways to encourage the growth of wood with desirable properties, and discovered the many contributions of trees to our ecosystem. Entire institutions have been set up dedicated to further our knowledge of wood.

The modern wood industry, is comprised of many complimentary parts. There are forestry companies to plant, maintain, and harvest trees. Sawmills and pulping plants to process the raw lumber into more usable forms. Specialist products like laminates, plywoods, and particle boards undergo further processing. Then there is the chain of exporters and distributors who transport wood products to where they are needed. This industry has international significance, involving billions of dollars worldwide.

The value of wood in a forest is determined by the quality and quantity of the trees. Properties of wood that can affect its value include the presence of defects, colour, grain, breaking strain and stiffness. Stiffness in particular is a characteristic that wood is regularly graded by because of its impact on the timber's suitability for construction. It is stiffness that this thesis focuses on in terms of acoustic investigation.

Research into acoustic waves also has a great depth of history, perhaps beginning in 5 BC with Pythagoras proof involving harmonics. This research has come a long way from its humble origins. Acoustic waves are now used for more than just music and communications. Modern applications include:

- Underwater sonar.
- Aircraft component inspection.
- Concrete inspection.

- Medical imaging and treatments.

Acoustic inspection of wood is not an entirely new field, but recent industrial application is sparking a renewed interest in further research. The ability to ascertain properties of trees and lumber using acoustic Non-Destructive Testing (NDT) opens up new possibilities for the forestry industry. Technologies for determining the stiffness of wood using acoustics are now a reality.

Stiffness is, of course, not the only wood property of interest to the forestry industry. The same technology is also being applied to tomographic systems for detecting rot in trees. Other properties of interest include:

- Aesthetic appeal. Colour and grain.
- Breaking strain.
- Stability. The amount of warp, twist, and bow exhibited as wood is cut and dried.
- Defects in the wood.
- Natural resistance to deterioration.

It is hoped that in the future acoustic NDT systems may be able to measure other properties of wood that effect its value.

This thesis is part of a larger project to further investigate the use of acoustics in timber inspection. In particular, the use of stress wave timers to ascertain wood stiffness. The Treetap device, developed at the University of Canterbury, has already been used successfully in many forestry research projects to determine stiffness in standing trees.

The Treetap 5.0 device is not, strictly speaking, a stress wave timer. Its aim is to provide a new research platform that can be used in a forest to sample and store acoustic signals. These signals can then be processed on a PC to investigate the acoustic properties of wood. The ultimate outcome of this will hopefully be to improve the accuracy of future stress wave timing devices. This processing, however, is beyond the scope of this thesis [Chauhan 2004] [Fife *et al.* 2004].

In addition to improving accuracy of stress wave timers it may be possible to begin drawing correlations between processed acoustic data and some of the other properties of wood listed above. An ideal NDT device would be small, cheap, quick and easy to use, and accurately report every characteristic of interest regarding a tree; but this is some way off yet.

1.1 OUTLINE

This thesis has three main chapters. The first two are a theoretical investigation of acoustic propagation. The last is a detailed description of the Treetap 5.0 hardware.

- **Chapter 2: Acoustics in solids.** This chapter is an introduction to the physics of waves travelling in solids. The concepts of stress and strain are introduced and their relationship to Hooke's law. Isotropic and orthotropic solids are discussed and how they effect wave propagation. The chapter ends with an introduction to dispersion and the development of a model that can be used to simulate its effect on acoustic waves.

- **Chapter 3: Acoustics in wood.** The second chapter builds on the concepts introduced in the first chapter and applies them specifically to wood. It begins with a discussion of wood's useful and necessary orthotropic and homogeneous properties. This is followed by a look at wood stiffness; how it effects the wood industry and a review of the two main technologies used to measure it acoustically. There is an examination of methods used to induce and receive acoustic pulses in wood with specific attention to the accelerometer used in the Treetap devices. Dispersion is revisited along with its implications on acoustic measurements. Finally some of the problems facing acoustic measurements are reviewed before a general summary of the strengths and weaknesses of both stiffness estimating technologies.

- **Chapter 4: System hardware.** This chapter is a description of the newly developed Treetap 5.0 hardware. After a general overview of the system specific of the probe and handunit devices are discussed. Most of the detail regards the physical hardware of the devices, but certain important software aspects are also covered. There is a review of the wireless communication system employed, including the software protocol. Lastly there is a look at the system employed to synchronise the data from the independent devices and how the stability of the independent clock sources effects this.

Chapter 2

ACOUSTICS IN SOLIDS

2.1 INTRODUCTION

This chapter introduces principles of acoustic propagation for solid materials. The results developed here are used in the following chapter to analyse acoustic propagation in wood.

Definitions of stress, strain, and the fundamentally important Hooke's law are dealt with in the first section. Section 2.3 is devoted to isotropic solids, leading onto orthotropic solids in Section 2.4. Orthotropic solids require significantly more investigation to develop an expression for Hooke's law and finally equations for wave velocity.

The last section of the chapter is devoted to dispersion. A mixture of electrical and mechanical models are used to look at how dispersion effects wave propagation in terms of attenuation as well as wave velocity.

The content of this chapter is in no way an exhaustive reference for all facets of acoustics. Its aim is to explain only the necessary elements for understanding the remainder of this thesis.

2.2 STRESS WAVES

Stress waves are a mechanical stress propagating through a solid medium in the form of a wave [Kolsky 1963]. The moving stress causes elastic displacement for the medium as it travels. There are two primary categories of stress waves in solids.

Longitudinal waves cause displacement in the direction of propagation. These waves are also called p-waves (primary or pressure), dilational waves, or compression waves.

Shear waves cause displacement perpendicular to the direction of propagation. These waves are also called s-waves (secondary or shear) or as transverse waves.

Other types of stress wave may also exist, like the Rayleigh wave. This wave occurs at the boundary of two media with different elastic properties (such as wood and air)

as a result of the superposition of shear and longitudinal waves [Kolsky 1963] [Emms and with assistance from Mike Andrews 2004]. By strict definition, acoustic waves are longitudinal [Kinsler *et al.* 1982]. However, when the medium of interest is a solid other waves types must also be considered, particularly shear waves.

In bounded solids superposition of waves can produce many other types of vibrational modes. Torsional and flexural modes have been observed and characterised in rods and bars [Kolsky 1963] [Bucur 2006]. Modelling trees as a hollow cylinder because of the changes in density across the growth rings leads poses even more vibrational modes [Wang and Williams 1996]. Circumferential modes have also been observed in spruces [Axmon and Hansson 1999].

2.2.1 Stress and strain

Stress waves take their name from the stress, or internal pressure, within a medium. Strain refers to the displacement caused by a stress. Consider an axial force applied to a uniform wire by pulling the ends apart. The stress σ is the pressure inside the wire and is equal to the force on the wire divided by its cross sectional area:

$$\sigma = \frac{F}{A} \tag{2.1}$$

Stress is measured in pascals (Pa). The strain ϵ is the displacement per unit length, and is equal to the distance the wire is stretched divided by its original length:

$$\epsilon = \frac{\Delta l}{l} \tag{2.2}$$

In this one dimensional case the stress and strain are related by a constant, C:

$$\sigma = C\,\epsilon \tag{2.3}$$

where the constant C is the stiffness. This relationship is Hooke's Law. In the one dimensional case, C can be denoted by E (Young's modulus).

2.2.2 Hooke's law

A more generalised version of Hooke's Law describes the elastic properties of three dimensional solids with three dimensional stresses and strains:

$$[\sigma_{ij}] = [C_{ijkl}] \cdot [\epsilon_{kl}] \tag{2.4}$$

There are three physical dimensions in Cartesian coordinates so: $(i, j, k, l \in x, y, z)$. σ_{ij} is the stress in direction i acting on the surface with its normal in the direction j.

The values of the fourth-rank tensor $[C_{ijkl}]$ are the elastic stiffnesses. ϵ_{kl} is the three dimensional strain component defined by

$$\epsilon_{kl} = \frac{1}{2}\left(\frac{\partial \xi_k}{\partial \mathrm{x}_l} + \frac{\partial \xi_l}{\partial \mathrm{x}_k}\right) \tag{2.5}$$

where ξ_i is the displacement in direction i, and x_i is the coordinate in direction i [Bucur 2006]. An alternative notation for Hooke's Law uses the compliance tensor $[S_{ijkl}]$:

$$[\epsilon_{kl}] = [S_{ijkl}] \cdot [\sigma_{ij}] \tag{2.6}$$

where the compliance tensor is the inverse of the stiffness tensor: $[S_{ijkl}] = [C_{ijkl}]^{-1}$. If the solid is homogeneous, the values of $[C_{ijkl}]$ and $[S_{ijkl}]$ are constants since they are independent of location. Since i, j, k, and l all belong to sets of size 3, the stiffness and compliance tensors have 81 terms.

2.2.2.1 Reduced notation

Hookes's law can be reduced into a vector and matrix notation. The stress and strain tensors are symmetric: $\sigma_{ij} = \sigma_{ji}, \epsilon_{ij} = \epsilon_{ji}$. As a result, the fourth order stiffness tensor is symmetric also: $C_{ijkl} = C_{jikl} = C_{ijlk} = C_{jilk}$. Only 21 of the 81 terms of $[C_{ijkl}]$ are independent. Hooke's Law can be restated using vectors for σ and ϵ, and a matrix for the stiffness:

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{21} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{31} & C_{32} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{41} & C_{42} & C_{43} & C_{44} & C_{45} & C_{46} \\ C_{51} & C_{52} & C_{53} & C_{54} & C_{55} & C_{56} \\ C_{61} & C_{62} & C_{63} & C_{64} & C_{65} & C_{66} \end{bmatrix} \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{zz} \\ \epsilon_{yz} \\ \epsilon_{zx} \\ \epsilon_{xy} \end{bmatrix} \tag{2.7}$$

In this reduced notation, the subscripts of the stiffness matrix $[C_{mn}]$ are indices of the matrix only: $(m, n \notin x, y, z)$. The same reduction can be applied to the compliance tensor $[S_{mn}]$. In this matrix notation the inverse property $[S_{mn}] = [C_{mn}]^{-1}$ still applies. See [Auld 1973] for a more detailed description of this transformation.

The $[C_{mn}]$ and $[S_{mn}]$ matricies are symmetric, so $C_{mn} = C_{nm}$. Thus only 21 of the 36 terms shown are independent. The number of independent terms can be further reduced by taking advantage of specific properties of the medium.

2.2.3 Velocity notation

A stress waves has two velocities associated with it. To reduce potential confusion created by this they are defined here with different symbols and names. One is the

velocity of the wave itself, the other is the velocity of the displaced solid. The speed at which the wave propagates through the solid will be referred to as the *phase velocity* of the wave [Kinsler *et al.* 1982]. The phase velocity will have the variable v.

The displacement of a point in the solid caused by the wave will use the notation ξ_i where i is the direction of the displacement. The derivative of the displacement will be referred to as the *point velocity* and will have the notation u_i. The point velocity is the derivative of the displacement [Kinsler *et al.* 1982]:

$$u_i = \frac{\mathrm{d}\xi_i}{\mathrm{d}t} \tag{2.8}$$

2.3 ISOTROPIC SOLIDS

The propagation of waves in isotropic materials is independent of the direction of travel. Isotropic solids have only two independent stiffness constants, the Lamé coefficients μ and λ. These coefficients are defined as [Bucur 2006]

$$\mu = \frac{E}{2(1+\nu)} \tag{2.9}$$

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)} \tag{2.10}$$

where E is Young's modulus and ν is Poisson's ratio.

Equation 2.7 can be rewritten using the Lamé coefficients for a three dimensional isotropic solid:

$$\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} \lambda+2\mu & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda+2\mu & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & \lambda+2\mu & 0 & 0 & 0 \\ 0 & 0 & 0 & 2\mu & 0 & 0 \\ 0 & 0 & 0 & 0 & 2\mu & 0 \\ 0 & 0 & 0 & 0 & 0 & 2\mu \end{bmatrix} \begin{bmatrix} \epsilon_{xx} \\ \epsilon_{yy} \\ \epsilon_{zz} \\ \epsilon_{yz} \\ \epsilon_{zx} \\ \epsilon_{xy} \end{bmatrix} \tag{2.11}$$

Young's modulus E is the ratio of the longitudinal stress to longitudinal strain. This is the same constant as in the wire example, Equation 2.3. For an isotropic solid, $E = \lambda + 2\mu$.

Poisson's ratio $\nu$is the ratio of transverse contraction to longitudinal extension under tensile stress. Pull the ends of some elastic material, like a stick of rubber, and the lateral dimensions decrease. Poisson's ratio measures this relationship.

The shear modulus, usually denoted as G, is the ratio of the transverse stress to trans-

verse strain. For isotropic solids $G = \mu$ [1]. Young's modulus, the Poisson's ratio and the shear modulus are commonly referred to as the engineering parameters [Bucur 2006].

The two moduli can be calculated by measuring the speeds of propagating waves if the density is known. The velocity of a longitudinal wave is given by Young's Modulus, while the velocity of a transverse (shear) wave is given by the shear modulus [Bucur 2006].

$$v_L = \sqrt{\frac{E}{\rho}} \tag{2.12}$$

$$v_S = \sqrt{\frac{G}{\rho}} \tag{2.13}$$

where ρ is the mass density of the solid.

2.4 ORTHOTROPIC SOLIDS

The propagation of a wave in an anisotropic material depends upon the direction it travels. Anisotropy is caused by the internal structure of the material having some directional dependence. Examples of anisotropic materials include crystal arrays, metals, ceramics, rocks, and layered plastics [Bucur 2006]. Orthotropic materials are a subset of anisotropic materials that have three mutually perpendicular planes of elastic symmetry. Wood is an anisotropic material that is usually modelled as being orthotropic [Bucur 2006] [Bucur 2001].

2.4.1 Hooke's law for orthotropic solids

In an orthotropic solid the stiffness matrix can be reduced to 12 terms:

$$[C_{mn}] = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ C_{21} & C_{22} & C_{23} & 0 & 0 & 0 \\ C_{31} & C_{32} & C_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{66} \end{bmatrix} \tag{2.14}$$

Only 9 of these are independent since $C_{mn} = C_{nm}$. The compliance matrix $[S_{mn}]$ has exactly the same structure as the stiffness matrix. $[S_{mn}]$ can be written in terms of the engineering parameters as shown in equation 2.15. By observing the engineering parameters, the physical significance of the elements in $[S_{mn}]$ is deduced [Bucur 2006].

[1] Note that some literature uses a factor of 1/2 in the shear moduli in the stiffness matrix. This results from a difference in definition of shear strain and *engineering shear strain*. $\gamma_{xy} = \epsilon_{xy} + \epsilon_{yx} = 2\epsilon_{xy}$.

$$\begin{bmatrix} S_{11} & S_{12} & S_{13} & 0 & 0 & 0 \\ S_{21} & S_{22} & S_{23} & 0 & 0 & 0 \\ S_{31} & S_{32} & S_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & S_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & S_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & S_{66} \end{bmatrix} = \begin{bmatrix} \frac{1}{E_x} & -\frac{\nu_{xy}}{E_y} & -\frac{\nu_{xz}}{E_z} & 0 & 0 & 0 \\ -\frac{\nu_{yx}}{E_x} & \frac{1}{E_y} & -\frac{\nu_{yz}}{E_z} & 0 & 0 & 0 \\ -\frac{\nu_{zx}}{E_x} & -\frac{\nu_{zy}}{E_y} & \frac{1}{E_z} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{G_{yz}} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{G_{xz}} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{G_{xy}} \end{bmatrix} \tag{2.15}$$

- S_{11}, S_{22}, S_{22} relate extensional strain in a particular direction to the extensional stress in the same direction.
- S_{12}, S_{13}, S_{23} relate extensional strain in a particular direction to the extensional stress in a perpendicular direction.
- S_{44}, S_{55}, S_{66} relate a shear strain to a shear stress in the same plane.

Note that the subscripts in the right hand matrix refer to the three physical dimensions, not the indices of the matrix, as on the left. Because $[S_{mn}]$ is symmetric it follows that

$$\frac{\nu_{xy}}{E_y} = \frac{\nu_{yx}}{E_x} \tag{2.16}$$

$$\frac{\nu_{xz}}{E_z} = \frac{\nu_{zx}}{E_x} \tag{2.17}$$

$$\frac{\nu_{yz}}{E_z} = \frac{\nu_{zy}}{E_y} \tag{2.18}$$

$[S_{mn}]$ can be inverted to give $[C_{mn}]$ in terms of the engineering parameters too.

$$[C_{mn}] = \begin{bmatrix} \frac{1-\nu_{yz}\nu_{zy}}{E_yE_z\Delta} & \frac{\nu_{yx}+\nu_{zx}\nu_{yz}}{E_yE_z\Delta} & \frac{\nu_{zx}+\nu_{yx}\nu_{zy}}{E_yE_z\Delta} & 0 & 0 & 0 \\ \frac{\nu_{xy}+\nu_{xz}\nu_{zy}}{E_xE_z\Delta} & \frac{1-\nu_{zx}\nu_{xz}}{E_xE_z\Delta} & \frac{\nu_{zy}+\nu_{zx}\nu_{xy}}{E_xE_z\Delta} & 0 & 0 & 0 \\ \frac{\nu_{xz}+\nu_{xy}\nu_{yz}}{E_xE_y\Delta} & \frac{\nu_{yz}+\nu_{xz}\nu_{yx}}{E_xE_y\Delta} & \frac{1-\nu_{xy}\nu_{yx}}{E_xE_y\Delta} & 0 & 0 & 0 \\ 0 & 0 & 0 & G_{yz} & 0 & 0 \\ 0 & 0 & 0 & 0 & G_{xz} & 0 \\ 0 & 0 & 0 & 0 & 0 & G_{xy} \end{bmatrix} \tag{2.19}$$

where

$$\Delta = \frac{1 - \nu_{xy}\nu_{yx} - \nu_{xz}\nu_{zx} - \nu_{yz}\nu_{zy} - y\nu_{xy}\nu_{yz}\nu_{zx}}{E_x E_y E_z}$$

Because $[C_{mn}]$ is symmetric it follows that

$$\frac{\nu_{yx} + \nu_{zx}\nu_{yz}}{E_y E_z \Delta} = \frac{\nu_{xy} + \nu_{xz}\nu_{zy}}{E_x E_z \Delta} \tag{2.20}$$

$$\frac{\nu_{zx} + \nu_{yx}\nu_{zy}}{E_y E_z \Delta} = \frac{\nu_{xz} + \nu_{xy}\nu_{yz}}{E_x E_y \Delta} \tag{2.21}$$

$$\frac{\nu_{zy} + \nu_{zx}\nu_{xy}}{E_x E_z \Delta} = \frac{\nu_{yz} + \nu_{xz}\nu_{yx}}{E_x E_y \Delta} \tag{2.22}$$

2.4.2 Wave propagation in an orthotropic solid

To look at wave propagation in a solid, the laws of motion are combined with Hooke's Law and the definition of stresses and strains. A restatement of Newton's second law ($F = ma$) can be made in terms of stress:

$$\frac{\partial \sigma_{ij}}{\partial \mathrm{x}_j} = \rho \frac{\partial^2 \xi_i}{\partial t^2} \tag{2.23}$$

where ξ_i is the displacement in direction i, x_j is the coordinate in direction j, ρ is the density of the solid, and t is time [Bucur 2006]. Combining equations 2.4, 2.5, and 2.23 yields the wave equation

$$\rho \frac{\partial^2 \xi_i}{\partial t^2} - C_{ijkl} \frac{\partial^2 \xi_k}{\partial \mathrm{x}_k \partial \mathrm{x}_j} = 0 \tag{2.24}$$

The direction of travel for a wave is defined by a unit *propagation vector* $\mathbf{n}$. The *wavenumber* (or wave vector) $\boldsymbol{\kappa}$ is defined as

$$\boldsymbol{\kappa} = \frac{2\pi}{\lambda}\mathbf{n} = \frac{\omega}{v}\mathbf{n} \tag{2.25}$$

where λ is the wavelength, ω is angular frequency, and v is the phase velocity [Bucur 2006]. The wavenumber is normally denoted by the symbol k, but to avoid a conflict with the subscript used prior, is denoted here by $\boldsymbol{\kappa}$. The displacement of a wave in direction i can then be defined as

$$\xi_i = A_i \exp\left(\sqrt{-1}(\kappa_j \mathrm{x}_j - \omega t)\right) \tag{2.26}$$

where A_i is the amplitude in direction i, κ_j is the component of the wave vector in direction j, and x_j is the coordinate in direction j [Bucur 2006]. The amplitude A_i can be rewritten as AP_m, where P_m are the components of a unit *polarisation vector* $\mathbf{P}$. The polarisation vector is a unit vector in the direction of the displacement. For a longitudinal wave the polarisation and propagation vectors are aligned: $\mathbf{P} = \mathbf{n}$. In a shear wave the displacement is orthogonal to the direction of travel: $\mathbf{P} \perp \mathbf{n}$.

Coefficients:	${n_x}^2$	${n_y}^2$	${n_z}^2$	$n_y n_z$	$n_x n_z$	$n_x n_y$
Γ_{xx}	C_{11}	C_{66}	C_{55}	$2C_{56}$	$2C_{15}$	$2C_{16}$
Γ_{yy}	C_{66}	C_{22}	C_{44}	$2C_{24}$	$2C_{46}$	$2C_{26}$
Γ_{zz}	C_{55}	C_{44}	C_{33}	$2C_{34}$	$2C_{35}$	$2C_{45}$
Γ_{xy}	C_{16}	C_{26}	C_{45}	$C_{25}+C_{46}$	$C_{14}+C_{56}$	$C_{12}+C_{66}$
Γ_{xz}	C_{15}	C_{46}	C_{35}	$C_{36}+C_{45}$	$C_{13}+C_{55}$	$C_{14}+C_{56}$
Γ_{yz}	C_{56}	C_{24}	C_{34}	$C_{23}+C_{44}$	$C_{36}+C_{45}$	$C_{25}+C_{46}$

Table 2.1 Components of the Kelvin-Christoffel tensor for a general anisotropic solid.

After substitution of Equation 2.26 into Equation 2.24, taking derivatives and rearranging:

$$(C_{ijkl}\, n_j\, n_k - \delta_{ik}\, \rho\, v^2)P_m = 0 \tag{2.27}$$

δ_{ik} is a Kronecker tensor; if $i = k$ then $\delta_{ik} = 1$, if $i \neq k$ then $\delta_{ik} = 0$.

2.4.2.1 Christoffel's equations

The Kelvin-Christoffel tensor Γ_{ik} is defined as [Bucur 2006]:

$$\Gamma_{ik} = C_{ijkl}\, n_j\, n_l \tag{2.28}$$

Since $[C_{ijkl}]$ is symmetric, the matrix $[\Gamma_{ik}]$ is also symmetric. Table 2.1 shows the components of each Kelvin-Christoffel term for an anisotropic solid using coefficients from the reduced form stiffness tensor $[C_{mn}]$. For an orthotropic solid, many of the coefficients of $[C_{mn}]$ are zero. Γ_{ik} values for an orthotropic solid are shown in Equation 2.29.

$$\begin{aligned}
\Gamma_{xx} &= {n_x}^2 C_{11} + {n_y}^2 C_{66} + {n_z}^2 C_{55} \\
\Gamma_{yy} &= {n_x}^2 C_{66} + {n_y}^2 C_{22} + {n_z}^2 C_{44} \\
\Gamma_{zz} &= {n_x}^2 C_{55} + {n_y}^2 C_{44} + {n_z}^2 C_{33} \\
\Gamma_{xy} &= n_x\, n_y (C_{12} + C_{66}) \\
\Gamma_{xz} &= n_x\, n_z (C_{13} + C_{55}) \\
\Gamma_{yz} &= n_y\, n_y (C_{23} + C_{44})
\end{aligned} \tag{2.29}$$

The Kelvin-Christoffel tensor substituted into Equation 2.27 gives

$$(\Gamma_{ik} - \delta_{ik}\, \rho\, v^2)P_m = 0 \tag{2.30}$$

or, expressed as a matrix:

$$\begin{bmatrix} \Gamma_{xx} - \rho v^2 & \Gamma_{xy} & \Gamma_{xz} \\ \Gamma_{xy} & \Gamma_{yy} - \rho v^2 & \Gamma_{yz} \\ \Gamma_{xz} & \Gamma_{yz} & \Gamma_{zz} - \rho v^2 \end{bmatrix} \begin{bmatrix} P_x \\ P_y \\ P_z \end{bmatrix} = 0 \tag{2.31}$$

The set of equations defined by Equation 2.30 or Equation 2.31 are Christoffel's equa-

tions. They define the relationships between the elastic constants C_{mn} and the phase velocity of the wave v. There are three different solutions for v, since the velocity depends on the direction of propagation [Bucur 2006].

2.4.2.2 Axial propagation solution

Consider a wave propagating along the x axis in an unbounded medium. The propagation vector has values $n_x = 1, n_y = 0, n_z = 0$. Equation 2.31 reduces to

$$\begin{bmatrix} C_{11} - \rho v^2 & 0 & 0 \\ 0 & C_{66} - \rho v^2 & 0 \\ 0 & 0 & C_{55} - \rho v^2 \end{bmatrix} \begin{bmatrix} P_x \\ P_y \\ P_z \end{bmatrix} = 0 \tag{2.32}$$

which has three solutions for v:

$$\begin{aligned} v_{\mathrm{P}} &= \sqrt{\frac{C_{11}}{\rho}} \qquad (P_x = 1) \\ v_{\mathrm{S1}} &= \sqrt{\frac{C_{66}}{\rho}} \qquad (P_y = 1) \\ v_{\mathrm{S2}} &= \sqrt{\frac{C_{55}}{\rho}} \qquad (P_z = 1) \end{aligned} \tag{2.33}$$

The solutions for v depend on the polarization vector. All three solutions are propagating along the x-axis. The first solution is a p-wave. The latter two are shear waves with displacements in the xy and xz planes respectively. All waves have independent phase velocities.

2.4.2.3 Off axis propagation solution

A similar, three part solution can be developed for an acoustic wave travelling at an angle to one of the principle axes of propagation. Let the direction of propagation be at some angle α to the x axis, in the x, y plane. The propagation vector $\mathbf{n}$ has components $n_x = \cos\alpha, n_y = \sin\alpha, n_z = 0$. Equation 2.31 can be reduced to

$$\begin{bmatrix} \Gamma_{xx} - \rho v^2 & \Gamma_{xy} & 0 \\ \Gamma_{xy} & \Gamma_{yy} - \rho v^2 & 0 \\ 0 & 0 & \Gamma_{zz} - \rho v^2 \end{bmatrix} \begin{bmatrix} P_x \\ P_y \\ P_z \end{bmatrix} = 0 \tag{2.34}$$

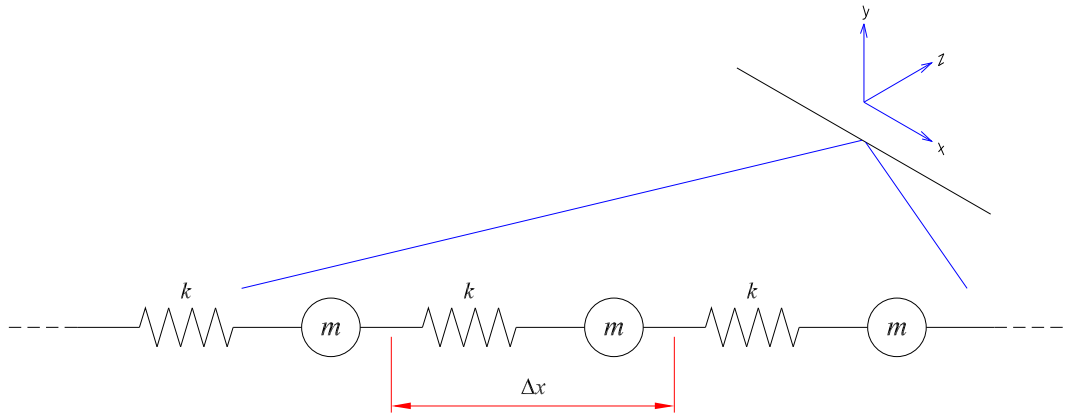

Figure 2.1 An infinitely thin rod from an unbounded, orthotropic solid. The rod is aligned with the x axis, so it can be modelled by a chain of mechanical elements.

Three solutions can be developed for v, as before, dependent on the polarisation:

$$\begin{aligned} v_{\mathrm{QP}} &= \sqrt{\frac{(\Gamma_{xx}+\Gamma_{yy})+\sqrt{(\Gamma_{xx}-\Gamma_{yy})^2+4(\Gamma_{xy})^2}}{2\rho}} \\ v_{\mathrm{QS}} &= \sqrt{\frac{(\Gamma_{xx}+\Gamma_{yy})-\sqrt{(\Gamma_{xx}-\Gamma_{yy})^2+4(\Gamma_{xy})^2}}{2\rho}} \\ v_{\mathrm{S}} &= \sqrt{\frac{\Gamma_{zz}}{\rho}} \end{aligned} \tag{2.35}$$

With these solutions, rather than clearly separated p-waves and shear waves, there is one shear wave, a quasi-shear wave, and a quasi-p-wave [Bucur 2006].

It is quite clear that the velocity solution for off-axis propagation is much more complex than when the propagation vector has only one non-zero component. Similarly, the general solution for propagation at an angle to all three axes is much more complex again. It is not necessary for this description of acoustic propagation and so is not presented.

2.5 DISPERSION

Dispersion is a phenomenon that accompanies attenuation in acoustic propagation. Dispersion causes the phase velocity of an acoustic wave to be dependent on the travelling wave's frequency.

Consider a line along one of the axes in an unbounded, orthotropic solid. This line can be modelled by a chain of mechanical elements of infinitely short length. Figure 2.1 depicts such a representation for a lossless medium. Note that this model can only be used to demonstrate p-waves, not shear as it has no lateral components.

Electrical	Mechanical Equivalent 1	Mechanical Equivalent 2
Voltage	Point velocity (u)	Force (F)
Current	Force (F)	Point velocity (u)
Resistance	Lubricity (inverse damping)	Damping coefficient
Capacitance	Mass	Compliance (inverse stiffness)
Inductance	Compliance (inverse stiffness)	Mass

Table 2.2 Electrical–mechanical conversions.

2.5.1 Model conversion

Rather than analyse the problem in terms of mechanical parameters, the model can be converted into an electrical transmission line. Each component of the mechanical model can be converted into an equivalent electrical component. Table 2.2 gives two possible systems for converting mechanical components into electrical components depending on the choice of mechanical equivalents for voltage and current. While either conversion system is equally valid, the voltage → velocity conversion makes for easier analysis of the electrical transmission line. This is the conversion used for the following analysis.

It should be noted that the velocity considered here is the point velocity, not the phase velocity. See Section 2.2.3 for velocity notation.

2.5.2 Transmission line model

Electrical transmission lines are conductive paths from one node to another in a circuit. Transmission line models are used to analyse or predict the voltage and current propagation in conductors ranging from circuit board tracks to high voltage power lines. Converting from a mechanical model into an electrical model allows the use of well defined results for determining electrical transmission along the line to analyse acoustic propagation.

Figure 2.2 shows a single element from a general electrical model in the Laplace domain. The series components are combined into an impedance per unit length $Z(s)$. The parallel components are combined into an admittance per unit length $Y(s)$. The length of each module is Δx. The s variable of the Laplace domain is

$$\begin{aligned} s &= \sigma + j2\pi f \\ &= \sigma + j\omega \end{aligned} \tag{2.36}$$

However, in the steady state $\sigma = 0$. For more information on the Laplace domain or Laplace transform see [Ambardar 1999], or [Cunningham and Miller 1995].

The output voltage and current from each module can be defined in terms of the input voltage and current:

$$V(x + \Delta x, s) = V(x, s) - \Delta x\, Z(s) I(x, s) \tag{2.37}$$

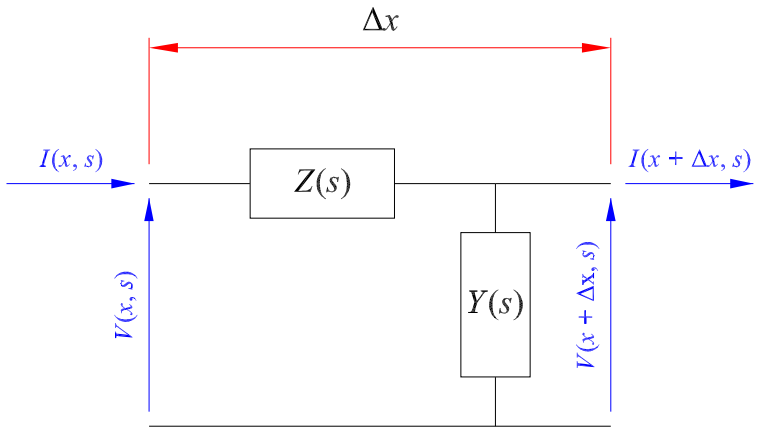

Figure 2.2 Generic electrical transmission line model (Laplace domain). $Z(s)$ is the series impedance per unit length and $Y(s)$ is the parallel admittance per unit length.

$$I(x+\Delta x,s) = I(x,s) - \Delta x\, Y(s) V(x+\Delta x,s) \tag{2.38}$$

Defining $\Delta V(x,s)$ and $\Delta I(x+\Delta x,s)$ allows the computation of the voltage and current gradients:

$$\Delta V(x,s) = -\Delta x\, Z(s) I(x,s) \tag{2.39}$$
$$\Delta I(x+\Delta x,s) = -\Delta x\, Y(s) V(x+\Delta x,s) \tag{2.40}$$

Each module is of infinitesimal length. Let Δx tend to 0 to produce differential equations:

$$\frac{\partial}{\partial x} V(x,s) = \lim_{\Delta x \to 0} \frac{V(x,s)}{\Delta x} \tag{2.41}$$
$$= -Z(s) I(x,s) \tag{2.42}$$
$$\frac{\partial}{\partial x} I(x,s) = \lim_{\Delta x \to 0} \frac{I(x,s)}{\Delta x} \tag{2.43}$$
$$= -Y(s) V(x,s) \tag{2.44}$$

To solve the differential equations, find the second derivative of the voltage:

$$\frac{\partial^2}{\partial x^2} V(x,s) = \frac{\partial}{\partial x}\left[\frac{\partial}{\partial x} V(x,s)\right] \tag{2.45}$$
$$= -Z(s) \frac{\partial}{\partial x} I(x,s) \tag{2.46}$$
$$= Z(s) Y(s) V(x,s) \tag{2.47}$$

This can be rearranged into

$$\frac{\partial^2}{\partial x^2} V(x,s) - \gamma^2(s) V(x,s) = 0 \tag{2.48}$$

where $\gamma(s)$ is the *propagation constant*:

$$\gamma(s) = \sqrt{Z(s)Y(s)} \tag{2.49}$$

2.5.2.1 General solutions

The general solution to Equation 2.48 is [Young 1999]:

$$V(x,s) = V_1(s)\exp(-\gamma(s)x) + V_2(s)\exp(\gamma(s)x) \tag{2.50}$$

where $V_1(s)$ and $V_2(s)$ are determined from the endpoint boundary conditions. The solution supports waves travelling in both directions.

The same type of solution exists for the current:

$$I(x,s) = I_1(s)\exp(-\gamma(s)x) + I_2(s)\exp(\gamma(s)x) \tag{2.51}$$

where

$$I_1(s) = \frac{V_1(s)}{Z_0(s)} \tag{2.52}$$

$$I_2(s) = -\frac{V_2(s)}{Z_0(s)} \tag{2.53}$$

and $Z_0(s)$ is the *characteristic impedance*, defined by

$$Z_0(s) = \frac{Z(s)}{\gamma(s)} = \sqrt{\frac{Z(s)}{Y(s)}} \tag{2.54}$$

2.5.2.2 Transfer function

To develop a transfer function for propagation along the line, consider the voltage input at some point x_0 and the output voltage at some point $x_0+\Delta x$. Only the wave travelling in the forward direction is considered. This sets the variable V_2 in Equation 2.50 to zero. The transfer function is the ratio of the output voltage to input voltage:

$$H(s,\Delta x) = \frac{V(s,x_0+\Delta x)}{V(s,x_0)} \tag{2.55}$$

$$= \frac{V_1(s)\exp(-\gamma(s)(x_0+\Delta x))}{V_1(s)\exp(-\gamma(s)(x_0))} \tag{2.56}$$

$$= \exp(-\gamma(s)\,\Delta x) \tag{2.57}$$

The transfer function for a wave propagating down the line is independent of the absolute position on the line and a function only of the distance traveled and frequency.

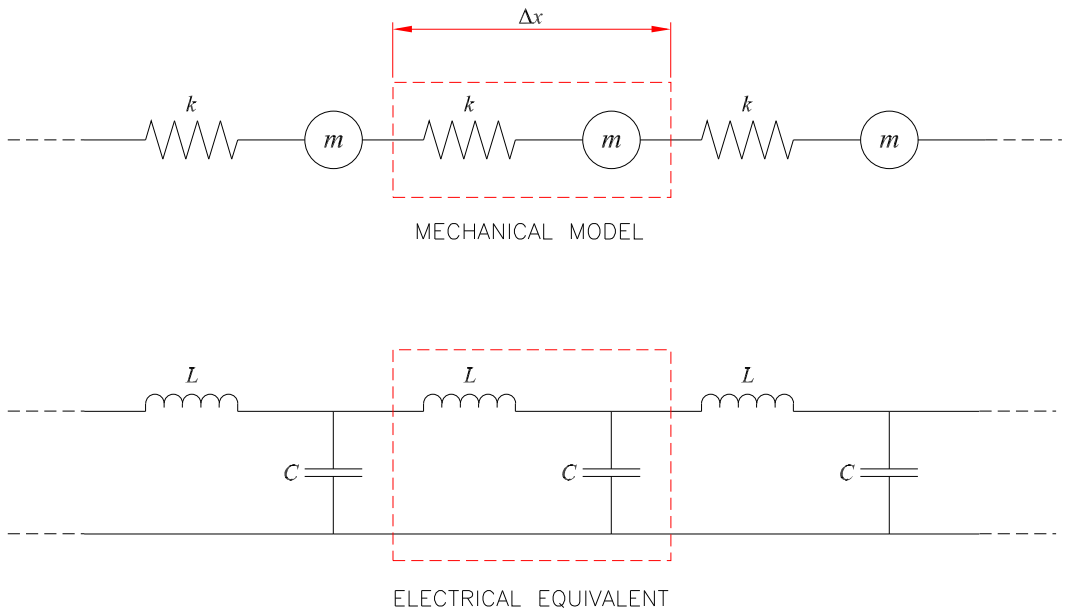

Figure 2.3 Lossless transmission line models.

2.5.2.3 Attenuation and velocity

The attenuation (absorption) and the phase velocity of the wave travelling on the transmission line are defined in terms of the propagation constant [Young 1999]. The propagation constant can be split into its real and imaginary components:

$$\gamma(s) = \alpha(s) + j\beta(s) \tag{2.58}$$

The real component is related to the absorption $a(s)$ in dB/m by

$$a(s) = 0.868\,\alpha(s) \tag{2.59}$$

The complex component is related to the phase velocity $v(s)$ by

$$v(s) = \frac{\omega}{\beta(s)} \tag{2.60}$$

where ω is the angular frequency.

2.5.3 Lossless model

Consider the case where transmission through an elastic solid is a lossless process. The mechanical model in Figure 2.3 depicts a mechanical model for propagation along an axis of an orthotropic medium. Since there are no mechanical dashpot components, the electrical equivalent model has no resistors. Because of this, there are no mechanisms for energy loss.

The series impedance of the electrical model is

$$Z(s) = Ls \tag{2.61}$$

The parallel admittance is

$$Y(s) = Cs \tag{2.62}$$

The propagation constant for the lossless electrical transmission line is then:

$$\gamma(s) = s\sqrt{L\,C} \tag{2.63}$$

The propagation constant has no real component so there is no attenuation, as defined by Equation 2.59.

The phase speed, as defined by Equation 2.60, is

$$v = \frac{1}{\sqrt{L\,C}} \tag{2.64}$$

The phase velocity is a constant. All frequencies travel at the same speed, so the lossless model is non-dispersive.

2.5.4 More complex models

Real solids, however, are generally not lossless; they absorb energy from acoustic waves. Absorption can be modelled in a mechanical transmission line by including one or more Maxwell elements parallel to the mechanical spring.

A Maxwell element consists of a spring and dashpot in series [Smith 1966]. Figure 2.4 shows a mechanical model with three Maxwell type loss mechanisms. A solid may have many loss mechanisms depending on its structure. Each mechanism has a *relaxation time* τ, which characterises the exponential rate of decay of the force on the mechanical element.

The component $Y(s)$ in this electrical model is the same as it was for the lossless case:

$$Y(s) = sC \tag{2.65}$$

The $Z(s)$ component has changed:

$$Z(s) = \frac{1}{\frac{1}{sL_0} + \sum_{n=1}^{N} \frac{1}{sL_n + R_n}} \tag{2.66}$$

$$= \frac{s}{\frac{1}{L_0} + \sum_{n=1}^{N} \frac{1}{L_n} \frac{s}{s + \frac{1}{\tau_n}}} \tag{2.67}$$

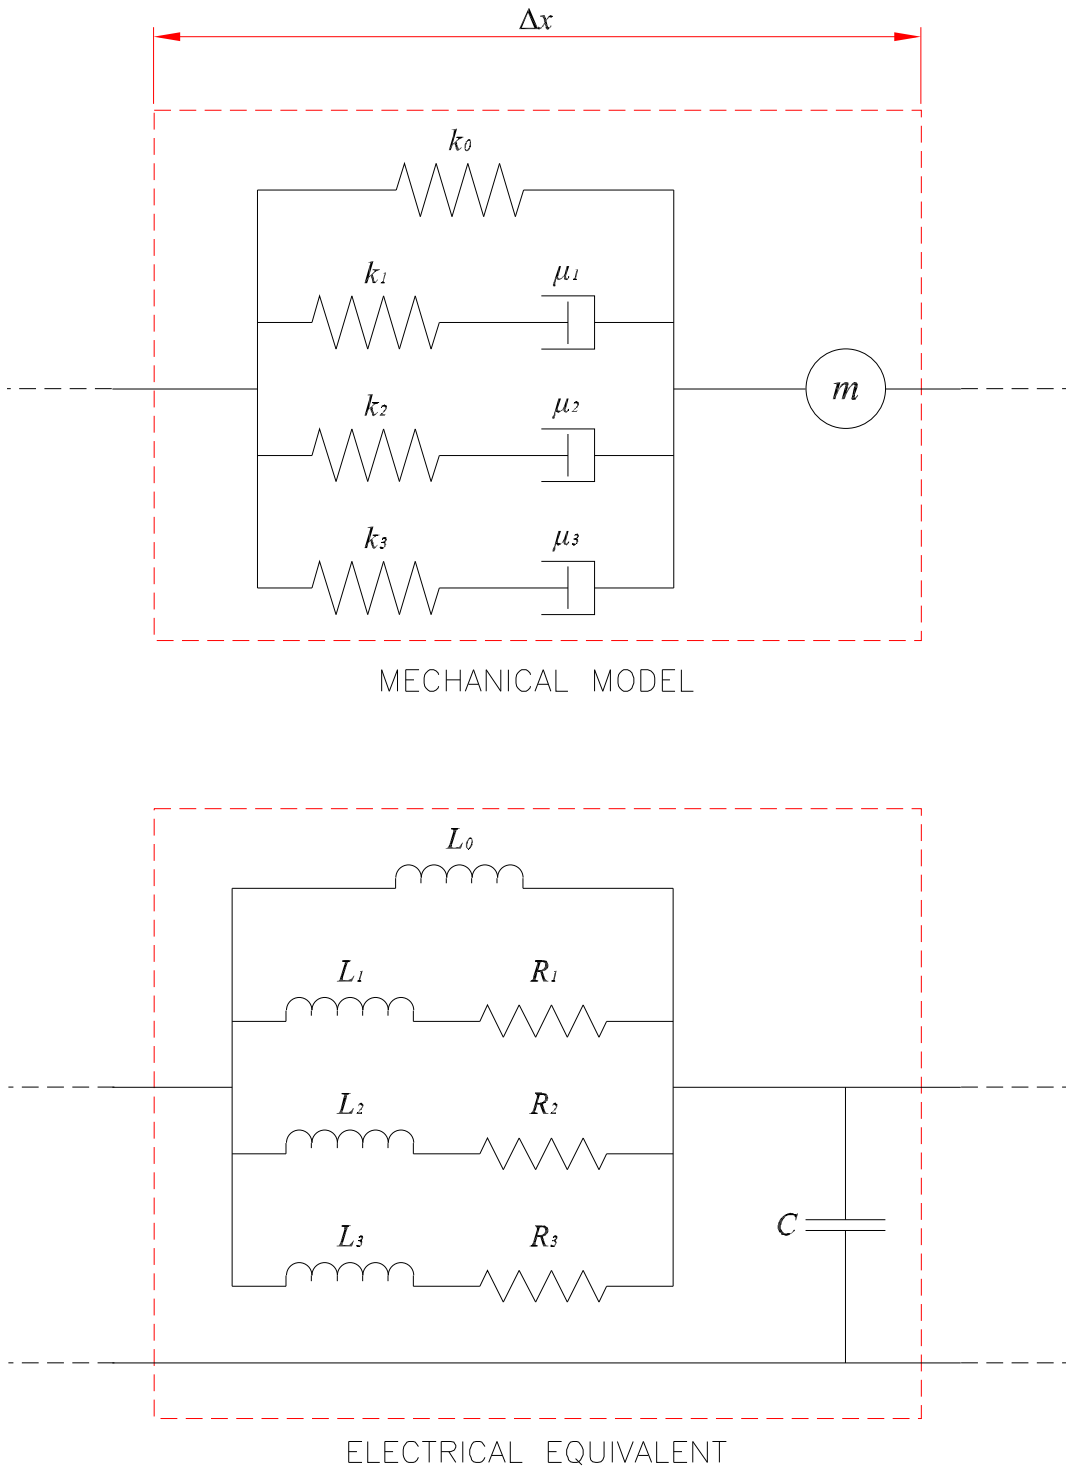

Figure 2.4 More complex transmission line models including loss mechanisms. The mechanical model, and equivalent electrical model both contain three loss mechanisms.

where the n^{th} relaxation time τ_n is

$$\tau_n = \frac{L_n}{R_n} \tag{2.68}$$

From Equation 2.49 the propagation constant for the complex model is

$$\gamma = s\sqrt{\frac{C}{\frac{1}{L_0} + \sum_{n=1}^{N} \frac{1}{L_n} \frac{s}{s + \frac{1}{\tau_n}}}} \tag{2.69}$$

2.5.5 Mechanical notation

To make useful simulations of acoustic propagation in solids the electrical model is converted into mechanical notation. The model is transformed by reversing the mechanical $\rightarrow$ electrical conversion used before. Under this notation the parallel admittance (per unit length) is

$$Y(s) = sm \tag{2.70}$$

where m is the mass per unit length. The series admittance (per unit length) is now defined in terms of the spring constants k_i:

$$Z(s) = \frac{s}{k_0 + \sum_{n=1}^{N} k_n \frac{s}{s + \frac{1}{\tau_n}}} \tag{2.71}$$

The resulting propagation constant is then

$$\gamma = s\sqrt{\frac{m}{k_0 + \sum_{n=1}^{N} k_n \frac{s}{s + \frac{1}{\tau_n}}}} \tag{2.72}$$

This notation, however, implies a bounded solid. The spring constants are dependent on the cross sectional area, as is the mass per unit length. The mass can be redefined in terms of the density ρ:

$$m = \rho A \tag{2.73}$$

where A is the cross sectional area. Spring constants can also be redefined in terms of the elastic moduli.

$$k_i = AE_i \tag{2.74}$$

Making these substitutions results in a propagation constant expressed in terms of

variables that describe an unbounded, homogeneous medium:

$$\gamma = s\sqrt{\frac{\rho}{E_0 + \sum_{n=1}^{N} E_n \frac{s}{s + \frac{1}{\tau_n}}}} \tag{2.75}$$

At low frequencies, the propagation constant is dominated by the E_0 term. As the frequency is increased, the sum of the E_n terms come to dominate the equation.

2.5.6 Using the model to simulate materials

The mechanical model developed through the used of electrical transmission line theory can now be used to simulate the dispersive effects of a wave propagating through a solid. Figures 2.5 and 2.6 show the phase speed and attenuation profiles of two contrived models. The simulated medium in Figure 2.5 has a single loss mechanism, while the simulated medium in Figure 2.6 has three loss mechanisms.

In both cases, the attenuation and phase speed are functions of frequency as opposed to the constant values derived for the lossless model. Higher frequencies have higher phase speeds and greater attenuation. Each successive loss mechanism increases the phase speed and attenuation for frequencies greater than the inverse of the relaxation time. The magnitude of the increase in phase speed and attenuation is defined by the modulus of elasticity of the loss mechanism.

The materials are contrived for demonstration purposes only and are not indicative of any real materials. The figures show frequencies far higher than those of acoustic interest.

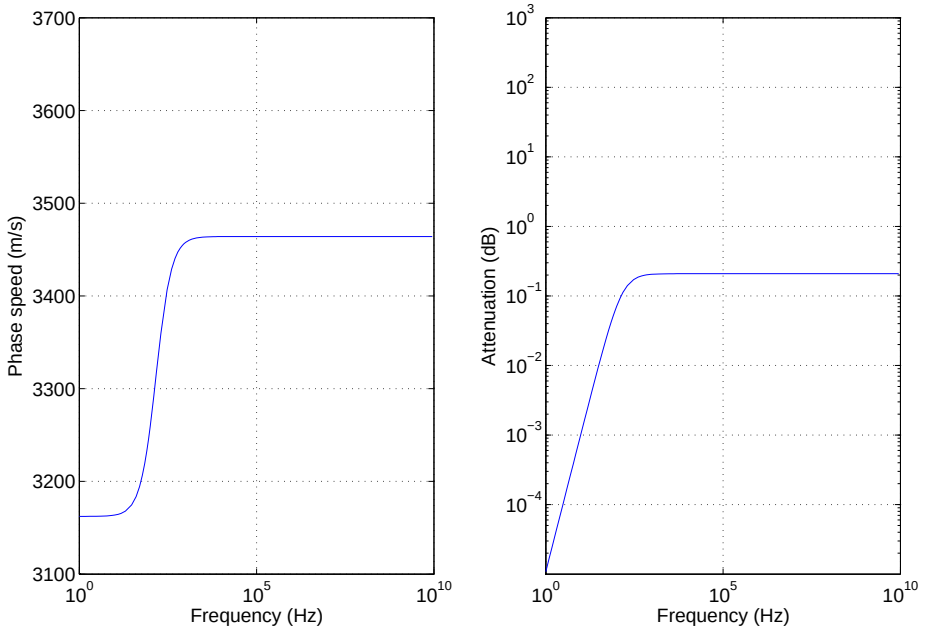

Figure 2.5 Phase speed and attenuation plots for a model with one relaxation time. $\rho = 1000\,\mathrm{kg/m^3}$, $E_0 = 10\,\mathrm{GPa}$, $E_1 = 2\,\mathrm{GPa}$, $\tau_1 = 1\,\mathrm{ms}$.

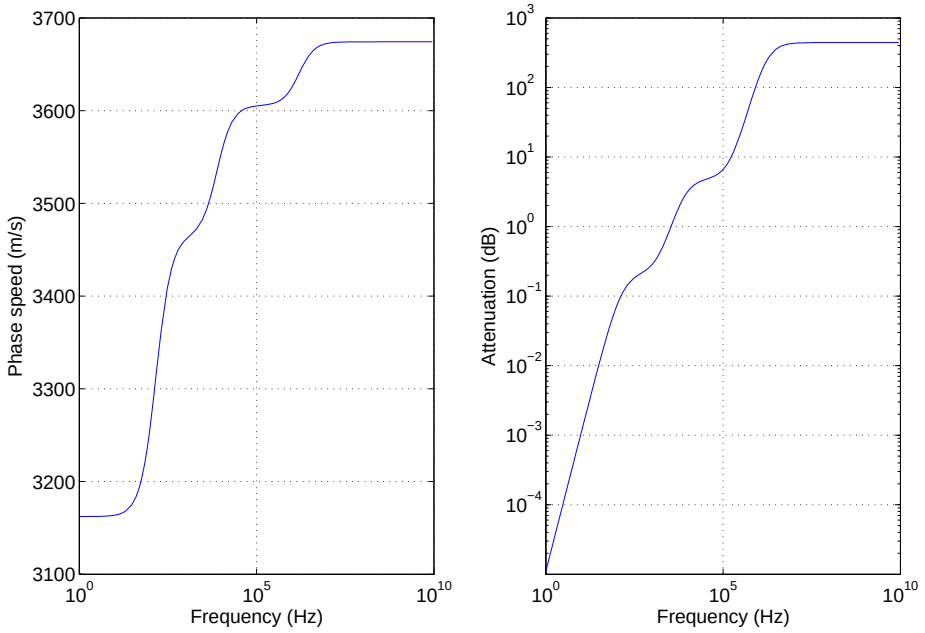

Figure 2.6 Phase speed and attenuation plots for a model with one relaxation time. $\rho = 1000\,\mathrm{kg/m^3}$, $E_0 = 10\,\mathrm{GPa}$, $E_1 = 2\,\mathrm{GPa}$, $\tau_1 = 1\,\mathrm{ms}$, $E_2 = 1\,\mathrm{GPa}$, $\tau_2 = 20\,\mu\mathrm{s}$, $E_1 = 0.5\,\mathrm{GPa}$, $\tau_3 = 0.1\,\mu\mathrm{s}$.

Chapter 3

ACOUSTIC MEASUREMENT OF WOOD

3.1 INTRODUCTION

This chapter applies the general stress wave fundamentals developed in the previous chapter to acoustic propagation in wood. This is done with the specific aim of finding the phase velocity of an acoustic wave in wood and relating this to stiffness.

Some assumptions must be made about the nature of wood as an acoustic medium. Section 3.2 shows that wood can be considered both orthotropic and homogeneous. The major provision for both of these assumptions to apply is that the wood be free from defects.

Section 3.3 begins the discussion of wood stiffness and relates it to the construction industry. The industrial test for stiffness is reviewed and compared to the more scientific definition developed for orthotropic solids in the previous chapter. Then follows a look at the two current systems for measuring stiffness using acoustics, stress wave timers and resonance testers.

A brief look at the creation of acoustic pulses in wood is taken in Section 3.5, followed by a review of the methods used to detect and record acoustic signals. Section 3.6 is an examination of an accelerometer, how it affects the measurement process, and suggests alterations and improvements to the design.

Dispersion is an important feature of acoustic propagation that has largely been ignored or marginalised in interpretations of stress wave timing measurements. In a dispersive media like wood, the phase velocity of an acoustic wave is a function of frequency. Dispersion is dealt with in Section 3.7 with specific consideration given to both stress wave timer and resonance testing measurements.

The chapter concludes with a look at some of the problems with acoustics in wood that have not already been discussed, and a summary and comparison of both velocity measuring systems.

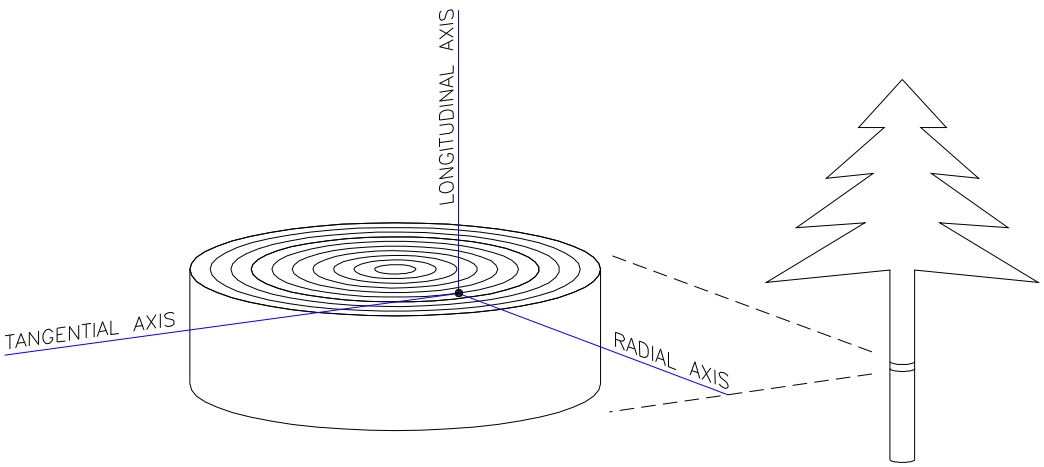

Figure 3.1 The three principle axes for a point in a standing tree.

3.2 WOOD AS AN ACOUSTIC MEDIUM

Wood is an extremely complex material. However, despite the many natural, non-uniform processes by which it is grown, wood can be approximated as an orthotropic, homogeneous solid [Bucur 2006]. Of course, certain provisions need to be made in order for the approximation to conform to the reality.

3.2.1 Orthogonality

Wood can be modelled as an orthotropic solid [Bucur 2006]. Wood has three approximately orthogonal planes of elastic symmetry, as required for orthotropy (see Section 2.4). These planes are normal to the principle axes of the timber. The three axes are:

Longitudinal: the axis in the vertical direction, aligned with the stem.

Radial: the axis in the horizontal plane from the center of the tree, and across the growth rings.

Tangential: the axis in the horizontal plane, tangential to the growth rings.

These axes are shown for a point the stem of a tree in Figure 3.1. It must be noted that the axes are dependent on location. The tangential axis is only *approximately* straight at the specified point. If the point is moved in the longitudinal or radial directions the axis definitions hold but not for movement in the tangential direction. Moving along the tangential line means that the axis is no longer tangential to the growth rings and the planes of elastic symmetry are no longer orthogonal. So, orthogonality holds only while considering a region where the tangential movement is much smaller than the radial coordinate so that the growth rings are approximately straight.

Using the subscripts L, R, and T to denote these axes, Equation 2.7 can be restated to define the reduced form of Hooke's law specifically for wood:

$$\begin{bmatrix} \sigma_{LL} \\ \sigma_{RR} \\ \sigma_{TT} \\ \sigma_{RT} \\ \sigma_{TL} \\ \sigma_{LR} \end{bmatrix} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & 0 & 0 & 0 \\ C_{12} & C_{22} & C_{23} & 0 & 0 & 0 \\ C_{13} & C_{23} & C_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{66} \end{bmatrix} \begin{bmatrix} \epsilon_{LL} \\ \epsilon_{RR} \\ \epsilon_{TT} \\ \epsilon_{RT} \\ \epsilon_{TL} \\ \epsilon_{LR} \end{bmatrix} \tag{3.1}$$

3.2.2 Homogeneity

Wood is approximately homogeneous [Bucur 2006]. This means that the properties of wood are independent of location and purely a function of the direction the wave is travelling. This assumption is accurate, provided that there are no major defects within the timber sample [Bucur 2006]. Section 3.8.3 details types of defects that affect acoustic propagation.

When looking at the structure of wood in microscopic detail it appears distinctly inhomogeneous. Figure 3.2 shows a microscopic view of a wood sample from the online *Wood anatomy of central European Species* [Schoch *et al.* 2004]. In this image, the density and spacings of the tracheid cell walls vary greatly depending on location within the growth ring. The three large white dots are resin canals. These are not considered to be defects as they naturally occur throughout the wood. Similarly inhomogeneous effects appear when observing wood samples in the radial and tangential directions. All these observations point to an strongly inhomogeneous material.

However, a solid can be approximated as being homogeneous for wave propagation if the inhomogeneous elements presented to the wavefront are smaller than a tenth of the wavelength in at least one direction [Hunter 2006]. Figure 3.3 shows two solids with inhomogeneous artefacts to illustrate this point. Ignore any wave-guiding effects caused by the bounding dimensions of the cylinder. If the wavelength of the transmitted wave is just over ten times larger than the diameter of the cylindrical artefact in the middle of the right hand solid, then this solid can be considered homogeneous. There will be some diffraction and reflection of energy at the boundary of the artefact, but the bulk of the wave's energy will travel around the artefact and continue on in the same direction of propagation. For the solid on the left, however, the size of the artefact is still much larger than a tenth of the wavelength so it is *not* homogeneous. When the wave hits the boundary of this artefact there will be significant reflections. The depth of the artefacts does not affect homogeneity.

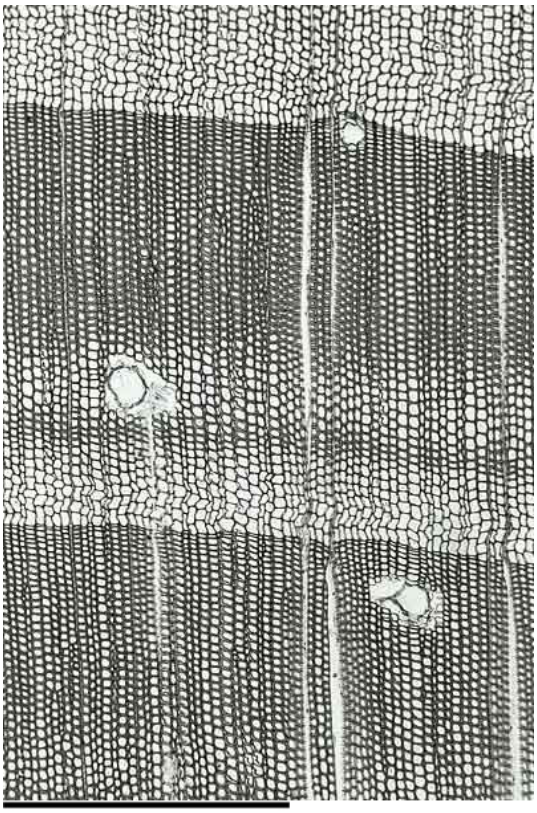

Figure 3.2 Microscopic view of a wood sample in the longitudinal direction. The black bar in the lower right is 1 mm long. The specimen is from the stem of a Pinus Silvestris (Common/Scots Pine) and is typical of most pine species.

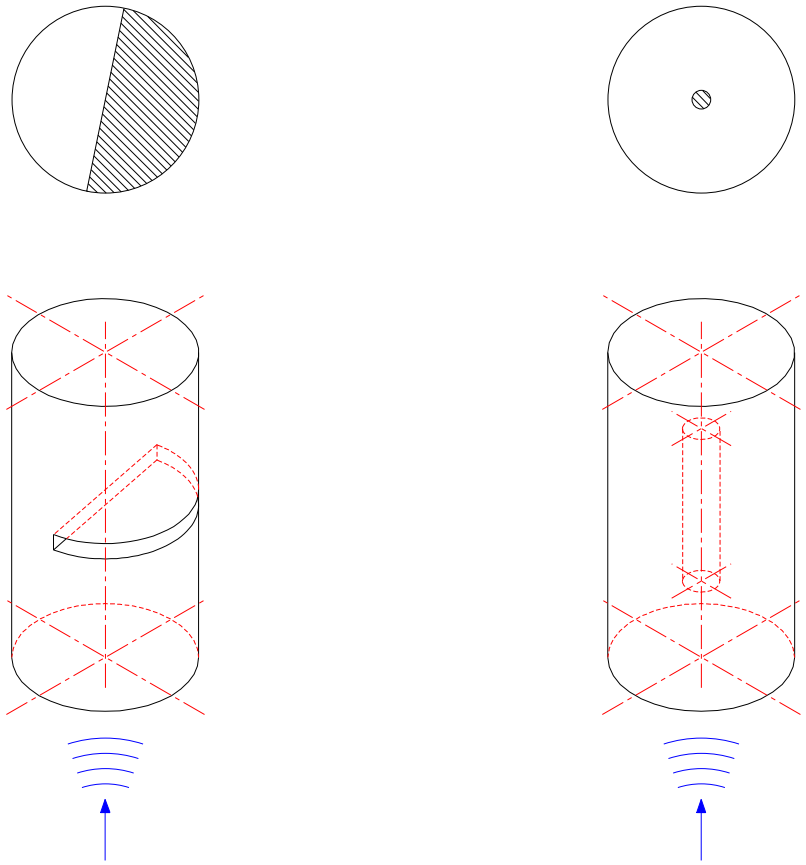

Figure 3.3 Two solids with inhomogenous artefacts. In both cases assume that the artefact has a different acoustic impedance to the solid. The direction of acoustic propagation is vertically along the cylinders. The hatched areas show the face of the artefacts presented to the wavefront.

v	f	λ	$\lambda/10$
1500 m/s	1 kHz	1.5 m	15 cm
1500 m/s	10 kHz	15 cm	1.5 cm
1500 m/s	100 kHz	1.5 cm	1.5 mm
1500 m/s	1 MHz	1.5 mm	0.15 mm
1500 m/s	10 MHz	0.15 mm	15 μm

Table 3.1 Acoustic wavelengths for a variety of frequencies.

The wavelength λ of an acoustic wave is defined as

$$\lambda = \frac{v}{f} \tag{3.2}$$

where v is the phase velocity of the wave and f is frequency. Note that, in general, v is a function of frequency. To find the frequency range for which homogeneity will hold in wood the wavelength is calculated for an low velocity acoustic wave. Usual velocity measurements report speeds of above 3000 m/s. Let v be an extreme low at 1500 m/s. Wavelengths for a variety of frequencies are shown in Table 3.1.

Returning to Figure 3.2; the largest features in the wood are approximately 0.25 mm across. Given that features need to be smaller than $\lambda/10$ to be homogeneous, the calculated wavelengths in Table 3.1 suggest that, even for extremely slow acoustic waves, wood will be approximately homogeneous for most frequencies below 1 MHz.

3.3 WOOD STIFFNESS

The stiffness of wood is an important characteristic of wood grown for construction. Stiffness relates strain to stress, as established in Section 2.2.1. Since wood is orthotropic, stiffness depends on the direction of the applied force. But for construction applications it is the longitudinal stiffness that is critical.

Wood is always sawn along the longitudinal axis to yield useful sized boards. There are a number of different cutting arrangements for a log. These include: through and through, plain sawn, rift sawn, quarter sawn, true quarter sawn, and commercially quarter sawn [Corbett 2004]. In every arrangement however the wood is still cut down the longitudinal axis, and it is this stiffness that is the most important.

There are two primary applications of wood used in construction; they are shown in Figure 3.4. Other applications of wood in construction are usually a reorientation or combination of these. In the first, the timber is aligned vertically (e.g. a stud in a wall). The timber is compressed in the vertical direction. In the second case, the timber is aligned in the horizontal (e.g. a floor joist). Although the force is applied orthogonally to the longitudinal direction, it is the longitudinal stiffness that determines how much the timber deforms. There is no vertical compression of the timber. Instead the top edge of the beam compresses in the longitudinal direction. Simultaneously,

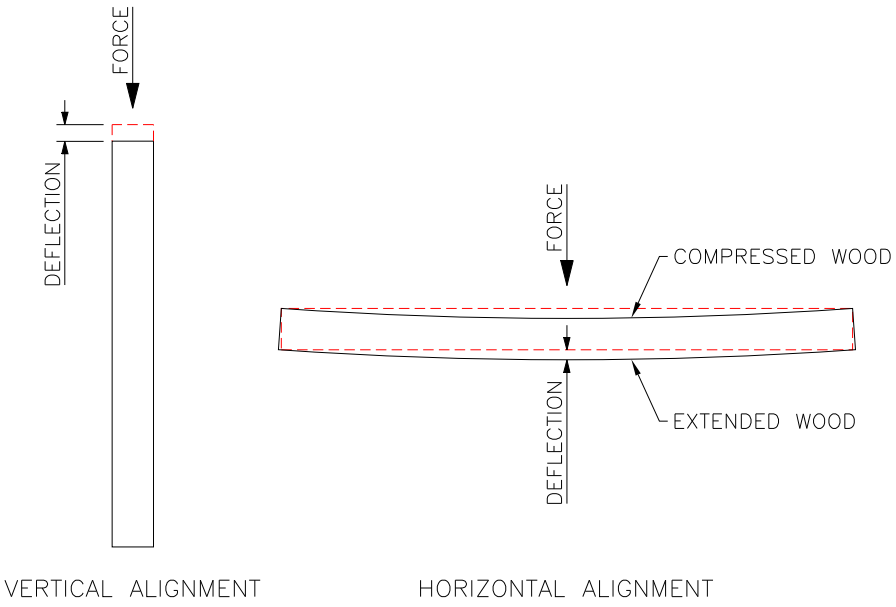

Figure 3.4 Two applications of wood used in construction.

the bottom edge of the beam stretches longtudinally. The opposed compression and tension produce a reactive force. Wood with higher longitudinal stiffness will suffer less deformation for the level of applied force in both cases.

3.3.1 Industry test

In order to grade wood stiffness at the sawmill a lateral force is applied to each plank of wood after they are sawn and the lateral deformation is measured[1] (the same principle as shown in Figure 3.4). As discussed above, this is a measurement of the longitudinal stiffness. Since this is the only significant stiffness for the forestry industry the term "stiffness" is used to refer to the longitudinal stiffness.

3.3.2 Stiffness using acoustics

The scientific definition of stiffness, as covered in Section 2.2.2, states that the stiffness of a solid is defined by a fourth rank stiffness tensor $[C_{ijkl}]$ with 81 terms. Wood is an orthotropic solid, allowing this tensor to be reduced to just 9 independent terms. If the stress on the timber if purely longitudinal, as discussed in Section 3.3, then the radial and tangential terms in the stress vector in Equation 3.1 are zero ($\sigma_{RR} = 0, \sigma_{TT} = 0$). Thus the stress in the longitudinal direction (σ_{LL}) is given by:

$$\sigma_{LL} = C_{11} \times \epsilon_{LL} \tag{3.3}$$

[1]See: ASTM D198/84 standard, Method of Static Tests of Timbers in Structural Sizes, of the American Society for Testing Materials

The industrially recognised stiffness of wood is therefore the C_{11} coefficient of the stiffness matrix.

In Section 2.4.2 it was established that, for an orthotropic medium, coefficients of the stiffness matrix $[C_{mn}]$ could be determined from the density of the medium and the velocity of particular travelling waves. In particular the velocity of a p-wave travelling along the longitudinal axis directly relates to the C_{11} term as shown in Equation 2.33. So, the results from the industrial test done at a sawmill to determine the stiffness of wood can also obtained by measuring the acoustic velocity of a p-wave travelling along the longitudinal axis of the wood.

There are two current technologies for measuring this acoustic velocity; stress wave timers and resonance testers. The systems are quite different but rely on similar methods of receiving acoustic data through an accelerometer. Some resonance testers use a microphone as an alternative to an accelerometer. Both systems have different, positive features in terms of operation and measurement accuracy.

3.3.3 Stress wave timers

Stress wave timers measure the time of flight of an acoustic wave between two points in a wood specimen. By measuring the time of flight and the distance between the points the phase velocity of the stress wave can be directly calculated ($v = d/t$). Previous generations of Treetap [Wang 2003], the Director ST300, the IML Impulse Hammer, and the Fakopp 2D systems are all stress wave timers.

Figure 3.5 shows a typical stress wave timer setup. An acoustic pulse is generated either with a piezoelectric device or by using a force hammer. The *time of flight* is measured between the generation of the pulse and detection by the receiving accelerometer.

Each part of the system causes a delay. t_w is the time of flight (the goal of the measurement). This is the time taken for the pulse to travel some distance d from the generator to the receiver. Delays are also introduced in the generation of the pulse and the receiving probe electronics. Let the delay from the start of timing (t_0) until the generated pulse (either electrically or with a hammer) enters the timber be t_g. Let the delay from when the acoustic pulse passes the tip of the accelerometer until the resulting voltage signal is interpreted as a pulse (t_1), be t_r. The length of t_r depends on the method used to interpret the voltage signal as well as the voltage signal itself. Interpolation methods, threshold levels, amplification, peak detection algorithms, pulse shape, and mechanical responses are all possible influences on this delay. t_g depends similarly on the method of pulse generation.

Assume that all delays are constants. The timer starts at t_0 when the pulse is generated and the signal is picked up some time later at t_1 when the pulse is detected.

$$t_1 - t_0 \quad = \quad t_g + t_w + t_r \tag{3.4}$$

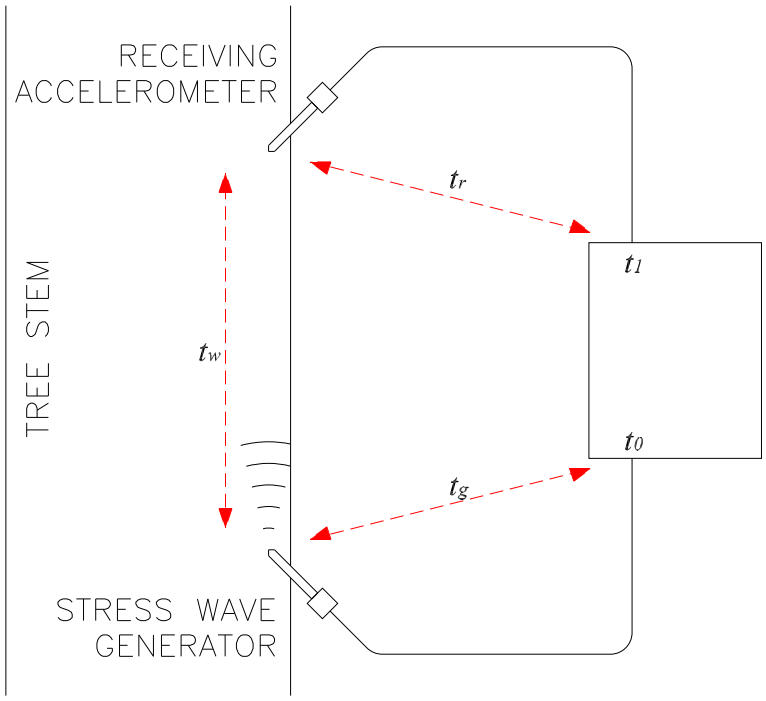

Figure 3.5 Typical stress wave timer setup.

$$t_w = t_1 - t_0 - t_g - t_r \tag{3.5}$$

To be able to estimate the time of flight through the wood accurately, t_g and t_r must be known.

3.3.3.1 Using two receivers

One of the key differences between the Treetap system and other stress wave timing devices is the use of two receivers. Figure 3.6 shows the setup of the Treetap system. In contrast to the typical setup shown in Figure 3.5, the time of flight is measured between when the first and second receivers to detect the pulse.

Again, assume that all the delays are constants. Let t_w and t_r be the same as the previous case and let t_1 and t_2 be the detection times of the two receivers.

$$t_2 - t_1 = t_w + (t_r - t_r) \tag{3.6}$$

$$t_w = t_2 - t_1 \tag{3.7}$$

In this case the generating delay (t_g) is removed completely and the receiver delays cancel. The time of flight through the wood is now dependent only on the difference between the detection time of receiver one and receiver two.

However, accurate measurement of the time of flight is not as simple as this argument suggests. The receiving delays cannot be relied upon to be identical. The reasons

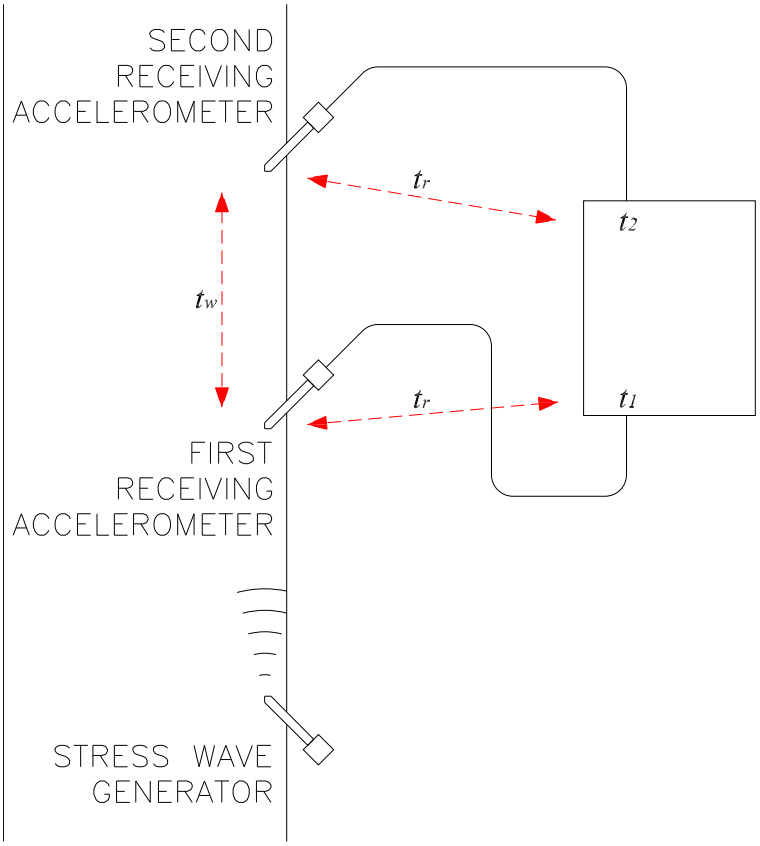

Figure 3.6 Dual-receiver stress waver timer setup.

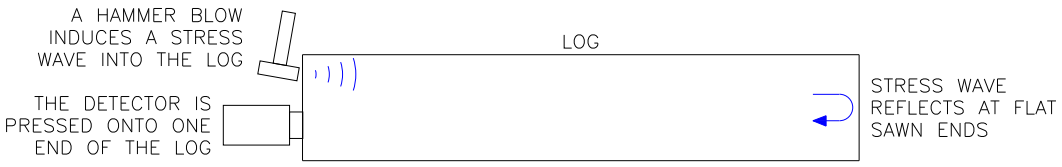

Figure 3.7 Resonance testing device setup.

for this are discussed in Section 3.7. Despite this, a two probe system still benefits from the removal of the generator delay and provides some (if not complete) cancellation of the receiver delays.

3.3.4 Resonance method

The alternative to a stress wave timer is a resonance tester, e.g., the Director HM200. This device can only be used on logs, since it relies on a stress wave reflecting at the ends of a sample [Andrews 2002]. Figure 3.7 shows how this type of tester is set up.

When the hammer strikes the log an acoustic pulse travels down the stem, reflects at the far end, travels back along the stem, reflects at the near end, and so on. The acoustic impedance match between air and wood is poor so most of the acoustic energy is reflected and only a small amount radiates into the air. The pulse continues to travel up and down the log until it finally dies away due to attenuation.

A single receiver is placed at the end of the log to collect acoustic data. A Fourier transform of the data gives a frequency domain representation of the result. The resonating signal is a series of peaks at harmonics of the frequency of the pulse's arrival at the accelerometer. These harmonics correspond to the longitudinal modes of resonance. The fundamental frequency of the resonance is an inverse measurement of the time of flight from the accelerometer to the far end of the log and back. By measuring the length of the log (L) and the fundamental frequency (f_0) the phase velocity of the pulse can be calculated:

$$\lambda = 2\,L \tag{3.8}$$

$$v = \lambda\, f_0 \tag{3.9}$$

$$v = 2\,L\, f_0 \tag{3.10}$$

3.4 INDUCING ACOUSTIC PULSES

There are two main ways of inducing an acoustic pulse into a timber sample; mechanical inputs like a hammer or a pendulum, or electro-mechanical inputs like a piezoelectric transducer. It is difficult to ascertain the exact profile of a pulse transmitted into a wood sample without accurately knowing the transfer functions of both the wood sample and the receiver. However, observation of a great number of trials suggests some strong correlations between particular methods of induction and the characteristics of the pulse.

Hammers are the most commonly seen system for inducing stress waves mechanically. Different types of hammers produce different pulses. The way an operator swings the hammer also affects the pulse. While there are a potentially endless variables in the type of hammer and how it is used, a few generalisations can be made:

- Hammers with a larger mass tend to impart more energy into the pulse as a result of having more momentum from the swing.
- Stiff construction material produces a higher bandwidth pulse, e.g., the pulse from a steel hammer has more high frequency components than a pulse from a soft rubber mallet.
- When a hammer strikes its target it is reflected sometime later by a rebounding compression wave. A relaxed grip and gentle swing reduces the probability that the hammer will be forced back onto the target, creating a second tap.

Alternatives to hammers include spring loaded devices or pendulums for imparting repeatable taps. For stress wave timers an object like a small metal spike or screw is usually embedded into the wood as a target. When the hammer makes contact a much higher bandwidth pulse is created, but the object can also create additional interference:

- Striking an embedded object with too much force can overcome the friction holding the object in place, moving it, and reducing the energy imparted to the pulse.
- Embedded objects can produce additional effects, e.g., a steel spike was observed adding a 30 kHz component to an acoustic signal. This was later attributed to the pulse resonating within the spike.
- The rebounding compression wave is stronger and faster when striking an embedded object, making it easier to create additional taps.

Resonance testers are able to take advantage of the flat end of the log and the preference for a high energy, low bandwidth pulse. Accordingly, the end of the log is usually struck directly with a hammer.

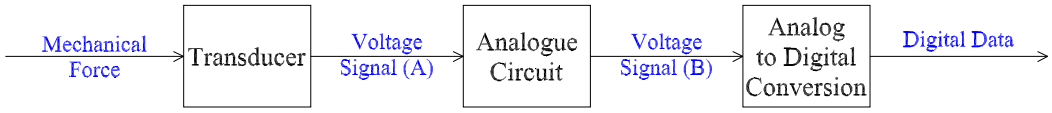

Figure 3.8 Diagram showing the conversion of a mechanical signal into recorded digital data.

Electro-mechanical transducers also create different pulse shapes depending on their design and the electrical input. The energy imparted by these devices is in general much, much less than even a small hammer is able to develop. They have two major advantages over a hammer tap however. The first is that the pulse shape is defined by the electrical input. By driving the device with different sources a variety of pulses can be created. The second advantage is repeatability. While a hammer induced pulse can change significantly for very small alterations in the operator's swing, a piezoelectric device transmits identical pulses each time.

3.5 RECEIVING ACOUSTIC PULSES

One of the difficulties of many applied technologies is in acquiring signals of interest in a useful format without altering either the system generating the signal, or the signal itself. For acoustic measurements in wood the signal of interest is the stress at some point as a result of an acoustic wave. The resulting measurement data, however, is several stages removed from this as shown in Figure 3.8. At each stage in this process there is potential for the acoustic signal can be distorted.

For the purposes of this section, specific reference will be given to the measurement process of the Treetap 5.0 system. This will not be an entirely accurate representation of all acoustic measurement systems for wood. However, the amplification, filtering, and digitisation process is typical of any sampling system and all the acoustic measuring systems discussed above use accelerometers of one type or another. While the accelerometers in each device are likely to be different, the analysis presented below could be easily applied to other accelerometers.

3.5.0.1 Force to voltage conversion

The conversion from mechanical acoustic wave to a voltage signal is done by an accelerometer. Accelerometers produce an electrical signal when they experience a mechanical force. An ideal accelerometer would convert force linearly into a voltage signal with some conveniently high gain. Unfortunately, real accelerometers are somewhat

more limited. At best we expect a device to perform the conversion linearly over some limited frequency range. See Section 3.6 for more details.

3.5.0.2 Filtering and amplification

The analogue circuit stage applies both filtering and amplification. The amplification applied to the signal is linear over the frequency range of interest. The signal content is unchanged by this process other than by scaling. The filtering applied at the analogue stage attenuates frequency components of the signal that are higher than 50 kHz. Provided that the received signal has a lower bandwidth than 50 kHz, the signal is not altered at all. See Section 4.2.2 for a more detailed description of the analogue circuitry.

3.5.0.3 Analog to digital conversion

The ADC stage of the conversion process digitises the signal, allowing it to be stored and processed. Sampling theory states that to avoid aliasing the sampling frequency must be greater than the Nyquist frequency. Aliasing is a form of signal distortion where spectral copies of a signal bleed into the original signal band [Ambardar 1999]. The Nyquist frequency is double the bandwidth of the signal being sampled. The filtering in the analog stage limits the bandwidth of the signal to 50 kHz so that with a sampling rate of 150 kHz the signal is digitised without being aliased.

As with any digitising process, a certain amount of quantisation noise is added. However, since the ADC being used has 16 bit quantisation, the signal to noise ratio is still high even when the signal amplitude is low.

So, provided that the signal bandwidth is lower than 50 kHz, the analogue and ADC stages of the conversion process are linear and do not alter the content of the signal. The possible effects of the accelerometer are more complex and are addressed in Section 3.6.

3.6 ACCELEROMETER

The accelerometer employed in the Treetap 5.0 system is the SD-02 from Fakopp Enterprise Bt., Hungary. The same accelerometer has been used successfully in previous generations of Treetap and the Fakopp 2D stress wave timers [Fife *et al.* 2004] [Chauhan 2004].

This accelerometer uses a piezoelectric ceramic element. Piezoelectricity is the ability of certain materials to generate charge under mechanical compression [Kinsler *et al.* 1982]. There are several ways of using this effect that are employed in a variety of accelerometers. Figure 3.9 shows the internal structure of the SD-02. A mass is bonded to the piezoelectric element. The ceramic and mass combination is held against a flat

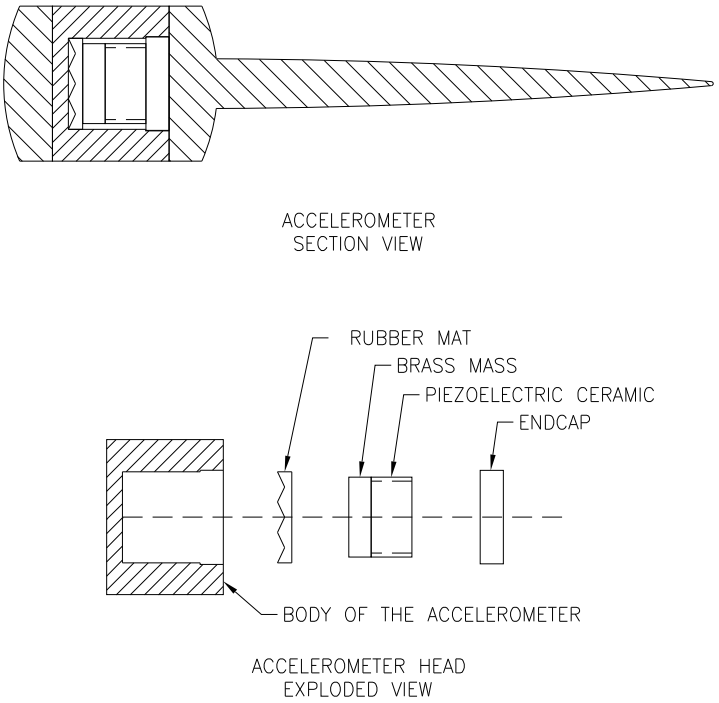

Figure 3.9 Diagram showing the internal components of the SD-02 accelerometer.

surface inside the accelerometer's head by static pressure from the rubber mat. When an acoustic wave displaces the accelerometer a force is imposed upon the ceramic, causing it to produce an electrical charge which is routed into the analogue circuits discussed above.

The accelerometer is not a linear converter of force into voltage. Its response can be characterised by a transfer function. Through good design, the transfer function is linear over a finite range. This range is the usable bandwidth of the device.

3.6.1 Accelerometer model

In order to find the limitations of the accelerometer an electro-mechanical model is created. The mechanical elements of the model are converted into electrical equivalents so that the model can be analysed. The component values for the electrical model are then measured or estimated so that a transfer function for the accelerometer can be determined.

Figure 3.10 shows a model of the mechanical components of the accelerometer. Relative to the other components, the rubber mat has approximately zero stiffness. This simplification makes the brass mass free except for its connection to the piezoelectric ceramic. Likewise, the body (m_2) and the brass mass (m_1) are assumed to have stiffness much grater than that of the ceramic and so are modelled as ideal masses. The mass of the piezoelectric ceramic is distributed by placing half at either end of the

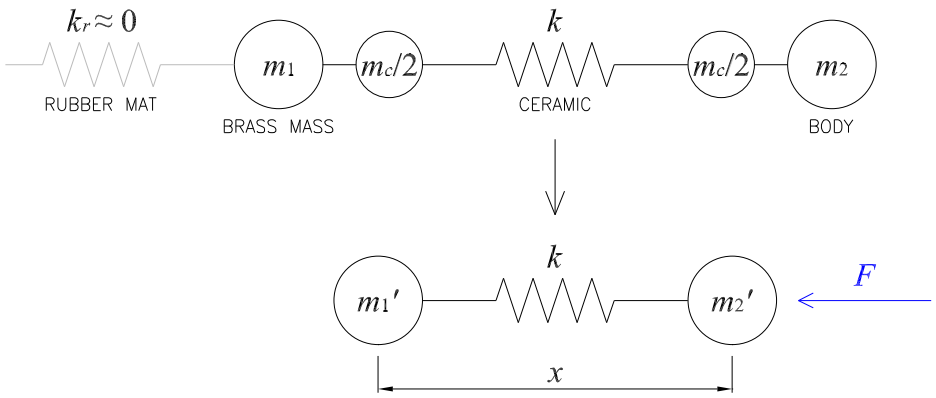

Figure 3.10 Mechanical model of the SD-02 accelerometer.

spring.

$$m_1' = m_1 + \frac{m_c}{2} \tag{3.11}$$

$$m_2' = m_2 + \frac{m_c}{2} \tag{3.12}$$

Thus m_1' and m_2' are the effective masses of the brass and body, inclusive of the mass of the ceramic (m_c).

3.6.1.1 Reciprocal transducer

Before developing a complete model for the accelerometer consider the piezoelectric element alone. The ceramic's conversion of force to voltage can be modelled by a generic two port model, like the one shown in Figure 3.11. The clamped force ($u(s) = 0$) and the short-circuit current ($V(s) = 0$) for this model require two conversion factors:

$$F(s)|_{u(s)=0} = \alpha_{ME} V(s) \tag{3.13}$$

$$I(s)|_{V(s)=0} = \alpha_{EM} u(s) \tag{3.14}$$

where $u(s)$ is the point velocity (see Section 2.2.3), $F(s)$ is force, $V(s)$ is voltage and $I(s)$ is current. For a reciprocal transducer $\alpha_{ME} = \alpha_{EM} = \alpha$. The force to voltage conversion uses the same factor as current to velocity conversion. Piezoelectric transducers are reciprocal [Kinsler *et al.* 1982]. As a result, they can be modelled by an ideal electrical transformer. Transducers that do not exhibit this property are anti-reciprocal.

3.6.1.2 Three port model

An accelerometer is a three port device. Two mechanical ports correspond to the external forces acting on the two masses, while a single electrical port corresponds to

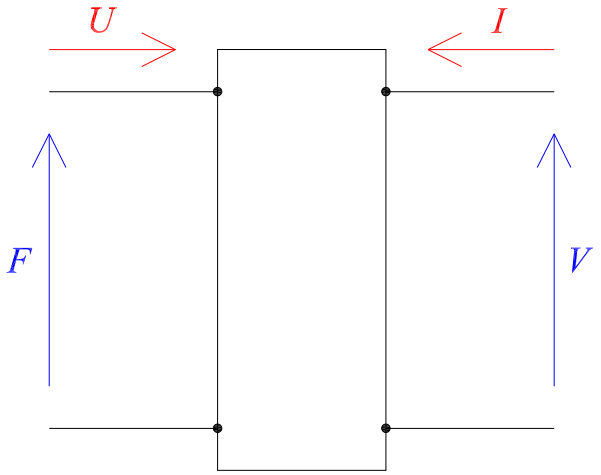

Figure 3.11 Generic, two port transducer model.

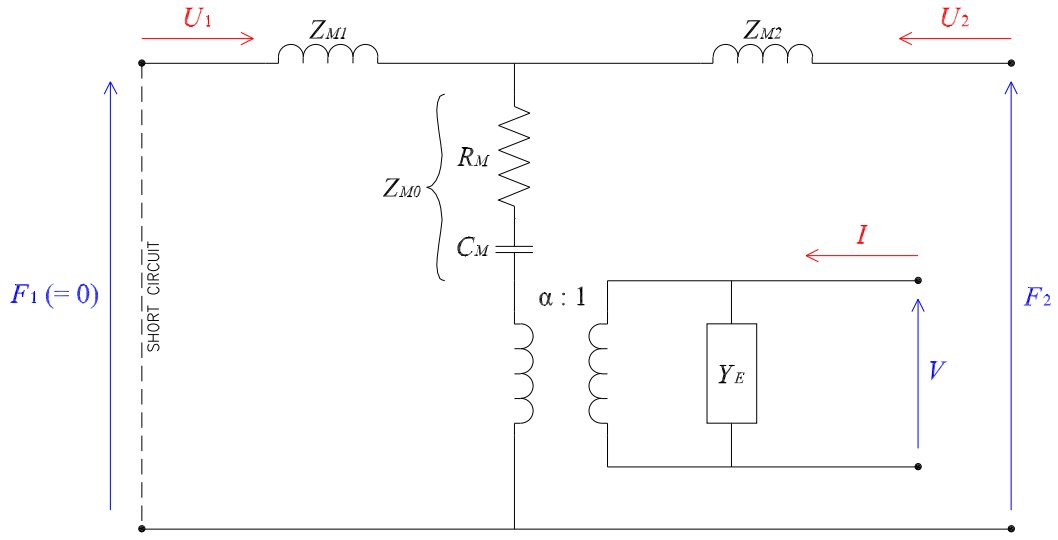

Figure 3.12 Electro-mechanical model of the SD-02 accelerometer with mechanical inputs.

the electrical output. The three ports Figure 3.12 are labeled F_1, F_2 and V. The input to the first mechanical port is the external force on the brass mass. However, since the rubber mat has approximately no stiffness the brass mass is free of external force, making it equivalent to an electrical short-circuit. The input to the second mechanical port is the external force on the accelerometer body. This is in effect the mechanical input to the device, since it is the body of the accelerometer that is affected by the passage of the acoustic wave. The electrical port corresponds to the output terminal voltage.

The model is shown in the electrical domain. The mechancial components have been converted into electrical equivalents using the force → voltage transformation (see Section 2.5.1 for details). Table 2.2 shows that in this transformation (Mechanical

Equivalent 2), masses convert into inductances and springs convert into capacitances. The mechanical impedances (in the electrical domain) are

$$Z_{M0}(s) = R_M + \frac{1}{sC_M} \tag{3.15}$$
$$Z_{M1}(s) = sm_1' \tag{3.16}$$
$$Z_{M2}(s) = sm_2' \tag{3.17}$$

To analyse how the input force and output voltage are related, start with circuit equations for the three port model in Laplace notation:

$$I(s) = \alpha U_1(s) + \alpha U_2(s) - Y_E(s)V(s) \tag{3.18}$$
$$F_1(s) = \alpha V(s) + Z_{M1}(s)U_1(s) + Z_{M0}(s)(U_1(s) + U_2(s)) \tag{3.19}$$
$$F_2(s) = \alpha V(s) + Z_{M2}(s)U_2(s) + Z_{M0}(s)(U_1(s) + U_2(s)) \tag{3.20}$$

where U_1 and U_2 are velocities, F_1 and F_2 are forces, I is the output current, and V is the output voltage. Assuming that the electrical load has a high impedance, the current flowing from the accelerometer model will be approximately 0. Equation 3.18 can be rearranged to

$$U_1(s) + U_2(s) = V(s)\frac{Y_E(s)}{\alpha} \tag{3.21}$$

The first mechanical port is equivalent to an electrical short circuit. Substituting $F_1 = 0$ and Equation 3.21 into Equation 3.19 and solving for $U_2(s)$ gives

$$Z_{M1}(s)U_2(s) = V(s)\left(\alpha + \frac{Y_E(s)}{\alpha}\left(Z_{M0}(s) + Z_{M1}(s)\right)\right) \tag{3.22}$$

Substituting Equations 3.21 and 3.22 into Equation 3.20 yields an expression for the input force $F_2(s)$ in terms of the output voltage $V(s)$.

$$F(s) = V(s)\left(\alpha\left(1 + \frac{Z_{M2}(s)}{Z_{M1}(s)}\right) + \frac{Y_E(s)}{\alpha}\left(Z_{M2} + Z_{M0}(s)\left(1 + \frac{Z_{M2}(s)}{Z_{M1}(s)}\right)\right)\right) \tag{3.23}$$

Substituting for the mechanical impedances and rearranging:

$$F(s) = V(s)\left(1 + \frac{m_2'}{m_1'}\right)\left(\alpha + \frac{Y_E(s)}{\alpha}\left(\frac{sm_1'm_2'}{m_1' + m_2'} + R_M + \frac{1}{sC_M}\right)\right) \tag{3.24}$$

Finally, since m_1' is much smaller than m_2' the equation can be simplified to

$$F(s) = V(s)\left(\frac{m_2'}{m_1'}\right)\left(\alpha + \frac{Y_E(s)}{\alpha}\left(sm_1' + R_M + \frac{1}{sC_M}\right)\right) \tag{3.25}$$

It is the transfer function for the accelerometer that relates the input force to the output voltage that will define the accelerometer's bandwidth. Define $H(s)$ as the ratio

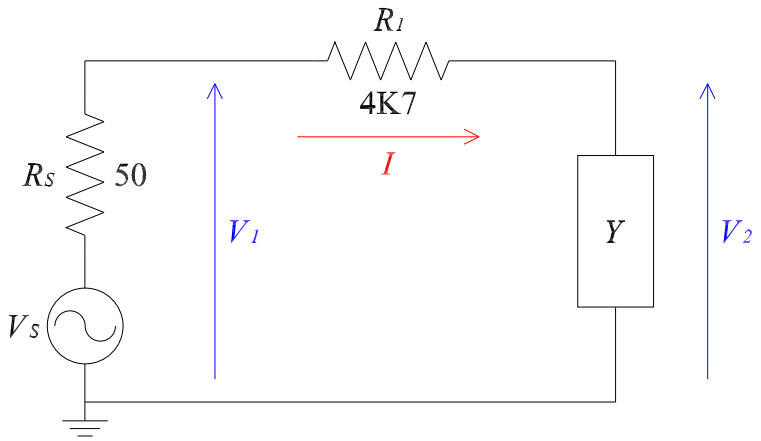

Figure 3.13 Circuit used to find the admittance of the accelerometer.

of the two:

$$H(s) = \frac{V(s)}{F(s)} \tag{3.26}$$

$$H(s) = \frac{1}{\left(\frac{m_2'}{m_1'}\right)\left(\alpha + \frac{Y_E(s)}{\alpha}\left(sm_1' + R_M + \frac{1}{sC_M}\right)\right)} \tag{3.27}$$

The transfer function could be considered to relate input *acceleration* to output voltage if scaled by the mass.

$$F(s) = m_2' A(s) \tag{3.28}$$

$$H(s) = \frac{V(s)}{m_2' A(s)} \tag{3.29}$$

3.6.1.3 Admittance characterisation

In order to find component values for the accelerometer model the admittance of the accelerometer is measured. The accelerometer is placed in series with a known resistor and a signal generator as illustrated in Figure 3.13. Sinusoidal voltages[2] of varying frequency are applied to the circuit at $V_S(f)$. The resulting voltages at $V_1(f)$ and $V_2(f)$ are measured. From this data the current $I(f)$ and the admittance $Y(f)$ are inferred.

The measurements are automated using an Agilent 54622D scope and an Agilent 33220A signal generator. Both are connected to a PC and controlled from a Matlab application. The scope measures the peak-peak voltages $V_{1(pp)}(f)$ and $V_{2(pp)}(f)$ for each frequency along with the phase angle between the two sinusoids, $\phi(f)$. The voltages

[2] 20 Volts peak-peak.

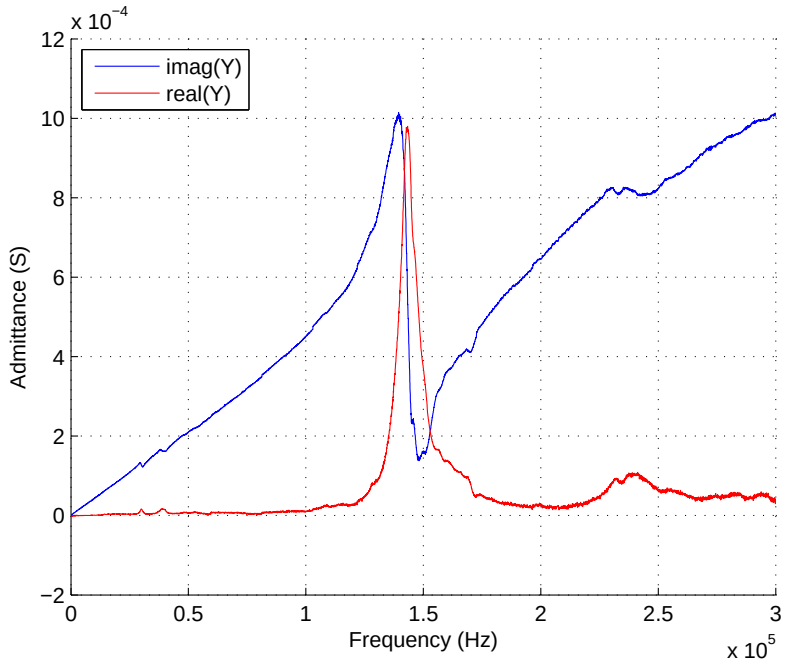

Figure 3.14 Plot of the measured admittance for an SD-02 accelerometer.

$V_1(f)$ and $V_2(f)$ are then

$$V_1(f) = V_{1(pp)}(f) \tag{3.30}$$
$$V_2(f) = V_{2(pp)}(f)\ \exp j\phi(f) \tag{3.31}$$

The admittance can now be defined in terms of the measured voltages:

$$Y(f) = \frac{I(f)}{V_2(f)} \tag{3.32}$$
$$= \frac{V_1(f) - V_2(f)}{R_1\, V_2(f)} \tag{3.33}$$

The measured $Y(f)$ is shown in Figure 3.14.

The impedance of the accelerometer model in Figure 3.12, as seen from the electrical port, can be expressed as

$$Y(s) = Y_E(s) + \frac{\alpha^2}{sm_1' + R_M + \frac{1}{sC_M}} \tag{3.34}$$

Where the electrical impedance $Y_E(s)$ (note that $s = j2\pi f$ in the steady state) for the ceramic is expected to be dominantly capacitive. The complex admittance data in Figure 3.14 shows a linear slope that suggests a capacitive component with a small superimposed resonance at $f_0 = 30\,\text{kHz}$. This small resonance can be attributed to the mechanical resonance defined by m_1', R_M, and C_M in Equation 3.34.

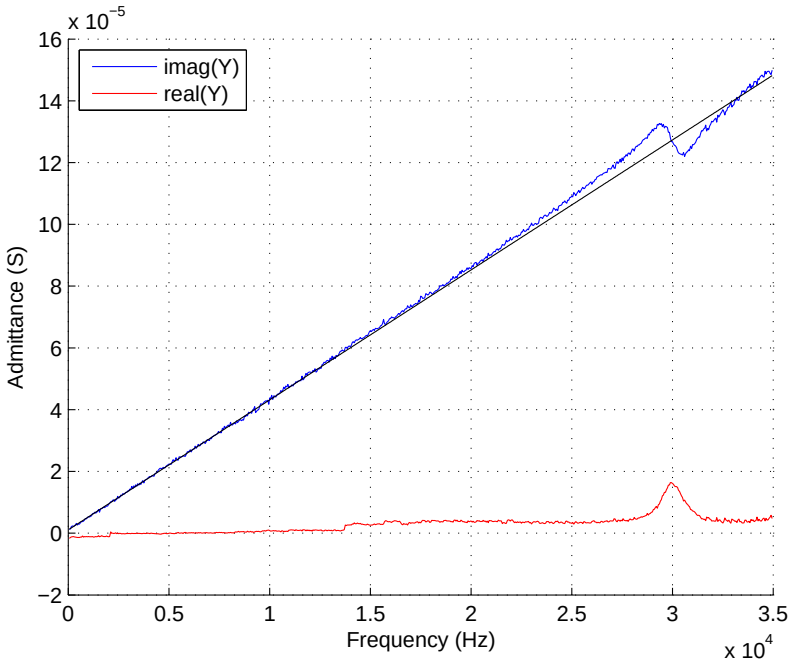

Figure 3.15 Plot of the measured admittance for an SD-02 accelerometer showing only frequencies up to 35 kHz. The straight line estimates the value of C_0.

A much stronger resonance appears around 130 kHz, and more small resonances occur at even higher frequencies. These higher order resonances are possibly caused by electrical resonance of the ceramic, but they are not of great significance as will become apparent in Section 3.6.2.

3.6.1.4 Estimating component values

If the electrical admittance is modelled as being purely capacitive it can be defined as

$$Y_E(s) = sC_0 \tag{3.35}$$

$$\frac{|Y_E|}{f} = 2\pi C_0 \tag{3.36}$$

C_0 can be estimated directly from the data in Figure 3.15. From the mean slope of the complex data, the estimated value of C_0 is 670 pF.

The mass m_1' is estimated from a calculated volume of the mass and the known density of brass. Taking the dimensions of the brass mass gives a volume of approximately 215 mm^3. Brass has a range of densities but values of around 8565 kg/m^3 are common. This calculates to a value of $m_1 = 1.8$ g for the brass mass. The total mass of the ceramic and brass mass are 4.8 g, as measured using electronic scales. The mass of

the ceramic is therefore 3.0 g. Substituting the known values into Equation 3.11 gives:

$$\begin{aligned} m_1' &= 1.8 + \frac{3.0}{2} \\ &= 3.3\,\text{g} \end{aligned} \tag{3.37}$$

The mass of the body of the accelerometer (m_2') is approximately 50 g. Accurate measurement of this is not required, since it acts only as a scaling factor, provided it is much greater than m_1'.

C_M can be determined from the resonant frequency of the admittance. The resonant frequency occurs when

$$sm_1' = \frac{1}{sC_M} \tag{3.38}$$

thus

$$f_0 = \frac{1}{2\pi\sqrt{m_1' C_M}} \tag{3.39}$$

$$C_M = \frac{1}{m_1'(2\pi f_0)^2} \tag{3.40}$$

Since f_0 and m_1' are known ($f_0 = 30\,\text{kHz}$), a value can be calculated for C_M. $C_M = 8.5 \times 10^{-9}\,\text{s}^2/\text{kg}$.

At resonance, the peak in the real data is proportional to the mechanical resistance R_M. The real data peak has a value of 16.4 μS.

$$\begin{aligned} \text{Re}[Y(s)|_{s=j2\pi f_0}] &= \frac{\alpha^2}{R_M} \\ R_M &= \frac{\alpha^2}{16.4 \times 10^{-6}} \end{aligned} \tag{3.41}$$

To acquire a value for α, the frequency spacing between the positive and negative peaks around resonance in the model is optimised to match the spacing in the measured admittance data. In the measured data, the spacing between the peaks is 1.2 kHz. α is approximately 0.02.

Figure 3.16 shows the modelled admittance in comparison to the measured data. The close fit of the model to the measured data confirms that the component values are accurate.

3.6.2 Transfer function

The transfer function (Equation 3.27) can now be populated with the component values derived above. The transfer function $H(f)$ is plotted in Figure 3.17[3]. This gives a complete picture of the accelerometer's response as a function of frequency.

[3] $H(s)$ and $H(f)$ are directly related by $s = j2\pi f$.

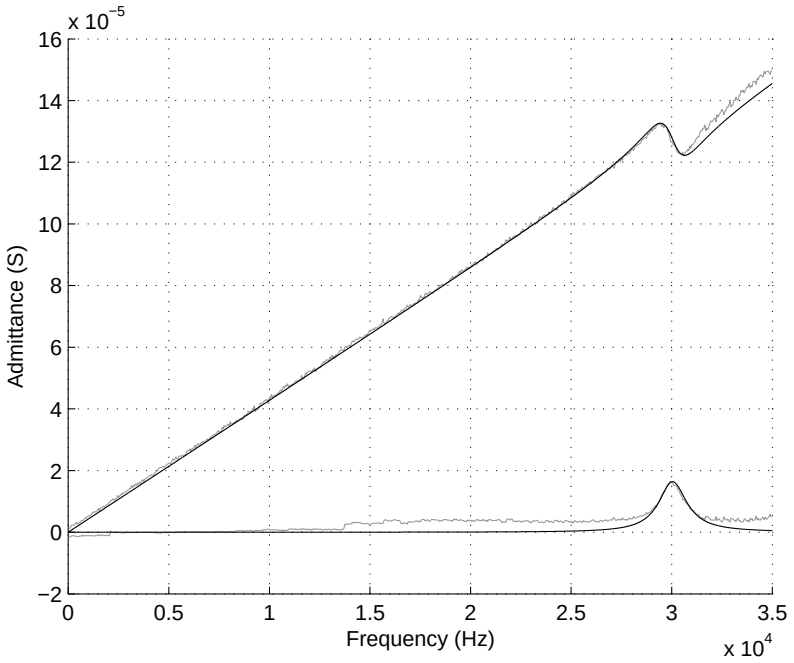

Figure 3.16 Plot of the modelled admittance for an SD-02 accelerometer. The measured admittance is also shown in grey.

The increased response magnitude at resonance shows a dramatic change of nearly 60 dB from the low frequency response. The frequency response drops away steeply for higher frequencies due to the capacitive nature of Y_E. This is why the higher order resonances of the attenuation profile can be ignored as stated in Section 3.6.1.3.

The bandwidth of the device is the range of frequencies where the response is a linear conversion from force to voltage. The bandwidth of the SD-02 is approximately 10 kHz as can be seen from Figure 3.17.

3.6.3 Redesigning the accelerometer

Given the relatively high price of the SD-02 for its acceptable, but not outstanding, performance it is possible that future prototypes may use an alternative accelerometer. Some brief design notes are presented here with particular reference to the Treetap 5.0 system and the components of the accelerometer affecting the transfer function derived in Equation 3.27.

- The mass of the accelerometer's body is a scaling factor of the transfer function. This makes some fundamental sense, since a heavier body will require more force to move it with the same acceleration. It is possible that the plastic housing and hardware connected to the device are adding to this mass, although they are not aligned with the spike.

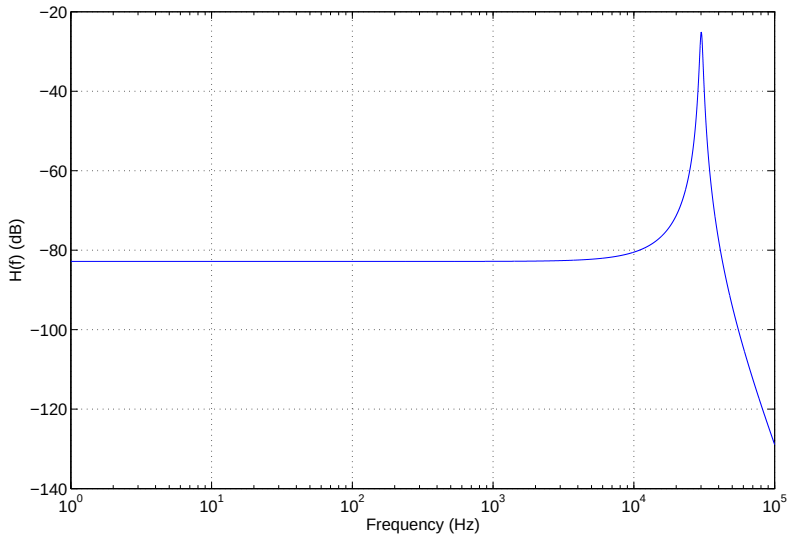

Figure 3.17 Transfer function for the modelled SD-02 accelerometer.

- The brass mass controls scaling as well as the resonant frequency of the device. A larger mass will increase the voltage signal amplitude. However, the resonant frequency will become lower, reducing the bandwidth.

- The piezoelectric ceramic is responsible for the values of α and C_M. Very little is known about the ceramic element used in the SD-02, but it may be possible to find a better replacement for this part.

3.7 DISPERSION

Section 2.5 demonstrated how, in a dispersive system, phase velocity and attenuation are functions of frequency. This section looks at how dispersion affects an acoustic pulse travelling in wood and how this in turn affects both the stress wave timing and resonance techniques for velocity measurement discussed in Section 3.3.2. Since both methods make phase velocity measurements along the longitudinal axis of the timber, the transmission line model developed in Section 2.5.2 can be used to model the dispersive nature of stress waves applicable to both these methods.

Propagation along the longitudinal axis in wood can be modelled by a transmission line with loss mechanisms. The number of mechanisms, their stiffness and relaxation times depend upon the composition and structure of the wood. Equation 2.57 defined the transfer function of a complex mechanical transmission line using the propagation constant[4]. Important equations are restated here for reference. Equations 2.75 and 2.58

[4]Recall that, in a dispersive media, the propagation constant is *not* constant, but a function of frequency.

Density	$\rho = 800\,\mathrm{kg/m^3}$
Young's Modulus	$E_0 = 10\,\mathrm{GPa}$
Relaxation 1 Young's Modulus	$E_1 = 6\,\mathrm{GPa}$
Relaxation time 1	$\tau_1 = 10\,\mu\mathrm{s}$
Relaxation 2 Young's Modulus	$E_2 = 4\,\mathrm{GPa}$
Relaxation time 2	$\tau_2 = 20\,\mu\mathrm{s}$
Relaxation 3 Young's Modulus	$E_3 = 1\,\mathrm{GPa}$
Relaxation time 3	$\tau_3 = 30\,\mu\mathrm{s}$

Table 3.2 Component values for a dispersive wood model with three relaxation processes. The values are not based on any real data and are contrived for illustration purposes only.

define the propagation constant:

$$\begin{aligned}\gamma(s) &= s\sqrt{\frac{\rho}{E_0 + \sum_{n=1}^{N} E_n \frac{s}{s+\frac{1}{\tau_n}}}} \\ &= \alpha(s) + j\beta(s)\end{aligned}$$

where ρ is density, E_n are the elastic moduli, and τ_n are the relaxation times. Equation 2.57,

$$H(s, \Delta x) = \exp(-\gamma(s)\,\Delta x)$$

defines the transfer function of the transmission line. Equation 2.59,

$$a(s) = 0.868\,\alpha(s)$$

defines its attenuation. Equation 2.60,

$$v(s) = \frac{\omega}{\beta(s)}$$

defines the phase velocity of the transmission line.

The high structural complexity of wood presents a potentially huge number of loss mechanisms. In order to demonstrate the effect of dispersion a model using just three is created arbitrarily. Component values of the wood model are given in Table 3.2. This develops a transfer function with attenuation and phase speed profiles as shown in Figures 3.18 and 3.19.

An idealised pulse is plotted in Figure 3.20. Let this pulse be defined as $p(t)$. The pulse, having traveled some distance Δx through the wood has a frequency domain representation:

$$P(s, \Delta x) = H(s, \Delta x)\ \mathcal{L}[p(t)] \tag{3.42}$$

and time domain representation:

$$p(t, \Delta x) = \mathcal{L}^{-1}[P(s, \Delta x)] \tag{3.43}$$

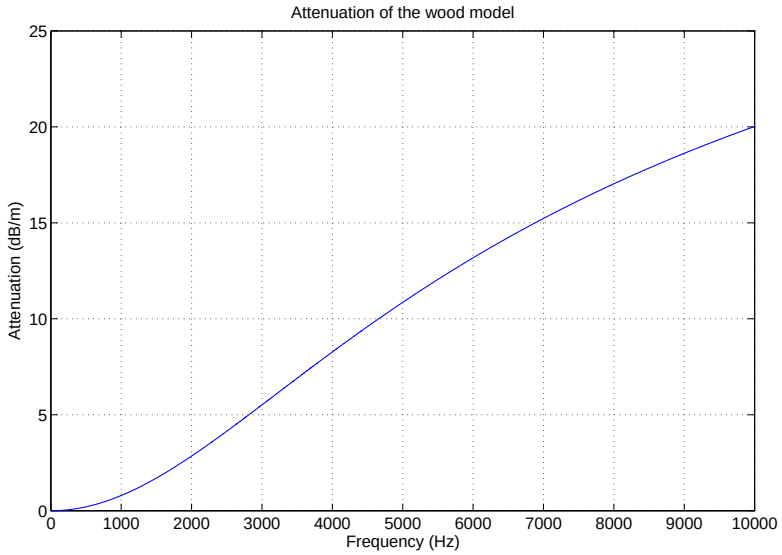

Figure 3.18 Attenuation profile of the wood model.

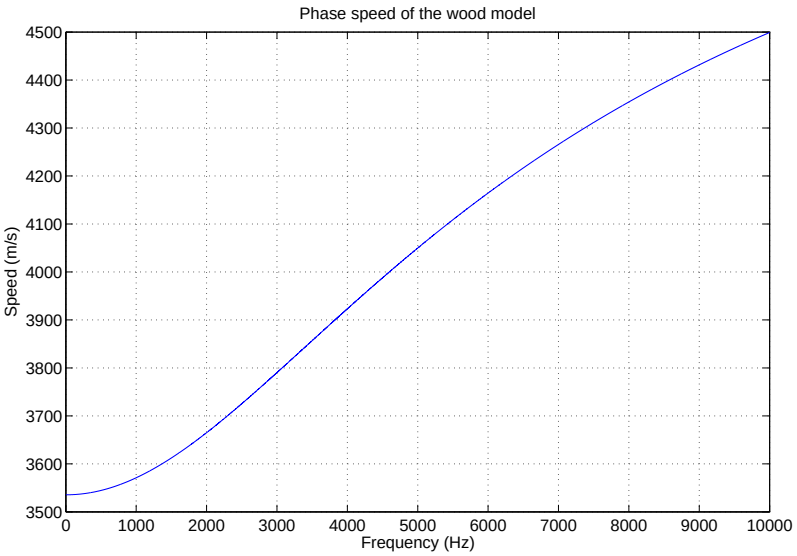

Figure 3.19 Phase speed profile of the wood model.

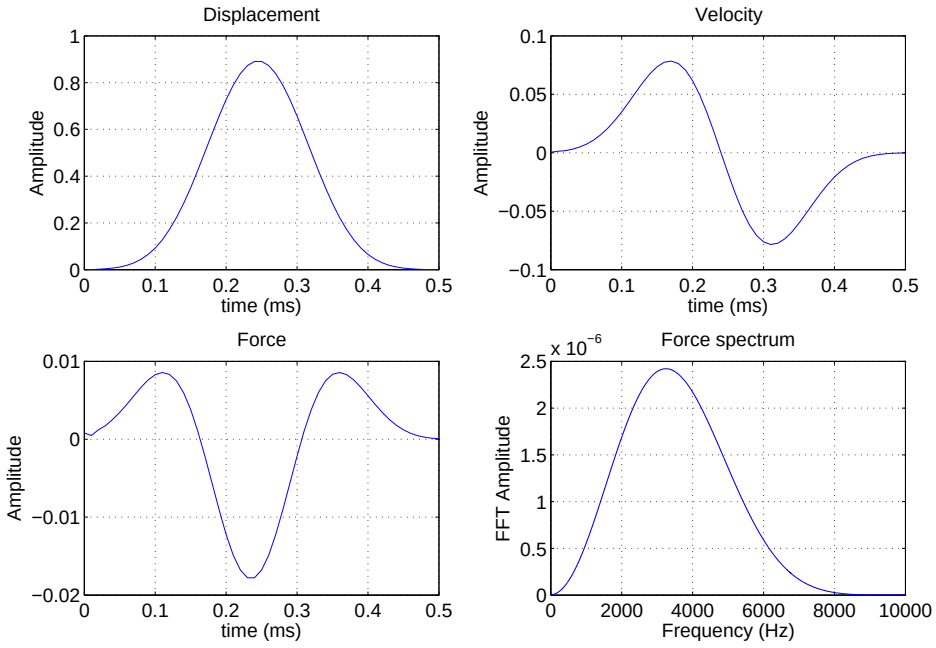

Figure 3.20 Broadband acoustic pulse: displacement, velocity, force and spectra plots.

where $\mathcal{L}$ and $\mathcal{L}^{-1}$ are the Laplace transform and inverse Laplace transform operators.

Figures 3.21 and 3.22 show the time and frequency domain representations of the pulse as it travels along the wood at 1 m intervals. It is quite apparent that both the size and shape of the pulse changes as it travels further through the wood.

The increased phase speed at higher frequencies causes a steeper leading edge of the wave and a smoother tail. Also, since components of the wave are travelling at different speeds there is an overall spreading effect. After 4 m the pulse approximately doubles in length.

The attenuation of different frequencies is most apparent in the frequency domain. While all frequencies are being attenuated, it is clear that the higher frequencies are attenuated more quickly. This is evidenced in particular by the shift of the frequency peak from over 3 kHz in the original pulse to just over 2 kHz after 4 m. The effect is also visible in the time domain, if somewhat more subtly. Note the increased smoothing of the pulse with distance. The removal of the higher frequencies reduces the sharpness of the peaks.

3.7.1 Stress wave timer measurements

The dispersive effects of the wood cause changes to both the shape and amplitude of the pulse as it travels. To consider how this effects stress wave measurements, basic threshold detection and peak detection algorithms are applied to the pulse from the previous example. The pulse is observed at 1 m from the source point and again at 4 m

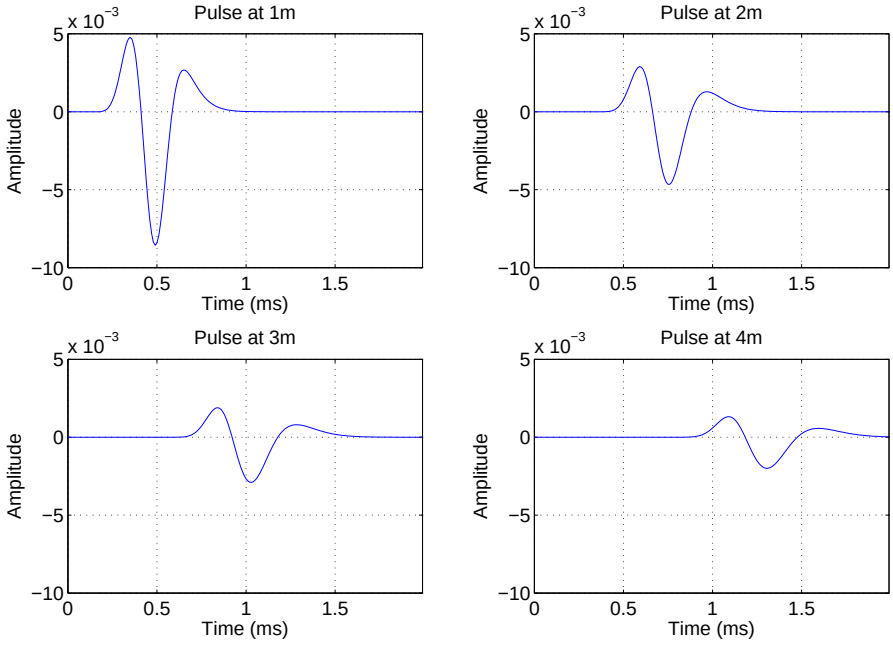

Figure 3.21 Four snapshots of the force pulse travelling through the wood model. Top left: $p(t, 1)$. Top right: $p(t, 2)$. Bottom left: $p(t, 3)$. Bottom right: $p(t, 4)$.

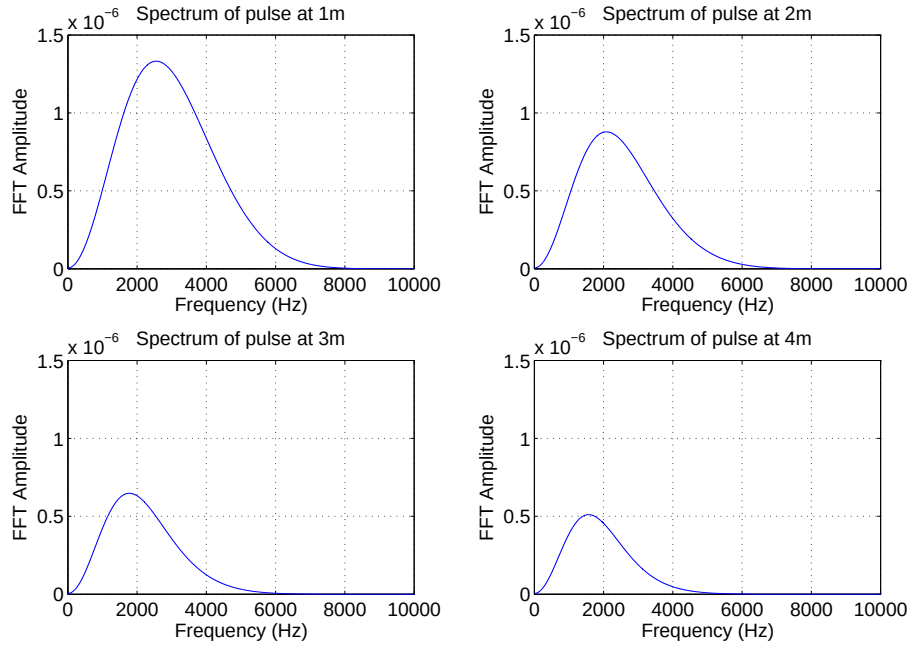

Figure 3.22 Four snapshots of the force pulse spectra. Top left: $P(f, 1)$. Top right: $P(f, 2)$. Bottom left: $P(f, 3)$. Bottom right: $P(f, 4)$.

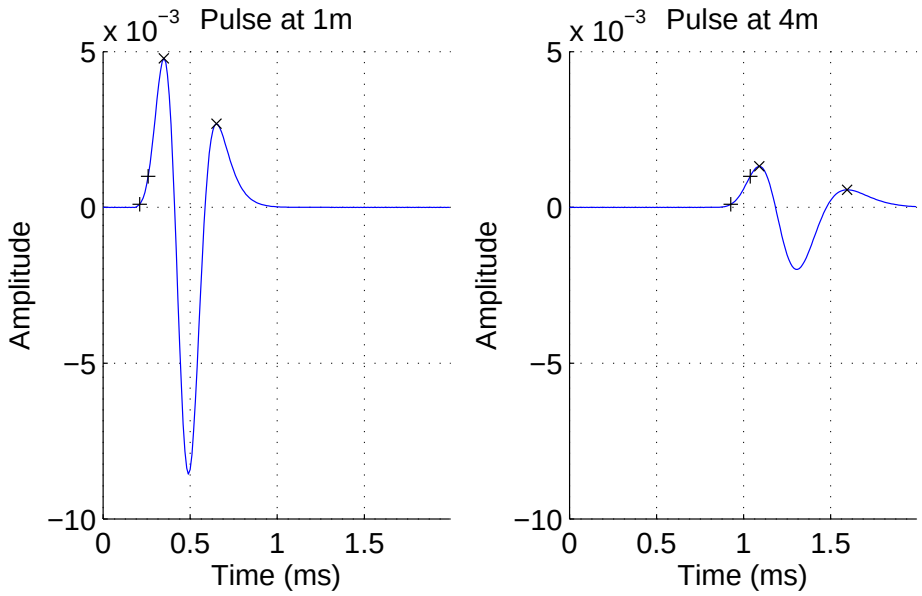

Figure 3.23 A pulse travelling through wood is measured at 1 m and 4 m from the pulse's starting point. "+" marks indicate where a threshold detection algorithm detects amplitudes of 10^{-4} and 10^{-3}. "×" marks indicate where the first and second positive peaks are detected.

Time between leading peaks $= 741\,\mu$s	Velocity $= 4048$ m/s
Time between secondary peaks $= 943\,\mu$s	Velocity $= 3180$ m/s
Time between lower thresholds $= 716\,\mu$s	Velocity $= 4191$ m/s
Time between higher thresholds $= 779\,\mu$s	Velocity $= 3849$ m/s

Table 3.3 Velocity measurements from the data in Figure 3.23.

as shown in Figure 3.23. The two positive peaks (marked ×) are found using quadratic interpolation. The threshold crossings (marked +) are interpolated linearly. Time of flight measurements are listed in Table 3.3 along with the resulting pulse velocity.

It is apparent from Table 3.3 that the detection scheme used dramatically affects the measured time of flight result and hence the velocity. The problem is that neither of these schemes, nor any other detection methods investigated, can be accurately related to the phase velocity in a meaningful, since the phase velocity is frequency dependent.

3.7.2 Resonance measurements

The resonance testing methods are affected by dispersion in a different manner to the stress wave timers. Figures 3.24 and 3.25 show an acoustic pulse resonating in a 10 m log, as seen from one end of the log where a resonance tester would be placed. The pulse clearly degrades with travel as seen in the time domain. However, in the frequency domain it is apparent that the signal degradation still leaves strong, measurable peaks at the low frequencies. The fundamental frequency peak occurs at 177 kHz, corresponding to a velocity of 3540 m/s. This is an accurate estimate of the low frequency phase speed of the wood model (see Figure 3.19).

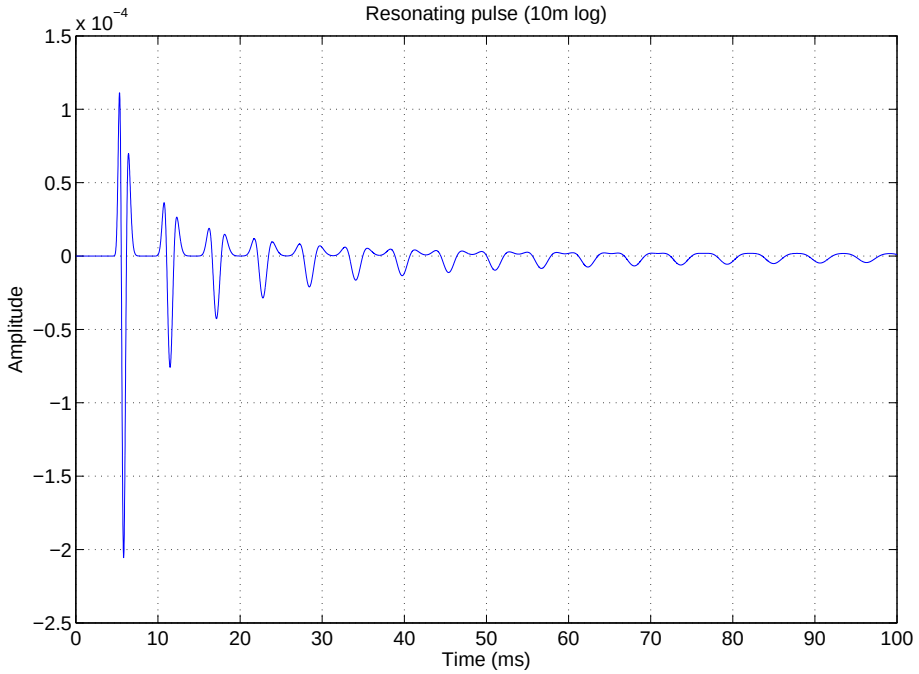

Figure 3.24 Simulated pulse observed by a resonance tester at one end of a 10 m log.

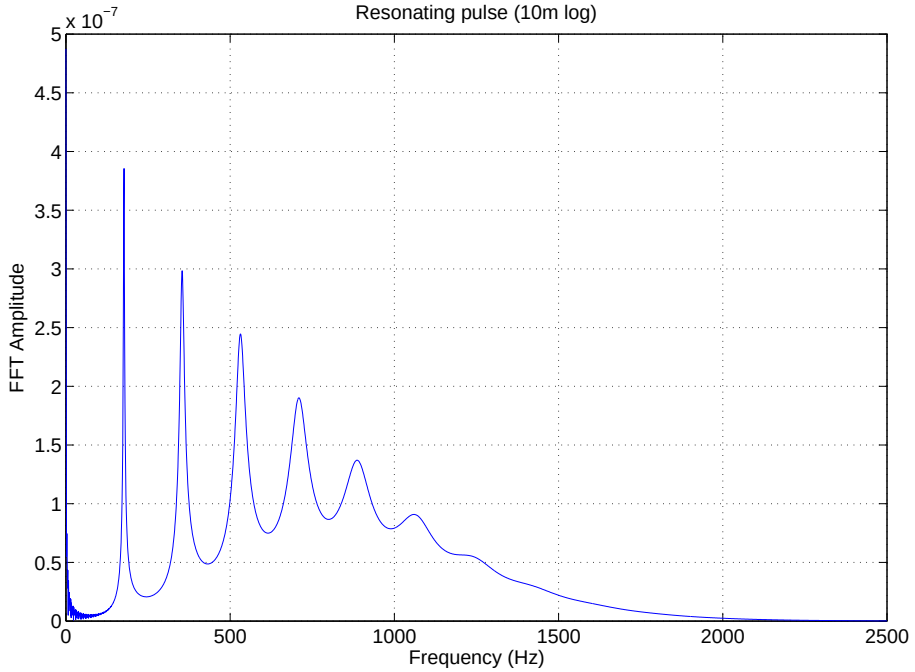

Figure 3.25 Fourier transform of the pulse observed in Figure 3.24.

3.7.3 Measurement comparison

Both measurement systems produce different results. The results from the resonance method produce an accurate estimate of the low frequency phase speed for the wood. The stress wave timer, by contrast, has difficulty determining any particular phase speed of the wood using the simple detection algorithms demonstrated above.

However, these simulations paint a bleaker picture for the stress wave timer than results from practical trials. The wood model and the ideal pulse used here are contrived for demonstration purposes only. Results from field trials show a high correlation of stiffness estimations from stress wave timers using threshold detection with stiffness estimations from a resonance tester and subsequent industrial testing [Fife *et al.* 2004].

3.8 PROBLEMS

The results developed so far all assume that the wood is an unbounded, homogeneous medium through which a purely longitudinal pulse is transmitted along a direct path. Unfortunately, in real wood there are a whole host of problems in transmitting and receiving an acoustic pulse without interference.

3.8.1 Shear waves

The transmission line type models established to model dispersion only deal with the longitudinal (compression) wave travelling along the longitudinal axis. In Section 2.4.2 the existance of three wave components with different polarisations was established for propagation of a wave in an unbounded orthotropic medium. So, there are two shear waves (one in each of the remaining planes) that are not considered by the transmission line model that arrive at the receiver sometime after the longitudinal wave. In addition to the three wave components, surface waves may also be detected.

Accelerometers like the SD-02 are only sensitive to aligned forces. Lateral forces produce negligible results. So for stress wave timers, where the accelerometers are mounted at an angle to the surface, it is not just the longitudinal wave that will be received. The resulting signal depends on the angle α (shown in Figure 3.26). The force aligned with the accelerometer that will be converted into the output signal has at least two components:

$$F_{\text{output}} = \sin\alpha\, F_{\text{longitudinal}} + \cos\alpha\, F_{\text{shear(LR)}} \tag{3.44}$$

Note that only the shear wave in the longitudinal-radial plane is included. The shear wave in the longitudinal-tangential plane acts in the direction perpendicular to the accelerometer, and so has no effect.

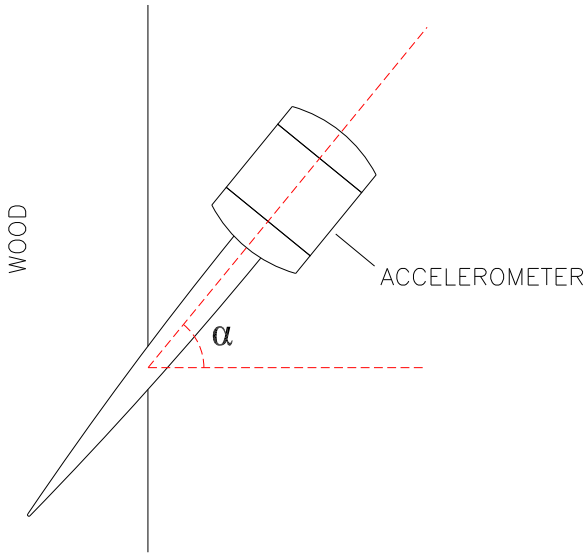

Figure 3.26 An accelerometer inserted at an angle to the longitudinal axis of a tree.

This analysis still assumes that there is only one wave being propagated, although the wave has three components with different polarisations.

The positive side to this is that the longitudinal component travels faster than the shear components [Beall 2002]. This makes is still possible to acquire the longitudinal velocity in stress wave timer measurements.

3.8.2 Multiple paths

Wood is not an unbounded medium. When an acoustic wave comes up against a boundary between two media the amount of reflection and refraction depends on the impedance match between the two media [Kolsky 1963]. Wood and air are poorly matched, so acoustic waves travelling inside wood are almost totally reflected at boundaries. Reflected waves present a huge number of paths from the generator to the receiver for a single pulse. Figure 3.27 shows in just two dimensions how a transmitted pulse may reflect from boundaries and arrive via several paths at the receiver. Only four possible paths shows are shown. There are many more possible paths with a higher number of reflections. If the medium is a cylinder, then the number of possible paths increases. If the medium is an imperfect cylinder with a rough internal surface (like a tree stem) then there are even more, unpredictable, paths. In short, there are an immense number of superimposed pulses arriving and the receiver via different paths.

Each path imposes different affects upon the acoustic wave. Due to the orthortopic nature of wood and dispersive effects, these different paths have phase velocities dependent on frequency and direction. To further complicate matters the wave travelling along each path has three components (one longitudinal and two shear). Trying to deconvolve the multitude of wave arrivals at the receiver is an extremely complex task

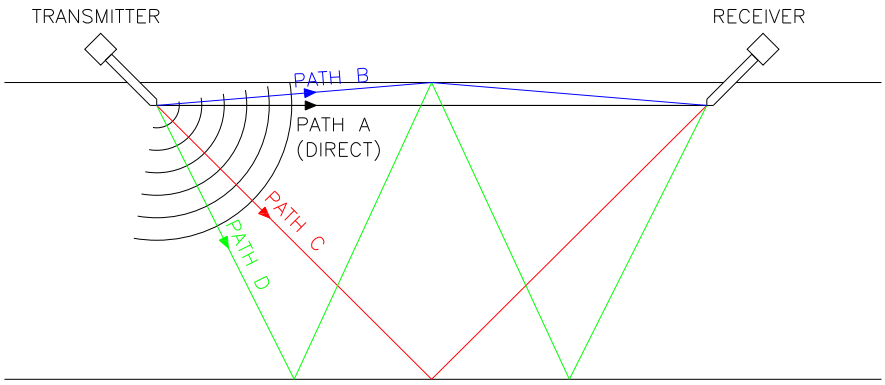

Figure 3.27 Diagram showing four possible paths from transmitter to receiver in a two dimensional plane of a log or tree. The transmitter is modelled as a point source.

and is beyond the scope of this thesis.

3.8.3 Defects

Defects in wood affect acoustic propagation. The size and type of a defect will determine the type and magnitude of the effect on acoustic propagation. Examples of defects include [Corbett 2004]

Shakes: cracks between and parallel to the growth rings.

Checks: cracks in the radial direction.

Knots: portion of a branch or limb of the tree.

Wormholes: caused by insects eating the wood.

Pitch pockets: space packed with an accumulation of resinous material.

Rot, or **decay:** disintegration of wood fiber.

Spiral grain: alignment of the fibers twists around the tree.

Some of these defects can have a major impact on stress wave timing measurements. Many defects may not be apparent without internal inspection.

Recently products have been developed that take use of the altered acoustic properties of defective wood to detect its presence. Acoustic tomography tools using multiple receivers and transmitters can be used to take a snapshot of a horizontal plane through a tree looking for areas of low acoustic velocity that may indicate rot [Divos and Divos 2005].

3.9 SUMMARY

Wood has both good and bad properties as an acoustic medium. On the positive side wood is:

- Orthotropic, provided that there is no significant movement in the tangential direction.
- Homogeneous, with the proviso that there are no significant defects in the sample.

On the negative side:

- Dispersion alters the shape of acoustic waves. Phase velocity is a function of frequency.
- Shear waves are also picked up by receivers attached at an angle.
- Multiple wavefronts arrive at a receiver simultaneously.
- Defects in the wood affect acoustic propagation.

However, both of the acoustic velocity measuring systems (stress wave timers and resonance testers) are able to minimise some or all of these negative features.

3.9.1 Resonance testers

The resonance testing system is able to handle most of the negative influences listed above because of the geometry of the setup and the large sample size. The results from stress wave timers has been shown to correlate with stiffness from bending tests [Harris *et al.* 2002] [Andrews 2002]. Because the sample is long compared to its diameter a wave travelling the entire length and back is so far removed from the source point that it is approximately a plane wave. Since the wave is planar the multiple path problem, which is created by having a point source, is removed.

Most defects will not affect the resonance tester measurements. The small, common defects like knots can be ignored as they affect only a small cross section of the log. Since the wave is approximately planar, most of the wave energy goes around the defect. Only large defects that cover significant sections of the log will affect the measurement results. Such defects usually mean that the log is worthless anyway so the stiffness measurement is no longer critical.

The accelerometer in a resonance tester is aligned with the longitudinal axis. Only a plane wave travelling in this direction needs to be considered, so the shear waves accompanying the p-wave can be ignored as they will not be detected by the accelerometer.

3.9.1.1 Low frequency measurements

As shown previously, dispersion causes the phase velocity to be a function of frequency. For construction timber it is the low frequency stiffness that is of interest. Section 3.7.2 showed that it is this low frequency stiffness that is found by measuring the resonance of a log. This has some additional advantages in terms of designing the device.

In Section 3.6.3 it was found that one of the key elements in the design of an accelerometer was the internal mass. There is a tradeoff between accelerometer bandwidth and sensitivity depending on the size of this mass. The frequency being measured by the resonance tester is low allowing the use of a low bandwidth, highly sensitivity accelerometer. This will provide an increase in SNR[5] and a decrease in the need for amplification. However, the distance traveled by a pulse in the resonance testing system is much greater than for a stress wave timer, so there is also more attenuation of the pulse.

Resonance testers need to sample for long time periods to make accurate estimates of the resonant frequency peak. When making a Fourier transform of the time domain data, the sample spacing between points in the frequency domain (Δf) is determined the length of time the signal is sampled for. Long sampling times give a smaller Δf and therefore better accuracy. Also, longer sampling times allow more reflected pulses to arrive, giving a better average result. A long sampling time means that there is more data to collect, requiring more memory for storage and longer processing times. These requirements can be minimised, however, simply by using a low sampling rate. With appropriate anti-aliasing filtering, this is permissible since the frequency being measured is low.

3.9.2 Stress wave timers

There are many reasons to expect that stress wave timers will produce inferior results to resonance testers. This type of device, however, has one huge advantage over resonance testers; they can be used on standing trees.

Where the resonance testers is basically immune to the problems presented by multiple wavefronts and the reception of shear waves along with p-waves, the stress wave timer must contend with both these problems.

Stress wave timers are able to avoid defects in the wood, provided that they are visible. Since they operate over relatively short distances the device can usually be setup in an area free from knots and other visible defects.

[5]Signal to Noise Ratio

3.9.2.1 Dispersion for the stress wave timer

In Section 3.7.1 it was established that dispersive effects potentially make velocity measurements hard to do accurately. The distorting transfer function of the wood is not known so the same features of a pulse at two points in time cannot be identified. Field trials have, however, indicated that there is still a high correlation between stress wave timer results and structural stiffness [Fife *et al.* 2004] [Wang *et al.* 2001] [Chauhan 2004]. There is still hope that a refinement of the detection algorithms may yield more accurate results from stress wave timing.

There is another possibility for stress wave timers that may prove dispersion to be positive feature. If it is possible to measure a single p-wave pulse at two points in a sample it may be possible to deconvolve the transfer function of the wood. Knowing the transfer function of the wood will reveal not only the low frequency stiffness, but also make it possible to approximate the loss mechanisms. Since the loss mechanisms are based on the composition and structure of the wood they may be used to estimate other properties of the timber. The isolation of a single p-wave pulse, however, would necessitate the removal of geometric effects like multiple wavefronts and shear waves detection. This is no problem.

Chapter 4

SYSTEM HARDWARE

4.1 INTRODUCTION

This chapter describes the design of new hardware for acoustic investigation. The new system serves two major purposes. The first, and perhaps primary, reason for its development is to further the investigation of acoustic inspection techniques for wood. The new system is to provide a platform for extraction of acoustic signals with higher sampling rates, longer sampling times, and a better Signal to Noise Ratio (SNR). With improvements in all three of these areas, the new hardware will provide a superior system of gathering data for processing later.

The second purpose of the new hardware is to prototype new design features for the Treetap stress wave timer series. New features of the hardware include:

- Sampling and storage of digital signals.
- Wireless connections.
- Integration of the accelerometer and hardware for each probe.
- Minimisation of unit size.
- Rechargeable batteries.

4.1.1 System overview

The Treetap 5 hardware is a three part system; one handunit and two (or more) probes connected by a wireless link. The system is shown in Figure 4.1.

The handunit is the central controller of the system. It is the arbitrator of the wireless link, provides the user interface, controls battery charging, uplinks to a PC, and has the potential to perform any necessary calculation of results. Details of the handunit are given in Section 4.3.

The purpose of each probe is to receive and store acoustic data. Probes do not perform any processing of the data. Each probe is completely independent so that the

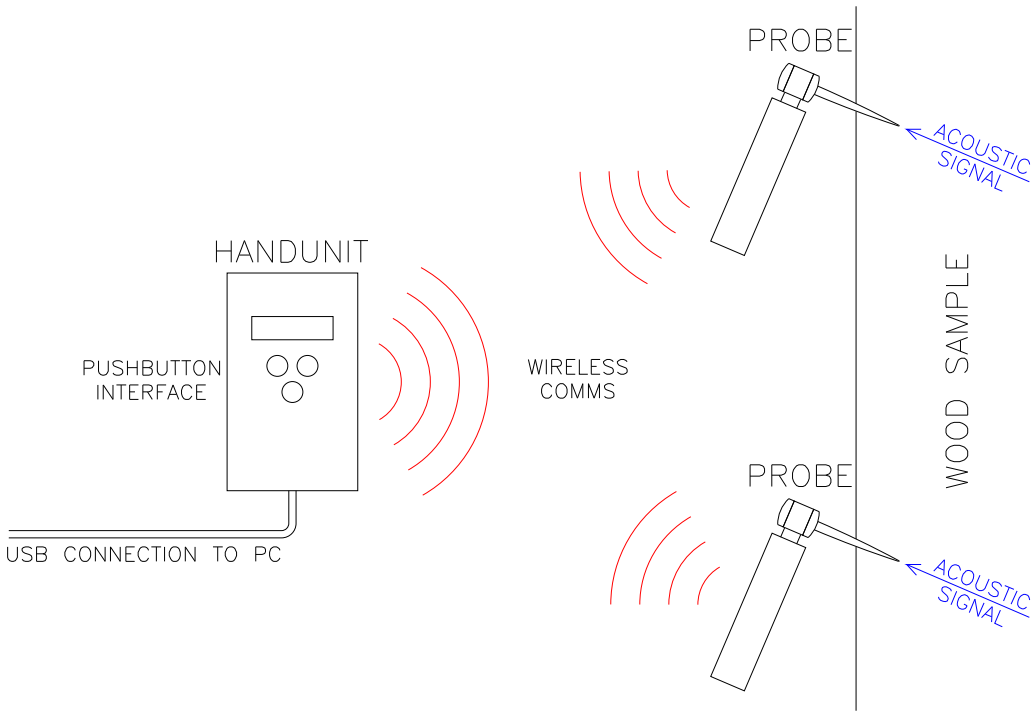

Figure 4.1 Overview of the Treetap 5.0 system.

system can be expanded to use more than two probes. Details of the probe design is given in the following section.

4.2 PROBES

The purpose of each probe is to record acoustic waves travelling through timber. An accelerometer is used to convert mechanical waves into a voltage signal. This signal is amplified by the front end circuitry and converted into a digital format. The memory on board the probes stores the digital data so that the handunit can later download it for processing. The core of each probe is an ATMega8 microcontroller. The main components of the probes are shown in Figure 4.2 and discussed in the following sections.

4.2.1 Accelerometer

The accelerometer used by the probes is an SD-02 vibration sensor with 60 mm spike[1]. This device has already been discussed in Section 3.6. To minimise noise the connecting wire is kept short and hardwired directly to the Printed Circuit Board (PCB).

[1]Supplier: FAKOPP Enterprise Bt., Hungary

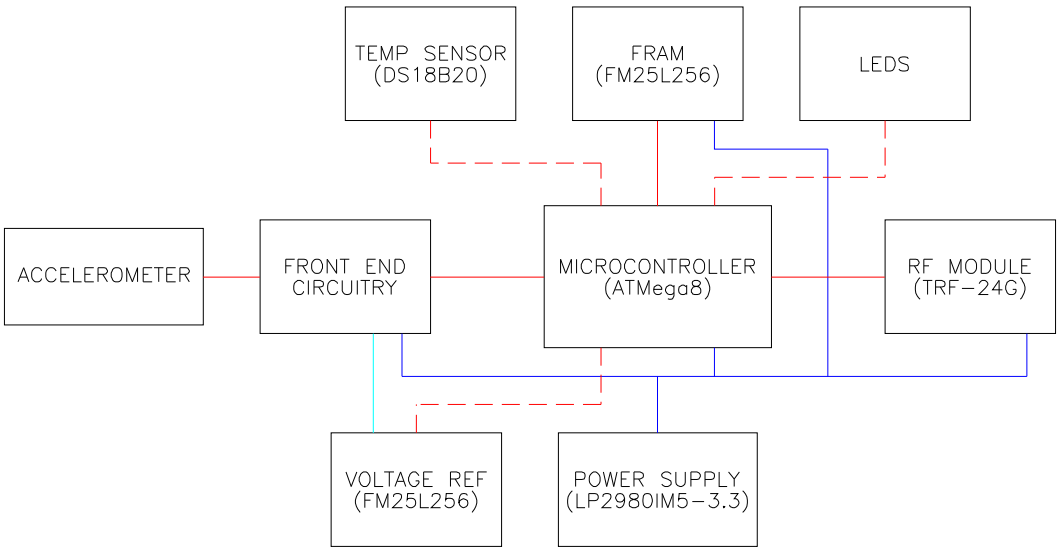

Figure 4.2 Block diagram of each probe's hardware components.

4.2.2 Front end

The analog front end electronics amplify and digitise the signal from the accelerometer. A signal detection circuit and a wakeup circuit are also included.

A blocking capacitor isolates the DC voltage of the accelerometer. This allows the signal to be biased to 1.6 V (midrail). Two safety diodes then limit the signal voltage to −0.7—4.0 V. This ensures that any large voltages, such as those produced while hammering the unit into a wood sample, do not damage the electronics. The mechanical acoustic signal is now represented by a analog voltage signal clipped to −0.7—4.0 V, with zero mechanical force equivalent to 1.6 V.

The weak signal from the accelerometer is then amplified by two Programmable Gain Amplifiers (PGAs) arranged in series. The PGAs are MCP6S21 integrated circuits. These low noise amplifiers offer a programmable gain ranging from 1—32 in discrete steps that can be set individually via the microcontroller's SPI bus. Thus the amplifiers in series have an amplification range of 1—1024. After each PGA is a simple RC, low-pass filter. These filters stop high frequency signals becoming aliased by the ADC. Aliasing, as discussed in Section 3.5, is a form of distortion that arises from the analog to digital conversion.

The simplicity of these filters means that they create little electrical noise, but the downside is that a they have a low roll-off rate (poor attenuation near the cutoff frequency). Nyquist's law states that to avoid aliasing in a sampled signal, the analog

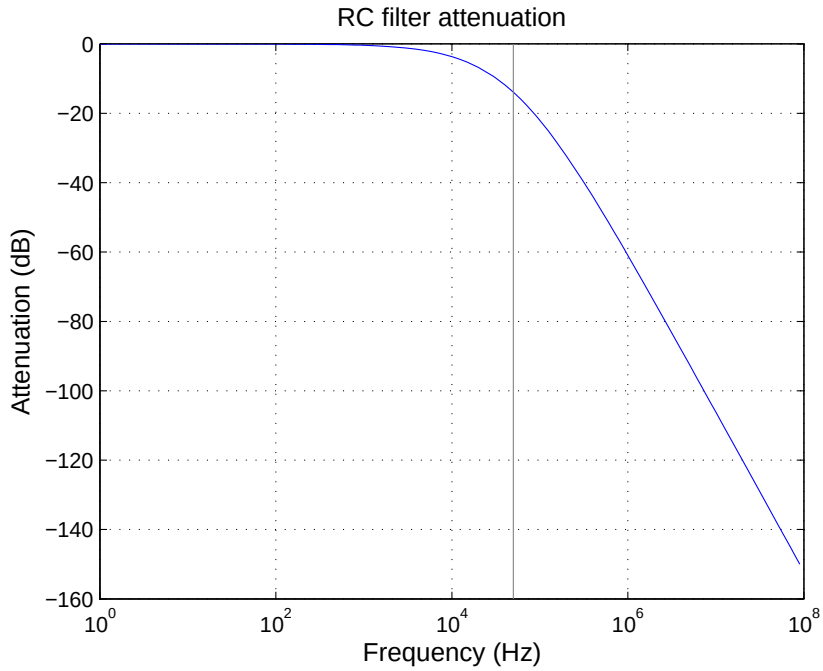

Figure 4.3 Attenuation profile of the anti-aliasing RC filters. The vertical line shows 50 kHz.

signal must be band limited to half the sampling frequency [Ambardar 1999].

$$f_{\text{Nyquist}} = \frac{f_{\text{sampling}}}{2} \tag{4.1}$$

where the Nyquist frequency is the maximum allowable signal bandwidth. The sampling frequency of the probes is 150 kHz, so the Nyquist frequency is 75 kHz. Since the filters have a low roll-off rate (shown in Figure 4.3) the cutoff frequency must be lower than the Nyquist frequency. The cutoff frequency of the filters is set at 50 kHz.

The LTC1864L 16 bit ADC digitises the amplified signal. The signal is sampled at 150 kHz, the maximum rate of the ADC. Data is clocked out over the SPI bus into the FRAM under the control of the microcontroller (see Section 4.2.5).

4.2.2.1 Signal detector

The signal detector observes the signal after the first amplifier. The OPA363 amplifier is configured with a gain of 100. This produces a 3.3 V output when the signal from the first PGA rises slightly above 1.6 V, indicating that an acoustic signal has been detected. The output is routed to an interrupt in the microcontroller, active high. The microcontroller polls this flag during the main measurement loop to determine when an acoustic signal should be recorded.

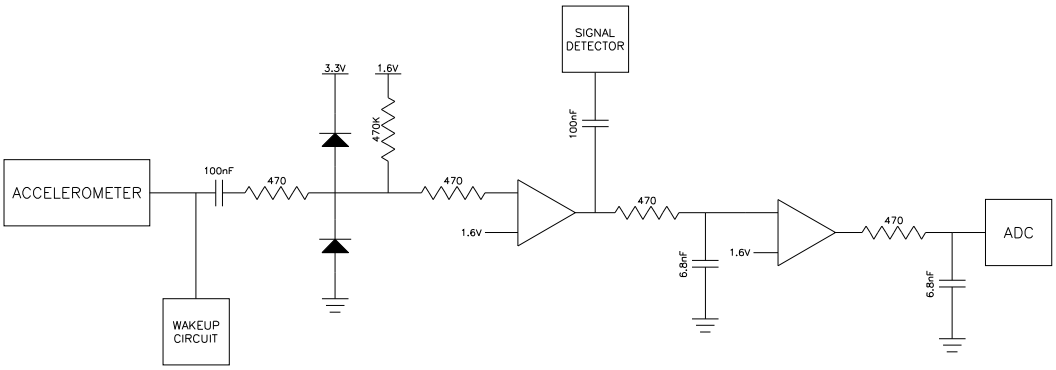

Figure 4.4 Simplified diagram of front end circuitry.

4.2.2.2 Wakeup circuit

The wakeup circuitry is connected directly to the input from the accelerometer. A MOSFET is used to assert an interrupt pin on the microcontroller when a large signal is detected. This interrupt is used to wake the microcontroller from sleep mode. Since there is no amplification a large signal is required to trigger a wakeup, e.g., when the probe is tapped directly with a hammer. The microcontroller wakes up and listens for instructions from the handunit. If a set period of inactive time elapses, the microcontroller powers down the unnecessary circuitry and enters sleep mode again. This conserves the life of the batteries.

4.2.3 Memory

The memory for storing data on the probes is a FM25L256 from Ramtron; a 32 kB Ferro-electric Random Access Memory (FRAM) chip. FRAM offers the best features of both Dynamic Random Access Memory (DRAM) and Electrically Erasable Programmable Read Only Memory (EEPROM). The memory has fast write times, necessary for the high speed data recording in the measurement loop. It is also non-volatile, so the data remains intact when the memory is powered down and can be retrieved later. FRAM manufacturers claim that their products have a virtually unlimited number of read/write cycles; at least 10 billion. Each data value for Treetap 5.0 requires 2 bytes of memory. At a sampling rate of 150 kHz the probe could continuously write to the FRAM at full speed, 24 hours per day for 35 years before exhausting the memory's non-volatility. Thus, in this application the read/write cycles can be considered infinite. Refer to the manufacturer's datasheet for technical details or to [Fox *et al.* 2001] for a description of the underlying technology.

4.2.4 Software - time

Before discussing the critical measurement loop code, it is necessary to look at how the probe software keeps track of time. The timer 1 module of the microcontroller is a 16 bit timer configured in free-running mode and resets at a value of 64000. With a system clock of 16 MHz this occurs every 4 ms. At each reset an interrupt increments a static unsigned, 16 bit variable. Note that if interrupts are disabled the static variable is not updated, so accurate time is not kept.

A timestamp `struct` contains two 16 bit unsigned integers. The first is called `ms_ticks`, the other `us_ticks`. When a `get_time` function is called, the returned timestamp's `ms_ticks` variable is set to the hidden, static variable and the `us_ticks` variable is set to the value of `TCNT1` (timer 1 main register). At a system clock rate of 16 MHz, the time in μs is then

$$t = \texttt{ms\_ticks} \times 4000 + \frac{\texttt{us\_ticks}}{16} \tag{4.2}$$

See the ATMega8 datasheet for additional information on the timer functionality.

4.2.5 Software - measurement loop

The most important part of the probe's software is the measurement loop. This loop captures the acoustic signal and stores the resulting data in the FRAM. When entering measure mode some initialisation tasks are executed before entering the actual loop. The loop controls the ADC's convert signal, reads data from the ADC and writes it to the FRAM, checks for signal detection, and keeps track of how many samples are still to be taken. Upon exiting the loop several more tasks are executed before returning control to the main function. The timing of this loop is critical to making accurate measurements.

4.2.5.1 Initialisation

The function `probe_measurement_loop` is called when the probe receives a command from the handunit to enter measurement mode. Various initial parameters are configured before entering the main looping code, including the disabling of all interrupts. With interrupts disabled the `get_time` functions do not return accurate time (see Section 4.2.4). Because of this, and to reduce execution cycles, the microcontroller's timer register `TCNT1` is accessed directly during the loop. Before entering the loop, the SPI speed is increased to from 1 MHz to 16 MHz to reduce the delays in reading and writing.

4.2.5.2 Loop functions

The loop begins by finishing the current ADC conversion and performing the read/write cycle. Data is read serially from the ADC and written to the FRAM at the same time. The ADC has 16 bit quantization, so two read/write cycles are needed to move each data element from the ADC to the FRAM. The process of writing data is one byte behind the read process. As the first byte of data is being clocked into the SPI Data Register (SPDR) from the ADC, via the `MISO` pin, the last byte of the previous measurement is being sent to the FRAM from the same register, via the `MOSI` pin. As the second byte of new data is read, the first byte is written. The data is not passed through the CPU, just shifted serially through the SPDR. This process is depicted in Figure 4.5. For convenience, one empty byte of data is written to the FRAM in the loop initialisation to compensate this delayed write process. The resulting data elements in memory then have even numbered addresses (0x00, 0x02, 0x04...).

Once the read/write cycle is complete the next conversion is begun by asserting the ADC's convert pin. While this conversion is taking place the microcontroller uses the remaining loop time for other required tasks:

- Triggered status is checked.
- FRAM address counter is incremented.
- Number of samples remaining is decremented, if necessary.

The `triggered_flag` status is checked to see if a signal has already been detected. If not, the signal detection flag (this flag is set by the signal detection circuit, see Section 4.2.2.1) is polled. If this flag is set, the value of `TCNT1` is written into the `us_ticks` variable of a *start timestamp* and the `triggered_flag` is set.

If the `triggered_flag` is already set, then the number of samples remaining to be collected is decremented. The loop will exit when this number drops to zero. The FRAM address must be incremented every loop, regardless of the trigger status.

4.2.5.3 Loop timing

It is essential that each iteration of the loop has exactly the same execution time to avoid sampling jitter. To achieve a sampling frequency of 148.148 kHz the loop must take exactly 108 clock cycles. The ideal sampling frequency is 150 kHz but the period is limited to an integer number of clock cycles at a frequency of 16 MHz. The original design of the loop used a timer to count the number of clock cycles and set a flag when it was time to restart the loop. Unfortunately, the code polling the flag could not be optimised to take less than four cycles of execution time. This presented a period jitter problem. On average the loop time was 108 cycles but any individual iteration might take anywhere between 104–112 cycles.

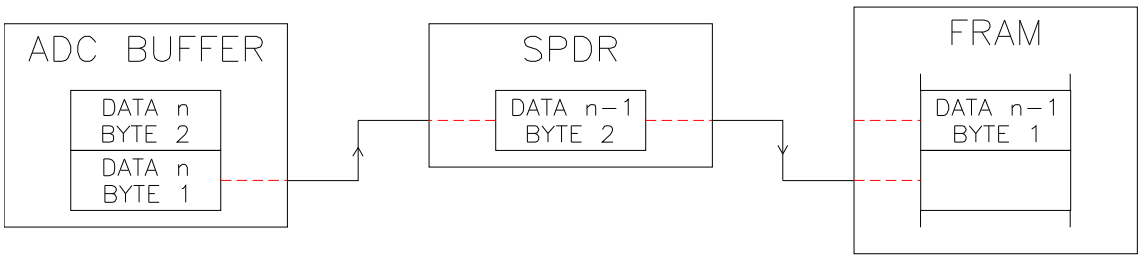

State of memory buffers at the beginning of the read/write cycle. An ADC conversion has just completed.

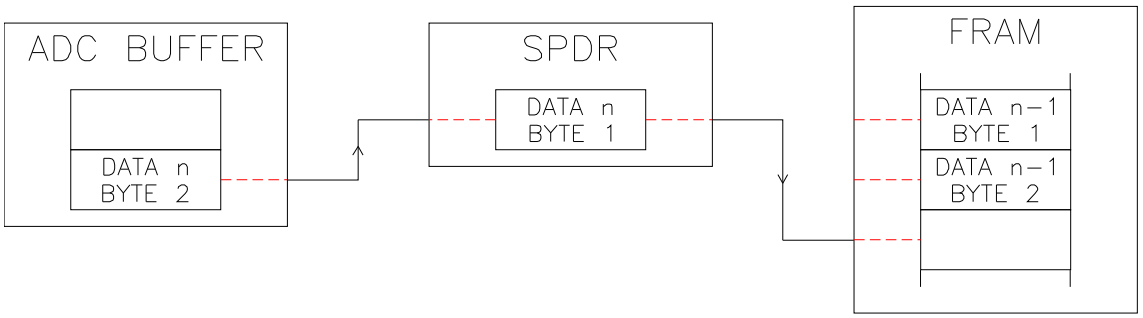

First half of the read/write cycle complete.

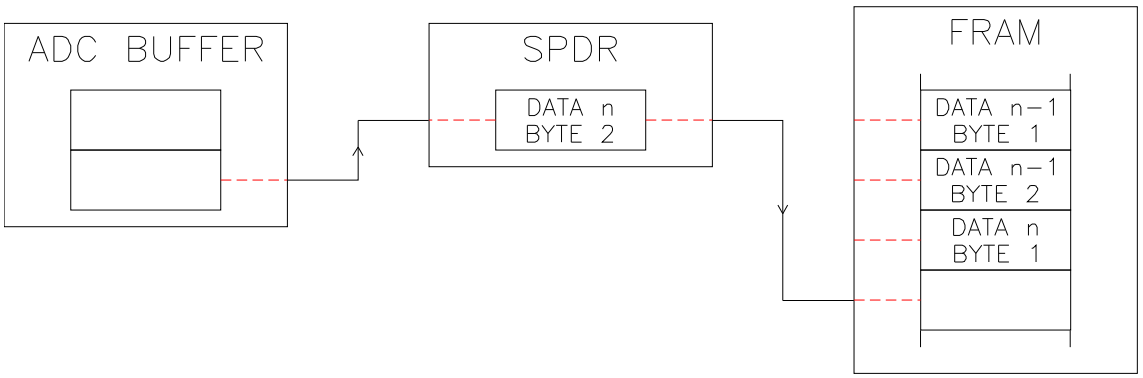

Read/write cycle complete. Note that the second byte of the new data is still in the SPDR.
This will be moved into the FRAM during the first half of the read/write cycle for the next data element.

Figure 4.5 The read/write process of the measurement loop.

To ensure that there was no period jitter, delays were added to each path of the code. `nop`[2] instructions throughout the code made every possible path a multiple of four clock cycles to execute. The correct locations of `nop`s could only be found by studying the assembler code and working out the number of clock cycles through each path of the code. Note that if the C code is modified, the location and number of `nop`s must be recalculated.

Since all the assembler code must be reviewed to balance the path lengths, the timer is now removed altogether and the total number of cycles for each iteration padded out to 108. Paths no longer need to take a multiple of four cycles but must be equal in length. This reduces the number of cycles needed by the code, allowing more time for extra instructions in future versions. It also allows more fine tuning of the sampling frequency, since there is no longer a restriction on the period being a multiple of four cycles.

4.2.5.4 Finishing

The loop exits when all the required samples have been collected. The last timing dependent function is to record a *finish timestamp*. A delay is added, equal to the number of cycles from the start of the loop to the collection of the start timestamps. The value of `TCNT1` is written into the `us_ticks` variable of a finish timestamp. With interrupts disabled, the start and finish timestamps have the same `ms_ticks` value. The start timestamp's `ms_ticks` variable is recalculated based on the finished timestamp and the number of samples taken. The two timestamps should now be separated by exactly the number of loop cycles multiplied by the number samples taken. This is a check for timing errors in the loop.

The final tasks of the `probe_measurement_loop` function are to:

- Write the last byte of the data to the FRAM, since the SPI write to FRAM is one byte behind the SPI read from ADC.
- Return the SPI clock back to the normal 1 MHz.
- Disable the FRAM.
- Re-enable interrupts.
- Fix the FRAM address so that it points to the start of the data, rather than the end.

Once these are completed the `probe_measurement_loop` function returns.

[2] No-operation. A blank instruction word taking one clock cycle to execute.

4.2.6 Power

The power source for each probe is a 3.6 V Lithium Polymer battery with 200 mAh capacity. This battery has significant capacity while being extremely small and light. Lithium batteries are rechargeable but require a charge monitoring device to do so safely [Moore and Schneider 2001]. They also require a charge protection circuit but this is usually integrated with the battery. Charge monitoring circuits are built into the handunit, rather than the probe, to save space. See Section 4.3.6.1 for details of the charging circuits.

The power source is routed through a linear regulator to ensure an even 3.3 V power supply. Since the supply and regulated voltages are so close the linear regulator works efficiently. An additional 1.6 V regulator supplies a bias point for the accelerometer's signal and a voltage reference for the PGAs (see 4.2.2). This second voltage regulator requires a load for voltage stability and is much less efficient so it is disabled when not in use.

The microcontroller extends the battery life of the probe by enabling peripherals only when they are needed. When active, the microcontroller sets all the peripheral devices to their working state. If the microcontroller receives no instructions for a set period of time (currently configured at 30 s) or receives a sleep command from the handunit, it sets all the peripheral devices into low power use mode. For some devices this involves setting or clearing an enable pin, for others an instruction is sent via the SPI. The microcontroller then puts itself into sleep mode. When the wakeup circuit (see Section 4.2.2.2) triggers the microcontroller's external interrupt it returns to its active state and re-enables peripheral devices. With such low power usage at standby there is no need for an on/off switch for the probes.

4.2.6.1 Future work

Two features of the power supply and management system could be handled better in future prototypes. Power management could be better optimised and the location of the charge circuit should be changed.

Currently there are only two modes of power management; everything on or everything off. An additional mode could be added specifically for measurements to reduce the amount of power wasted on the 1.6 V regulator. Then only the devices required at any one time would be powered up:

Normal mode: RF circuit powered. FRAM powered on demand. No power supplied to front end circuits or voltage reference.

Measure mode: Front end circuits and FRAM powered. No power supplied to RF circuit.

Sleep mode: Everything powered down, including microcontroller.

The other future change would be to the location of the charge controller. Placing the charging circuit on the probe unit, rather than the handunit would make the unit more independent. More probes could be charged simultaneously and charged from a greater variety of sources. Originally, space on the PCB was conserved in order to fit the unit into its housing. The housing is now longer than originally intended (to give the operator more grip) allowing more space for PCB expansion and removes the need to be so conservative.

4.2.7 Temperature sensor

The primary role of the temperature sensor on the probe is not to measure temperature. The DS18B20 is a three pin module that connects to the microcontroller via a one wire bus. These modules each have a unique 64 bit serial number. This number is read by the microcontroller and used as the probe's device ID. In a future software version it may be useful to sample the temperature each time a measurement is made and send this data back to the handunit along with the measurement data.

4.2.8 Housing

The probe hardware is housed in a sealed, hollow tube with the accelerometer at one end. This provides a safe, sealed environment for the hardware and a handle for the operator to hold while inserting and removing the probe from timber. The body of the housing is made from clear polycarbonate tubing. Plastic plugs at both ends are sealed with an o-ring and have a single Allen bolt (outside the o-ring) to hold them in place. The plug at the top is bored out to jamb fit the accelerometer head. It also has a smaller diameter standoff to keep the accelerometer head clear of the main casing, reducing the likelihood of the casing being hit by a hammer when driven into timber. The accelerometer spike points perpendicular to the tube. The bottom plug has a recess for the battery charging socket. The hardware PCB fits snugly inside the tubing, along with the battery without the need for any packing material. The photo in Figure 4.6 shows the probe in its casing. The result is a light, durable, water resistant package.

4.3 HANDUNIT

The handunit acts as an interface between a PC and the probes so that acoustic data can be recorded and processed on a PC. However, there is sufficient program space on the handunit's microcontroller to implement data processing, allowing the device to operate as a stand alone stress wave timer.

Figure 4.6 Image of the probe hardware in its casing.

The handunit has a central microcontroller governing all the device's systems. A set of pushbuttons provide input from an operator, and a Liquid Crystal Display (LCD) provides feedback. The unit has an RF module identical to those on the probes for wireless communication. The microcontroller's memory is supplemented with a FRAM unit, also identical to those on the probes. The uplink to a PC is via USB.

4.3.1 Microcontroller

The core of the handunit is an ATMega32 microcontroller. This is from the same series of microcontrollers as those used on the probes but is a larger model. This larger microcontroller has more program memory to store the greater quantity of software as well as significantly more external I/O pins. Since they are from the same series, the microcontrollers on the probes and handunit are able to share common software, e.g., RF communication software.

4.3.2 Display

The display module is a two-line LCD with a built in hardware driver. The hardware driver simplifies the I/O to the device by allowing the microcontroller to send text in ASCII format. The microcontroller communicates with the LCD on a six wire bus.

The display needs a +5 V supply. To supply this from the available 3.3 V, one of the Pulse Width Modulation (PWM) ports of the microcontroller is set up as a charge pump. This produces a -1.3 V node, to supply the LCD's V_0 pin. This gives a relative

supply voltage of up to 4.6 V; a sufficient supply for bright operation. The contrast can be altered by changing the PWM duty cycle. The LCD operates with 3.3 V logic levels, but I/O values still lie within the required ranges.

4.3.3 Pushbuttons

A set of three pushbuttons provides the main input from the user when disconnected from a PC. They are de-bounced in software. The buttons are also connected to hardware interrupts within the microcontroller to wake the unit from sleep mode.

4.3.4 Memory

The memory in the handunit is the same FRAM module as is used in the probes. See 4.2.3 for details. The memory in the handunit is not used when it is connected directly to a PC, since the data can be offloaded directly.

4.3.4.1 Future work

To make full use of the 32 kB memories on each probe, it may be useful to have more memory modules on the handunit. To operate stand alone the current prototype must partition its memory into data space and result space. Data downloaded from each probe can be written to data space, processed and the result stored in result space. The maximum size of probe data space is then

$$\text{probe data space} = 32768 - \text{bytes per result} \times \text{number of results} \tag{4.3}$$

If an algorithm comparing data from both probes is implemented the space for each probes data is halved again. Ultimately, the type of data processing used dictates the size of the results and how much probe data needed. If the exact nature of this processing is not known at the next prototype stage additional FRAM should be considered.

4.3.5 USB

The USB connection allows the handunit to interface to any PC. A driver needs to be installed for Microsoft Windows to allow this, but one is freely downloadable from the FTDI website[3]. Most Linux platforms include the required driver. The PC creates a virtual serial port, connected to one of its physical USB ports. On the handunit, the FT232 chip connects to a Universal Asynchronous Receiver Transmitter (UART) on the microcontroller. Data is converted from serial to USB format by the FT232, is send from the handunit to the PC via a USB cable, and is converted back to serial format by the virtual serial port driver. The PC can then communicate directly with

[3] http://www.ftdichip.com/FTDrivers.htm

Syntax	Description
e	Enumerate probes.
cp n	Set probe p to RF channel n.
u	Make a full measurement and download to USB.
m	Put probes into measurement mode.
dp	Download data from probe p.
fp	Download probe p's entire FRAM
ip	Download measurement info from probe p.
b n	Set the number of samples n recorded before a signal is detected on all probes.
a n	Set the number of samples n recorded after a signal is detected on all probes.
ga n	Change the gain of PGA a to n on all probes.
t n	Timestamps test with n timestamps.

Table 4.1 Handunit commands.

the microcontroller through this serial port. A simple terminal program can be used for this.

Single characters, sometimes with an additional numeric value, followed by a carriage return are recognized by the handunit as commands. The available commands and syntax is shown in Table 4.1.

When downloading data from a probe, the data is printed directly to the serial port. An example of the format of the data is shown in Figure 4.7.

4.3.6 Power

The power source for the handunit is the USB 5 V supply with an Lithium Polymer battery when operating away from a PC. The handunit has no power switch; it relies on power management to conserve battery power. The power source is routed to a linear regulator supplying 3.3 V directly to the microcontroller. The regulator also supplies the other devices on the PCB through a high-side MOSFET. The gate of the MOSFET is controlled by the microcontroller. When a 90 s period of inactivity occurs, the microcontroller removes power to all other devices on the PCB by toggling the gate value. It then puts itself into sleep mode. It will return from sleep when one of the buttons is pushed causing an interrupt (see Section 4.3.3).

4.3.6.1 Charge circuits

The battery charging circuits are powered from the USB 5 V supply and are governed by three MAX1501 charge controller chips. These charge controllers ensure safe charging of the lithium batteries. They also control a green LED indicating that charging has

```
Info P0                          ← Start of info. Specifies probe number.
id 10 13 d8 ac 00 08 00 54       ← Device ID number (hex).

sb 500                           ← Number of samples before edge.
sa 800                           ← Number of samples after edge.
ad 22636                         ← FRAM address where the data starts.

ts 270610031                     ← Probe time at start (edge).
tf 270632428                     ← Probe time at finish.

270933984, 123280407             ← Timestamp taken after measurement.
4085596195, 245506600            ← Subsequent timestamps for synchronisation.
4101698317, 261611683
...

Data P0 L1300                    ← Start of data. Specifies probe number
31792                            and length of data.
31796
31782                            16 bit data is represented in decimal ASCII.
31792
31784
31780
...
```

Figure 4.7 Example of the output format appearing on a terminal program.

Figure 4.8 Image of the handunit casing.

been completed and the units are ready to go[4]. The handunit automatically charges its own battery whenever it is plugged into the USB port. The probes must be connected to the handunit via a charge cable, while it is connected to the USB port, for their batteries to be recharged.

4.3.7 Housing

The handunit hardware is housed in a plastic, rectangular box with a clear window for the LCD. The three pushbuttons are mounted on the front and have sealing o-rings. Sockets are fitted into the bottom of the unit to connect the charging cables and the USB cable. The overall package is much smaller that the previous Treetap units, making it easier to hold and control in one hand. The unit is not completely sealed due to gaps around the cable sockets but is fairly rugged and splash-proof. An image of the handunit's casing is shown in Figure 4.8.

[4]Note that there is still an unresolved issue with the charge light not functioning properly for both charging chips.

4.4 RF COMMUNICATIONS

The communications link between the devices in the system is wireless. Probes can be placed anywhere on a tree or timber sample without the inconvenience and limitations of cables. It removes the possibility of damaging the cables or connectors. By having the amplification and sampling electronics close to the accelerometer, the SNR of the system improved. Since there is not a limited number of connectors, the system can be expanded to use any number of probes.

The advantages of a wireless system, however, come at a price. There is a significant increase in software complexity of the devices. The data from each unit has to be synchronized, since there is no common clock signal. It takes longer to upload data because of the additional communication overheads. There is an increase in production cost since many of the components (RF modules and probe batteries particularly) are not necessary in a wired system.

4.4.1 RF hardware

A TRF24G module from Spark Fun[5] is mounted on each probe and handunit. The TRF24G uses a Nordic nRF2401 chip; a 2.4 GHz band RF transceiver, designed to communicate with other identical chips. The nRF2401 chip requires soldering of hidden pins, making it impractical for direct use in this application. The TRF24G module can be easily and directly soldered to the main PCB. The TRF24G comes complete with the handful of additional components needed by the nRF2401, a metal shield to reduce interference, and a flat PCB loop antenna. The module has many configuration options including channel selection, address checking, CRC check, and shockburst mode. For details any of these features, see the nRF2401 datasheet (available online at the Spark Fun website[6].

The module has six pins connected to the microcontroller. An SPI interface (`SCK`, `MISO` and `MOSI`) is the main data interface. The chip enable (`RF_CE`) and chip select (`RF_CS`) pins allow the microcontroller to define the mode of operation of the device. The data ready (`RF_DR`) pin alerts the microcontroller that there is data ready to be read from the device. The First-in, First-out (FIFO) buffer in the RF module allows the microcontroller more flexible access to access the RF module.

4.4.1.1 RF band

The 2.4 GHz band is an unlicensed RF band in most countries. While every country has its own legislation governing the use of this band, most follow the regulations set out by the International Telecommunication Union (ITU), the Geneva based, specialised

[5] http://www.sparkfun.com
[6] http://www.sparkfun.com/datasheets/RF/nRF2401rev1_1.pdf

agency of the United Nations. The Radio Regulations[7] is an intergovernmental treaty specifying the allocation of frequency bands and technical parameters for radio use.

4.4.1.2 RF channels

The TRF24G divides the 2.4 GHz band into 125 RF channels. Channels are centered at 1 MHz spacings, starting at 2.4 GHz. The channel number selects which channel the TRF24G is tuned to:

$$f_{RF} = 2400 + \text{channel number} \times 1.0 \quad \text{MHz}$$

All devices are tuned to channel 0 on startup but can be shifted to another channel later by sending an RF instruction from the handunit to the probes demanding the change. The microcontrollers then reconfigure their TRF24G modules appropriately.

4.4.1.3 RF addresses

The TRF24G uses a 5 byte address to identify units operating on the same channel. For the Treetap5 devices, the last four bytes of this address are fixed to the sequence `CA 7E A9 52`. This arbitrary number set was chosen because it has a high number of transitions. The first byte is set to the `RF_ID` of the unit. This number must be known to the other devices on the wireless network that wish to communicate with it. See the section on probe enumeration (4.4.3) for details on how the `RF_ID`s are set.

4.4.2 RF protocol

The protocol used by Treetap5 has been designed specifically for this system. It ensures reliable transmission from the handunit to all probes. A single master node (the handunit) with a number of slave nodes (probes) are connected by wireless links. The master node is the arbiter of the wireless network and is the only node able to initiate communication. The slave nodes are only able to respond to instructions from the master node.

The high levels of the protocol are implemented in the software running on each device's microcontroller. These include packet acknowledgment, command packets, timestamps, and special commands. The TRF24G hardware handles the lowest levels of the protocol.

[7]The Radio Regulations, edition of 2004, contain the complete texts of the Radio Regulations as adopted by the World Radiocommunication Conference (Geneva, 1995) (WRC-95) and subsequently revised and adopted by the World Radiocommunication Conference (Geneva, 1997) (WRC-97), the World Radiocommunication Conference (Istanbul, 2000) (WRC-2000), and the World Radiocommunication Conference (Geneva, 2003) (WRC-03) including all Appendices, Resolutions, Recommendations and ITU-R Recommendations incorporated by reference.

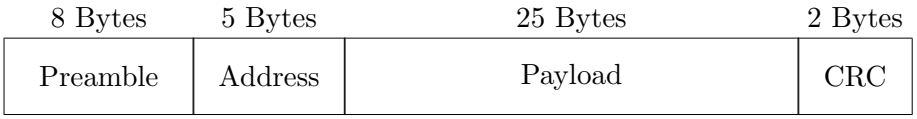

Figure 4.9 Packet layout diagram showing the number of bytes allocated to each section.

Most of the communication across the network takes the form of single packet commands from the handunit to probes. There are two other forms of communication; the special enumeration command and download requests.

4.4.2.1 Hardware protocol

The hardware of the RF module handles the lowest levels of the communications protocol. Once the device is configured, packet framing, address and CRC checking are all done by the TRF24G without intervention from the microcontroller.

Each packet contains a preamble, address, payload, and CRC check. Figure 4.4.2.1 shows the number of bytes in each section. 25 is the maximum number of bytes allowed in a packet payload using the current configuration. To increase this number bytes from the address or CRC would have to be sacrificed. To send a packet, the microcontroller configures the TRF24G for transmit then clocks the receiving address and payload into the buffer via the SPI. The preamble and CRC check are added by the TRF24G hardware.

When a packet is received by a TRF24G it checks the preamble for errors, compares the packet address to its own, and performs a CRC check. If the preamble is correct, the address matches, and the error check is passed, the packet is considered valid. The preamble, address, and CRC are stripped from the valid packet. The payload is inserted into the TRF24G's outgoing buffer and the data ready (`RF_DR`) pin to the microcontroller is asserted. If the packet is not valid for any reason then it is discarded completely. Invalid packets are those with incorrect addresses (intended for another unit) and those with incorrect preambles, CRC checks, or packet length (corrupted packets).

4.4.2.2 Commands and acknowledgments

Reliable command transmission is ensured using a packet acknowledgment system. In each command packet the first byte of the payload is a command header. The value of this byte determines the type of command. Subsequent bytes contain additional command parameters where necessary. When the master node sends a packet to a slave, the slave must return an acknowledgment packet. If no such packet is received by the master unit within a certain time, the request is re-transmitted. Figure 4.10

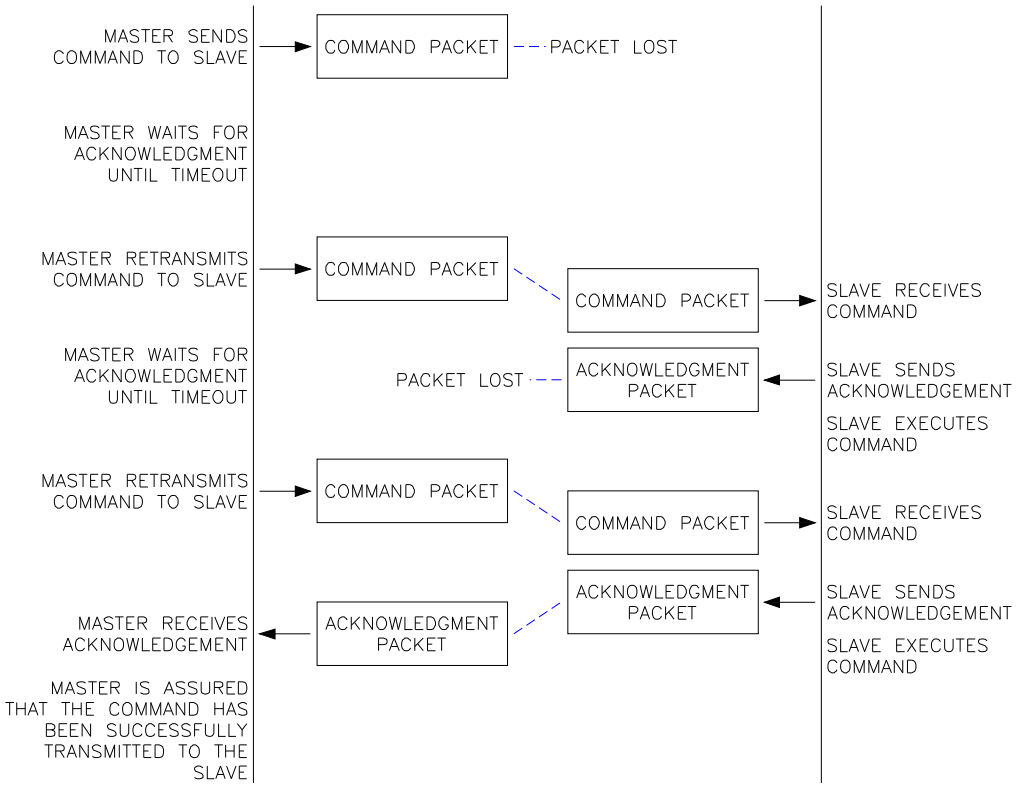

Figure 4.10 Recovery from lost packets. Diagram shows recovery from loss of both commands and acknowledgments.

shows how the protocol copes with packets lost. Note that the slave may re-execute a command if an acknowledgment packet is lost, but commands are designed in such a way that this does not present a problem.

4.4.2.3 Timestamp generation

All acknowledgment packets contains a timestamp. When the `RF_DR` pin is driven high by the RF module an interrupt is generated on the microcontroller. This interrupt captures and stores the current local time as a timestamp. See Section 4.2.4 for details on timestamps. These timestamps are later used to synchronise the data from both probes (see Section 4.5).

4.4.3 Probe enumeration

When the units are first powered up all the slave devices have a common `RF_ID` and hence the same address. The probes must be assigned individual `RF_ID`s so that they may be communicated with independently. The special *enumeration* command handles this.

On startup the handunit and probe RF_IDs are set to default master and slave values. To start an enumeration command the handunit broadcasts a DEVICE_ID_GET command to the default slave address. All probes currently listening on that address hear the command but *do not* immediately send an acknowledgment. This is an exception to the rule of acknowledging all command packets. If every probe was to respond immediately the packets would collide and be rejected by the master.

Each probe uses its device ID (see Section 4.2.7) to determine a random delay before responding. Each time a probe receives a DEVICE_ID_GET it chooses a byte from its device ID, takes the value of the lower four bits and adds two. This is the number of milliseconds the probe will wait before responding to the handunit. Consecutive retries use consecutive bytes from the device ID, thus generating different wait times. Since the probes have unique device IDs, each probe waits a different amount of time on each retry.

Once this delay has elapsed the probe transmits a DEVICE_ID_RESPONSE packet. This packet contains the device ID and the version number of the probe software. Currently the software version number is not checked, but this may be implemented in a future version. The probe's RF module is returned to the receive mode and waits for a response from the handunit.

When the handunit receives a DEVICE_ID_RESPONSE packet it responds by sending an ENUMERATE_DEVICE command. This is still transmitted to the default slave address so all probes may receive this packet. The device ID of the probe intended to receive the ENUMERATE_DEVICE command is inserted into the command packet so that even though all probes receive the command, only the intended device will act upon it. The new RF_ID that the handunit is assigning the probe is also included in the command packet. No acknowledgment is expected for this command so it is repeated a number of times to guard against channel loss. If an acknowledgment was to be sent by the probe and was lost, the handunit would have to repeat the enumeration command. However, since the probe has already received the enumeration command it is no longer listening to the broadcast address and would never hear the request. The handunit then adds the probe with new RF_ID and device ID to its list of enumerated probes.

When the correct probe receives this packet it sets its RF_ID to the new one given by the handunit. Having completed its enumeration it returns to normal operation. Other probes receiving this packet identify that it is intended for another probe and stop trying to send DEVICE_ID_RESPONSE packets. They return to listening on the broadcast address.

This process is continued until the handunit is satisfied that all possible probes have been enumerated. The handunit can now communicate selectively with each probe by using the specific addresses defined by the new RF_IDs.

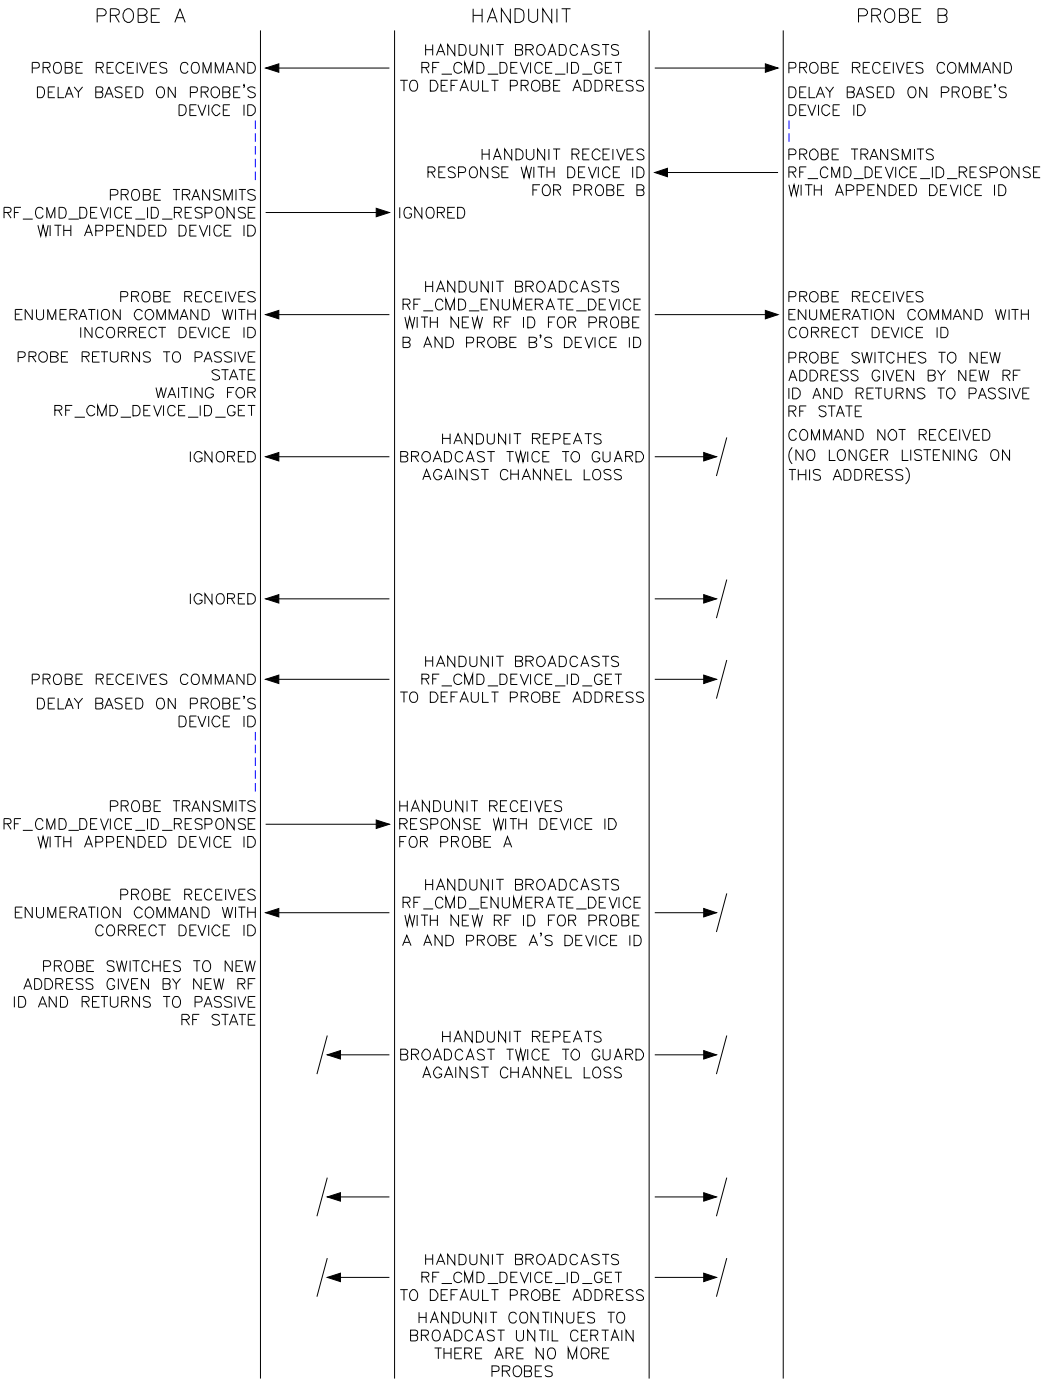

Figure 4.11 Enumeration Process. Diagram shows the communications sequence for two probes being enumerated.

4.4.3.1 Alternative scheme

The enumeration scheme currently used works well. However, there is a possible problem for two or more complete systems working within range of each other. The systems enumerate their probes every time they power up. If two systems both power down, the first unit to power back up may enumerate one or more of the other system's probes. An alternative scheme for a future prototype could bind the probes to their handunit, eliminating this problem. This would also make it more feasible to have each system operate on a separate channel to avoid collisions. This option is available on the current system but must be reset every time the unit powers up. Figure 4.12 shows a possible scheme for enumeration that involves having a separate mode for the handunit where it is able to bind probes permanently.

4.4.3.2 Data download

There are two download requests that the handunit makes from each probe for each measurement. The first acquires the measurement information from the probe. This includes timestamps for the start and finish of the measurement and the memory address where the measurement data starts. Once this information has been checked the handunit requests a download of the probe's data memory. The handunit specifies which section of the probe's memory it wants to download.

The request for information is relatively simple. The handunit sends a `CMD_MEAS_INFO_GET` packet and the probe acknowledges. The probe then sends an additional packet with a `CMD_MEAS_INFO_SEND` header. In the payload of this second packet is

- The start timestamp for the last measurement.
- The finish timestamp for the last measurement.
- The address in memory where the data from the last measurement starts.

This information is received and stored by the handunit.

The data download is more complex because of the greater volume of data. The handunit sends a `CMD_DATA_GET` packet which includes the start address and the length of data to be downloaded. The probe acknowledges this packet then transmits a `RF_CMD_DATA_START` packet that states the available length of data. The available length will limited by the size of the probe's data buffer. This is currently configured to 240 bytes.

The probe fills the buffer with the appropriate block of data from its FRAM and transmits this in a series of data packets. The first byte of each packet contains a packet number. The remaining 24 bytes are data. The handunit receives this stream of packets and monitors for packet loss using a worklist algorithm. If one or more packets are

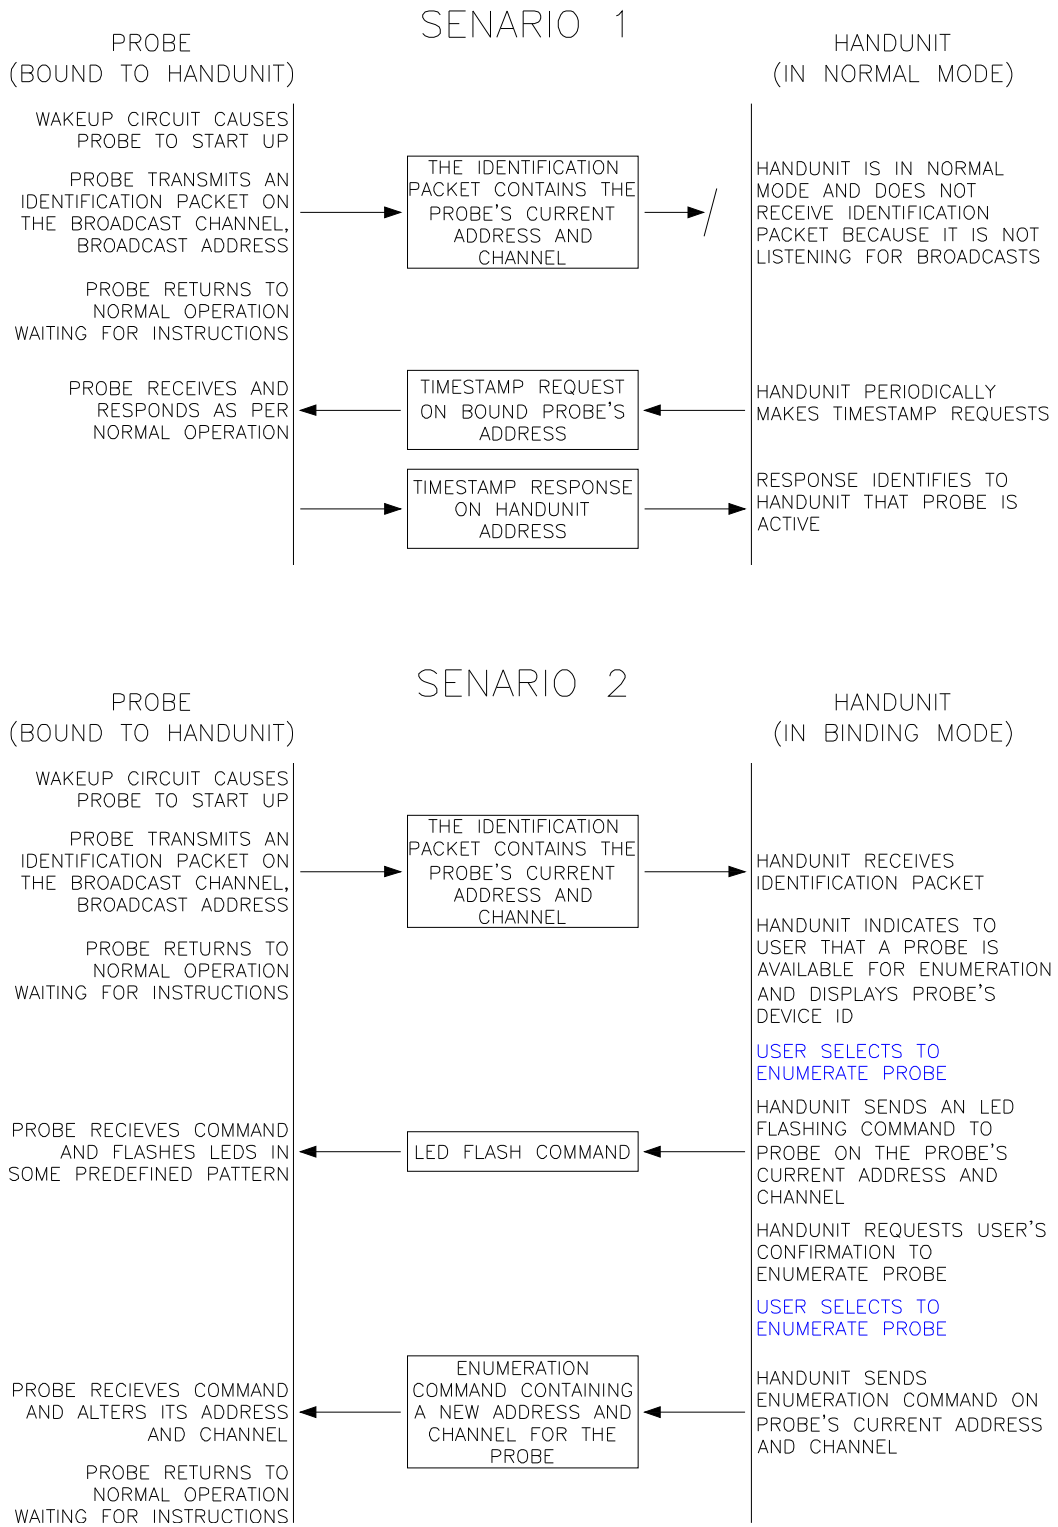

Figure 4.12 Alternative enumeration process.

missing they will be requested individually with an RF_CMD_RESEND_PACKET command once the stream of packets finishes. When all the packets have been correctly received, the handunit sends an RF_CMD_ALL_PACKETS_RECEIVED command to return the probe to its normal state. The handunit continues requesting data in this way until it has downloaded all the data it requires.

By not having to acknowledge each packet the speed of the download is increased while the worklist algorithm maintains a reliable transfer. To improve the speed of the system further the buffer size needs to be expanded so that the process is repeated fewer times. This requires a microcontroller for the probe with more memory or better optimisation of existing code.

4.5 SYNCHRONISATION

The wireless connection of devices means that the measurements taken by each probe are completely independent of each other; they have no common timescale. The timescale of a measurement made by a particular probe is known only relative to the local time of that probe. The relationship between the local time of each device is unknown, so the relationship of the measurement timescales from multiple probes is also unknown.

The synchronisation process uses timestamps from the handunit and the probes to relate the local time of each probe to the local time of the handunit. It follows that the timescales of the measurements can also be related to the local time of the handunit and thus to each other.

4.5.1 Local time and timestamps

Local time refers to the time kept by the timestamping software on a particular device. The details of this are discussed in Section 4.2.4. With a 16 MHz clock source, the timestamps measure local time with a resolution of 62.5 ns. Timestamps have a maximum value of just over four minutes, at which point they reset to zero.

4.5.1.1 Clock sources

Because the devices are not connected in any physical way, the system clocks on each piece of hardware have to be created by independent oscillators. The clock source on each device has a nominal frequency of 16 MHz. However, the accuracy of these oscillators is imperfect. The particular crystal oscillators used in the probes[8] have an accuracy of ±50ppm[9], while the handunit uses a quartz crystal oscillator with the same accuracy. By using Rubidium or Cesium crystals (these are the kinds used in

[8] ABM3B - Ultra miniature surface mount microprocessor crystals.
[9] Parts Per Million.

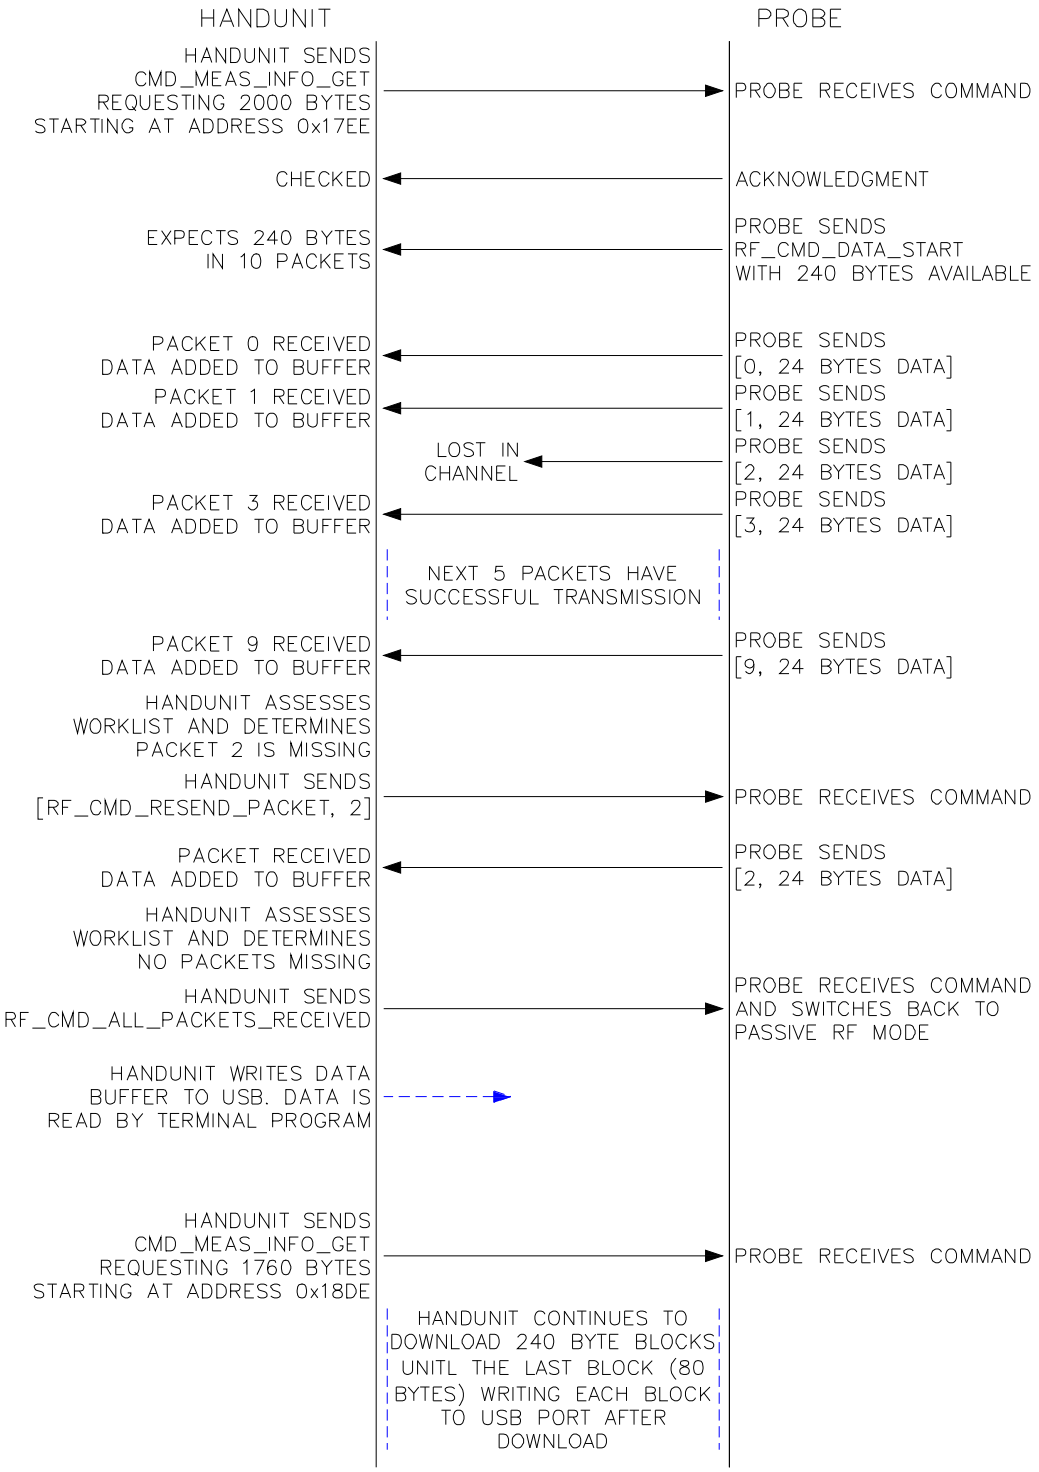

Figure 4.13 Data download process showing the handling of a lost packet.

atomic clocks), the accuracy would be improved to around 25×10^{-6}—375×10^{-9}ppm, however, the cost of such items is much, much greater.

The crystal oscillator is the most accurate of the standard methods for clock generation [Gadre 2001]. There are other, more accurate alternatives for clock sources (at slightly greater expense) that may be able to improve future prototypes. These include more expensive better crystal oscillators (XO) or temperature controlled crystal oscillators (TCXO) [Filler and Vig 1993].

There are three primary causes of varying crystal frequency:

- Temperature
- Supply voltage
- Aging

Aging has a much smaller effect than the other two causes and is ignored in this application. Identical power supplies are used on each device and they are operated in the same environment, minimising the difference in temperature and supply voltage. Despite this, the system clocks on each unit persist in having noticeably different frequencies. This difference is enough to cause added difficulty in the synchronisation of the measurement timescales.

4.5.1.2 Relating local time on two units

A first order estimation is used to relate the two local time on two devices. The local time difference of the units (α) and the relative frequency of their clocks (β) are the only parameters needed to do this. If the local time on the handunit is t_H and the local time on probe 1 is t_{P1} then they are related by:

$$t_H = \beta t_{P1} + \alpha$$

When α and β are known it is a simple process to convert the timescale of each probe's measurement into handunit local time. By applying the same process to the second probe the timescales of both sets of measurement data can be made relative to the handunit's local time and thus to each other.

4.5.1.3 Timestamps

The only timing data available to estimate α and β are the timestamps from acknowledgment packets (see Section 4.4.2.3). A complete timestamp contains the local time of the handunit and the local time of a probe. Figure 4.14 shows the process of acquiring a complete timestamp.

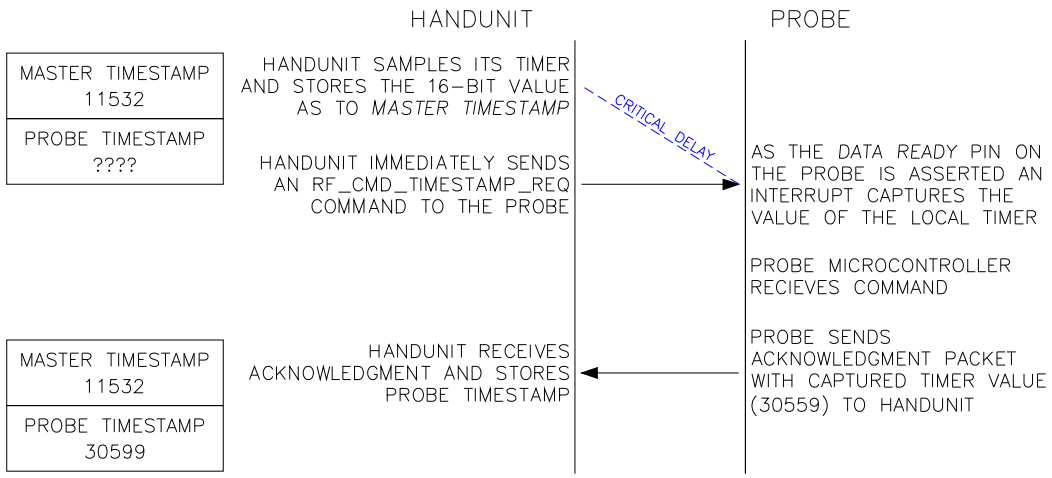

Figure 4.14 Timestamping process.

	Probe 1	**Probe 2**
Points Analysed	1320304	1320270
Max Jitter Element	0.25894	0.26738
Mean (μ)	-1.2342×10^{-6}	-6.6961×10^{-6}
Std Dev (σ)	0.063353	0.066148
99th Percentile (3σ)	0.19006	0.19844

Table 4.2 Jitter measurements.

The two timestamps are not taken simultaneously because there is a delay (labeled *critical delay* in Figure 4.14) from when the handunit makes its own timestamp to the execution of the interrupt on the probe, capturing the probe's timestamp. The section of code that sends the RF packet uses only low level instructions to ensure that the software based part of this delay is constant. The hardware delays of the RF modules are unfortunately not controllable.

4.5.1.4 Timestamp jitter

There is a still a certain amount of jitter that occurs in this critical delay due to the variable delays of the RF module. The jitter is analysed by taking a long sequence of timestamps (1.3×10^6) on both probes and calculating a moving average of the difference between the handunit and each probe. The jitter is then approximated as the difference between the moving average and the real data points. The histogram in Figure 4.15 suggests a zero mean, Gaussian type distribution for the jitter. Measurements of the jitter data are found in Table 4.2. The table gives a maximum timestamp error of $0.25\,\mu$s.

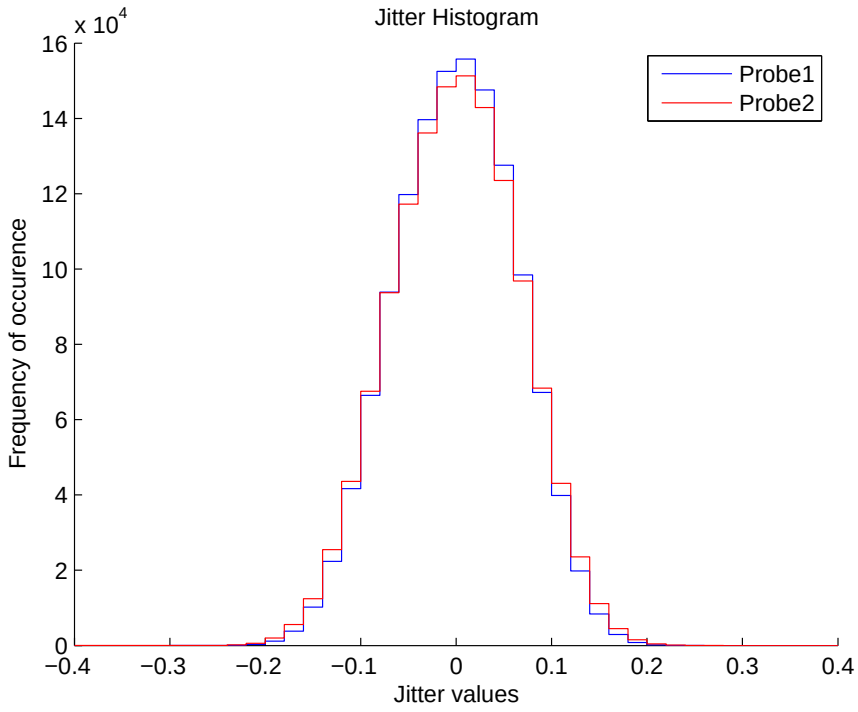

Figure 4.15 Histogram of timestamp jitter.

4.5.1.5 Estimating parameters

Parameters α and β are estimated from a series of complete timestamps. The timestamps can be plotted with the handunit time on the x-axis and the probe time on the y-axis. α and β are then estimated as the offset and slope of a straight line fit through then plotted points. A minimum mean square error calculation is used to determine the optimal values of that line.

4.5.2 Clock stability

Being able to estimate values for α and β (see above) gives a fix on the slight differences in frequency between the clock sources for each unit. Clock *stability* refers to the rate of change of the value for β. Figure 4.16 shows a typical set of relative timestamps for the handunit and two probes over a period of approximately 3 minutes. 10000 timestamps were collected for each probe at 18 μs spacings. Applying a first order approximation to both data sets gives an estimate of α and β for each probe. At first glance the straight line appears to be an excellent fit.

The approximation error, shown in Figure 4.17, is the difference between the straight line fit and the actual timestamp values. The frequency of the clock sources on both units is changing over the three minute window indicating significant clock instability. The sharp change in direction of the errors in the approximation for probe 1

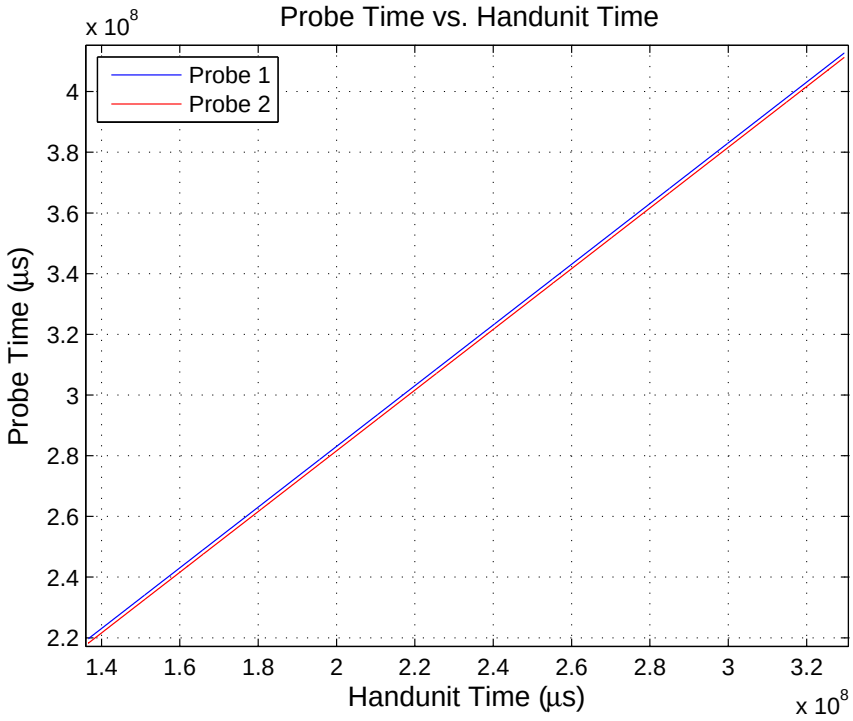

Figure 4.16 Timestamp values for local probe time plotted against local handunit time. 10 000 samples collected over approximately 3 minutes.

indicates an ability for the source frequency to change rapidly.

The problem of clock instability is exaggerated by the fact that timestamps can only be requested once the measurement loop has finished. The subsequently requested timestamps can only be used to extrapolate the conversion parameters of the timescale during the measurement cycle. More complex models for time conversion will not produce lower errors without additional information to predict the clock behavior. One possible improvement to this may be to use temperature monitoring for more accurate prediction. The most simple solution for future prototypes will be to use more stable clock sources.

4.5.2.1 Period of sampling and period of stability

If the period of sampling is defined as the interval when the timestamps are taken, then there is a period of stability either side of this. The period of stability is the interval when the error of the local time conversion is acceptably low. The length of this period is affected by the maximum tolerable error and by the length of the timestamp sampling period t_{est}. For measurement timescales to be converted to a common timescale with confidence, the data must be measured within the period of stability.

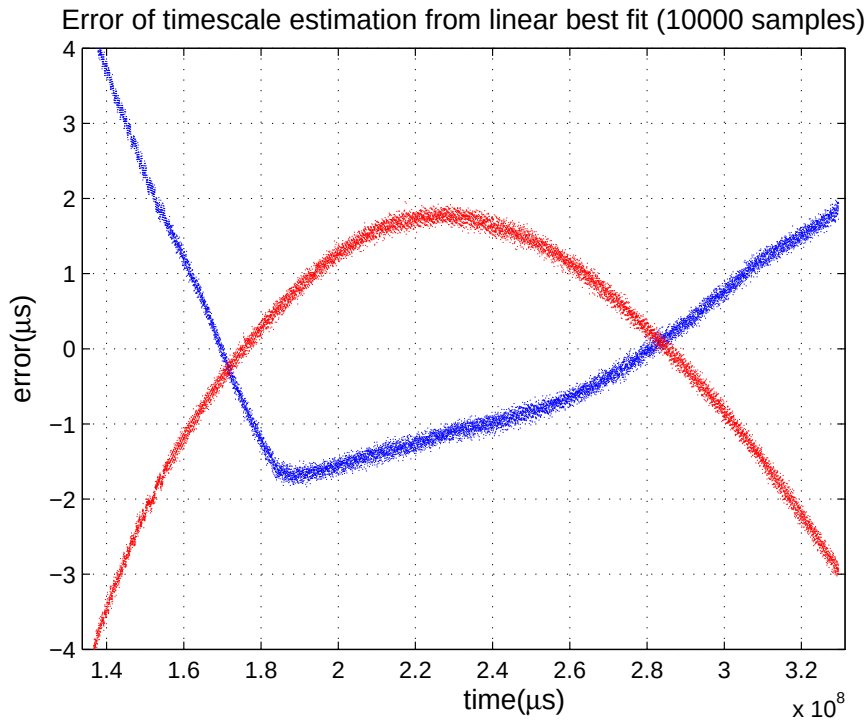

Figure 4.17 Approximation error between probe timestamp values from Figure 4.16 and a straight line fit. Timestamp jitter causes the vertical scattering of the points.

4.5.2.2 Optimal length of timestamp sampling

A high number of timestamps sampled gives the best averaging of the timestamp jitter error. A short timestamp sampling period is less likely to run into clock instability. To find the optimal length of t_{est}, a brute force investigation involving a huge number of timestamps was undertaken. Approximately 1.3 million timestamps were collected from each probe over a period of 6 hours. For 100 000 of these timestamps several parameter estimations are made using different periods of sampling. The estimation is compared to the timestamp values prior to the period of timestamp sampling. A series of error values are plotted on the horizontal axis of Figures 4.18 and 4.19. The vertical axis shows the minimum time from the period of timestamp sampling until the error values are reached. The legend shows the different t_{est} values used.

Figures 4.18 and 4.19 illustrate that the optimal length for t_{est} is between 2—10 seconds, but depending on the maximum acceptable error, it may not be necessary to sample this long.

4.5.2.3 Minimum length of timestamp sampling

An optimal value for t_{est} of around 5 s would involve collecting over 250 timestamps for each probe. This requires more processing power, time, and RF power. The operator

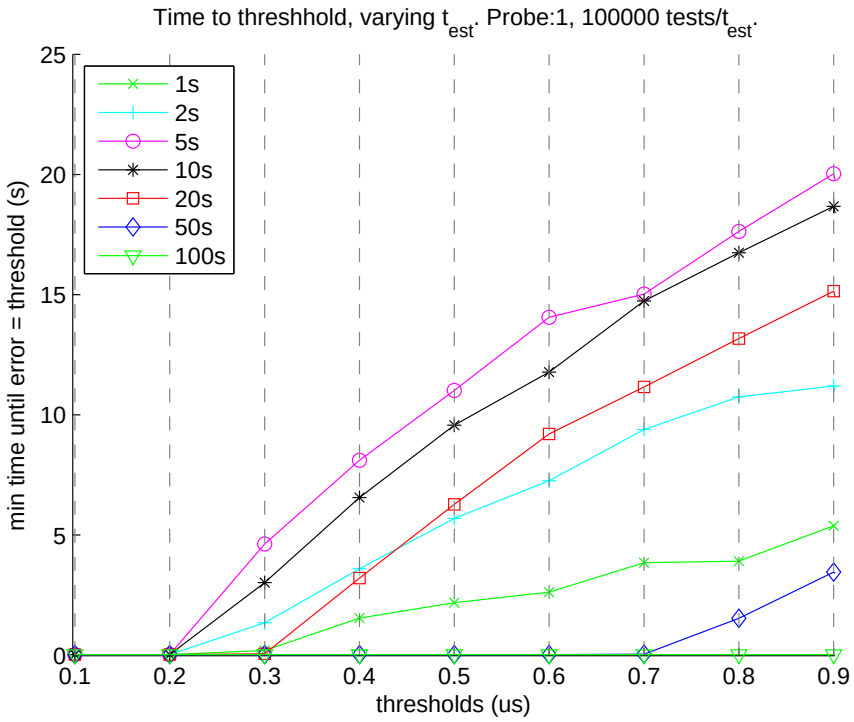

Figure 4.18 Minimum time until error values - Probe 1.

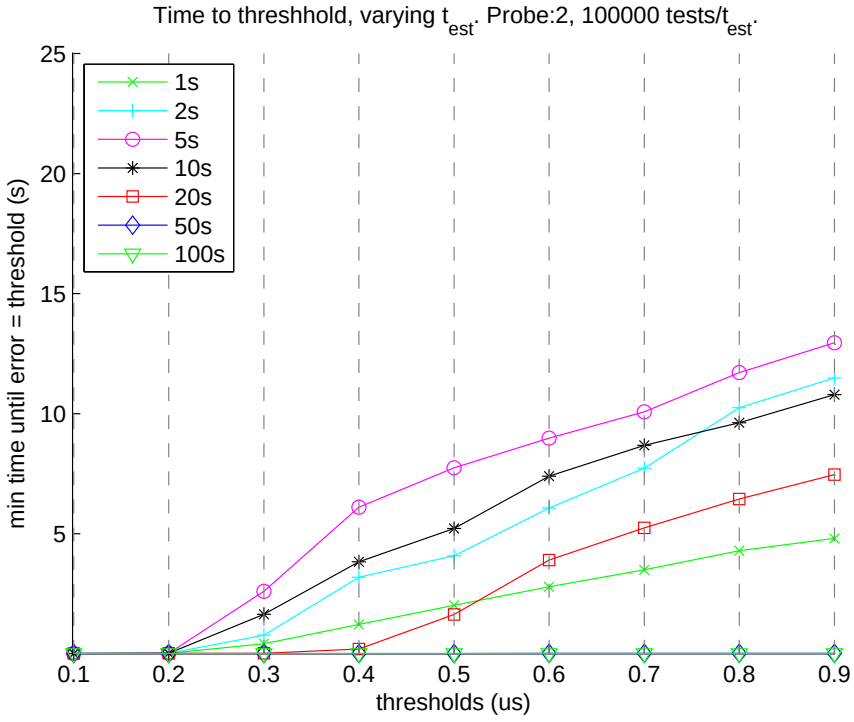

Figure 4.19 Minimum time until error values - Probe 2.

also has to wait the 5 s for the timestamps to be collected. In this case, it is more appropriate to seek the minimum acceptable value for t_{est}.

The longest possible measurement time is approximately 100 ms (limited by the size of the probes' FRAM). On top of this it may take around 50 ms for the handunit to begin collecting timestamps. This makes the required region of stability only 150 ms. The minimum acceptable value for t_{est} is found by deciding on a maximum tolerable error. Note that errors values less than 0.2 μs have minimum times near zero due to timestamp jitter. Not until clear of the timestamp jitter (see Section 4.5.1.4) do the minimum times become significant. Selecting a maximum tolerable error of 0.4 μs, a t_{est} of around 1 s is acceptable. This requires only around 50 timestamps per probe and results in much shorter wait times for the user.

4.6 SUMMARY

The Treetap 5 prototype is successful. It fulfills both its mandates to provide a new acoustic measurement system and to test new possible design features. With the benefit of hindsight, however, it is apparent that aspects of the design could certainly be improved.

4.6.1 Positive features

Review first some of the particularly positive features of the design:

- High SNR
- 150 kHz sampling rate
- Flexibility

By having the signal conditioning and digitisation electronics in close proximity to the accelerometer very little interference enters the measurement, giving a high SNR. The programmable amplification system provides a wide range of gain (from 1—1024) with a minimum of added noise. With the use of a 16-bit ADC quantisation noise is minimal. The high sampling rate gives a high resolution view of the detected acoustic signal. However, if a longer duration of sampling is required delays can be added into the software timing loop to decrease the sampling frequency.

4.6.2 Suggested future work

There are some very small alterations to some features that could be made to the device. Most of these have been discussed at the appropriate sections within the chapter. These include:

- Changes to the power management system. See Section 4.2.6.1.
- Inclusion of additional memory for the handunit. See Section 4.3.4.1.
- An alternative scheme for probe enumeration. See Section 4.4.3.1.

Other suggested minor changes include:

- Removal of the signal detection circuit on the probe. This is superfluous because there is enough time during the measurement loop to observe the digitised data in software and detect the beginning of a signal. It was added for comparison with previous Treetap systems [Wang 2003].
- Using a larger data buffer for data downloads. This will increase the efficiency of the transfer and reduce the wait time. Alternatively, the buffer might be done away with altogether and the FRAM accessed directly.

The above suggestions for future prototypes, however, are not the main limiting factors of the system. The accelerometer bandwidth and the wireless communication network are much more significant problems.

4.6.2.1 Accelerometer bandwidth

The processing of the voltage signal from the accelerometer has high SNR, flexible amplification, high quality digitisation, and a high sampling rate. None of these features, however, are able to make up for inadequacies of the accelerometer itself.

The initial design of the system was based on an assumption that the accelerometer was a high bandwidth device. A suitable sampling rate (150 kHz) was selected accordingly and the anti-aliasing filters cutoff frequency set at 50 kHz. Investigation of the SD-02 accelerometer (detailed in Section 3.6) revealed that the resonant frequency of the device is a much lower than expected 23 kHz and the usable bandwidth of the device approximately 15 kHz. As a result, signals detected by the accelerometer that have higher frequency components than the accelerometers bandwidth will be misrepresented. Frequency components near the resonant frequency will increase in magnitude, while higher frequency components will be heavily attenuated.

Two suggestions for fixing this problem:

- Redesign the accelerometer to have a higher bandwidth. Increasing the bandwidth of the accelerometer will result in a less sensitive transducer (see Section 3.6.3). But, since there is ample signal amplification available in the front end circuitry this can be compensated for.

- Alternatively, by reducing the cutoff frequency of the anti-aliasing filters to the 15 kHz bandwidth of the accelerometers the magnitude increase at resonance can be removed or even reduced. This does not increase the bandwidth of the detection system, but does ensure that the signal recorded is truly band-limited to the frequency range or which the accelerometer is accurate.

4.6.2.2 Wireless communications

The wireless communication system connecting the handunit to the probes causes multiple problems. The convenience of having no cables is at the expense of high complexity, slow data transfer, extra processing time, extra program space, and timing accuracy.

The speed of data transfer is lowered dramatically by the wireless system. For commands and timestamps the slow rate of transfer is not noticeable, but when uploading measurement data from the probes it is. The wait times are proportional to the amount of data being downloaded, but even for modest quantities of data significant wait times are involved. To upload the whole 32 kByte memory from both probes the delay is over two minutes.

Once data has been uploaded, timestamps must be collected for synchronisation and the constants to estimate relative clock speeds evaluated. The clock rate adjustment must then be applied to the result of whatever detection or processing algorithm is being used. Although the synchronisation process is able to reduce the error caused by the clock rate to less than half a microsecond, there is a high price in terms of complexity, processing time, and program space to accomplish this.

Ultimately, an unstable or inaccurate clock source limits the accuracy of the system regardless of whether the system is wireless or not. The problem with a wireless system is that instability errors are compounded by having multiple clock sources. The half microsecond error should then be considered an additional error (per probe) on top of a potentially inaccurate master clock source.

The simplest solution to the problems presented by the wireless system is to do away with it entirely. It causes too many major problems, including fundamental timing accuracy issues, to justify the slight extra convenience of not having cables in the field. The advantages of having electronics close to the accelerometer could be maintained by having a common bus like the possible system depicted in Figure 4.20. A CAN bus may be a suitable solution [Tindell *et al.* 1995]. Advantages to this type of system over a wireless one include:

- Amplification and sampling hardware is located close to the accelerometer for optimal SNR.

- Data transfer speeds will be much higher than for the wireless system.

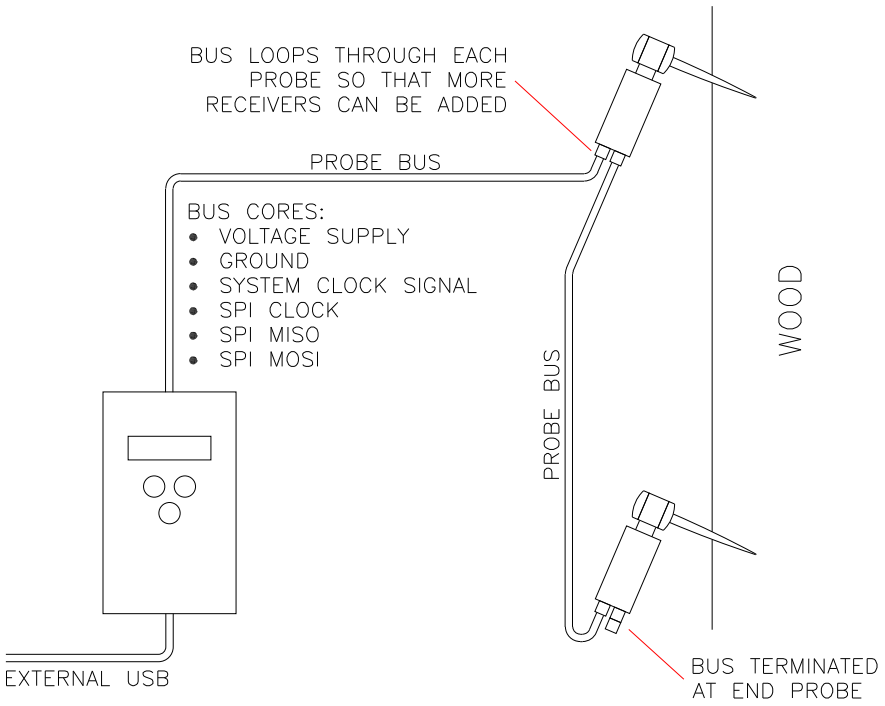

Figure 4.20 Alternative bus driven communication system

- A common system clock eliminates the need to synchronise data and all associated problems.
- The probes are chain linked. As a result, the number of probes is not restricted.

4.7 CONCLUSIONS

The conclusion of this thesis is in many ways more of a beginning than an end. The research and design described here is only a stepping stone to making the most of acoustics in timber inspection. Stress wave timers already being used for time of flight measurements use simple algorithms, like threshold detection, that ignore the fundamental problem of dispersion. As shown in Section 3.7.1, this has the potential to produce results that cannot be related to the true stiffness.

Resonance testing systems are able to cope much better with problems like dispersion, multiple paths, and defects. While this makes them the obvious choice between the two methods of stress wave timing they still have a major, fundamental flaw: they will always be limited to use on logs. Although still a useful tool in some respects, making available the stiffness of logs already committed to a sawmill is somewhat redundant. For any forestry company that wants to ascertain the properties of their trees without them being felled, a stress wave timer is currently the only real option.

However, as stated above the current suite of stress wave timers have some major inadequacies of their own. To understand these problems more fully, more investigation is required. A lack of tools tailored to acoustic inspection of wood has made it difficult to study in detail the effect of superimposed waves, dispersion, and defects. To overcome this the Treetap 5.0 prototype was designed and constructed.

4.7.1 Suggestions for future experiments

With a better system for gathering acoustic data experiments can now be undertaken to attempt to quantify the effects of dispersion, multiple paths, defects, and transducer coupling. Some suggested experiments include:

- Investigation of the effects of multiple signal paths through changing sample geometry.
- Use multiple accelerometers, at a common point, in different alignments to detect the presence of p-waves and shear waves to determine how an angled accelerometer may be affected by both types of wave.
- Use specific sample geometry to investigate the effect of dispersion. In a small diameter wooden rod the propagating wave should be approximately planar. Alignment of the accelerometer with the axis should detect only the p-wave component. The sample could now be considered a transmission line and it may be possible to characterise the loss mechanisms.
- Better modelling of boundary conditions to assess geometric dispersion.

4.7.2 Treetap 5.0

The Treetap 5.0 system, while a successful prototype, has some major room for improvement (as discussed in Section 4.6).

Many of the fundamentals, and details, of the design are good. In particular:

- The transfer of the voltage signal from the accelerometer to the digitally stored data is done at a high sampling rate and with a minimum of added noise.
- The system is deliberately not limited to two receivers to allow a range of experiments to be conducted.
- The USB connection provides a fast, universal interface to a PC.

The areas of the design requiring the most attention are the wireless link and the transducer. The wireless communication system simply does not add enough value to the system to justify the huge increase in complexity, synchronisation issues, and loss of transfer speed. It may be possible to overcome these obstacles without creating more problems, but the simplest solution to improving the communication between the devices is to use a physically connected system.

As the first step in recording acoustic signals, the limitations of the accelerometer are perhaps the most serious. The old adage: garbage in–garbage out, is true of the electronics in the device. If the electrical signal produced by the accelerometer is a poor representation of the acoustic signal in the timber, then the recorded digital signal is likewise a poor representation. In fairness to the SD-02 device it *is* a good, linear accelerometer up until 12—15 kHz. However, to analyse accurately signals with higher frequency components, it will be necessary to find a substitute for the accelerometer. One option to cover all the possible uses of the device may be to make the accelerometer detachable in future prototypes, so that transducers tailored to specific experiments can be used.

REFERENCES

Ambardar, A. (1999), *Analog and Digital Signal Processing*, Brooks/Cole Publishing Company, 2nd ed.

Andrews, M. (2002), 'Wood quality measurement', *NZ Journal of Forestry*, November.

Auld, B.A. (1973), *Acoustic Fields and Waves in Solids, Volume I*, John Wiley & Sons.

Axmon, J. and Hansson, M. (1999), *Nondestructive Detection of Decay in Spruces Using Acoustic Signals: Evaluation of Circumferential Modes*, Technical Report, Lund Institute of Technology, May.

Beall, F.C. (2002), 'Overview of the use of ultrasonic technologies in research on wood properties', *Wood Science and Technology*, Vol. 36, pp. 197–212.

Bucur, V. (2001), 'Ultrasonic energy flux deviation and off-diagonal elastic constants of wood', *IEEE Ultrasonics Symposium*, pp. 697–700.

Bucur, V. (2006), *Acoustics of Wood*, Springer-Verlag, 2nd ed.

Chauhan, S. (2004), *Selecting and/or Processing Wood According to its Processing Characteristics*, PhD thesis, Department of Forestry Engineering, University of Canterbury.

Corbett, S. (2004), *The Practical Woodworker*, Hermes House.

Cunningham, D.R. and Miller, J.A. (1995), *Circuit Analysis*, Houghton Mifflin Company, 2nd ed.

Divos, F. and Divos, P. (2005), 'Resolution of stress wave based acoustic tomography', In *14th International Symposium on Nondestructive Testing of Wood*, University of Applied Sciences, Germany, May.

Emms, G. and with assistance from Mike Andrews, J.I. (2004), *Background Review of Dynamic Stress-Strain Relation and Elastic Wave Propagation as Applied to Wood or Similar Materials*, Technical Report, WQI Report No. Str 7, February.

Fife, D., Ilic, J., M, K., Lausberg, M. and McConchie, D. (2004), *Assessment of Standing Tree Stiffness in Australian Radiata Pine*, Technical Report, WQI Report No. Str 10, September.

Filler, R. and Vig, J. (1993), 'Long-term aging of oscillators', *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, Vol. 40, No. 4, July, pp. 387–394.

Fox, G.R., Chu, F. and Davenport, T. (2001), 'Microelectronics and nanometer structures', *Journal of Vacuum Science and Technology*, Vol. 19, No. 5, September, pp. 1967–1971.

Gadre, D.V. (2001), *Programming and Customizing the AVR Microcontroller*, McGraw-Hill.

Harris, P., Petherick, R. and Andrews, M. (2002), 'Acoustic resonance tools', In *13th International Symposium on Nondestructive Testing of Wood*, pp. 195–201.

Hunter, A.J. (2006), *Underwater acoustic modelling for synthetic aperture sonar*, PhD thesis, Department of Electrical and Computer Engineering, University of Canterbury, New Zealand, June.

Kinsler, L.E., Frey, A.R., Coppens, A.B. and Saunders, J.V. (1982), *Fundamentals of Acoustics*, John Wiley & Sons, Inc., 3rd ed.

Kolsky, H. (1963), *Stress Waves in Solids*, Dover Publications, Inc.

Moore, S.W. and Schneider, P.J. (2001), *A Review of Cell Equalization Methods for Lithium Ion and Lithium Polymer Battery Systems*, Technical Report, Delphi Automotive Systems.

Schoch, W., Heller, I., Schweingruber, F. and Kienast, F. (2004), 'Wood anatomy of central european species'. Online version: www.woodanatomy.ch.

Smith, J.C. (1966), 'Wave propagation in a three-element linear spring and dashpot model filament', *Journal of Applied Physics*, Vol. 37, No. 4.

Tindell, K., Burns, A. and Wellings, A. (1995), 'Calculating controller area network (CAN) message response times', *Control Engineering Practice*, Vol. 3, No. 8, pp. 1163–1169.

Wang, M. (2003), 'Treetap project progress report 1', June. Internal Report: University of Canterbury, Department of Electrical Engineering.

Wang, H. and Williams, K. (1996), 'Vibrational modes of thick cylinders of infinite length', *Journal of Sound and Vibration*, pp. 955–971.

Wang, X., Ross, R.J., McClellan, M., Barbour, R.J., Erickson, J.R., Forsman, J.H. and McGinnis, G.D. (2001), 'Nondestructive evaluation of standing trees with a stress wave method', *Wood and Fiber Science*, pp. 522–533.

Young, P.H. (1999), *Electronic Communication Techniques*, Prentice Hall, 4th ed.