INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

Ultrasonic Arterial Vibrometry
With Wavelet Based Detection and Estimation

by

Melani Irene Plett

A dissertation submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2000

Department of Electrical Engineering

Doctoral Dissertation

In presenting this dissertation in partial fulfillment of the requirements for the Doctoral degree at the University of Washington,
I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of the dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to Bell and Howell Information and Learning,
300 North Zeeb Road, P.O. Box 1346, Ann Arbor, MI 48106-1346, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature Melanie Irene Platt

Date March 2, 2000

University of Washington
Graduate School

This is to certify that I have examined this copy of a doctoral dissertation by

Melani Irene Plett

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Chair of Supervisory Committee:

Kirk W. Beach

Reading Committee:

Donald B. Percival

James A. Ritcey

Date: March 2, 2000

University of Washington

Abstract

Ultrasonic Arterial Vibrometry
With Wavelet Based Detection and Estimation

by Melani Irene Plett

Chairperson of the Supervisory Committee

Research Professor Kirk W. Beach
Department of Surgery

Since the invention of the stethoscope, detection of vibrations and sounds from the body has been a touchstone of diagnosis. However, the method is limited to sounds transmitted to the skin with no means to determine the anatomic and physiological source of the sounds save the cunning of the examiner.

To aid diagnosis, researchers have attempted to correlate the characteristics of the sounds to the type and significance of the disorder, and they have sought to understand the mechanisms that generate the vibrations producing the sound. Despite three decades of research these goals have met limited success. Important obstacles have been the inability to non-invasively and efficiently 1) detect the vibrations and 2) fully measure the vibration characteristics.

This document introduces an extension to current techniques of detecting and measuring arterial vibrations. Using ultrasound quadrature phase (complex) demodulation methods similar to those of ultrasonic color flow imaging, we have developed a system to detect and measure tissue vibrations with amplitude excursions as small as 30 nanometers. Morlet wavelet analysis, combined with non-parametric binary hypothesis testing, permits sensitive and specific detection and measurement of short duration vibrations amidst clutter and time-varying, colored noise. Specifically, the wavelet spectra will consist of Gaussian noise under the null hypothesis that vibration and blood flow are absent. Determination of the presence of vibration is then based on the expected distributions of the novel normalized wavelet power

spectra or normalized cross wavelet spectra associated with the Gaussian noise, as derived in this document. Receiver operating curves for the detector were generated from simulated vibrations. These curves reveal expected detection rates in excess of 99.5% at false alarm rates below 1% for signal to noise ratios as small as 0.5. Subsequent to detection, the Morlet wavelet is also used for characterizing the detected vibration. Vibrations and their quantified characteristics are displayed on conventional ultrasound images to aid the clinician's diagnosis.

The detector and estimator have been successfully applied to phantom (physical model) and *in vivo* vibrations. A few examples are shown, including an arterial bleed. The current system is offline, though a real-time implementation will be straightforward.

TABLE OF CONTENTS

LIST OF FIGURES

Number *Page*

LIST OF TABLES

GLOSSARY

Analytic: A complex representation of a real signal in which all negative frequency magnitudes are set to zero.

Aortic valve: A valve in the heart lying between the left ventricle and the aorta.

Aneurysm: A region of dilatation (enlargement) in a blood vessel wall.

Arterio-venous shunt or fistula: An abnormal communication between an artery and a vein caused by defect, surgery or trauma.

Atherosclerosis: Thickening, loss of elasticity and calcification of an arterial wall.

Auscultation: The process of listening, usually with the aid of a stethoscope, to sounds produced by movement of gas or liquid in the body.

AV shunt: See arterio-venous shunt.

Broadband: A signal with a low quality factor*. If the signal were in the audio range, it would have a harsh or hissing sound.

Colored noise: Unlike white noise which has a constant variance across all frequencies, the variance of colored noise changes with frequency.

Diagnostic peritoneal lavage: A test to detect hemorrhaging in the abdominal cavity. It involves insertion of a catheter into the peritoneum*. If no blood, bile, urine or intestinal fluid can be aspirated, the abdominal fluid is irrigated with 250 to 1000 ml of warmed saline. If feasible, the patient is moved from side to side to assure the saline reaches all areas of the abdominal cavity. The fluid is then allowed to drain via gravity. The resulting fluid is tested for the presence of blood.

Diastole (diastolic): The period of the cardiac cycle in which the heart chambers refill.

Dissecting aneurysm: A separation of the medial layer of a vessel wall from the outer (adventitial) layer, creating two vessel lumens*.

Filter bank: A series of filters with different frequency passbands arranged to span a desired large frequency band.

Hidden Markov model: This is a pattern recognition method based on some training set. A model of sequences of events is built where the probability of an event occurring depends upon the fact that a preceding event occurred, with the probabilities determined from the training set. The term hidden implies that the individual events cannot be directly observed.

Innocent bruit or murmur: Bruits or murmurs that are not associated with a disorder.

Intrachamber septa: The partition between the left and right halves of the heart.

Lumen: Open channel of a vessel.

---

* Elsewhere in glossary

Matched filter: A filter whose time domain shape is the same as the signal to be detected, except reversed in time. This filter will yield the best signal-to-noise ratio to enhance detection.

Mitral valve: A valve in the heart at the opening between the left atrium and the left ventricle. It allows blood to pass from the atrium to the ventricle.

Narrowband: A signal with a high quality factor*. If the signal is in the audio range, it would have a pure tone.

Neural networks: A type of artificial intelligence that attempts to imitate the way a human brain works. A neural network works by creating connections between processing elements, the computer equivalent of neurons. The organization and weights of the connections determine the output. Neural networks are particularly effective for predicting events when the networks have a large database of prior examples to draw on.

Palpation: The process of examining the body by careful feeling with the hands and fingertips.

Pansystolic: A period of the cardiac cycle encompassing both systole and diastole.

Patent ductus arteriosis: The ductus arteriosis is a blood vessel which connects the pulmonary artery directly to the ascending aorta during fetal development. Patent ductus arteriosis is the condition in which this vessel remains open after birth.

Periodogram: The magnitude squared of the discrete Fourier transform of a time series tapered by a rectangular window.

Peritoneum: The membrane surrounding the abdominal cavity.

Pseudo-aneurysm: A pseudo-aneurysm differs from a true aneurysm* in that its wall does not contain the components of an artery. The wall of a pseudo-aneurysm consists only of fibrous tissue and surrounding structures

Pulmonic valve: A valve in the heart lying between the pulmonary artery and the right ventricle.

Ray line: One ultrasound beam location.

Quality factor: The ratio of a signal's center frequency to its bandwidth.

Stenosis: A region of narrowing or constriction in a blood vessel, often caused by the build-up of atherosclerotic plaque.

Sternum (sternal): The breastbone.

Subcutaneous: Occurring below the skin.

Systole (systolic): The period in the cardiac cycle in which the cardiac muscle of the left ventricle contracts, forcing blood into the aorta.

Tricuspid valve: The valve in the heart between the right atrium and the right ventricle.

Vasospasm: Sustained constriction of vascular smooth muscle causing significant vessel narrowing.

---

* Elsewhere in glossary

Vein graft: A section of vein is placed in lieu of a section of diseased artery.

Zero padding: The process of adding extra zeros to the end of a time series before applying a discrete Fourier transform in order to achieve a finer frequency spacing.

LIST OF VARIABLES

A	amplitude of an ultrasonic RF echo (Volts)
A_v	vibration displacement amplitude, or amplitude estimate
a_k	an autoregression coefficient
C_δ	amplitude correction factor for wavelet estimation
c	speed of sound
D_δ	intensity correction factor for wavelet estimation
F	frequency of ultrasound transmission (Hz)
$F(2,\nu)$	a statistical F distribution with 2 and ν degrees of freedom
f_D	Doppler frequency of moving tissue (Hz)
f_v	vibration center frequency (Hz)
I	in-phase representation of a signal
I_{avg}	average intensity of acoustic thermal noise
I_{inst}	instantaneous intensity of acoustic thermal noise
J	the number of depths sampled along an RF line
j	sqrt(-1)
K	the number of pulses in a color Doppler ensemble
k	a time index
k_B	Boltzmann's constant
L	the number of frames collected in a color Doppler sequence
M	a point under test for Fourier-based adaptive filtering
m	multiplicative factor for a chi-square approximation to the noise estimate distribution
N	the length of a time series
n	a depth index
n_t	samples of a noise process
PRF	ultrasound pulse repetition frequency
P	Power estimate
p	acoustic pressure
Q	quadrature-phase representation of a signal
Q_v	vibration quality factor*
Q_w	wavelet quality factor*
R	reflection coefficient
RF	radio frequency ultrasound (MHz)
S	received ultrasonic RF echo
s	a wavelet scale
T	time between consecutive ultrasound pulses
T_K	absolute temperature (Kelvin)
T_0	temporal length of a segment of time samples

---

* See glossary.

t	"slow" time – time scale for changes at a given depth
u	the parametric model order
v	noise estimate for a detection test
v_L	initial noise estimate left of the test point
v_R	initial noise estimate right of the test point
vel	velocity of solid tissue
$W(t,s)$	wavelet transform coefficients at time t and scale s
x	displacement
X	depth in tissue
Z	acoustic impedance of tissue
α	decay factor for Morlet wavelet
ϕ	phase of quadrature representation of a signal
η	degree of freedom measure of initial noise estimate
φ	the phase between vibrations at neighboring depths
λ	wavelength of ultrasound
ν	degree of freedom measure of noise estimate
θ	the angle between the ultrasound beam line and the artery axis
ρ	density
σ	standard deviation of a time series
τ	"fast" time – time scale for ultrasound pulse propagation
ω_0	the center frequency of a wavelet (radians)

ACKNOWLEDGEMENTS

First, I acknowledge and thank Dr. Kirk Beach for generously offering his time and expertise in numerous areas ranging from ultrasound to research to teaching. Thank you, Kirk, for taking the time to bring me up to speed, for valuing my input, and for caring about me as a person. It has been a true privilege working under your guidance for the past few years.

I thank, too, the members of my supervisory committee: Drs. Don Percival, Jim Ritcey, John Sahr, and Joan Sanders. Thank you for your time and input at very important stages of my research.

Unlike many graduates students, I have had the opportunity to be a part of a large research group. Drs. Larry Crum, Kirk Beach, Steve Carter and Roy Martin are the principal investigators of the Acoustic Hemostasis project that uses High Intensity Focused Ultrasound (HIFU) to stop internal bleeding. They have graciously included me in their meetings and help to guide my research. I thank each of you PIs. I am also grateful to DARPA for funding this research under grant #N00014-96-0630.

Thank you Barbrina Dunmire for your countless hours of theoretical discussion, design and troubleshooting of data collection instruments, and puzzling over data analysis. Barb, I am indebted to you for your skillful editing of this dissertation and other papers. Thank you!

To Marla Paun, I thank you for lending your sonographic skills, graceful interaction with patients, and anatomic knowledge. Thank you also for your willingness to handle administrative tasks! You are appreciated.

I also thank the numerous engineers, students, and medical professionals that have contributed important counsel and expertise to this work: George Barrett, Mike Bertoglio, Susannah Bloch, Don Brocha, Kathy Brown, Matt Bruce, Kim Cantwell-Gab, Ron Daigle, Boyd Edmondson, John Flynn, Gene Hokanson, Peter Kaczkowski, John Kucewicz, Ravi Managuli, Chuck Matson, Mark Moehring, Tim Meyers, Curt

Plett, Jean Primozich, Helen Routh, Xuegong Shi, Frank Starr, Sherry Stemen, D. Eugene Strandness, Jr., Shahram Vaezy, and JongTae Yuk

Finally, I thank my family. In particular, I thank my husband Curt Plett, my parents and my siblings for their support and encouragement in this entire process. I could not have done this without you.

DEDICATION

This dissertation is dedicated to the one true God: the Father, Christ the Son, and the Holy Spirit. He has developed in me whatever skills and gifts have been used for this work. Along the way He as created the necessary opportunities to accomplish this research, and He has given me the strength, perseverance and wisdom to take advantage of them. May this research bring Him glory.

1 Introduction

1.1 Clinical overview

For centuries, clinicians have been listening to the human body as a means to detect and diagnose cardiovascular disorders. In 1816, René Laennec dignified this technique with the invention of the stethoscope (1819, De l'Auscultation Mediate). Today, auscultation* with a stethoscope has become routine. Clinicians listen for sounds emanating from all dynamic organs: heart, lungs, gut, arteries, and veins. Cardiovascular sounds are known as bruits or murmurs. These sounds indicate vibrations in the tissue which are likely produced by disturbances in the blood flow. Such disturbances typically arise from stenoses*, aneurysms* and arterio-venous shunts* (see Figure 1.1), although some occur in the absence of a disorder. Vibrations have also been found in punctured arteries. The sounds which are used to detect a possible cardiovascular disorder also aid in the diagnosis of the disorder type based on several characteristics of the sound, including: spatial location, temporal duration, timing in the heart cycle, frequency, loudness (intensity or amplitude), quality, and propagation distance. When these characteristics are quantified, they will potentially provide insight into the significance and natural progression of the disorder, as well as the mechanism of vibration.

Quantification of arterial vibration parameters requires the ability to capture the vibration for subsequent analysis. To date, vibrations can be captured non-invasively at the skin surface with stethoscopes and microphones, and below the skin with conventional Doppler ultrasound. None of these methods, though, has the ability to measure all of an arterial vibration's characteristics at the site of the disorder below the skin surface. The stethoscope provides the qualitative aspects of the vibration. Its utility is limited to vibrations that propagate to the skin surface with sufficient intensity to be discerned by the clinician. Phonoangiography involves analyzing

---

* See glossary.

signals received by a microphone placed at the surface of the skin. This method has succeeded in quantifying some, but not all, important vibration characteristics, and is rarely used in the clinic. Doppler ultrasound can examine regions below the skin surface, but current techniques do not promote efficient and effective vibration measurement. Researchers, therefore, have developed invasive techniques that are only applied to animals and alter the region of interest. Thus, the research has been limited by the absence of a non-invasive means of completely quantifying arterial vibrations. The fact that a lot of effort has been made to understand and quantify these vibrations over the past thirty years underlines both their importance and the continued need for an efficient means of accomplishing this task.

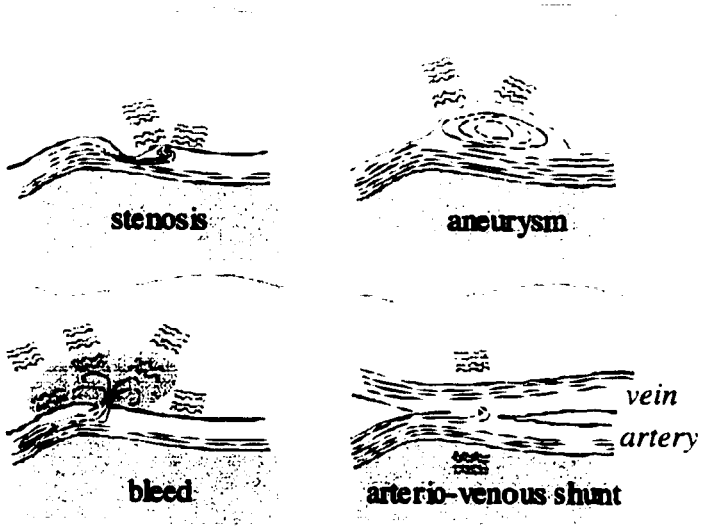

Figure 1.1. Illustration of arterial disorders.
Four arterial disorders are known to sometimes generate vibrations. The vibrations originate at the disorder site and propagate into surrounding tissue. They are thought to arise due to turbulence in the blood flow.

1.2 Approach

Detection and measurement of tissue vibrations require two steps: data collection and data analysis. The data to be collected are tissue displacements, including vibrations. Ultrasound is uniquely suited for this. Specifically, it is able to non-invasively examine subcutaneous* physiology in real time, capture signals in the vibration frequency range, and measure sub-microscopic displacements. Also important, it accomplishes these at a relatively low cost. The ultrasound method of choice, called (complex) quadrature phase demodulation (described later in section 4.1.3), is a variant of Doppler ultrasound that was shown to measure tissue displacements as early as 1968 (Baker and Simmons 1968). When appropriately processed, this method can measure the displacements associated with arterial vibrations, which can be as small as 30 nm.

The second issue in vibration detection and measurement is data analysis. Even with the myriad of modern signal processing techniques available, detection and measurement of transient oscillations with unknown onset, duration, frequency, bandwidth, amplitude and phase has remained elusive. This is true especially when the oscillation occurs in colored, non-stationary noise (i.e., frequency dependent and time varying noise) and in large, low frequency clutter, both of which are present in arterial vibrations.

We have developed a means to accomplish this detection and measurement based on the continuous wavelet transform. Use of the Morlet wavelet causes the transform to approximate a filter bank* of matched filters* for the vibration detection. The transformed data, then, lends itself to binary hypothesis testing for the presence of a signal in Gaussian distributed noise, as well as to estimation (measurement) of the vibration parameters. To further enhance detection and estimation, a cross wavelet spectrum is used to emphasize a correlation between successive ultrasound depth

---

* See glossary.

gates. To evaluate its detection ability, receiver operating characteristic (ROC) curves are generated using computer simulated vibrations. The method is then applied to *in vivo* data as well as data from physical models.

1.3 Applications

The methods outlined here for arterial vibrometry have applications in the clinic, the biomedical research laboratory and in the signal processing laboratory. For the clinician, the system developed here provides a sensitive, automatic detection and measurement method that will enhance diagnosis and treatment of arterial disorders that generate vibrations. The ability to non-invasively measure subcutaneous vibrations will also afford long-awaited measurements in the research of arterial disorders. From a signal processing perspective, this detection system expands existing Morlet wavelet binary hypothesis testing methods to a class of signals where traditional detection methods fail: transient, oscillatory signals in additive Gaussian noise with unknown, non-stationary noise variance in low frequency clutter.

In addition to arterial disorders, other physiology can be studied with this technique: muscle tremor, speech, respiration, and digestion are all known to produce similar vibrations.

Finally, the vibration signals of interest here resemble signals that occur in other settings: underwater acoustics, radar, sonar, seismology and vibrating mechanical systems. In addition, if the temporal coordinate is replaced with an appropriate spatial coordinate, the analysis can also be applied to identifying transient periodic structures in images, provided they have the appropriate noise characteristics.

1.4 Project Statement

The long term goal of this research, which reaches beyond this thesis, is to develop an automatic, real time system to detect, measure and locate the vibrations that arise from stenoses*, aneurysms*, arterio-venous shunts* and bleeds. The system will obtain vibrational amplitude, frequency, and duration measurements, as well as locate the

---

* See glossary.

vibration source temporally in the cardiac cycle. Analysis is applied simultaneously to multiple spatial locations within a two dimensional region of the body (~ 10 cm^2). This provides information about the spatial distribution of the vibrations, as well as fast, efficient detection and measurement of the vibration. Combined, this information will aid understanding, diagnosis, prognosis and management of the arterial disorder. For ease of interpretation, the detected vibration parameters will be superimposed on a conventional ultrasound image. Also, to minimize the efforts needed for commercial implementation and United States Federal Drug Administration approval, the design uses straightforward modifications to the transmit sequences of modern ultrasound instruments.

Currently, a one-dimensional, offline system is in place and the success of the system has been demonstrated in humans and animals. A two-dimensional real-time system has been outlined, but has not yet been implemented.

This document will introduce the technique that will aid clinical diagnosis, enable further research of vascular disorders that generate vibrations, and provide a novel detector and estimator for an important class of signals. Specifically, I will provide background, theory, results, and plans for clinical implementation of non-invasive, arterial vibrometry.

2 The clinical problem: a patient that presents with potential vibrations

A patient who has an arterial disorder may or may not be symptomatic upon presentation to a clinician. Possible presenting symptoms depend on the type of disorder, and may include a history of trauma, loss of consciousness, pain, numbness, weakness, or an audible noise emanating from a vibration.

Often vibrations are discovered upon auscultation* during a physical exam: the clinician listens over certain regions of the body for sounds that are not typical "heart sounds" (see Figure 2.1). If present, the sound would indicate a likely abnormality, the clinical significance of which depends on its frequency (pitch), quality, intensity (loudness), duration, radiation (propagation distance), timing in the heart cycle, spatial location, patient symptoms, and patient history. These sounds are referred to as murmurs or bruits.

While the term murmur often refers to sounds heard in the chest, and bruit to sounds heard elsewhere, the terminology is not uniform. The term murmur is also used in peripheral locations in place of bruit. In addition, the generic terms bruit and murmur often generally refer to a low quality harsh or wind-like sound; whereas, use of the terms "musical bruit" or "musical murmur" indicates a pure tone. Another vibration encountered upon physical examination is the arterial thrill. A thrill is any vibration at the skin surface that can be palpated* by the examiner. When thrills are discovered upon physical examination, their presence is noted, but their characteristics are not.

In general, the presence of a bruit warrants notice and further consideration, regardless of the location in the body. Determining that a bruit is innocent (i.e., not due to a disorder) is important, yet not clear-cut. For instance, although the majority of stenoses are due to atherosclerosis, stenoses can occur in normal arteries due to vessel kinks or compression from adjacent organs.

---

* See glossary.

To summarize, in spite of the fact that the terms bruit and murmur are common in the literature and have been around as long as the stethoscope (Laennec, 1819) they are problematic for clinical use for three reasons. 1) They are not well-defined. 2) They are often anecdotal. 3) Their clinical significance is poorly understood. They remain, however, an indicator of many significant vascular disorders, as outlined in the next sections and summarized in Figure 2.1.

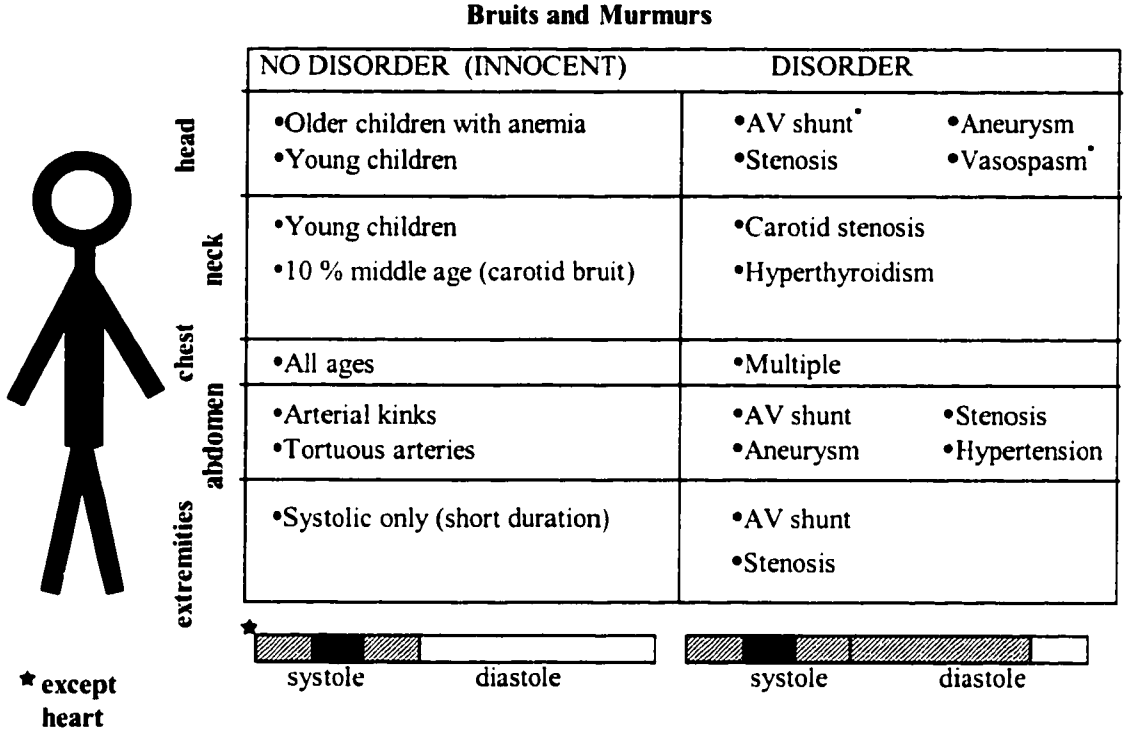

Figure 2.1. Bruit locations and associated arterial disorders.
Bruits, which are caused by vibrations, are heard throughout the body. Other than in the heart, vibrations arising innocently* occur only during systole; whereas, vibrations arising due to arterial disorders can extend into diastole.

2.1 The head

Cranial bruits are fairly common in healthy young children. The presence of a cranial bruit, however, in an adult or non-anemic older child can indicate aneurysm, stenosis,

arterio-venous shunts, vasospasm* or increased intracranial pressure (Ferguson 1972, Aaslid and Nornes 1984, Bates 1991, Hasegawa and Kobayashi 1995, Young et al. 1996), though its absence does not exclude them. Both general bruits and musical bruits arise in intracranial arterial disorders. Headache has been noted as a symptom of these disorders, though few signs arise before the situation is critical.

2.2 The neck

In younger people innocent bruits* are common. An innocent continuous hum may be found in the jugular vein of children and young adults. (Bates 1991). Additionally, about ten percent of middle age people will have an innocent carotid bruit (Bates 1991).

Bruits in the carotid arteries of older people, however, suggest partial arterial obstruction. Carotid bruits tend to occur during systole and may extend into diastole with any range of quality, frequency and intensity. Cardiac murmurs may also radiate to the carotids and mimic a carotid bruit. The frequency of these referred cardiac murmurs will decrease with distance from the heart.

A bruit at the thyroid indicates hyperthyroidism. This would be accompanied by such symptoms as nervousness, weight-loss with increased appetite, sweating and heat intolerance, frequent bowel movements, and muscular weakness or tremor. (Bates 1991, Greenberger and Hinthorn 1993).

2.3 The chest

There are multiple causes and types of cardiac murmurs. The majority are caused by either valvular stenosis or regurgitation. The murmurs can be further differentiated into other binary classifications: 1) right sternal* or left sternal, 2) systolic* or diastolic*, 3) short duration (i.e. ejection) or long duration (pansystolic*) (see Table 2.1). Murmurs can also be caused by patent ductus arteriosis* or by defects in the cardiac intrachamber septa*. Cardiac murmurs may occur in systole, diastole, or both with any range of quality, frequency and intensity. Murmur intensity (or loudness) at

---

* see glossary

the disorder site is proportional to the pressure difference between the involved chambers. Cardiac murmurs do arise in normal patients, but more commonly in abnormal patients; this distinction can be difficult.

Innocent bruits occurring in the chest include a "mammary souffle", which is heard over the base of the heart due to increased blood flow during pregnancy and lactation, and venous hum, which is heard over the base of the heart or above the clavicles. Both of these bruits are continuous.

Table 2.1. Heart murmur locations, associated disorders, and timing.

Heart Side	Ventricle Function	Valve		Murmur		
		Name	Defect	Location	Timing	Duration
Right	Inflow	Tricuspid*	Stenosis	Left	Diastolic	Mid Diastolic
Right	Inflow	Tricuspid	Regurg-itation	Left	Systolic	Pansystolic*
Right	Outflow	Pulmonic*	Stenosis	Left	Systolic	Ejection
Right	Outflow	Pulmonic	Regurg-itation	Left	Diastolic	Early Diastolic
Left	Inflow	Mitral*	Stenosis	Apex	Diastolic	Mid Diastolic
Left	Inflow	Mitral	Regurg-itation	Apex	Systolic	Pansystolic
Left	Outflow	Aortic*	Stenosis	Right	Systolic	Ejection
Left	Outflow	Aortic	Regurg-itation	Right	Diastolic	Early Diastolic

2.4 The abdomen

Abdominal bruits may indicate arterio-venous shunts, stenoses or aneurysms. Specifically, bruits that occur in both systole and diastole strongly indicate stenoses. Arterio-venous shunts would be expected only as a complication of biopsy (Bates 1991, Okamoto et al. 1992, Plainfosse et al. 1992, Greenberger and Hinthorn 1993, Strassburg et al. et al. 1996, Young et al. 1996). A dissecting aneurysm may be accompanied by a sudden appearance of an aortic bruit (Young et al. 1996).

---

* See glossary

2.5 Extremities

Innocent bruits occurring only during systole are relatively common (Bates 1991, Young et al. 1996), but a bruit that extends from systole into diastole is highly indicative of a disorder, likely stenosis or arterio-venous fistula*.

2.6 Trauma

Although injuries of large arteries are usually identified readily, injuries to the smaller blood vessels may be more difficult to ascertain, especially when the resultant bleeding is into an expandable closed space such as the peritoneum*. Currently, diagnostic peritoneal lavage (DPL)*, computed tomography (CT), magnetic resonance imaging (MRI), angiography, and surgery are used to find evidence of bleeding and to locate the bleeding site (Perry 1996, Young et al. 1996).

Several clinicians, though, have documented the use of auscultation and palpation to discover various vascular disorders induced by trauma (Meagher et al. 1979, Reichard et al. 1994). Khoury et al. (1994) and Ilijevski et al. (1998) report murmurs and thrills specifically associated with traumatic arterio-venous shunts. Islam et al. (1996) found that all 38 patients whom they studied with ruptured aneurysm (of sinus of Valsalva) had a continuous "machinery" murmur (i.e., rumbling) and 34 had an associated thrill. In addition, gunshot wounds have generated murmurs in two reported cases by Myers and Morgan (1998), and Kaklikkaya et al. (1999).

---

* See glossary.

3 Arterial vibration in the literature and clinic

Since 1970, researchers have studied vibrations associated with arterial disorders to develop understanding of their source mechanism and to hopefully improve their diagnosis and management. Although they have met with limited success, one result of their studies has been the development of a knowledge base of the general characteristics of the vibrations, a few of which are quantitative.

3.1 Vibrations from the various disorders

3.1.1 Stenoses and vasospasm

With invasive techniques, researchers have shown that the sound emanating from an arterial stenosis corresponds to the vibration of the arterial wall (Foreman and Hutchison 1970, Clark 1977, Miller et al. 1980), but its presence or absence from a particular stenosis is not yet completely understood. It is known, though, that bruits do not occur from arteries that either are barely narrowed or are nearly occluded. For the stenoses that lie between these two extremes, information about the frequency, amplitude and locations of the vibrations may provide additional valuable insight into the clinical significance of the stenosis (Lees and Dewey 1970). For instance, Roach (1963) and Vito et al. (1975) have demonstrated that vibrations can lead to arterial dilatation and fatigue in *in vitro* human iliac arteries and *in vivo* dog carotid arteries. Perhaps there is a link between the vibration parameters and the degree of dilatation or fatigue.

In the 1970's and 1980's, microphones were used to study arterial wall vibrations from stenoses. *In vivo* studies were carried out non-invasively in humans, and invasively in animals. Additional studies used excised arteries. (Foreman and Hutchison 1970; Lees and Dewey 1970; Kirkeeide et al. 1977; Roach and Stockley 1980; Knox, Breslau, and Strandness 1981). The composite results indicate somewhat narrowband* arterial vibrations in the range of 100-750 Hz with quality factors*, Q_v, between 1.4 and 15.

---

* See glossary

Foreman and Hutchison (1970) also found multiple resonant frequencies, but not at exact harmonic frequencies.

Instead of the vibration characteristics, some researchers have focused on the mechanism of vibration. Collectively, their work indicates that broadband blood turbulence generates vessel wall resonance (i.e., vibrations). Theories began to be published in 1970 when Lees and Dewey theorized that turbulence in the blood flow distal to a stenosis produces pressure fluctuations at the vessel wall. They expected a broadband signal to radiate from the artery. However, Foreman and Hutchison (1970) studied isolated arteries and found that the vibration frequencies distal to stenoses corresponded to the resonant frequencies of the arterial walls. These researchers also found narrowband vibrations in isolated arteries distal to artificial stenoses when exposed to wideband stimulus, which is consistent with the resonance theory. Sabbah and Stein (1976) found similar results in their invasive *in vivo* studies of cardiac ejection murmurs, a disorder similar to the arterial stenosis. The theory that the vibration frequency is a resonant frequency of the arterial wall is further corroborated by Chapman and Charlesworth (1981). These researchers used a tubing model with various degrees of stenosis (i.e., degrees of narrowing) and found that vibration frequencies were fairly constant and independent of the size of stenosis and Reynold's number. Hence, many researchers, working independently and with varying aims, have provided evidence that bruits heard with stenosis are likely the result of vessel wall resonance induced by blood turbulence.

Another arterial disorder closely related to stenosis is vasospasm in which the arterial muscle has a sustained contraction that narrows the artery. Olinger and Wasserman (1977), as well as Aaslid and Nornes (1984), have recorded both broadband* and narrowband bruits from intra-cranial vasospasm which fall in the same frequency range as bruits associated with stenosis. Aaslid and Nornes hypothesized that vasospasm caused increased blood velocities which, in turn, generated the vibrations

---

* See glossary.

in a manner similar to stenosis. They found that the vibration frequency increased with increased vasospasm (i.e., decreased vessel diameter).

3.1.2 Bleeds

I have found weak, localized vibrations in both exposed and *in situ* porcine arterial bleeds (see sections 4.4.1.2, 4.5, and 7.4.6). The vibrations are likely generated by turbulence in the blood as it exits the high-pressure vessel through a puncture or tear in the vessel wall. In a manner similar to vibrations from stenoses, the blood invokes resonance of the arterial walls as it flows past an orifice with a significant pressure differential. One paper in the literature discusses murmurs found in patients with ruptured aneurysm (Islam et al. 1996).

3.1.3 Aneurysms

Several researchers have studied bruits arising from intra-cranial aneurysms (Ferguson 1972; Olinger and Wasserman, 1977; Sekhar and Heros, 1981). They reported narrowband bruits with frequencies in the range between 140 and 840 Hz with high quality factors. Detection and location of these aneurysms with an electronic stethoscope has alerted caregivers to life-saving intervention, as ruptured cranial aneurysms can be lethal. Although the cause of aneurysmal rupture is not fully understood, it is thought that vibrations may weaken the arterial walls (Ferguson 1972, Sekhar and Heros 1981).

In addition to the cranium, vibration due to aneurysm has also been noted in the abdomen (Okamoto et al. 1992).

As with stenoses, the cause of aneurysmal vibrations is still unknown. Simkins and Stehbens (1973) studied various aneurysm shapes in sheep and tubing models with mixed results in terms of reproducing the vibrations. In their review, Mast and Pierce (1995) conclude that the prevailing thought seems to be that a fluid resonator is responsible and that possibly only certain aneurysmal configurations will facilitate this phenomenon. Aaslid and Nornes (1984) disagree, however. Their studies of patients following subanachroid hemorrhage revealed that vibrations were not co-located with the aneurysm, but rather appeared at the site of the associated vasospasm.

It should be mentioned that a cousin to the aneurysm, the pseudo-aneurysm (or false aneurysm), is also known to generate bruits (Agarwal et al. 1993). A pseudo-aneurysm differs from a true aneurysm in that its wall does not contain the components of an artery. The wall of a pseudo-aneurysm consists only of fibrous tissue and surrounding structures. Pseudo-aneurysms are readily detected with conventional ultrasound imaging, though vibrometry would potentially aid their detection as well.

3.1.4 Arterio-venous shunts

Direct communication between an artery and a vein causes turbulence in the blood as it leaves a high pressure artery for the lower pressure vein. This turbulence can cause vibrations in the surrounding tissue. These well-known vibrations are typically discovered as thrills upon palpation, but occasionally upon color Doppler ultrasound examination (3.3.3.2). They have been specifically documented by numerous clinicians and researchers (Shu and Hwang 1988, Middleton et al. 1989a and b, Fillinger et al. 1990, Plainfosse et al. 1992, Khoury et al. 1994, L'opez et al. 1995, and Ilijevski et al. 1998). To quantitate the vibration frequency, Stehbens et al. (1995) formed experimental arterio-venous shunts in rabbits. Measuring the vessel wall vibrations with laser vibrometry, they discovered vibrations up to 2200 Hz.

3.2 Vibration characterization

The various vibration characteristics have been shown to have clinical relevance or are thought to be potentially clinically relevant. The degree of significance of a disorder varies according to the disorder type as listed in Table 3.1.

Table 3.1. Key factors in degree of significance according to disorder type.

Disorder Type	Item indicating degree of significance
Aneurysm	Artery diameter
Stenosis and vasospasm	Degree of narrowing
Bleed	Volume flow rate of artery leak
Arterio-venous shunt	Volume flow rate out of artery

- **Anatomic location** is clinically related to the disorder type. The spatial area over which the vibration is present is occasionally related to the disorder's significance.

- **Duration and Timing in heart cycle** are clinically recorded jointly as systolic, diastolic, pansystolic or continuous, and are clinically related to the type of disorder. They may be clinically related to the significance of the disorder for stenoses and vasospasm.
- **Frequency** is often clinically related to the type of disorder. It has been related to the disorder's significance for stenoses and vasospasm.
- **Amplitude and Intensity** are clinically recorded in terms of "loudness", and are possibly related to the disorder's significance.
- **Bandwidth and quality** are clinically recorded in terms of the degree to which the sound is "harsh" or "musical", and are occasionally clinically related to the type of disorder.
- **Phase propagation** (the relative phase relationship between vibrations at successive depths) is potentially related to distinguishing vibration from blood flow during vibrometry. Phase propagation is closely related to the time delay measure discussed in section 5.3.2.3.

3.2.1 Quantitative

In vivo arterial vibration characteristics quantified in the literature are limited to frequency, quality factor and duration, with ranges listed in Table 3.2. Since quantitative information on palpable thrills is unavailable, the values in Table 3.2 do not pertain to them.

Table 3.2. Vibration characteristics from the literature.

Vibration characteristic	Minimum value	Maximum value
Frequency	100	>1000[#]
Quality factor (Q_v)	1 (broadband)	15 (narrowband)
Duration (ms)	10	>100

[#] Vibrations up to 2200 Hz have been found in a gortex graft in a rabbit (Stehbens et al. 1995).

3.2.2 Qualitative

3.2.2.1 Location

Vibrations occur at arterial vessel walls and propagate radially into surrounding solid tissue away from the vessel. Vibrations may result from, or induce, oscillations in the vessel's blood flow. Propagating vibrations or induced oscillations are expected to be spatially coherent over some region. The vibration propagation may be extremely limited to within a millimeter of the vibration source, or it may extend for many centimeters through a homogeneous organ.

Other than in the cardiac arteries, arterial vibrations will occur primarily during systole, when the blood velocities and local power are at a maximum. Strong vibrations, however, often extend into diastole. Although heart beats are not completely uniform, a disorder that causes a vibration in one heart beat is expected to generate a vibration in each consecutive heart beat at roughly the same point in the cardiac cycle.

3.2.2.2 Bi-directional

Vibrations can be distinguished from other signals in a Doppler ultrasound examination (see section 3.3.3.1) in that they are motions that occur both toward and away from the transducer at a frequency greater than typical wall motion (see Figure 3.1, the signal is both above and below the baseline). Typical wall motion (shown in Figure 4.16) has a fundamental frequency of 60 to 90 beats per minute (1- 1.5 Hz) with up to approximately 30 harmonics.

3.2.2.3 Bandwidth

There are primarily two types of bruits: narrowband and broadband. A narrowband bruit may be called "musical" or "seagull" based on its audible sound. The sound heard from a broadband vibration is similar to that heard from blowing on a microphone or hissing.

Vibrations can also consist of multiple narrowband frequencies occurring at a single spatial location simultaneously.

3.2.2.4 Signal to noise ratio

Qualitatively, *in vivo* vibration signal to noise ratios vary widely from imperceptible to obvious.

3.2.2.5 Variability

The vibration amplitude, frequency and time duration can change from heart beat to heart beat. The vibrations do not remain fixed in space (with respect to the transducer over a ~1 mm^3 sample volume), but rather change in spatial location. This variability is possibly due to physiological changes in the vibration mechanism, movement of the artery, or movement of the transducer.

3.3 Existing methods and their ability to accomplish vibrometry

Ultrasound, stethoscopes, and phonoangiography are the existing methods that can be used to examine arterial vibrations. Each method, however, is limited in its ability to offer efficient and thorough detection and characterization. The following table presents a comparison of their abilities and limitations in efficiently and sensitively quantifying arterial vibrations.

Table 3.3. Existing methods to measure arterial vibrations.

	Stethoscope	**Phono-angiography**	**Pulsed Doppler**	**Color Doppler**
Acquisition Time	***Fast***	***Fast***	Slow	***Fast***
Spatial location	General	General	***Specific***	***Specific***
SNR	Low	Low	***Higher***	***Higher***
Timing in cardiac cycle	Qualitative	***Quantitative***	***Quantitative***	Qualitative
Duration	Qualitative	***Quantitative***	***Quantitative***	Qualitative
Frequency	Qualitative	***Quantitative***	***Quantitative***	None
Amplitude/ Intensity	Qualitative	Qualitative	None (unless >10μm)	None (unknown frequency)
Phase	None	None	None	None

3.3.1 Stethoscope

The stethoscope allows the clinician to examine a large region of tissue very quickly. In that respect, it is very efficient. However, it does not offer quantitative measures of the vibrational amplitude, phase, frequency, intensity, or specific anatomic location. Use of the stethoscope to detect a bruit is also very dependent on the ability of the examiner to hear the bruit and differentiate it from other sounds. In addition, due to the stethoscope's large volume of interrogation, the vibration signal strength may be very low with respect to other sounds resulting in a low signal to noise ratio.

3.3.2 Phonoangiograpy: qualitative and quantitative

Phonoangiography, called phonocardiography when applied to the heart, involves placing a microphone at the skin surface to record the acoustic waveform of a bruit. In qualitative phonoangiography, the digitized signal is displayed in time for viewing the vibration's duration, its timing in the cardiac cycle and its amplitude at the skin

surface. In quantitative phonoangiography, a Fourier transform is applied to the recorded pressure time series to obtain the bruit's frequency content (Lees and Dewey 1970; Pansini et al. 1980; Knox, Breslau and Strandness 1981; Semmlow et al. 1983, Pennestri et al. 1989). As with the stethoscope, phonoangiography is efficient because the microphone captures sounds from a wide volume, but it does not provide measurement of the amplitude, phase, frequency or intensity of the vibration at the site of the disorder. Furthermore, it fails to even locate the disorder site. Finally, like the stethoscope, phonoangiography interrogates a large spatial volume all at once. Therefore, a vibration's signal strength may be very low with respect to other sounds.

3.3.3 Ultrasound

[Readers not familiar with medical ultrasound may wish to read section 4.1, before continuing.]

3.3.3.1 Pulsed Doppler ultrasound

The single gate pulsed Doppler ultrasound mode provides one dimensional velocity measurements (scaled from frequency according to the Doppler equation) within a site called the sample volume (~ mm^3). The "velocities" are plotted versus time to form a spectral "velocity" waveform display. Although pulsed Doppler was not designed for measuring vibrations, when they are present, vibrations appear as a two-sided waveform indicating simultaneous motion toward and away from the transducer (Pennestri 1989). In current practice, these waveforms are considered "noise" by the examiners and by the ultrasound manufacturers. Such signals are avoided by the examiners, so manufacturers attempt to reject them.

Some researchers, however, have studied pulsed Doppler processing of vibrations (Holen, Waag and Gramiak 1985; Nowicki and Marciniak 1989; Wang, Bone and Hossack 1992; Heimdal and Torp 1997). They have shown that harmonics arise when the Doppler processing utilizes quadrature phase (complex) demodulation (described in section 4.1.3). With this demodulation technique, the voltage time series of the two resulting channels, in-phase (I) and quadrature-phase (Q), will each fit a Bessel function (Couch 1990). The Bessel function has a fundamental frequency at the

frequency of the vibration and additional harmonics. In the absence of a wall (clutter) filter, the magnitudes of the harmonics relate to the amplitude of the vibration. At low amplitudes, the magnitudes of the harmonics are negligible, and only the fundamental is seen. As the amplitude of the oscillations increases above ~$\lambda/30$, however, the first Bessel harmonic rises to within 20 dB of the fundamental, and begins to show up in the Doppler spectral waveform. Larger amplitudes will display multiple harmonics (e.g., Figure 3.1). For 5 MHz ultrasound, λ is 300 μm, and the harmonic's 20 dB threshold is ~ 10 μm. Vibrations larger than 10 μm, therefore, can be measured roughly by the number and/or amplitude ratio of Bessel harmonics seen on a Doppler spectral waveform (Shinozuka and Yamakoshi, 1993). Arterial wall vibrations, however, are often not large enough to generate the harmonics. When a wall filter is used, the vibration amplitude is not preserved, and the presence of harmonics cannot be related to the vibration amplitude.

The pulsed Doppler method quantifies the same vibration characteristics as phonoangiography (see Table 3.3), while also specifying the vibration's anatomic location. In addition, because pulsed Doppler can interrogate a very small volume of tissue, other sounds will minimally impact the vibration's signal. Exploring a region of tissue with single gate Doppler ultrasound, however, would be very laborious and would still fail to provide amplitude, intensity and phase measurements.

Figure 3.1. Strong, narrowband bruit appearing in pulsed Doppler ultrasound at a human carotid artery.
Bruits, which are caused by vibrations, have a unique signature on a Doppler spectral waveform: they appear symmetrically about the 0 Hz (or 0 velocity) axis. This example includes harmonics of the fundamental vibration frequency. The harmonics appear as the multiple lines seen on both sides of the center line.

3.3.3.2 Color Doppler ultrasound

The color Doppler ultrasound mode (also called Color Flow Imaging [CFI] or Color Flow Doppler (CFD)) is designed to colorize a region (~ 10 cm^2) of an ultrasound image (called the color box) with the spatial distribution of blood velocities where blood flow is present. Typically, flow in one direction is colored red and the opposite direction blue (see Figure 4.11). Blood velocity oscillations generate a checkerboard appearance in Color Doppler mode (Middleton, 1989a and b; Beach, 1993) as shown in Figure 3.2. [Overgaining, aliasing and blood turbulence will also render a chaotic color Doppler display. Overgaining should result in a wide region of non-pulsatile color. Aliasing can be eliminated with the adjustment of the pulse repetition frequency (PRF). Blood turbulence will arise only within a vessel.]

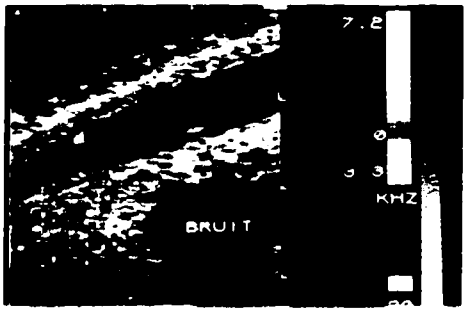

Figure 3.2. A bruit appearing in color Doppler ultrasound at a carotid artery. Vibrations cause "color bruits" which appear as a mosaic of the colors associated with motion both toward and away from the transducer.

The checkerboard appearance arising from vibrations is likely the result of two aspects of the color Doppler mode. 1) Consider that the there is a fixed time window over which the color Doppler method interrogates a given sample volume: 0.5 to 10 ms. The color Doppler method expects the velocity to be constant over this temporal window for its measurement, but vibrations often will not have a constant velocity over this time window. When the velocity is not constant, the color Doppler method will simply choose a best fit to the distribution of the velocities. In frame 1 of Figure 3.3, for instance, the eight samples at the high PRF will yield a velocity associated with the slope of the vibration during the short time window. The mid PRF and low PRF settings, though, would compute values closer to zero since the average slope of their points is smaller. 2) The second contributor to the checkerboard pattern is that the time at which the sample volume is revisited (the frame rate) is not synchronous to the vibration. Therefore, at the high PRF, the velocity estimate at frame 1 is very different from frame 2: frame 1's velocity is positive and frame 2's velocity is approximately zero. The same is true at the mid PRF: frame 1's velocity is positive and frame 2's is negative. Since the estimated velocity is different for each frame, and for similar reasons each spatial beam line, the display takes on the checkerboard appearance. Thus, the value determined by color Doppler processing will greatly depend on the frequency of the vibration and the PRF settings.

There are three additional factors affecting the display of vibrations with color Doppler ultrasound. 1) This mode considers solid tissue movement to be "noise" and attempts to avoid colorizing these regions by setting a threshold below the expected backscatter signal level of solid tissue. 2) Since the vibrations typically are very short in duration (10-100 ms) and recur just once per heart beat (e.g. once per second), and since the color Doppler mode interrogates each ultrasound beam line for less than 10 ms (10 pulse echo cycles at 1000 Hz PRF) recurring at 6-30 Hz, the vibration is easily missed by the interrogation. 3) The color Doppler mode is designed to measure blood velocities greater than approximately 1 cm/s. In fact, it will not display velocities below some adjustable fraction of the PRF, known as the wall filter setting. When wall filters are in place, this limitation is dependent on the nature of the wall filtering. In the absence of wall filtering, for a sinusoidal vibration at 1000 Hz, a peak velocity of 1 cm/s is achieved with a vibration amplitude of 1.7 μm corresponding to an acoustic intensity of 8.6 mW/cm^2 (see equation (6-20), with $Z = 1.5\ 10^6\ kg/(m^2s)$). If the vibration frequency (f_v) were known, then the vibration's minimum amplitude (A_v) could be determined according to:

$$A_v > v_{peak} / 2\pi f_v. \qquad (3\text{-}1)$$

By adjusting the velocity threshold, the vibration's amplitude could be deduced. The frequency of arterial vibrations, however, is not known *a priori*, and wall filters are typically in place, so this amplitude measurement method does not apply.
In summary, Color Doppler, offers possibly only one of the arterial vibration parameters of interest, the anatomical location of the vibration, and may provide vague, qualitative information about the duration and timing in the heart cycle with a time resolution equal to the frame rate (e.g. 10 frames/second).

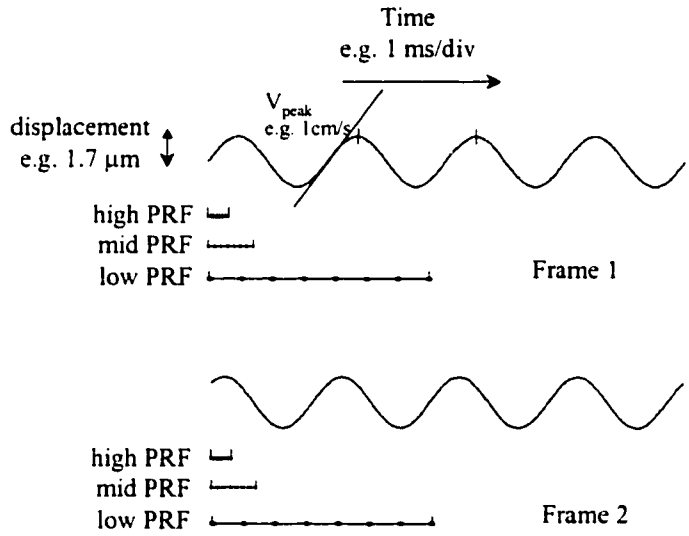

Figure 3.3. Color Doppler sampling of a vibration.
The samples obtained in one ensemble of the color Doppler acquisition vary greatly depending on the relationship between the PRF and the vibration frequency. A relatively high PRF will sample a region of the vibration that is approximately a constant velocity. A relatively low PRF will sample several cycles of the vibration.

3.3.3.3 Tissue Doppler ultrasound

Tissue Doppler is a recent outgrowth of the color Doppler technique, designed primarily for echocardiography (Sutherland et al. 1994, Miyatake et al. 1995, Schmidt-Trucksass et al. 1998). In tissue Doppler, the thresholding of both backscatter signal level and velocity values are eliminated. Velocities are calculated directly from the low frequency tissue movement.

Despite their emphasis on displaying tissue motion, the modifications to color Doppler processing will prevent the display of small vibrations. The primary reason is that vibrations typically arise in conjunction with large low frequency wall motion. Tissue Doppler would measure the low frequency wall motion (i.e., clutter) and be completely insensitive to the much smaller vibrations (see Table 4.3, in section 4.5.1).

For the rare case that the vibration amplitude approaches the clutter amplitude, tissue Doppler methods would suffer from the same factors as color Doppler. Again, the velocity measurement would vary with the number of vibration periods encompassed in the temporal window, if indeed the vibration site were interrogated during the short time that the vibration was present.

There is also a variant of tissue Doppler that is used in conjunction with M-mode ultrasound, in which only one spatial line is interrogated. This mode would suffer from the clutter as described above, but, because it does not have a limited time window, would not suffer the other timing issues.

3.3.4 Novel method by Huang, Lerner and Parker (1992)

In 1992, Huang, Lerner and Parker suggested several time domain techniques to calculate the vibration amplitude from as little as two pulse echo cycles (using samples of the quadrature signals described in section 4.1.3) when the sinusoidal frequency is known. These methods, however, were designed for purely sinusoidal (very narrowband) signals with no clutter. For the case of unknown frequency, the researchers go on to suggest a means of determining the frequency provided that the signal is extremely narrowband. This frequency measurement is also accomplished with a small number of samples. Because it does not allow for clutter, or broadband vibrations, though, this method is not suitable for arterial vibrometry.

3.4 Other ultrasound processing techniques with potential for measuring arterial vibrations

Ultrasound has been mentioned as the method of choice for capturing the tissue displacements associated with arterial vibrations, yet each of the ultrasound methods discussed previously failed in this regard. Several ultrasound methods do exist that have been used to measure tissue displacements, which is the first step in measuring vibrations. These displacement measurement methods are listed below in approximately chronological order of their implementation. One of these, quadrature phase (complex) demodulation, is well suited for vibrometry and is the chosen method (section 4.2).

3.4.1 M-mode

M-Mode ultrasound is intended for imaging large solid tissue displacements (see section 4.1.4.3). It interrogates a fixed spatial line in tissue and displays displacement versus time. Because it is based on amplitude demodulation of the ultrasonic echo, M-mode offers only course resolution (~ 10 μm) limited first by the wavelength of the ultrasound, and subsequently by either the digitization rate of the received echo or the display pixelation.

3.4.2 Phase techniques

In 1968, Baker and Simmons (1968) developed an analog phase-locked echo technique to measure low frequency (< 50 Hz) arterial wall displacements. His innovation recognized that the phase trajectory, $\phi(t)$, of the radio frequency (RF) ultrasonic echo corresponding to a specific depth in tissue would be proportional to the displacement, x(t), of that tissue according to:

$$x(t) = \frac{\phi(t)}{2\pi} \lambda/2 \qquad (3\text{--}2)$$

where λ is the wavelength of the ultrasound. This technique was furthered by Hokanson et al. in 1970, McLeod et al. in 1977, and Reneman in 1981.

Later, in 1982, Wilson and Robinson computed the phase digitally from an analytic* signal of a broadband, M-mode pulse. Their results had an error of +/- 38 μm due to asynchronous clocks.

Using related methods, Hartley et al. (1978), Kansky (1988) and Phillips, Hoeks et al. (1990), Shinozuka and Yamakoshi (1993), and Kanai et al. (1993) individually developed similar systems to track arterial wall and heart movements. [Kanai's group recently applied this technique to arterial wall vibrations (Sunagawa et al. 1999).] In these systems, $\phi(t)$ is determined from quadrature phase (complex) demodulation of RF ultrasound echoes received from narrowband pulses as described in section 4.1.3.

---

* see glossary

Specifically, the later systems use the quadrature components (I for in-phase and Q for quadrature-phase) to compute $\phi(t)$ according to:

$$\phi = \text{atan}(Q / I). \qquad (3\text{-}3)$$

These authors report their result in microns, which, in most cases, was limited by their digitization resolution. Displacements on the order of nanometers can be measured with this technique if 12 bit resolution is used, allowing the phase measurement to be limited only by signal-to-noise ratio (see section 4.6.1).

3.4.3 cross-correlation

Perhaps a more straightforward means of measuring tissue displacement involves cross-correlating the echoes received from one ultrasound pulse with the echoes received from a subsequent pulse. This technique, described in detail in a review by Hein and O'Brien (1993), was introduced by Dickinson and Hill in 1982, furthered by Bonnefous and Pesque in 1986, and continues to be a subject of active research (Brands et al. 1997). This method is extremely computationally expensive, and its resolution is limited primarily by computational power. Cross-correlation techniques are rapidly gaining acceptance as computational power becomes more available.

4 The Vibration Data

This chapter will introduce the vibration data, including: relevant ultrasound concepts and techniques, the instruments used to collect data in this project, the data components, vibration models, a presentation of some vibration examples with a summary of the experimental vibration characteristics, a discussion of the displacement resolution, and a validation of the accuracy of the data collection method.

4.1 Relevant ultrasound concepts and techniques

Medical ultrasound instruments are routinely used for vascular diagnosis. The information is obtained in the following manner. First ultrasound waves are generated in a piezoelectric transducer and transmitted into the body. Subsequently, the transducer receives the ultrasound waves that are scattered back to the transducer from subcutaneous tissue. The information of interest is then extracted from the returned sound waves and presented back to the operator. The instruments can thus be considered to have four primary steps: ultrasound transmission, ultrasound reception, signal processing, and display. Each step operates in various modes in order to accomplish both anatomical imaging and velocity measurements. We will consider the first three steps separately. The output displays will be shown where appropriate.

4.1.1 Pulsed ultrasound transmission

Ultrasound is transmitted in pulses along lines that are determined by the beam pattern of the ultrasound transducer. The duration of the transmission pulse is inversely proportional to the depth resolution. Therefore, relatively short pulses are used (~0.1 to 10 μs) to provide acceptable spatial information (resolution of ~ 0.15 to 15 mm).

4.1.1.1 Pulse types and duration

Two basic types of pulses are commonly implemented in ultrasound instruments: broadband and narrowband. Broadband pulses can be very short (~ 0.1 μs) and yield superior depth resolution. These are primarily used for anatomic imaging. The longer, narrowband pulses are used to obtain velocity and displacement information.

The difference is due to the analysis techniques employed for each of the modes, and will become clearer in section 4.1.4. Figure 4.1 illustrates the two pulse types. The length of each is affected by the wavelength of the selected ultrasound frequency.

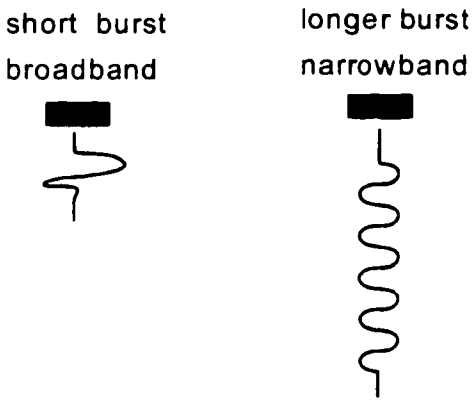

Figure 4.1. Ultrasound pulse types.
The two pulse types used in conventional diagnostic ultrasound are: broadband and narrowband. The broadband pulse is typically very short. The narrowband pulse usually contains at least several cycles of a sine wave at the chosen transmission frequency.

4.1.1.2 RF frequencies

Three factors affect the selection of the ultrasound transmission frequency (F). 1) It must be low enough to minimize attenuation from the depths of interest. 2) It must be selected to induce scatter from the tissue of interest. 3) The highest acceptable frequency should be used to minimize the wavelength for the best possible depth and phase resolution. Since solid tissue attenuates ultrasound at a rate of ~ 1 dB/MHz/cm (which is ~ 2 dB/MHz/cm including round trip travel), the choice of transmission frequency is constrained at the upper limit by the depth at which tissues are to be interrogated. For example, with 5 MHz ultrasound, echoes from tissues that are 5 cm deep would be attenuated by a factor of approximately 50 dB (2 dB/MHz/cm * 5 cm * 5 MHz) compared to those from the skin's surface. [Attenuation up to possibly ~ 90 dB can be mitigated with sufficient digitization and the introduction of a post-processing gain. See sections 4.2.1 and 4.3.1.1.2.]

In addition to attenuation concerns, the transmission frequency must be selected so that the tissue of interest has sufficient scatter. Since the transmission frequency choice is bounded from above by attenuation (typically a maximum of 10 MHz is used), the ultrasound wavelength (> 150 μm) is rather large with respect to the size of connective tissue, cells and fibers (~ 10 μm, Jensen 1996). Thus, Rayleigh scattering dominates in blood and soft tissues. The strength of the scatter varies with tissue type and increases with transmission frequency (Jensen 1996). Scatter returned from solid tissue have power levels larger than that from blood by a factor of 10 to 500 (Jensen 1996). Therefore, at a given depth, higher transmission frequencies that produce sufficient backscatter from solid tissue may not produce sufficient backscatter from blood.

As a tradeoff between depth resolution, small particle attenuation, and backscatter, diagnostic ultrasound frequencies typically lie between 1 and 10 MHz (radio frequency). Because of attenuation, the higher frequencies are used for superficial tissues, and the lower frequencies for deeper tissues and blood velocities, at a loss of spatial resolution. Most tissues of interest lie within a depth of 15 cm.

4.1.1.3 PRF values

In order to determine the depth from which an echo returns, sequential pulses must not overlap. That is, all echoes that have not been completely attenuated from a given pulse must have returned to the transducer before a subsequent pulse is transmitted. Hence the rate at which pulses are transmitted, called the pulse repetition frequency (PRF), is limited by the speed of sound in tissue, c, and the depth, X, at which the ultrasound is sufficiently attenuated as to be negligible compared to the echoes of interest. After accounting for round trip travel of the ultrasound, the maximum PRF is:

$$PRF = \frac{c}{2X}. \qquad (4\text{-}1)$$

4.1.1.4 One and two dimensional pulsed ultrasound

If echoes are desired along a single spatial line, then pulses are transmitted at a chosen PRF with the ultrasound beam at a fixed location. This mode is used for displacement and velocity measurement. To interrogate an area rather than a single line, the following sequence of events occurs: 1) some integer number (K) of pulses are transmitted and received along a given spatial line, 2) the ultrasound beam is moved electronically to a new spatial line, where K pulses are transmitted and received, and 3) the ultrasound beam is moved yet again until the entire region of interest is interrogated. For velocity measures K is greater than one, and for imaging, K is set to one.

4.1.2 Pulsed Ultrasound reception

Conventionally the same transducer aperture is used for both transmission and reception. For this configuration, the returned echoes are referred to as backscatter. The ultrasound echoes travel through the same beam pattern as the transmitted pulse, but the scatter from the tissue alters the original shape. For instance, the perpendicular boundary between two anatomic structures with differing acoustic impedances, Z, whose cross sections are larger than the ultrasound wavelength, will reflect the ultrasound pulse. This type of boundary is called a specular reflector. The returned echo is a replica of the transmitted pulse, but with a lower amplitude. The strength and sign of the reflection is determined by the reflection coefficient (R) for the interface between the two tissue types:

$$R = \frac{Z_2 - Z_1}{Z_2 + Z_1}. \qquad (4\text{-}2)$$

Since acoustic impedance is related to the density and speed of sound of the tissues, the ultrasound reflection is dependent on the density and speed of sound differences between them.

Impedance mismatches arising from clusters of particles which are smaller than the ultrasound wavelength, on the other hand, will exhibit Rayleigh scattering rather than

reflection. The returned ultrasound from Rayleigh scatterers will have considerably smaller, non-uniform amplitude compared to reflectors, roughly up to 60 dB smaller. They will also have a longer duration compared to the transmit pulse. The received signal from both types of scatterers, though, can be consider to be an amplitude modulated (AM) version of the transmitted ultrasound (Figure 4.2 and Figure 4.3).

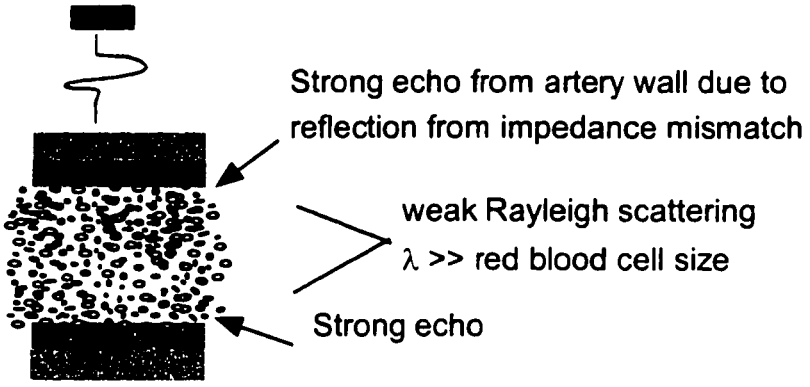

Figure 4.2. Illustration of backscatter from tissue.
Rayleigh scattering occurs in blood and within solid tissue. Reflection can occur at solid tissue boundaries.

In addition to amplitude modulation, tissue also introduces a phase shift on the returned echo. The ultrasound echoes, then, can be considered to be both amplitude and phase modulated, which mathematically takes the form:

$$S(\tau) = A(\tau)\sin(2\pi F\tau + \phi(\tau)), \qquad (4\text{-}3)$$

where S is a returned echo, A is the amplitude of the echo, F is the ultrasound transmission frequency, ϕ is the phase of the ultrasound echo, and τ is the time corresponding to the ultrasound propagation. The value of τ is proportional to the depth (X) of the tissue whose scatter produced the amplitude $A(\tau)$ and phase $\phi(\tau)$:

$$\tau = X/c. \qquad (4\text{-}4)$$

This expression dictates how a particular depth location along the beam pattern can be determined from the received signal based on the time elapsed between transmission and reception.

Equation (4-3) represents just one echo. For multiple transmit pulses, it should be recognized that the received echoes will be time varying:

$$S(\tau,kT) = A(\tau,kT)\sin(2\pi F\tau + \phi(\tau,kT)), \qquad (4\text{-}5)$$

where k is an integer representing discrete pulse-echo cycles and T in the time period between pulses.

For this discussion, the variable τ will be referred to as "fast" time with units of microseconds. The variable product kT will be used to represent "slow" time which corresponds to the phase or displacement changes at a given depth, across multiple pulse echo cycles. "Slow" time is on the order of milliseconds or seconds.

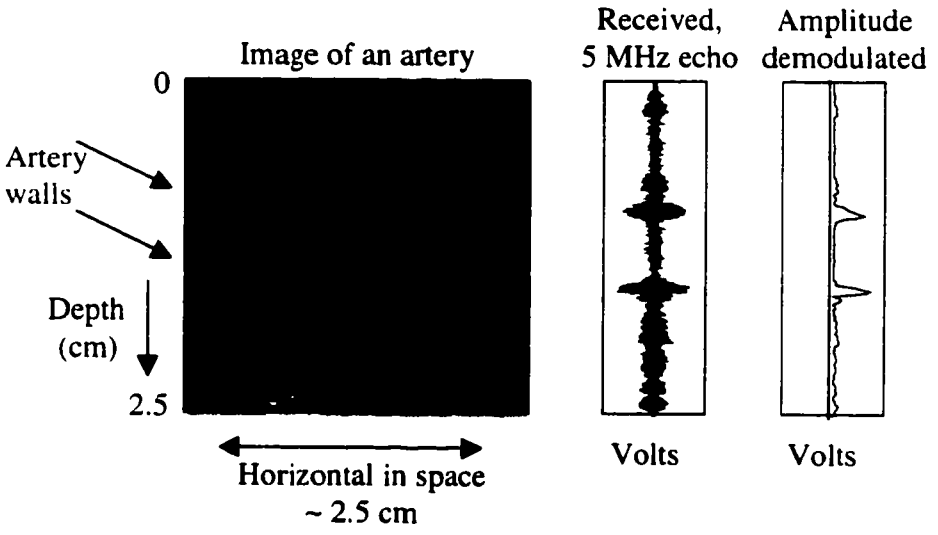

Figure 4.3. Ultrasound image and digitized echoes from an artery.
Left: B-mode image of a human carotid artery. Center: Received ultrasound echo from the center line of the B-mode image. Right: Demodulated amplitude of the echo from the center line of the B-mode image.

4.1.3 Quadrature (complex) demodulation

In engineering practice, an expression such as equation (4-3) is commonly used in complex notation for ease of computation:

$$S(\tau) = A(\tau)\, e^{j(2\pi F\tau + \phi(\tau))}, \qquad (4\text{-}6)$$

where $j = \sqrt{-1}$, and it is acknowledged that the true echo is the real part of $S(\tau)$. This notation is often simplified into a "phasor" form:

$$S(\tau) = A(\tau)\, e^{j\phi(\tau)}, \qquad (4\text{-}7)$$

where the term $e^{j2\pi F\tau}$ is noted as a fixed multiplier which can be reintroduced for computing the final result. $S(\tau)$ can be separated into its real, $I(\tau)$, and imaginary, $Q(\tau)$, rectangular components:

$$I(\tau) = A(\tau)\sin(\phi(\tau)) \qquad (4\text{-}8)$$
$$Q(\tau) = A(\tau)\cos(\phi(\tau)).$$

$A(\tau)$ and $\phi(\tau)$ can be determined from these according to:

$$A(\tau) = \sqrt{I^2(\tau) + Q^2(\tau)} \qquad (4\text{-}9)$$
$$\phi(\tau) = \text{atan}(Q(\tau)/I(\tau)).$$

Note from equation (4-8), that $I(\tau)$ and $Q(\tau)$ will be 90° out of phase, which is a quarter of one period of the ultrasound. $I(\tau)$ and $Q(\tau)$ are therefore referred to as the quadrature representation of the received echo.

Because of their dual role as both quadrature and complex representations, determination of $I(\tau)$ and $Q(\tau)$ is called both quadrature demodulation and complex

demodulation. This demodulation can be accomplished in either an analog or digital manner.

Figure 4.4 illustrates an example analog process. For transmission, a 5 MHz transmit sinusoid is time limited to a 1 μs duration, and transmitted into the body through the transducer. The received echoes are mixed with two versions of the original transmit sinusoid, in phase and 90° shifted, to create a pair of quadrature signals. Since an analog mixer effectively produces the sum and difference of the two input frequencies, the resulting signals will have significant energy around 0 Hz as well as around 10 MHz. After passing through a 5 MHz low pass filter, $I_{analog}(\tau)$ and $Q_{analog}(\tau)$ are the resulting in-phase and quadrature-phase baseband versions of the received signal, called the quadrature pair.

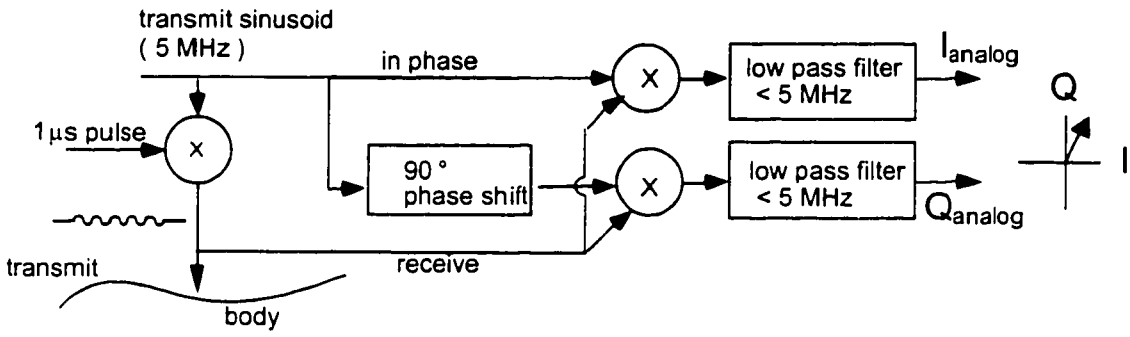

Figure 4.4. Quadrature demodulation.
The received echo is mixed directly with the transmit clock and also with the transmit clock shifted by 90°. The two resulting signals are low pass filtered to yield the in-phase and quadrature base-band representation of the original echo.

Many modern instruments obtain the quadrature pair through digital means, instead (Figure 4.5). When the received echo is sampled synchronously to the transmit clock at a rate of four times the transmit frequency, the quadrature pairs can be computed from the samples according to:

$$I = \tfrac{1}{2}(V_1 - V_3) \qquad (4\text{-}10)$$
$$Q = \tfrac{1}{2}(V_2 - V_4).$$

For noise filtering, multiple cycles can be averaged together or weighted in some other manner:

$$I = \tfrac{1}{6}(V_1 - V_3 + V_5 - V_7 + V_9 - V_{11}) \qquad (4\text{-}11)$$
$$Q = \tfrac{1}{6}(V_2 - V_4 + V_6 - V_8 + V_{10} - V_{12}).$$

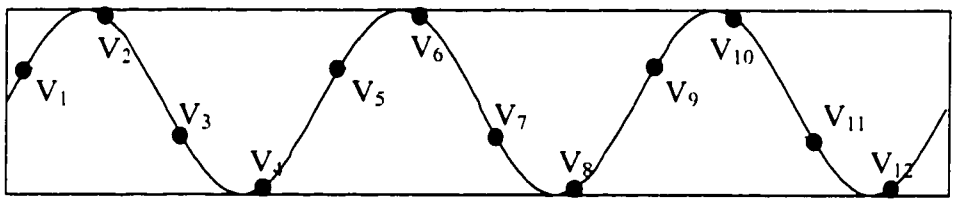

Figure 4.5. Quadrature demodulation by synchronous sampling.
A received echo is sampled at four times the transmit clock frequency. Alternating samples (such as samples 1, 3, 5, 7, etc.) are used to form one of the quadrature channels. The remaining samples are used to form the other of the quadrature channels.

4.1.4 Signal processing of pulsed ultrasound

Ultrasound instruments process received echoes to display anatomic images, velocity measurements and tissue displacements. Although vibration detection and estimation involve only displacement measurement from pulsed Doppler transmission, other ultrasound modes (B-mode, M-mode, and color Doppler) are included here for two reasons. 1) Vibrometry results are best displayed on a conventional ultrasound image for ease of integrating the new information with other familiar and also useful information. 2) It is helpful to understand the other modes in order to gain insight into their abilities and limitations in detecting and measuring vibrations (summarized in section 3.3). We will discuss displacement measurement first, followed by the conventional ultrasound modes.

4.1.4.1 Displacement measurement with pulsed ultrasound

When consecutive ultrasound pulses impact an object that is moving toward the transducer, the backscatter from the object will occur at progressively shallower depths and result in the signal reaching the transducer at progressively smaller time intervals. If the transmitted pulse is narrowband, then the delayed echo returning to the transducer will have a fixed frequency and some phase with respect to the transmitter clock. This phase will change proportionately with the object's displacement.

Figure 4.6 illustrates the process of measuring the phase changes corresponding to a moving object. At time t_0, a one microsecond transmit pulse leaves the transducer. At time X/c later it impacts an interface located at depth X. Only microseconds later, at time $t_0 + \tau_{return}$, the receiver compares one microsecond of the returned echo with one microsecond of the transmitter clock and determines that the received echo has some relative phase, ϕ_0. When the next pulse reaches the object, the object has moved Δx microns toward the transducer. At $t_1 + \tau_{return}$, when the receiver compares the second returned echo with the transmitted pulse, it determines the new phase relationship, ϕ_1. This change in phase between pulse 0 and pulse 1 is related to the object's displacement according to:

$$\Delta x = (\lambda/2)(\Delta\phi / 2\pi), \qquad (4\text{-}12)$$

where $\Delta\phi = \phi_1 - \phi_0$, λ is the wavelength of the ultrasound, the first factor of 2 accounts for round trip travel of the ultrasound pulse, and the factor of 2π converts from wavelength to phase. Subsequent displacements can be measured with respect to the initial location according to:

$$x(kT) = (\lambda/2)(\phi(kT) - \phi_0)/ 2\pi, \qquad (4\text{-}13)$$

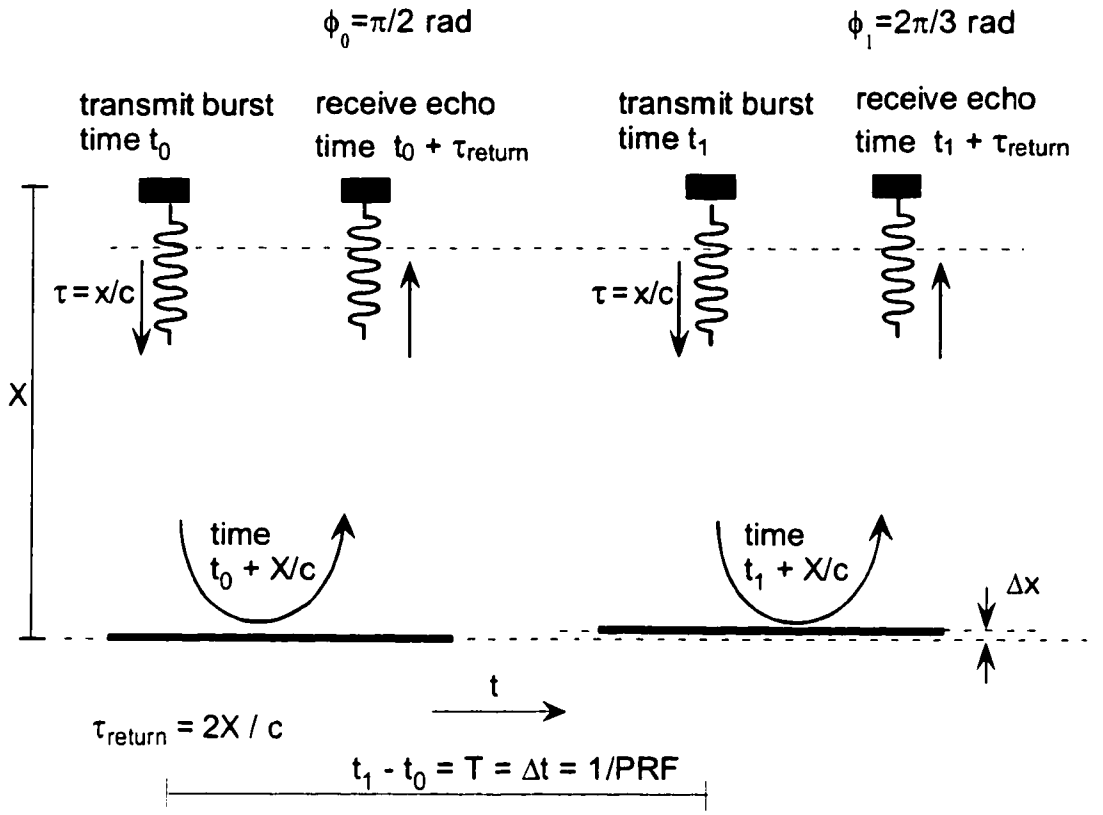

Figure 4.6. Phase measurement from narrowband ultrasound.
At time t_0, the first echo is transmitted and received while a reflective object is located at a distance X from the transducer. The received echo is compared with the transmit clock and the phase between the two is computed. At time t_1, the reflector has moved a distance Δx nearer to the transducer. A second echo is transmitted and received. Again, the received echo is compared with the transmit clock and the phase between the two is computed. The difference of the two computed phases is proportional to the distance Δx with proportionality constant λ, the wavelength of the ultrasound in the tissue.

4.1.4.2 B-mode

In anatomical imaging, motion is not a concern, so images are made simply from the amplitude demodulation of the received echoes along multiple adjacent beam lines. That is, $A(\tau)$ from equation (4-3) is plotted for each pulse echo line to form an ultrasound image. The amplitude values are displayed in shades of gray. The

standard anatomic imaging mode, therefore, is called B-mode, which is short for "brightness mode". Since phase information is ignored, a narrowband pulse is not needed, so a broadband transmit pulse is used in order to achieve fine depth resolution. An example of a B-mode image is found on the left of Figure 4.3 in section 4.1.2.

4.1.4.3 M-mode

M-mode, short for motion mode, is commonly used to visualize movements of tissues located along a single beam line. M-mode displays $A(\tau,kT)$, the AM demodulation for a specific beam line as a function of time (Figure 4.7). As with B-mode, the broad band transmit pulse is used to achieve fine depth resolution.

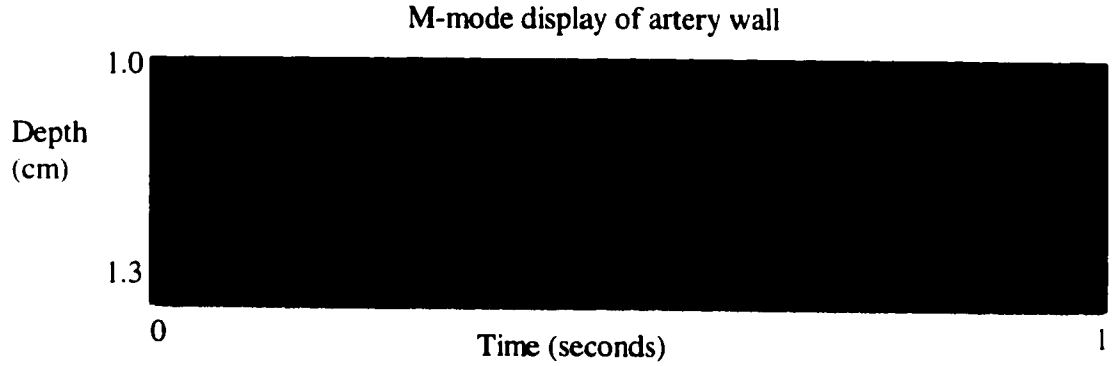

Figure 4.7. M-mode display.
This M-mode display of human artery wall motion shows tissue displacement at multiple depths versus time as computed from amplitude demodulation of the received ultrasonic echoes.

4.1.4.4 Pulsed Doppler

In pulsed Doppler mode, the goal is to measure the velocity in a certain small region of tissue called a sample volume (~ mm^3) (marked with a white rectangle in Figure 4.8). The velocity measurement is accomplished with the use of the narrowband pulse and the associated phase shift measurements. Since displacement is proportional to the phase changes, velocity, *vel*, can be measured as the differential displacement (Δx from equation (4-12)) divided by the time between position measurements ($\Delta t = T = t_1 - t_0$ in Figure 4.6):

$$vel = \Delta x / \Delta t = \lambda/2 \, (\Delta\phi / 2\pi) / \Delta t. \qquad (4\text{-}14)$$

This would be the velocity along the ultrasound beam line. When the velocity along the artery axis is desired, it can be calculated as (*vel*) cos(θ) where θ is the angle between the beam line and the artery axis. Care must be taken to ensure that the absolute value of the phase change between pulses does not exceed π, otherwise aliasing will occur and the measurement will be in error by some multiple of $\lambda/(2\Delta t)$. The time between pulses, Δt or T, is called the pulse repetition interval, and its inverse, $1/\Delta t$, is called the pulsed repetition frequency (PRF). (Figure 4.6).

Equation (4-14) illustrates how the velocity can be computed from just two phase measurements that form Δφ. This calculation is called the lambda process. To reduce noise effects, more than two phase measurements can be used in linear combination in what is called the autocorrelation calculation of the velocity. [A nice review of this method is found in Jensen (1996).]

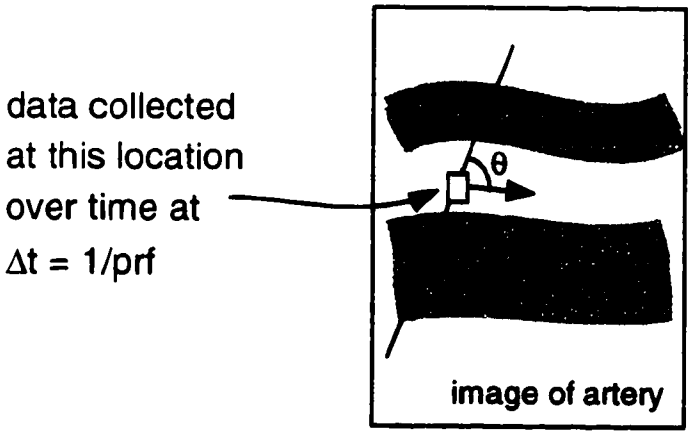

Figure 4.8. Illustration of pulsed Doppler data collection.
The pulsed Doppler mode repeatedly interrogates a single location (called a sample volume, denoted with a white box) along a fixed spatial line. The location is interrogated at a pulse rate called the pulse repetition frequency (PRF). Blood velocities associated with the artery axis are estimated by adjusting the computed Doppler frequencies by the cosine of the angle between the Doppler line and the artery axis, θ.

Typically a variety of blood velocities are present within a given sample volume. In order to measure the full range of velocities, pulsed Doppler ultrasound modes commonly compute blood velocities using a series of fast Fourier transforms (FFTs) applied to a time series of quadrature pairs. To understand this process, consider that changes in phase over time can be represented as frequency. Therefore, the velocity measurement can be calculated from the frequency content of the quadrature pair time series. The velocity values (*vel*) are proportional to the frequency values (f_D) and the transmit frequency (F) according to the Doppler equation for objects whose velocities are much lower than the speed of sound:

$$vel = (\lambda/2)\,(\Delta\phi / 2\pi) / \Delta t = (\lambda/2)\,(\Delta\phi / 2\pi / \Delta t\ ) = (\lambda/2)\, f_D = ((c / F) / 2) f_D, \quad (4\text{-}15)$$

where

$$\lambda = c / F. \quad (4\text{-}16)$$

The resulting Fourier spectra generate a spectral waveform in velocity (computed from frequency) and time such as that displayed in Figure 4.9 when the sample volume is located in the lumen* of an artery.

---

* See glossary.

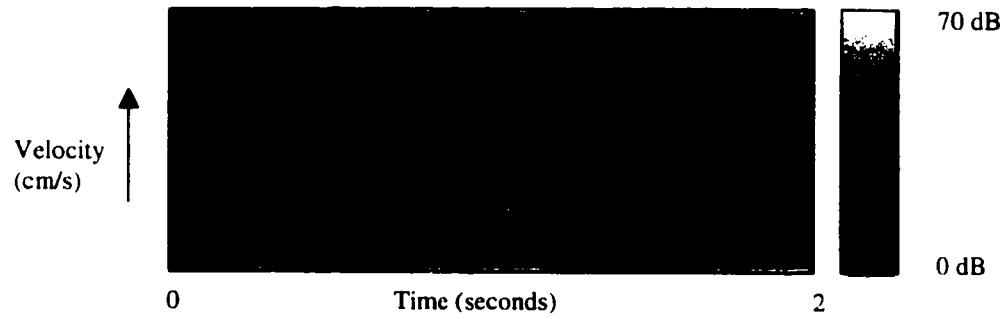

Figure 4.9. A Doppler spectral waveform.
A Doppler spectral waveform displays the velocities estimated for the Doppler sample volume. The velocity estimates are computed using the Doppler equation applied to the frequency content of the quadrature time series associated with the sample volume, as computed from the short-time Fourier transform (see section 5.1.2.1).

Note that the processing in pulsed Doppler mode does not involve measurement of a Doppler shift in the ultrasonic echo itself. Rather, it uses changes across multiple pulse echo cycles. This technique is called "time-of-flight" measurement. Despite this fact, the term "Doppler" has remained in the title of this ultrasound mode.

4.1.4.5 Color Doppler

Color Doppler mode essentially applies pulsed Doppler methods to a grid of spatial locations. Like pulsed Doppler, this mode also uses quadrature pairs to measure velocity, but it uses only 4 - 16 pairs for its calculation as opposed to pulsed Doppler's continuous interrogation of one spatial location. Since it has a limited number of values with which to compute the velocities, color Doppler modes typically use the autocorrelation method described in section 4.1.4.4 rather than the FFT. Figure 4.10 illustrates the color Doppler data collection method. Beginning at frame one, ray one, K pulses (an ensemble) are transmitted and received at the PRF rate. Upon completion of the ensemble, the instrument transmits and receives K pulses along ray 2 and repeats the process until all M ray lines have been collected. Subsequently, pulses resume along ray 1 and a new frame is begun. For each returned echo,

quadrature pairs are obtained at multiple depth locations. At each depth location, K quadrature pairs are used to determine the velocity.

The velocity information is commonly displayed by a color scheme in which motion toward the transducer is marked in red (sometimes extending to yellow) and motion away from the transducer is marked in blue (sometimes extending to green) with shades varying according to the magnitude of the velocity (Figure 4.11).

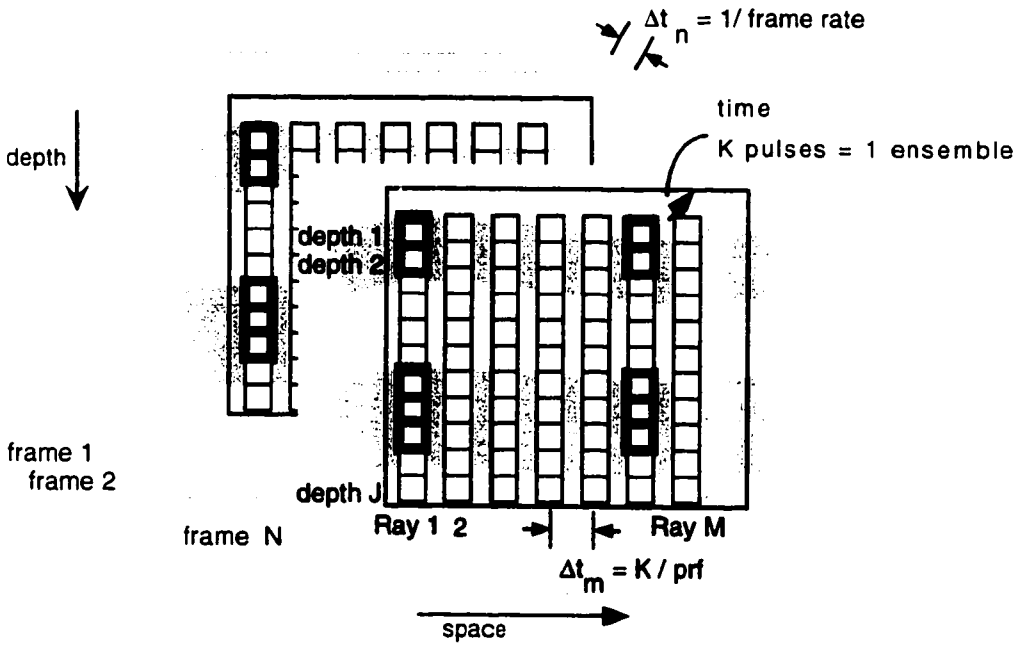

Figure 4.10. Illustration of color Doppler data collection.
Color Doppler data collection samples repeated received echoes from a given Doppler line (called a ray line) at multiple times (1 through K) corresponding to depths 1 through J. This group of echoes is called an ensemble. Next, the Doppler line is advanced to an adjacent ray line and the process just described is repeated. When M ray lines have been collected, one complete frame has been captured. The Doppler line is then returned to its original location at ray line number 1 for capturing subsequent frames.

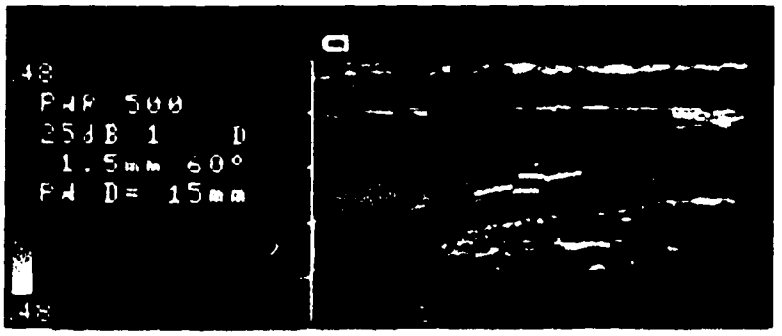

Figure 4.11. Example of color Doppler display.
A color Doppler display of an artery and neighboring vein. This mode colors motion toward the transducer as red and motion away from the transducer as blue. In this case the artery is colored red and the vein, blue.

4.2 Instruments used to collect data in this project

Three commercial ultrasound instruments have been modified so that received, quadrature demodulated ultrasound signals can be captured and processed off-line: Hokanson CP-1B, ATL UM8 and ATL HDI 1000. The instruments operate and transmit in the standard pulsed Doppler mode, but signal acquisition and processing of the returned echoes for vibrometry is different from the method used for traditional single gate pulsed Doppler. Like pulsed Doppler, data is collected along a single ultrasound beam line. Rather than acquiring signals from one sample volume, however, we use multiple sample volumes, called depth gates (Figure 4.12). Thus, data is collected from multiple locations along a line fixed in space using pulsed Doppler transmit sequences.

Figure 4.13 illustrates the initial analog and digital manipulation of the received signal from the received analog echo to the time series associated with a particular depth. Column a) plots one received RF echo, $S(\tau)$. Column b) shows the corresponding basebanded analog quadrature representation of the RF echo, I_{analog} and Q_{analog} (which represent $I(\tau)$ and $Q(\tau)$ of equation (4-8) in an analog form). The columns in c) plot the digitized quadrature data (marked with x's) from three different pulses. Processing, which occurs off-line, separates the digitized quadrature values (I and Q)

into individual depth gates (I_{nk} and Q_{nk}), where n represents the depth of interest and k represents the time of interest). At each depth gate, the phase (ϕ) of the received signal is calculated from the arctangent of the quadrature pair for each pulse echo cycle. Displacement can then be computed based on equation (4-13).

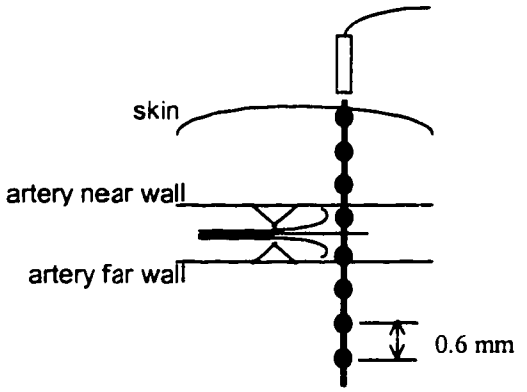

Figure 4.12. Illustration of multi-gate data collection.
For multi-gate data collection, repeated echoes from one Doppler line are sampled at fixed, multiple points in "fast" time to obtain quadrature time series corresponding to multiple depths. In this example, the time series correspond to depths separated by 0.6 mm.

Of the three ultrasound instruments that have been obtained that allow access to the quadrature signals, two of the instruments derive the quadrature signals in analog circuitry, ATL UM8 and Hokanson CP-1B. The other, ATL's HDI 1000, uses the digital sampling method of Figure 4.5. Table 4.1 summarizes their features.

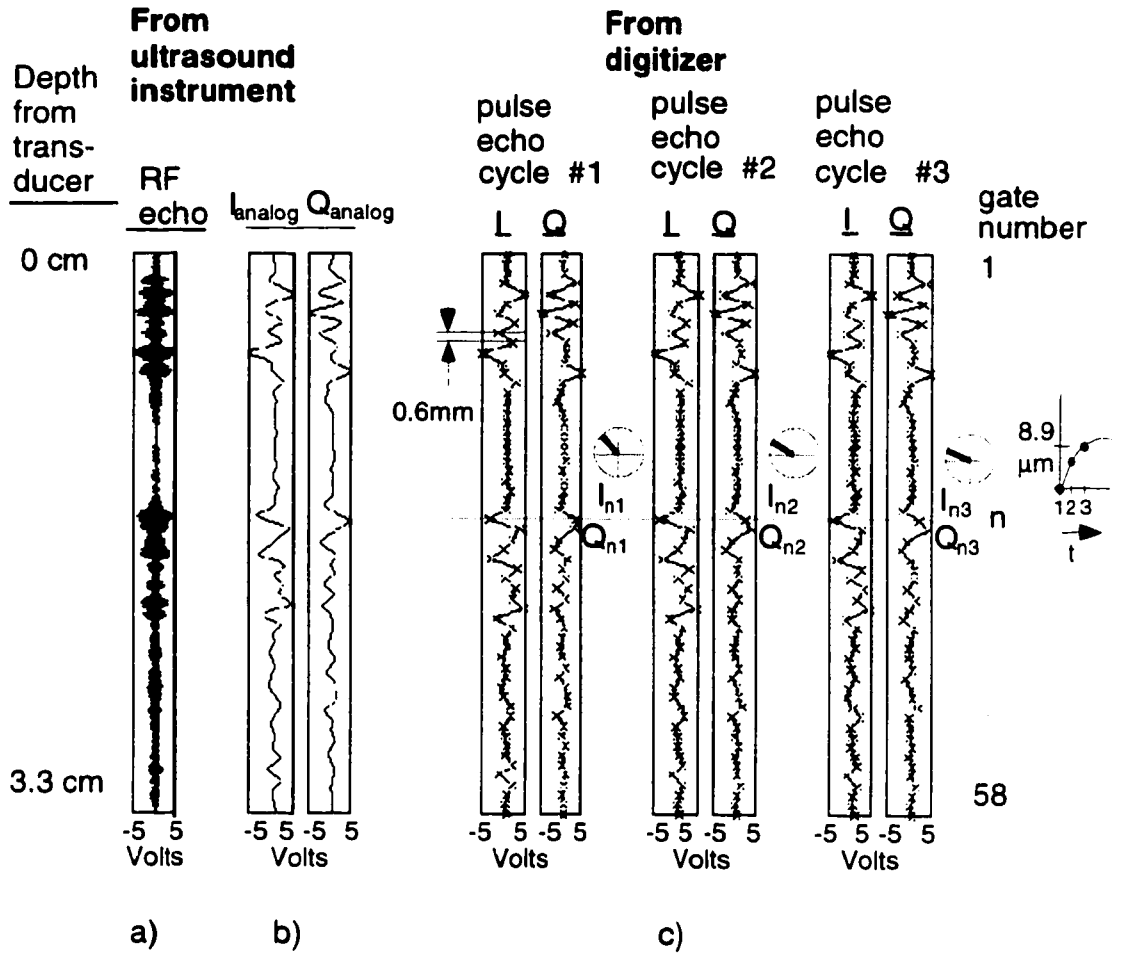

Figure 4.13. Multi-gate quadrature phase demodulation.
Column a): A received echo from a human vein graft. Column b): The base-banded quadrature representation of this echo. Column c): Digitized, base-banded quadrature representations of several echoes. The horizontal line highlights the quadrature time series associated with depth n for the pulses numbered 1 through 3. The magnitude and phase for each pulse is displayed in the complex plane. To the right of column c) is the displacement trajectory for the tissue at depth n for these three pulses.

Table 4.1. Multi-gate data collection instruments.

	B-mode for imaging	**Quadrature Channels**	**PRF (Hz)**	**Transmit frequency**	**Digitization Resolution**
Hokanson CP-1B	No	Analog	8000 Hz	5 MHz	12 bits
ATL UM8	Yes	Analog	8000 Hz	3 or 5 MHz	12 bits
ATL 1000	Yes	Digital	variable	variable	>12 bits

4.2.1 Analog instruments: Hokanson CP-1B and ATL UM8

The primary difference between the Hokanson CP-1B and ATL UM8 ultrasound instruments is that the former only operates in the pulsed Doppler mode, and the latter operates in both pulsed Doppler and B-mode. [B-mode is useful for imaging the tissue region of interest to aid placement of the ultrasound beam line used for data collection.] In the Doppler mode, both of these instruments transmit a fixed frequency tone burst of 1 μs duration. The resulting signal, therefore, has a bandwidth of approximately 1 MHz in the depth dimension. However, at each depth, or sample gate, the quadrature signals' frequency bandwidths are set by the phase modulation induced by tissue movement, which is typically less than 8 kHz. A pulse repetition frequency of 8 kHz was chosen to adequately sample typical tissue displacements and blood velocities while allowing interrogation of depths up to approximately 9 cm (see equation (4-1)).

The block diagram of Figure 4.14 shows the data collection method for these two systems. After transmission and reception, the instruments output analog quadrature representations of the received echo (I_{analog}, Q_{analog}). These are sampled at 1.25 MHz in a manner synchronous to the transmit burst to provide samples at fixed locations along the ultrasound beam. Accounting for round trip travel at a speed of sound of 1.5 mm/μs, the 1 μs pulse duration corresponds to 0.75 mm spread in the depth dimension. Similarly, the sampling rate of 1.25 MHz corresponds to 0.6 mm spacing. This provides overlap in the samples along the depth dimension. A 12-bit digitizer (Gage Applied Sciences CS512-8M) yields a 72 dB dynamic range to exceed the ~ 60

dB range required for capturing tissue and blood signals simultaneously (see section 4.1.2) with values sufficiently above the noise floor.

To capture the arterial vibrations, it is desirable to collect data from a 1 cm region on either side of the artery. Since the arteries of interest have diameters on the order of 1 cm, data along at least 3 cm depth range is deemed sufficient. With 0.6 mm spacing, 62 samples cover a depth spread of approximately 3.5 cm which is an appropriate tradeoff between depth range, sample depth spacing and data storage requirements. A large digitizer memory provides adequate storage for the data from several cardiac cycles and several centimeters in depth.

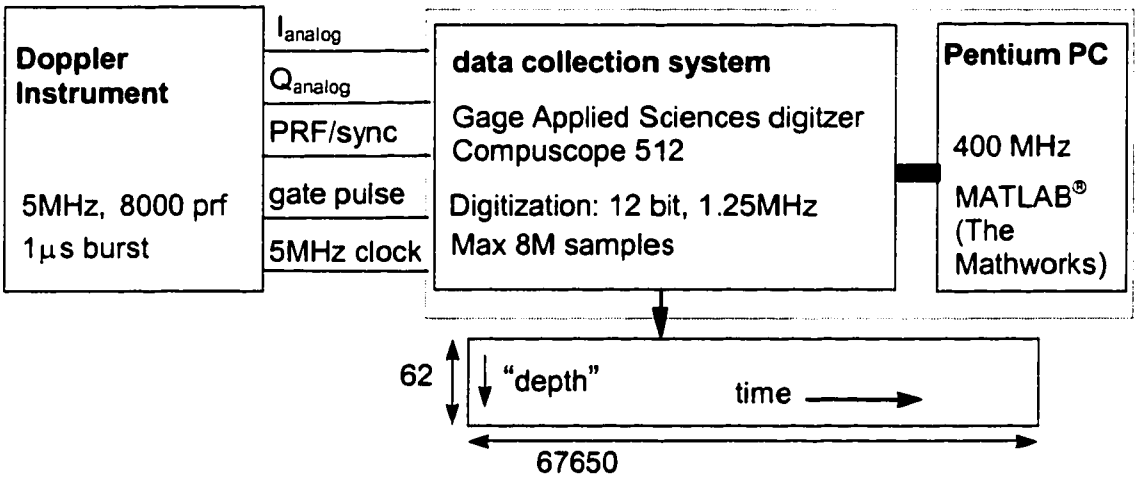

Figure 4.14. Block diagram of data collection from analog systems.
Data is collected by interfacing a Doppler instrument with an external digitizer. Digitization must be implemented completely synchronously with the ultrasound instrument's transmit sequences. The digitizer samples each received echo at 1.25 MHz which corresponds to the 0.6 mm spacing in the depth dimension. Samples are stored for 62 depths centered about the ultrasound instrument's gate pulse, which indicates the location of the selected Doppler sample volume. Data from several seconds are first stored on the digitizer's internal memory and later downloaded to a host PC. Off-line analysis is carried out in MATLAB®.

4.2.2 Digital instrument: HDI 1000

The third instrument, ATL's HDI 1000, differs from the others in that it digitizes the quadrature signals directly by sampling the returned radio frequency echo at four times its center frequency and performing the appropriate sample subtractions to obtain the quadrature samples, as described in section 4.1.3. It operates in pulsed Doppler, B-mode and color Doppler. This instrument receives each ultrasound echo from multiple transducer elements simultaneously and samples each element individually. Coherent summation of these elements yields an effective sampling resolution of greater than 12 bits. The resulting quadrature values, $I(\tau)$ and $Q(\tau)$, are provided once per transmit clock cycle, so the effective depth spacing is half the wavelength of the transmitted ultrasound. Unlike the other two instruments, data can be collected at various instrument settings, such as PRF and transmit frequency. A custom program allows storage to RAM of the multi-gate quadrature data in a depth range up to 8 cm for up to 750,000 quadrature pairs. This allows for sampling between the following extremes: many seconds of data for a short depth range at a low PRF and low transmit frequency, or a very short time duration for a long depth range at a high PRF and high transmit frequency.

4.3 Data components

The received data from solid tissue consists of four parts: vibration signal, noise, clutter and sometimes blood flow. Noise will be considered any contribution to the data that does not arise directly from blood or solid tissue within the ultrasound's sample volume. Clutter, then, consists of any movement in the tissue (solid tissue or blood) within the sample volume that is not vibration or blood flow, such as movement associated with the heart cycle. Clutter also includes the affects of stationary solid tissue. The vibration component consists of possibly multiple transient sinusoids at unknown onsets, durations, and amplitudes with unknown, time-varying frequencies and phases. Vibration examples and quantified characteristics will be discussed later in this chapter. The non-vibration data components will be discussed here.

4.3.1 Noise

The noise is assumed to be additive and consists of two parts: broadband and narrowband. It is thought to arise from the instruments' electronics, acoustic thermal noise of the tissue, and reverberation of the ultrasound between specular reflectors.

4.3.1.1 Broadband noise

4.3.1.1.1 Electronic

There are three primary types of random electronic noise: thermal (Johnson), shot and "1/f" ("excess"). Thermal noise arises from thermally excited charge carriers moving randomly in a conductor. The amplitude of thermal noise obeys a Gaussian distribution (Horowitz and Hill 1989). Shot noise, on the other hand, arises from the discrete electric charges that make up electric current. It is described by Horowitz and Hill as "rain on a tin roof. The finiteness of the charge quantum results in statistical fluctuations of the current," which are collectively call shot noise. Again, its distribution is Gaussian. A third common noise type is called by the names "1/f", "excess" and "flicker". Its name, "1/f", describes its frequency nature in that it has equal power per decade of frequency. This noise is present throughout nature. It arises due to such electronic sources as fluctuation in a resistor's resistance value and fluctuations in the base current noise in a transistor. There are also non-electronic examples of "1/f' signals such as the yearly flow of the Nile over 2000 years, loudness of classical music, and flow of sand through an hourglass. (Horowitz and Hill 1989).

4.3.1.1.2 Acoustic thermal noise from tissue

It is known that thermally excited atoms or molecules in tissue exhibit vibrations which produce random acoustic waves. The mean pressure fluctuation of these waves corresponds to the absolute temperature of the medium and the frequency spectrum of interrogation (Bowen 1987). The Rayleigh-Jeans' law describes this mean pressure fluctuation:

$$< p^2 >= \frac{4}{3}\pi k_B T_K \frac{c\rho}{c^2}(F_2^3 - F_1^3), \qquad (4\text{-}17)$$

where p is the acoustic pressure, k_B is Boltzmann's constant, T_K is the absolute temperature of the medium, ρ is the density of the medium, c is the speed of sound in the medium, F_1 is the low end of the interrogated frequency range, and F_2 is the upper end of the interrogated frequency range. This acoustic equivalent of thermal noise or blackbody radiation (Mellen 1952, Ezrow 1962) is also Gaussian in nature.

Due to acoustic attenuation in tissue, the noise associated with a particular depth in tissue will have its higher frequency content more attenuated than its lower frequency content. Thus the noise received will differ slightly from that predicted by the Rayleigh Jeans' law. Furthermore, due to inhomogeneities in tissue, thermal noise from some locations may be reflected or refracted toward or away from the ultrasound beam resulting in non-uniform thermal noise even at locations with the same absolute temperature and distance from the transducer. Therefore, the noise level may vary by depth, and if the transducer moves even slightly during data acquisition, the noise level may vary in time.

To understand the relevance of acoustic thermal noise, consider a 1 μs, 5 MHz ultrasound pulse transmitted into human tissue at normal body temperature. The pulse has approximately a 1 MHz bandwidth, causing $F_1 \approx 4.5$ MHz, and $F_2 \approx 5.5$ MHz. With $k_B = 1.38\ 10^{-23}$ J/K, $T_K \approx 310$ K, $\rho c \approx 1.5\ 10^6$ kg/(m^{2}s), and c = 1500 m/s, the pressure variance is approximately 0.90 (N/m^2)^2. Since instantaneous acoustic intensity is proportional to the square of the pressure and inversely proportional to the acoustic impedance of the medium (Z = ρc = acoustic impedance),

$$I_{inst} = p^2 / Z, \qquad (4\text{-}18)$$

the associated average acoustic intensity is found by,

$$I_{avg} = \langle p^2 \rangle / Z \approx 6\ 10^{-11}\ W/cm^2 \approx -102\ dB. \qquad (4\text{-}19)$$

Now consider the intensity of received ultrasound from a depth of 2 cm in solid tissue. Let the intensity at the skin surface be 0.3 W/cm^2 spatial peak temporal average which falls just below 0.32 W/cm^2, the maximum allowed in peripheral vessels according to the FDA (United States Food and Drug Administration) (Jensen 1996). If the pulse duration is 1 μs, then the peak intensity would be 30 W/cm^2 spatial peak temporal peak for a PRF of 10 kHz (i.e., for a 1% duty cycle, 0.3 temporal average corresponds to 30 temporal peak.). With 2 dB/cm/MHz round trip attenuation in solid tissue (see section 4.1.1.2), there is a 20 dB attenuation for 5 MHz ultrasound at the depth of 2 cm. The maximum possible peak received intensity, therefore, would be $3 * 10^{-1}$ W/cm^2 (-5 dB) which is approximately 97 dB above the average intensity of acoustic thermal noise. This value would only be achieved, however, with 100% reflection from tissue. The maximum reflection intensity, well below 100%, would occur from a boundary between body tissues with the greatest difference in acoustic impedance: fat ($1.7 * 10^6$ $kg/(m^2s)$) and muscle ($1.38 * 10^6$ $kg/(m^2s)$), excluding bone. The resulting intensity reflection coefficient, the square of R (equation (4-2)), is 0.011 or approximately -20 dB. Therefore, the more realistic resulting peak received intensity is 77 dB above the average intensity of acoustic noise, which lies just below the floor of the 12 bit digitization (72 dB) used for arterial vibrometry. Based on these assumptions, acoustic thermal noise will affect vibrometry measurements at depths below 2 cm when the transmit intensity is at the maximum level allowed. When a weaker transmit intensity is used, the ratio between the ultrasound intensity in the sample volume and the acoustic thermal noise intensity will decrease, causing the acoustic thermal noise to have a larger effect at the deeper locations and begin to affect the shallower ones.

4.3.1.2 Narrowband noise

Narrowband noise is present in many data sets (e.g. see Figure 4.15 at +/- 200 Hz and +/- 300 Hz throughout the time series). This noise can vary from depth to depth in frequency and in strength. As seen in the figure, the noise is not completely stationary. There are two likely sources of narrowband noise: 1) the instrument's electronics or 2) reverberation between two specular reflectors such as the ultrasound transducer and a bone or air interface. When the instrument's electronics are the source, the noise frequency should be constant for all sample depths (unless there are significant non-linearities), and the noise should be reasonably stationary. Narrowband noise arising from reverberation should be weak, depth dependent, and sensitive to changes in the transducer's orientation.

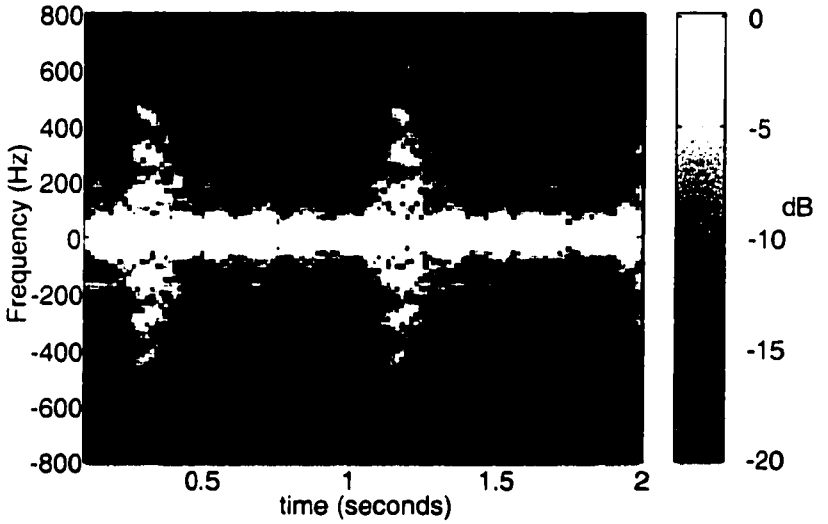

Figure 4.15. Example of narrowband vibrations in a Doppler spectral waveform. The narrowband noise appears throughout the time series at +/- 200 Hz, and less strongly at +/- 300 Hz. The data between -100 Hz and + 100 Hz is clutter.

4.3.2 Clutter

Since the vibration is the signal of interest, movement of the arterial walls at the heart rate is considered clutter. An example of this movement, common to all arteries, is

shown in Figure 4.16. When the artery moves three-dimensionally or the ultrasound beam is not located near the diameter of the vessel, the wall movement may show a considerable degree of variation from the highly regular motion shown in the figure. For discussion of the clutter in terms of frequency and filtering, the term clutter will refer to pulsatile arterial wall movement that occurs normally due to the heart beat, as well as a DC term from echoes from stationary tissue. A rapid change in vessel diameter (> 15 μm/ms) will have a Doppler shift (> 100 Hz, for 5 MHz ultrasound) that encroaches upon the range of frequencies expected for vibration. Unlike vibration which is bi-directional, this maximum velocity would occur in only one direction over a time duration on the order of 10ms: radially outward from the vessel. Displacements of aortic wall motion would resemble that of Figure 4.16 if the displacement range were multiplied by a factor of 10 (-600 to +200 μm). If this were the case, then the dotted box of Figure 4.16 would represent clutter that would achieve the 100 Hz Doppler shift.

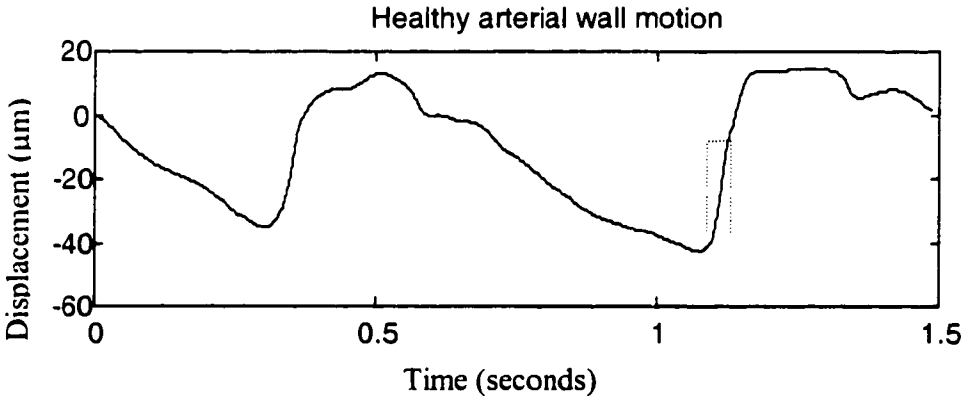

Figure 4.16. Typical arterial wall motion.
This displacement at a wall of a human carotid artery is computed as described in section 4.1.4.1. The dotted box marks the region that generates the highest clutter frequencies of this time series.

The quadrature pairs (I and Q) contain valuable information for the differentiation of stationary clutter from signal. Motion from arterial walls have a distinct appearance

when the in-phase data is plotted against the quadrature-phase data in what can be called the IQ plane as shown in Figure 4.17. The left side of the figure shows a plot of the I values versus Q values versus time. The middle view shows a plot of I and Q versus time. The right side shows a plot of I versus Q over the entire time segment. In this view, note that 1) the signal rotates about the origin, 2) the magnitude of the vectors is large corresponding to the strong echoes received from solid tissue, and 3) the trajectory encircles the origin corresponding to the large displacement of the wall (greater than a half wavelength).

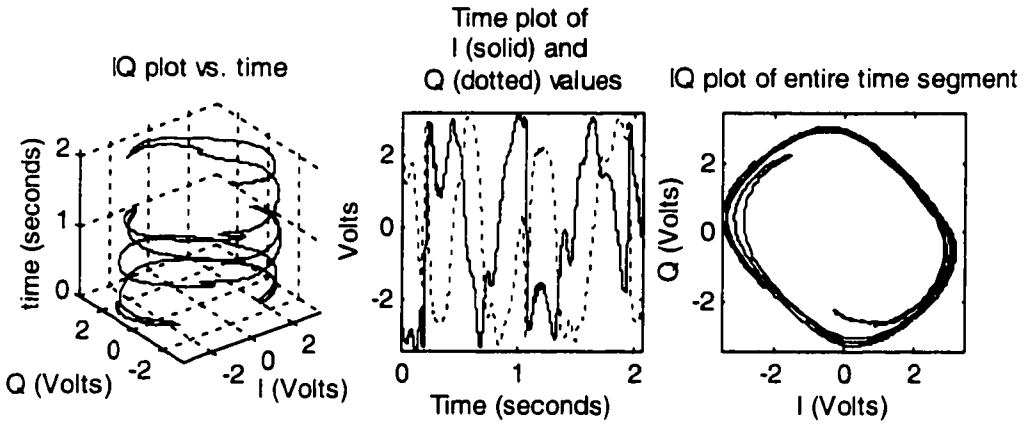

Figure 4.17. Plots of quadrature pairs from healthy arterial wall motion. **Left**: A plot of quadrature channel I vs. quadrature channel Q vs. time. **Middle**: A plot of I vs. time and Q vs. time. **Right**: A plot of I vs. Q over the entire two seconds.

Occasionally only a portion of the sample volume is in motion and another portion is stationary. In this case, the IQ trajectory will rotate about some point other than the origin (e.g., Figure 4.18). Another scenario that would produce this same phenomenon in the IQ plane would be reflection of echoes from stationary tissue not located within the chosen sample volume, but elsewhere. Two possible sources for this common effect are the following. 1) The stationary tissue coincides with a sidelobe of the ultrasound beam. 2) The transmit pulse is reflected from a depth shallower than the desired sample volume to an unexpected location, is subsequently

reflected back to the original reflector, and then backscattered to the transducer (Figure 4.19, and Beach (1993)). Regardless of the mechanism, when an IQ trajectory such as Figure 4.18 occurs, the true displacement cannot be computed as the phase of the raw quadrature pairs; rather, it must be computed from the quadrature pairs offset by the new center of the trajectory, (I_0,Q_0):

$$x(t) = \left(\frac{\lambda}{2}\right)\frac{\operatorname{atan}\left(\frac{Q(t) - Q_0}{I(t) - I_0}\right)}{2\pi}. \qquad (4\text{-}20)$$

Unfortunately, determining the values of I_0 and Q_0 can be complicated and sometimes impossible.

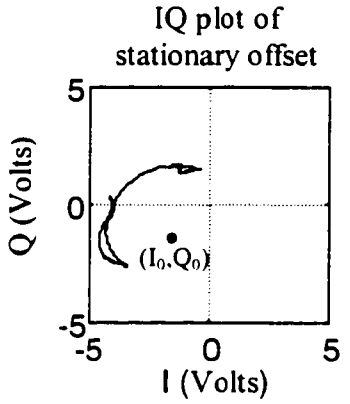

Figure 4.18. Plot of quadrature pairs for healthy arterial motion with stationary offset.
The time duration for this data is roughly 0.5 seconds.

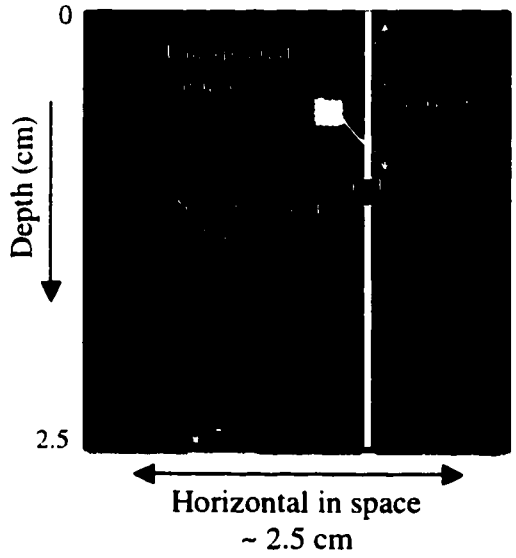

Figure 4.19. Reflected sample volume.
An ultrasound pulse propagates into the tissue along the thick, solid, white line. At the reflector, the pulse is both reflected toward the unexpected sample volume, and continues along the original path deeper into tissue. Backscatter from the unexpected sample volume propagates along the path of the reflected pulse until it reaches the reflector. The reflector, then, reflects this backscatter so that it combines with backscatter from the expected sample volume. The combined echoes finally propagate back to the transducer along the original ultrasound beam path.

4.3.3 Blood flow

Laminar blood flow commonly moves in one uniform direction within the sample volume over a period of time smaller than 10 ms, resulting in net motion either toward or away from the transducer. This will not be the case if the blood flow is turbulent, or if the ultrasound beam is nearly perpendicular to the blood flow direction.
Blood velocity can be computed from the phase of quadrature data in the same manner as solid tissue velocity (equation (4-14)). The integration of blood velocity, however, does not yield the displacement because a bolus of blood passes through the sample volume, never to return. Instead, the result is called the "time velocity integral" which represents the distance traveled by the blood if the ultrasound beam and the blood velocity are co-linear.

4.4 Vibration models and experimental setups

Several models have been used to collect vibrations that resemble true arterial disorders to varying degrees. In general, a few observations can be made about the various vibration models. *In vivo* data sets from patients and animals with arterial disorders are difficult to obtain and full of uncertainties, but represent the actual vibrations of interest. Phantom (i.e., physical model) data, on the other hand, provide vibrations from the desired mechanism, but do not accurately model the *in vivo* environment. Two compromises provide reasonable results: humming by a human volunteer and external vibration of *in vivo* tissue. Finally, computer simulations are used because they are intrinsically repeatable and controllable, albeit, simplistic. Additional comments related to collecting data from human patients and animals are given in Appendix A.

4.4.1 *In vivo*

4.4.1.1 Patients

Human subjects approval has been granted by the University of Washington for ultrasound vibrometry studies on patients and normal volunteers (Appendix B). Data has been collected from human patients with stenoses, arterio-venous shunts and pseudo-aneurysms. Since human bleeds arise only in emergent patients, the bleeding model will not be explored until vibrometry detection has been implemented on a commercial instrument and shown to be clinically useful.

Advantages:

- These are exactly the vibrations that need to be detected.

Disadvantages:

- Patients with bruits emanating from stenoses, arterio-venous shunts, or aneurysms are available only rarely and sporadically.
- The patient's physiology is constantly changing so that recruitment of a patient with a history of bruit does not guarantee that the patient will have a bruit at the time of a future exam.
- The ground truth (i.e., true frequency, amplitude, duration, presence of vibration, etc.) is not available.
- The vibration parameters are not controllable.
- These data sets are difficult to obtain due to logistics described in Appendix A.

4.4.1.2 Animals

Data has also been collected on arterial punctures in anaesthetized pigs, under approval from the University of Washington Animal Care department. Arterial punctures were performed in two ways. The simplest approach was the exposed bleed. In this case, the skin and fat layers are pulled away from the artery and the resulting cavity is filled with de-gassed water. The water serves as a coupling medium between the ultrasound transducer and the artery. B-mode ultrasound is then used to image the artery. With the artery in sight, the ultrasound instrument is switched to Doppler mode, the artery is punctured with a biopsy needle, and echoes are recorded from ultrasound beams placed in the region of the puncture. The advantage of this method is its simplicity and ease of implementation. Its disadvantage is that the region of interest is greatly modified by the exposure of the artery. Pressures and boundaries that would be present in a real trauma situation are absent in this model. A more realistic setting, then, is a puncture that is minimally invasive. I call this model the "in situ" puncture. To capture vibrations from "in situ" bleeds, the artery of interest is first

imaged as with the exposed bleed. The puncture, however, comes either from a needle within the artery guided by a catheter visualized with fluoroscopy (X-ray), or from a skilled surgeon guiding the needle through the skin and fat layers visualized with ultrasound imaging.

Advantages:

- This environment is more controlled than that of human patients. It is possible to obtain data both before and after the puncture.
- Noise, clutter and blood flow signals should highly resemble that expected from a human patient.

Disadvantages:

- The animal's physiology may change over the course of several data collections as the experimental animal loses blood volume or anesthesia changes.
- It is difficult to induce a disorder, such as bleeding, without otherwise affecting the artery and its surrounding tissue.
- It is difficult to place the ultrasound beam directly over the supposed disorder without a real-time method of detecting the vibration.
- The vibration parameters are not controllable.
- The ground truth is not available.
- These data sets are difficult to obtain due to logistics described in Appendix A.

4.4.1.3 Humming

Vascular vibrations can be simulated in the neck by short duration humming. This technique serves as an excellent model in that it preserves the expected clutter and noise while producing a signal very similar to that of arterial vibration.

Advantages:

- It is possible to obtain data both before and after the expected vibration.
- Noise, clutter and blood flow signals are taken from a human patient.
- The vibration parameters can be varied.
- It is possible to obtain data both before and after the generated vibration.
- The vibrations will propagate in a manner somewhat similar to an arterial disorder.
- These data sets are easily obtained.

Disadvantages:

- The vibration mechanism is different than from an arterial disorder.
- This method is limited to regions of the body that are in proximity to the vocal chords.
- Vibration parameters cannot be precisely varied or precisely repeated.

4.4.1.4 Externally vibrated tissue

External vibrations can be generated with an audio speaker that can be set for sequential tone bursts. Vibration amplitudes are varied by changing the amplification between the signal generator and the audio speaker.

Advantages:

- It is possible to obtain data both before and after the generated vibration.
- Noise, clutter and blood flow signals are taken from a human patient.
- The frequency, and timing of the pulsatile vibrations can be controlled exactly.
- The ground truth is known other than amplitude and intensity.
- The vibration parameters are controllable.
- The data is repeatable.
- These data sets are easily obtained.

Disadvantages:

- The ground truth for vibration amplitude, and therefore intensity, is not available.
- The vibration mechanism is different than from an arterial disorder.
- The vibrations will propagate in a more structured manner than they would from an arterial disorder.

4.4.2 Physical model: stenosis flow phantom

A pulsatile fluid flow system has been built for use as a model of an arterial stenosis (Figure 4.20). It consists of a cellulose seeded water-glycerol mixture circulated through flexible silicon tubing of 15 mm inner diameter and 2 mm wall thickness using a Harvard pulsatile pump. A magnetic stirrer was used to maintain mixture consistency. The axisymmetric stenoses were manufactured from acrylic with a graded entrance and abrupt exit. A variety of stenosis degrees and fluid flow rates were systematically applied to the system and single line, multi-gate, pulsed ultrasound data were recorded at various locations from the stenosis.

Advantages:

- The vibration's source mechanism should be similar to human physiological arterial vibrations.
- The data are repeatable.
- Some "physiological" parameters can be varied such as flow rate and degree of stenosis.

Disadvantages:

- The ground truth is not available.
- *In vivo* clutter is not included.
- The vibration parameters are not controllable.

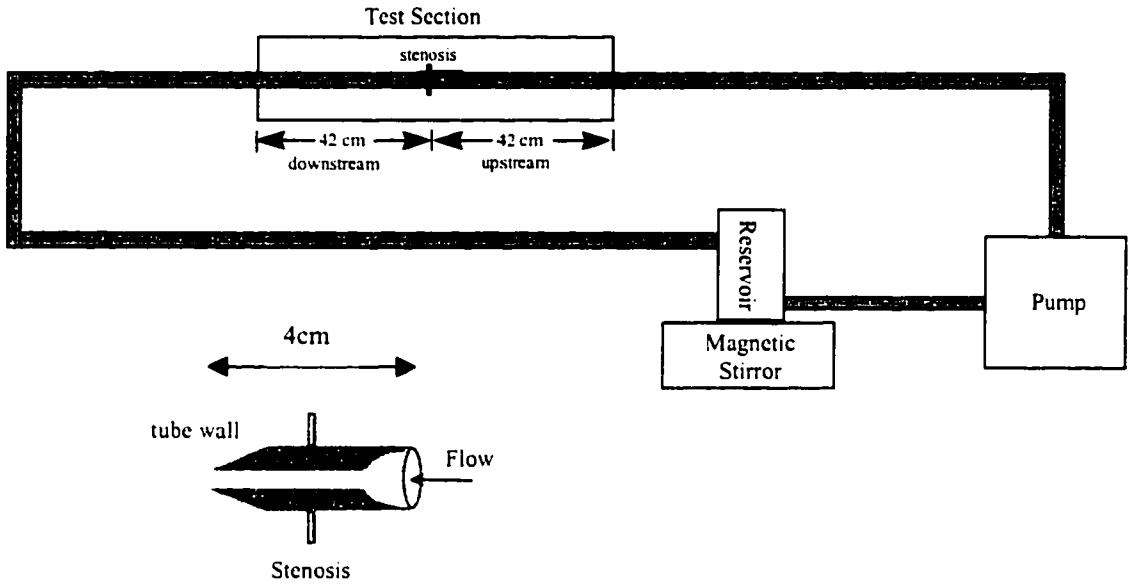

Figure 4.20. The stenosis phantom (physical model).
A stenosis resides at the center of a test section within flexible silicon tubing. A cellulose seeded water-glycerol mixture circulates through the tubing using a Harvard pump and a magnetic stirrer to maintain the mixture consistency.

4.4.3 Computer simulation:

Simulated vibrations have been generated via computer programs. Details are given in section 7.1.1.1.

Advantages:

- The vibration parameters can be completely controlled.
- The data is exactly repeatable.
- Large numbers of simulated vibrations can be generated.

Disadvantages:

- The noise and clutter can be included in a simplified manner, but will not replicate true values.
- The simulated vibrations are simpler and more structured than vibrations from arterial disorders.

4.5 In vivo vibration examples and quantified characteristics

4.5.1 Vibrations from arterial disorders

Before a system can be designed to automatically detect and estimate arterial vibrations, preliminary work must be done to collect and analyze them in order to customize the automated system, as well as predict what assumptions can be made about the data. This section will display vibration displacements from *in vivo* arterial disorders and summarize quantified characteristics found in preliminary analysis of *in vivo* arterial vibrations.

The following time series illustrates the displacement of a narrowband, *in vivo* vibration after high pass clutter filtering. This vibration has a very large signal to noise ratio (>10).

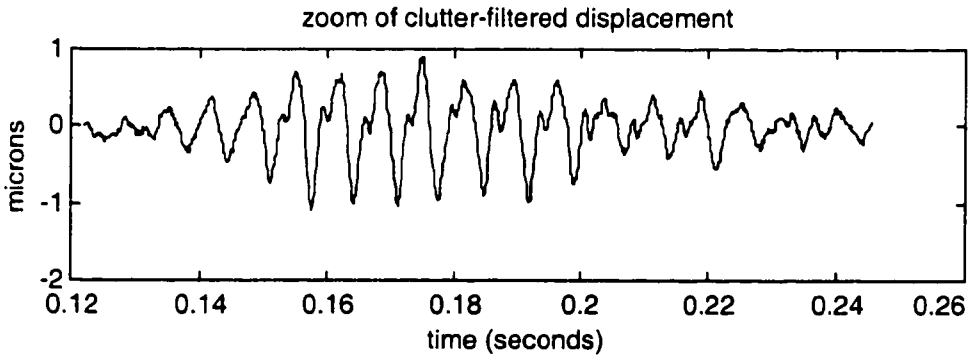

Figure 4.21. Example of narrowband vibration.
This time series was collected from the wall of a stenosed human infrainguinal vein bypass graft. The displacement is computed as described in section 4.1.4.1. Clutter filtering has been applied for displaying the vibration.

A lower SNR, broadband, *in vivo* vibration is shown in Figure 4.22 after clutter filtering.

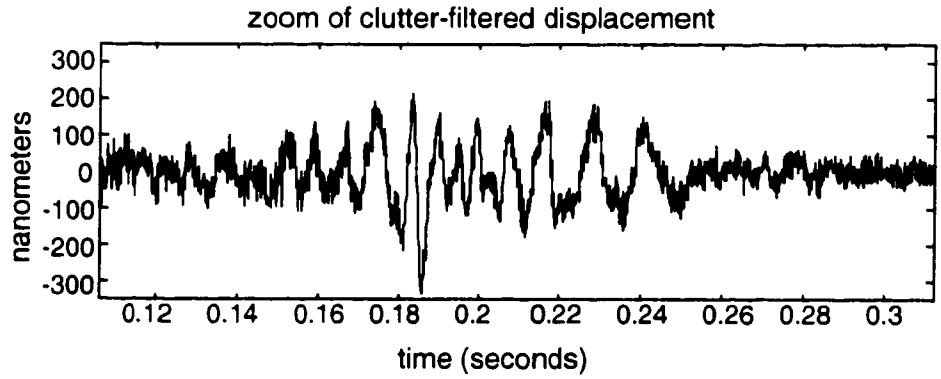

Figure 4.22. Example of broadband vibration.
This was collected from solid tissue near a porcine bleed created transcutaneously. The displacement is computed as described in section 4.1.4.1. Clutter filtering has been applied for displaying the vibration.

The vibration frequencies and amplitudes found from various *in vivo* cases are summarized in the Table 4.2 below. In addition to the quantities listed in the table, when a vibration was present in a data set at a certain time and frequency location at one depth, its presence at multiple consecutive depths ranged from 2 to 8 (0.6 mm to 4.8 mm). Since collecting *in vivo* data is time consuming and complicated, the amount of *in vivo* data available for analysis is limited. However, vibration values reported in the literature of the various disorders are completely consistent with the data described below.

In vivo arterial vibration quantities found in my experience, combined with those from the literature, are summarized in Table 4.3. Because ratios of the signal (vibration), clutter (low frequency wall motion) and noise are important in considering signal detection and estimation, they have been quantified for the *in vivo* vibrations in addition to the other vibration parameters. Specific ratios of interest are: the signal to

noise ratio (SNR), signal to clutter ratio, and the ratio of the lowest center vibration frequency to the maximum clutter frequency. Ultrasound signal to clutter ratios are typically close to the minimum values and are quite small. The lowest signal frequency approaches the highest clutter frequency when the solid tissue moves rapidly in one direction, causing a large Doppler shift.

Table 4.2. *In vivo* vibration measurements.

Data type	**Subject**		**# of data sets with vibration**	**Vibration Type**	**Frequ-ency Range (Hz)**	**Vibration amplitude (nm)**
***In vivo* Bleeds:**	Porcine	Exposed	at least 3	Narrowband Broadband	100-950 100-500	100-200 100-5000
		in situ	at least 1	Narrowband	500-700	30-1000
Stenoses/ Aneur-ysms:	Human	Pseudo-aneur-ysm	at least 1	Broadband	0-1000	~2000
		Stenosis / AV shunt	at least 4	Narrowband Broadband	150-600 0-500	30-4000 200-600

Table 4.3. Vibration characteristics summary.

Vibration characteristic	**Minimum value**	**Maximum value**
Duration (ms)	9	>120
Frequency	100	1000
Quality factor (Q_v)	1 (broadband)	10 (narrowband)
Amplitude (nm)	30	5000
Signal to clutter ratio	0.001	1
Signal to noise ratio (SNR)	1	>10
Lowest signal frequency to highest clutter frequency ratio	1	10

As noted in Table 4.3 some vibrations have durations greater than 100 ms, and some are even continuous. These long vibrations are more easily detected and characterized

with conventional methods, so this document will concern itself with the shorter duration vibrations.

The SNR is calculated as the ratio of the square of the amplitude of the vibration to the variance of the noise during the presence of the signal. The signal to clutter ratio is calculated as the ratio of the square of the amplitude of the vibration to the square of the clutter displacement only during the time that the vibration is present. Quality factor is defined as the center frequency of the vibration divided by the 3 dB bandwidth of the vibration.

It should be emphasized that the SNR values correspond to vibrations that have been detected via inspection of frequency spectra. It is fully expected that vibrations occur with much lower SNRs: this primitive detection method was simply not sensitive enough to reveal them.

4.5.2 Healthy vibrations

Healthy arterial motion occasionally exhibits a few cycles of vibration as shown in Figure 4.23, marked by the arrows. Since healthy arteries can be quite elastic it is likely that some oscillate very briefly simply due to strong pulsatile flow. These brief vibrations may fall within or below the frequency range of vibrations shown in Table 4.3.

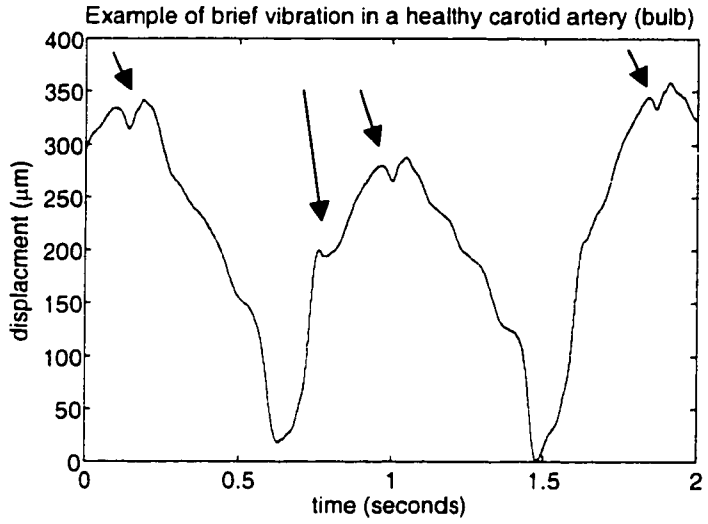

Figure 4.23 Vibration in a healthy human carotid artery wall.
The displacement from the wall of this carotid artery bulb is computed as described in section 4.1.4.1. The arrows mark the presence of short duration, low frequency vibration in this healthy arterial wall.

4.6 Vibration measurement validation

Two issues are involved in vibration measurement: displacement resolution and measurement accuracy. Displacement resolution is limited by the digitization resolution and the noise level. Displacement measurement accuracy is limited primarily by the accuracy of the assumed value for the speed of sound in the tissue. Since an alternate means of measuring displacement subcutaneously *in vivo* is not available, subcutaneous measures cannot be validated. Validation can be made *in vitro*, however, by comparing the ultrasonic measure with that of a fiber optic system. The following sections explore the resolution and accuracy of the displacement measurement methods discussed earlier in this chapter.

4.6.1 Displacement resolution

As described in equations (4-13) and (4-9), displacement is computed via a phase measurement which in turn is computed from the quadrature components. The displacement resolution, therefore, is dependent on the measurable changes in I and Q. To determine the minimum possible displacement resolution, consider a very strong

echo scattering from solid tissue such that the echo is greater than 72 dB above the noise floor. Consider further that the received echo is amplified until it just reaches the maximum voltage range of a 12-bit digitizer (+/- 2048 bins). Suppose that the received echo is very nearly in phase with the transmit clock so that most of the echo strength projects to the I channel. That is, the value of I is 2047 digitizer bins and the value of Q is 1 digitizer bin. If the previous echo had I equal to 2048 and Q equal to 0, then this would correspond to a displacement of 11.7 nm for 5 MHz ultrasound ($(300\mu m/2)(\tan(1/2047))/(2\pi)$). Thus, the best possible displacement resolution would be twice this value, or 23.3 nm. When the noise floor is within 72 dB of the echo strength, or the echo is not amplified to the maximum digitization range, the displacement resolution will exceed 23.3 nm, and thus worsen.

4.6.2 Measurement validation

Displacement measurement accuracy was verified with a fiber optic displacement measurement system on latex tubing that vibrated due to room noise (Figure 4.24). The fiber optic system is described in Appendix C. Thin-walled latex tubing was used with ¼" inner diameter and $^{1}/_{32}$" wall thickness, which is sufficiently elastic and flexible for vibrations to be induced and propagated. Also important, latex is both sufficiently echo-lucent (transparent to ultrasound) and echo-genic (reflects ultrasound). The tubing was filled with water and left otherwise untouched for measurements of room vibration effects on the tubing.

First the vibration displacement was collected with the Hokanson CP-1B ultrasound instrument and measured in the manner described in section 4.1.4.1. The left plot of Figure 4.25 shows the resulting displacement whose standard deviation was 1.8011. For this measurement, the speed of sound was simply assumed to be 1500 m/s, and the ultrasound radiation pressure was assumed to be negligible. Subsequently, the fiber optic system was used to measure the tubing vibrations (right plot of Figure 4.25). The standard deviation of the fiber optic measurement, 1.9787, is 9% greater than that for the ultrasound measurement. An exact agreement is not expected since the two measurements were not made simultaneously, and the tubing vibration amplitude is

somewhat time varying. In fact another ultrasound collection of 2.5 seconds duration had a standard deviation of 2.067, and was taken just after the one displayed. The agreement between the two measurement systems is deemed acceptable, and effort has not been made to improve the comparison, for the following reasons. 1) There is not a gold standard for *in vivo* measurements of this nature, so an error less than 9% is a tremendous improvement over the state of the art. 2) There are significant sources for error in *in vivo* measurement beyond any error sources present in this comparison experiment (see section 6.3.2).

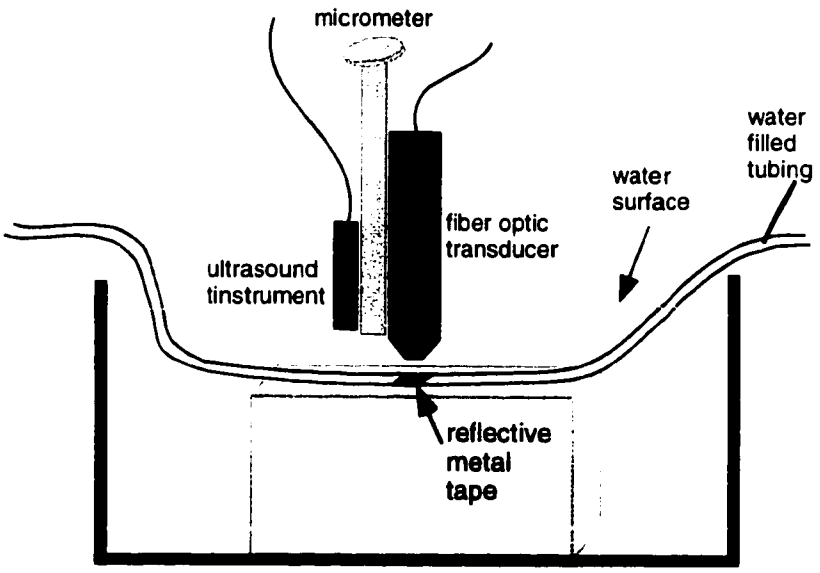

Figure 4.24. The fiber optic vibration measurement system.
A fiber optic transducer and an ultrasound transducer are both secured to a micrometer and placed above thin-walled latex tubing in a water bath. Bright, reflective metal tape is placed at the surface of the tubing just below the fiber optic sensor to reflect the fiber optic beam.

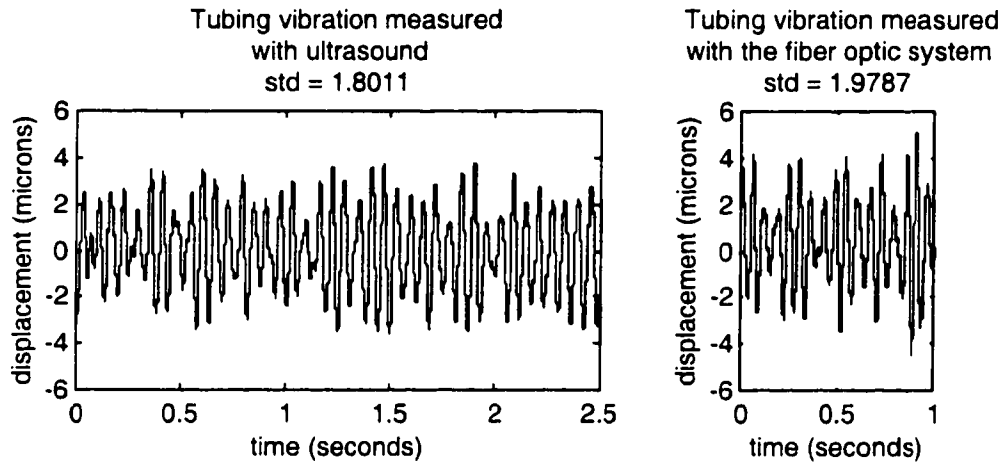

Figure 4.25. Vibration measurement comparison.
Vibrations were generated by ambient room noise only. **Left**: Vibration displacement of latex tubing as measured by the ultrasound method of section 4.1.4.1. **Right**: Vibration displacement of the same latex tubing at a subsequent time segment as measured by the fiber optic system.

5 Background on detection and estimation theory

To a first approximation arterial vibrations can be thought of as transient sinusoids. Each sinusoid can be parameterized by the following: frequency, amplitude, phase, arrival time and duration. Detection and estimation of transient sinusoids in noise is a common problem encountered in signal processing that has been solved for a single stationary sinusoid with unknown parameters in stationary additive white Gaussian noise and when there are many samples of the signal. However, detection remains a problem when the signal is transient, with an unknown number of sinusoids, with unknown parameters, in colored noise, and when the sample size is small. All of these complications are present in time series of arterial vibrations.

Detection and estimation are separate topics and will be addressed sequentially, with detection discussed first since the estimates would be erroneous in the absence of a sinusoid. Detection of a signal in a noisy time series can also be separated into two stages: feature extraction, and decision. When necessary, these are preceded by a signal conditioning stage which would eliminate any unwanted components that are known to be in the data.

The goal of the feature extraction stage is to transform the data into a form that facilitates the detection decision. The decision stage examines the features and outputs the decision that the target signal is either present or absent. The feature extraction and decision stages are extremely important, and greatly dependent on each other.

5.1 Feature extraction overview

When the features of interest in a non-stationary (i.e., time-varying) time series can be described by their frequency content, feature extraction is generally concerned with analysis in the joint time and frequency domains, since both temporal and frequency changes contain diagnostic information. The two approaches for time-frequency analysis of non-stationary data are: parametric and non-parametric. Parametric methods involve determining the parameters of a particular model, such as a sum of sinusoids, that fit the model to the data in a given segment. Non-parametric methods,

alternatively, make few assumptions about the data and can be applied directly to a broad range of time series.

5.1.1 Parametric

Important issues in parametric analysis include the selection of the model and its order (e.g., the number of sinusoids). Inappropriate model and order selection will cause the data to be underfit or overfit, and it may be catastrophic to subsequent analysis. Because a concise model of arterial vibrations is not available, parametric methods are not appropriate for vibrometry. Discussion of parametric methods is included here because these are common feature extractors for sinusoidal signals, and their inutility for vibrometry should be more carefully explained.

All of the following parametric methods assume stationary data; thus, non-stationary time series must be separated into segments that can be considered stationary.

5.1.1.1 Autoregressive methods (AR)

The most common parametric method employs autoregressive models (AR) in which it is assumed that a data value at a given time can be predicted from the preceding p data values and a noise term (Oppenheim and Schafer 1989). An advantage of this method is that any power spectrum can be modelled by an AR process of some order p; however, the value of p may exceed the length of the time series (Percival and Walden 1993). The AR model is written as:

$$x_t = \sum_{l=1}^{p} a_l x_{t-l} + n_t; t \geq l \,, \qquad (5\text{-}1)$$

where x_t represents time samples, a_k are the coefficients of the AR process, p is the model order, and n_t are samples of a stationary white noise process. AR systems can also be described by the following power spectrum:

$$\frac{\sigma_p^2 \Delta t}{\left|1 - \sum_{l=1}^{p} a_l e^{-j2\pi f l \Delta t}\right|^2}, \qquad (5\text{-}2)$$

where σ_p^2 is the variance of the noise term, n_t; f is frequency; Δt is the time between samples; and j = sqrt(-1). The AR process is an all-pole system where all of the roots lie within the unit circle.

Criteria have been developed for determining the model order, p, with various degrees of computation. Once the model order is chosen, linear predictive, maximum entropy and maximum likelihood methods are used to determine the parameters that fit the chosen model to the data. These, too, have varying degrees of computational cost. Unfortunately, the value of p is crucial in modeling the data, and the methods for determining the model order are known to be imperfect. For a more in-depth consideration of autoregressive parametric methods, the reader is referred to Percival and Walden (1993).

5.1.1.2 Statistical inverse theory

When the appropriate model for the desired signal is known, statistical inverse theory can be used to determine the model parameters. This technique makes intensive use of least squares methods. For the case of stationary, purely sinusoidal signals in white noise, the computation is minimal and straightforward; however, it can be extremely computationally intensive as the number of model parameters rises.

5.1.1.3 Singular value decomposition (SVD) for sinusoidal signals

If the signal can be modeled purely as a sum of stationary sinusoids, then the problem of computing the model parameters can be formulated as a matrix equation. The model parameters are directly computed from the eigenvalues, or singular values, of this matrix equation. Hence, the method is called the singular value decomposition. The reader is referred to Percival and Walden (1993) for more information.

5.1.2 Non-parametric

Non-parametric techniques include the short-time Fourier transform (STFT), the wavelet transform, the Cohen kernel time-frequency transforms, and the Karhunen-Loève expansion. For non-stationary data, a primary consideration is the tradeoff between the frequency resolution and time resolution. Each of these methods addresses this tradeoff differently.

5.1.2.1 Short-Time Fourier transform

The short-time Fourier transform (Oppenheim and Schafer 1989) involves the segmentation of a time series into possibly overlapping blocks of equal length that are assumed to be stationary. The fast Fourier transform (FFT) is applied to each block of data sequentially. The FFT projects the frequency content of each segment into equally spaced frequency bins (Δf) spaced according to the temporal length (T_0) of the segment:

$$\Delta f = \frac{1}{T_0}. \qquad (5\text{-}3)$$

This fixed frequency spacing implies that the bandwidth of each segment is constant. The value of T_0, should be chosen small enough so that the time series can be considered stationary over the time segment. Because it treats the data as if it were stationary over each segment, the STFT computes only one amplitude and phase value for each time-frequency bin. A plot of the modulus squared (i.e., the square of the magnitude) of the sequential FFT ordinates versus time is commonly called the spectrogram.

The STFT is extremely common and computationally efficient. A significant drawback of the STFT arises, though, when the data has a large dynamic range, as is the case for arterial vibrations in strong, low frequency clutter. In such a case, high frequency noise bias can mask the high frequency signal structure (Oppenheim and

Schafer 1989, Percival and Walden 1993). Tapers (temporal windows that weight the time samples unevenly) can be applied to each segment before transformation to limit the bias, but not eliminate it. One effective means of reducing the bias, called the multi-taper method, is the application of multiple tapers to each time segment, followed by an averaging of the resulting spectra (Percival and Walden 1993). The optimal choice of the type and number of tapers varies with the dynamic range of the segment and the frequency resolution necessary to fully analyze the data set. Therefore, to best analyze a time series, either knowledge of the dynamic range and frequency separation of the signals must be available *a priori*, or various tapers must be applied and compared. For sufficiently long stationary segments, the bias can be further mitigated by averaging smaller, overlapping segments using methods such as Welch's overlapping segment averaging (WOSA) technique.

In 1946, Gabor demonstrated that a Gaussian-shaped taper simultaneously maximizes the time and frequency resolution of the short time Fourier transform (Gabor, 1946). The use of a STFT with a Gaussian window is now called the Gabor transform.

5.1.2.2 Wavelet transform

The wavelet transform is another non-parametric technique with a growing number of applications. This technique begins with a mother wavelet that is a zero mean, short duration time signal, w(t), which is convolved with the time series of interest, x(t), in order to filter the time series. The mother wavelet is then dilated (expanded in time) at fixed dilation factors to form daughter wavelets, called wavelet scales. Each of these is convolved with x(t) as well. Thus, the mother wavelet, and its dilations, form a bank of possibly overlapping filters (scales) in which each segment of the bank has the same quality factor, Q_w (bandwidth/center frequency).

Figure 5.1 illustrates the filter bank concept of the Morlet wavelet. At the top of the figure, x(t) is displayed in both time and frequency:

$$x(t) = [\ \sin(2\pi\ 250t + \pi/3)+\sin(2\pi\ 700t)\ ]\ [0.5\ (1 - \cos(2\pi\ t\ /0.032))], \quad (5\text{-}4)$$

where the latter factor is a Hanning taper. This illustration uses the real part of the complex Morlet wavelet, which is a Gaussian tapered complex sinusoid (Figure 5.2):

$$w(t) = \pi^{-1/4} e^{-j\omega_0 t} e^{-\alpha t^2/2} \qquad (5\text{-}5)$$

$$W(\omega) = \pi^{-1/4} e^{-(\omega-\omega_0)^2/(2\alpha)}, \qquad (5\text{-}6)$$

where ω_0 represents the radian frequency of the sinusoid; α, the rate of decay; and $W(\omega)$, the Fourier transform of $w(t)$.

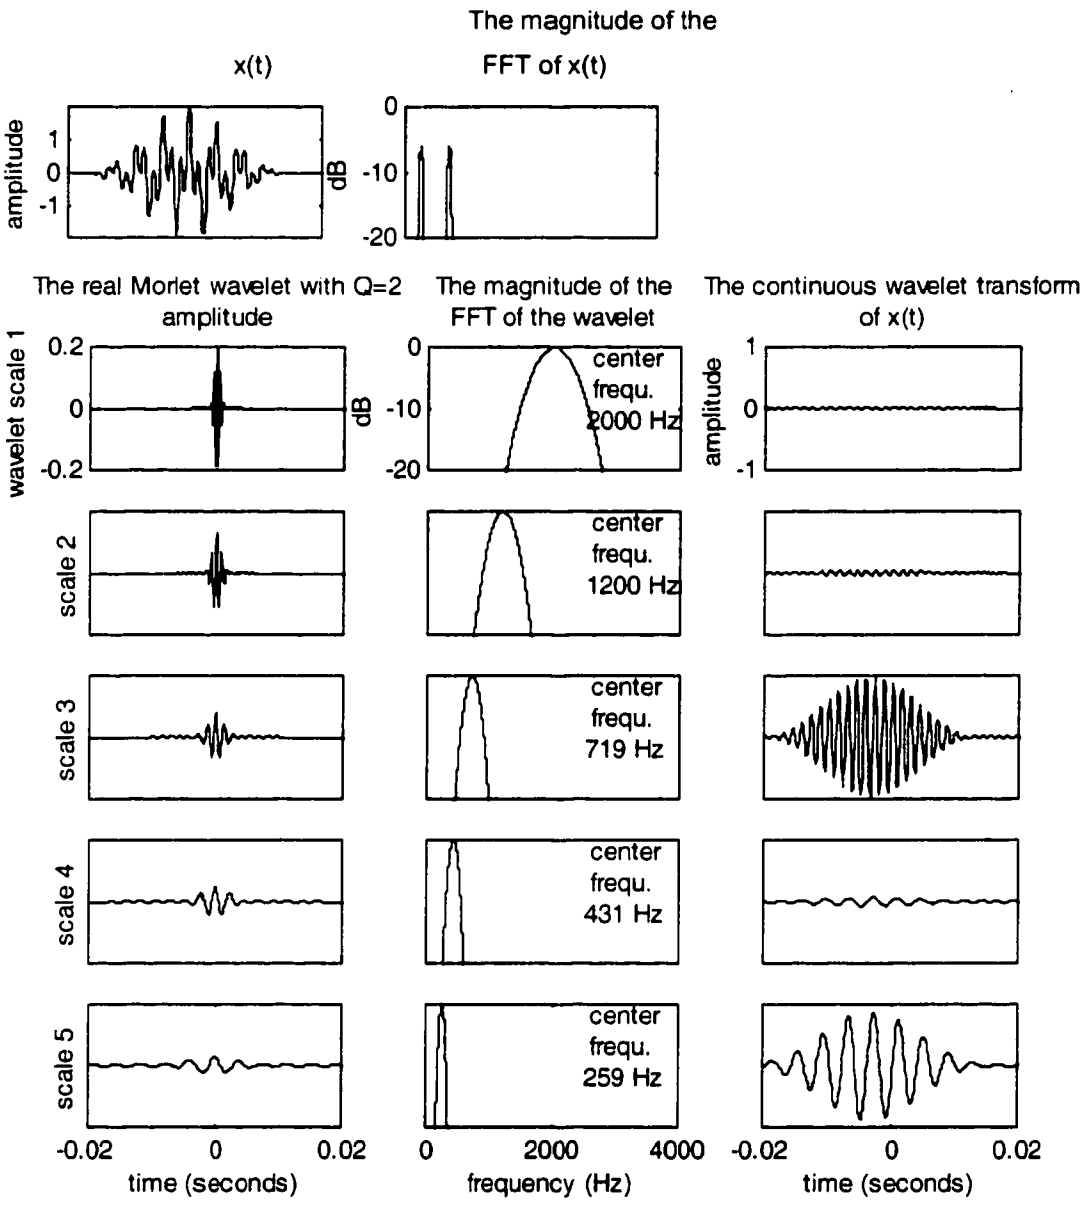

Figure 5.1. Example of a wavelet as a filter bank.
Top row: The example signal, x(t). Left: x(t) in the time domain. It is a Hanning tapered sum of two sinusoids with frequencies at 250 and 700 Hz, and amplitudes of one. Right: The frequency domain representation of x(t). **Left column:** (rows 2 through 6) The Morlet wavelet at various temporal dilation factors (wavelet scales). **Middle column:** (rows 2 through 6) The frequency domain representations of the various wavelet scales. **Right column:** (rows 2 through 6) The continuous wavelet transform of x(t) as computed from the convolution of the wavelet scales with x(t).

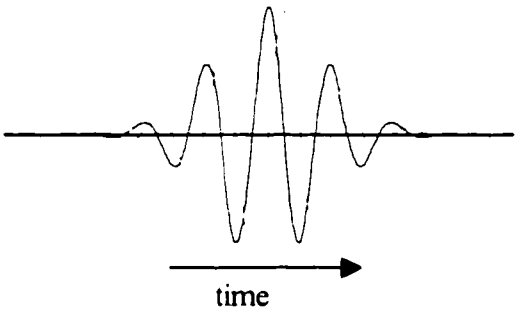

Figure 5.2. The complex Morlet wavelet.
Real{w(t)} is black. Imag{w(t)} is gray.

The quality factor, Q_w, of the Morlet wavelet is determined by the two parameters ω_0 and α (Teich et al. 1995):

$$Q_w = \frac{\omega_0}{2\sqrt{2\alpha}}. \qquad (5\text{-}7)$$

The mother wavelet resides at scale 1, shown in the second row of Figure 5.1. The center frequency, $\omega_0 / 2\pi$, for the mother wavelet has been chosen to be 2000 Hz, and α to be 4.93 10^6 for a fast decay resulting in a Q_w of 2. The daughter wavelets are dilations of the mother wavelet with a dilation factor of $^5/_3$.

At scale 1, the wavelet acts as a bandpass filter centered around 2000 Hz. Since the frequency content of x(t) falls well below this range, the wavelet transform coefficients at scale 1, resulting from the convolution of this mother wavelet with x(t), will be zero. Although the wavelet at scale 2 has a center frequency at 1200 Hz, its passband extends down to 700 Hz, albeit attenuated by nearly 20 dB. Thus the 700 Hz component of x(t) results in non-zero wavelet coefficients at scale 2, though they are small. At scale 3 the wavelet center frequency (719 Hz), nearly matches the 700 Hz component of x(t), resulting in large wavelet coefficients. Scales 4 and 5 act on x(t)'s

component at 250 Hz in a manner similar to scales 2 and 3 acting on x(t)'s 700 Hz component, respectively.

In this example, the wavelets have been normalized so that the amplitude of the wavelet coefficients will equal the amplitude of x(t) when the wavelet frequency exactly matches a component of x(t), provided that the wavelet is much shorter than x(t). This concept is discussed in greater detail in section 6.3.1.1, and Appendix D.

Figure 5.3 compares the wavelet transform to the STFT. Both transforms separate a signal by time and frequency, with a constant value of the product $\Delta f \Delta t$, where Δf is the width of the bins in the frequency dimension, and Δt is the width in the time dimension. The left side of the figure illustrates that the STFT parses a signal into equally spaced time and frequency bins (i.e., fixed Δf, and Δt). At the lowest, non-zero frequency bin, the time resolution is equal to one cycle at that frequency. At the higher frequencies, the time resolution degrades until one time-frequency bin represents the average of numerous cycles.

The wavelet transform is analogous to a STFT in which the duration of the temporal segment is not fixed, but varies with frequency. That is, Δf increases as Δt decreases, as shown in the right side of the figure. Unlike the STFT, there are a fixed number of oscillations per time-frequency bin. A significant advantage is that the wavelet itself is time-limited, so non-stationary data does not need to be artificially separated into stationary segments. The time segments in the right side of Figure 5.3 are shown to illustrate that the time duration of a wavelet at a specific scale (or frequency range) is proportional to the center frequency of the scale, and also to the bandwidth of the scale since the quality factor is a constant across scales. The data are not actually separated into segments before applying the wavelet transform. At the high frequency scales, the wavelet transform yields superior time resolution compared to the STFT without the stationarity assumption. As the wavelet center frequency decreases, the frequency resolution increases, but time resolution decreases. This tradeoff resembles the

vibration data quite well in that lower frequencies have longer time periods and thus they have inherently poorer time resolution.

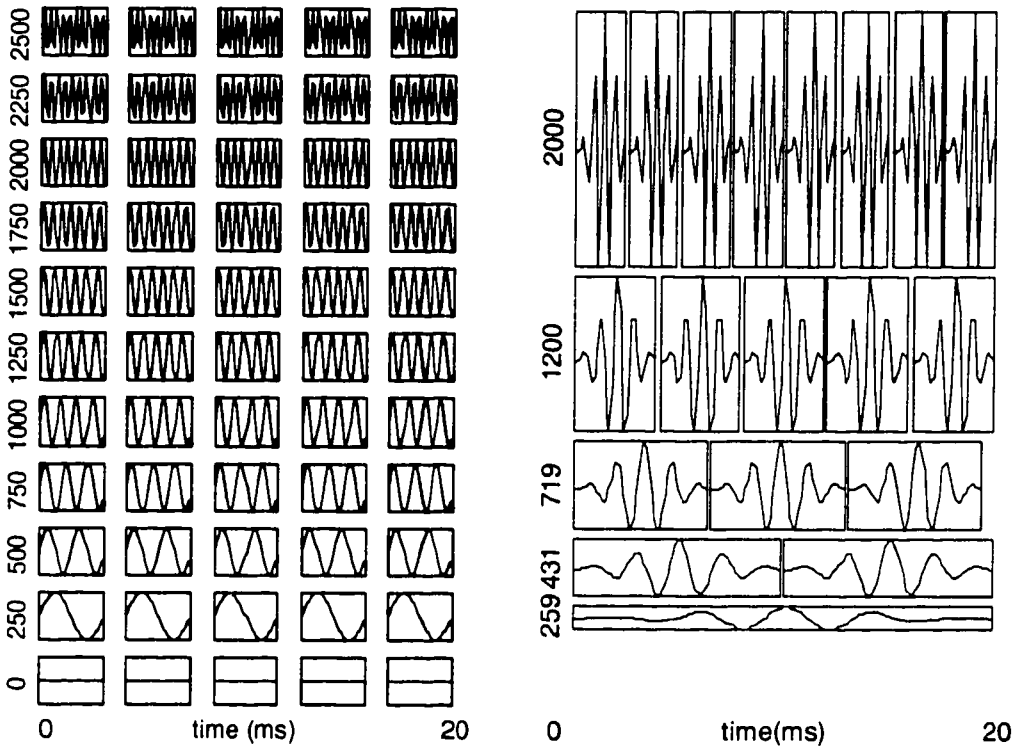

Figure 5.3. STFT (left) and wavelet (right) time frequency bins.
The same digitization rate is used for both transforms. **Left:** The bins of the STFT have fixed bandwidth and time duration. The time and frequency boundaries of the bins represent the boundaries of the resulting transform. **Right:** The bins of the wavelet transform have a fixed quality factor (bandwidth / center frequency). The frequency boundaries represent true boundaries of the resulting transform. The time boundaries, however, illustrate the time duration of the wavelets at the various frequency ranges, but the resulting transform does not segment the data into fixed time blocks.

Finally, it makes intuitive sense to employ a wavelet transform, rather than a Fourier transform, to extract features of a vibration because the wavelet more closely resembles a vibration than does a continuous sinusoid (Figure 5.4). This intuition is validated in a concept known as the matched filter, which is commonly used in the

feature extraction stage. The theory states that when a signal is known and is in zero-mean additive white Gaussian noise, the feature extractor offering the maximum signal to noise ratio (SNR) will have an impulse response equal to the time-reversal of the signal itself. [If the signal is symmetric in time about its center point, then its time-reversal is simply the original signal.] In other words, correlating the received data with the expected signal will yield a result that is best suited for the detection decision in terms of SNR. Since a set of wavelets resemble transient vibrations of many different frequency ranges, the wavelet transform can be roughly considered a bank of matched filters for use with vibration detection.

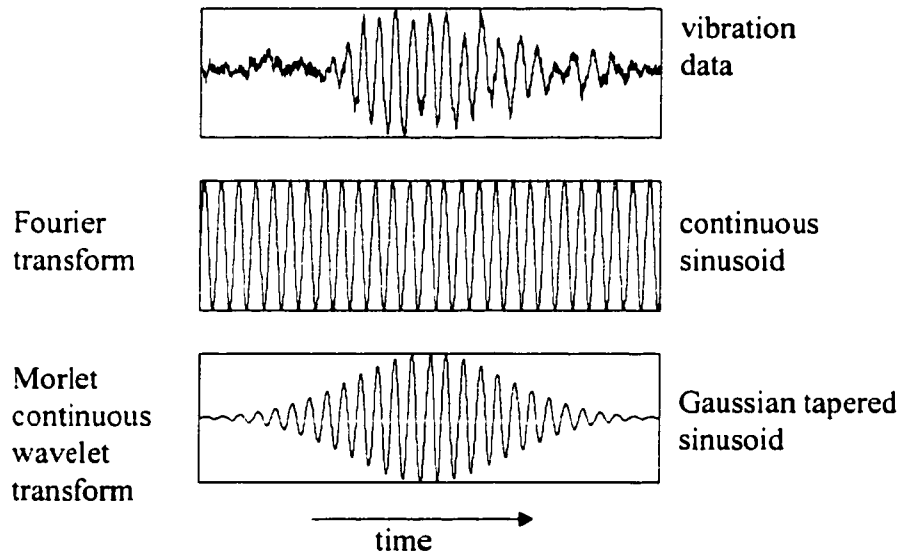

Figure 5.4. Comparison of a vibration to a continuous sinusoid and a Morlet wavelet.
Upper plot: A vibration from a stenosed human femoral vein bypass graft. **Middle plot:** A continuous sinusoid at a frequency near to that of the vibration. **Lower plot:** A Morlet wavelet (Gaussian tapered sinusoid) with a center frequency near that of the vibration.

The wavelet transform is commonly applied both in a "continuous" (CWT) and discrete (DWT) manner. The CWT incorporates redundancy into the transform in

order to facilitate analysis of the signal of interest, at the expense of computational cost. The DWT, alternatively, can be computed very efficiently, but at a significant loss of flexibility for signal analysis. The need for this flexibility is discussed in Chapter 6. Because the CWT is used for vibrometry, it will be discussed in detail; whereas, the DWT will only be summarized for completeness.

5.1.2.2.1 Continuous wavelet transform (CWT)

The wavelet, w(t),used for the continuous wavelet transform can be customized for the signal features of interest provided it meets two criteria (Farge 1992).

1) The integral of the wavelet is zero: $\int_{-\infty}^{\infty} w(t)dt = 0$

This requires that the wavelet has a mean of zero, but is not uniquely zero.

The other criteria, called the admissibility condition, ensures that the original signal can be constructed from its continuous wavelet transform:

2) $\int_{-\infty}^{\infty} \frac{1}{f} |W(f)|^2 \, df < \infty$, where W(f) is the Fourier transform of w(t).

This is satisfied when the wavelet is a true bandpass filter in that it removes components at both low and high frequencies. Reconstruction will be discussed later in sections 5.3.2.2 and 6.3.1.1.

The Morlet wavelet, mentioned earlier (equation (5-5)), is a continuous wavelet commonly used for analysis of signals containing sinusoids, and it is the wavelet that is used for arterial vibrometry. One advantage of the Morlet wavelet is that its quality factor is adjustable according to equation (5-7) so that it can be customized to the signals of interest. Another significant advantage of the Morlet wavelet is its superb clutter filtering ability. Figure 5.5 illustrates that the continuous wavelet transform with the Morlet wavelet will pass 135 Hz vibration signal unattenuated while applying 60 dB attenuation to 75 Hz clutter for a sampling rate of 8000 Hz.

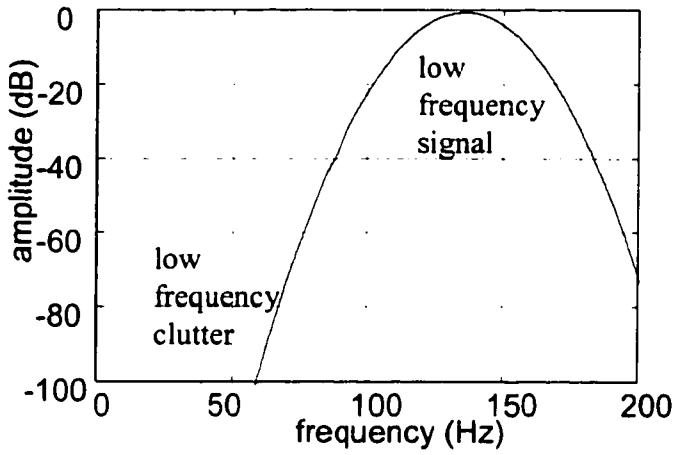

Figure 5.5. Illustration of the clutter filtering ability of the Morlet wavelet. The frequency domain representation of a Morlet wavelet with a quality factor of 3 with a center frequency at ~ 135 Hz. This wavelet will pass a 135 Hz vibration unattenuated, while attenuating 75 Hz clutter by 60 dB (a factor of 1000) and 100 Hz clutter by 20 dB (a factor of 10).

As with the Gabor transform, the Morlet wavelet uses a Gaussian window and thus provides simultaneously maximal time and frequency resolution.

The Morlet wavelet is certainly not the only continuous wavelet. Some others include the Mexican hat, Sinc and Meyer wavelets shown in Figure 5.6. All of them have the property that their integral is zero.

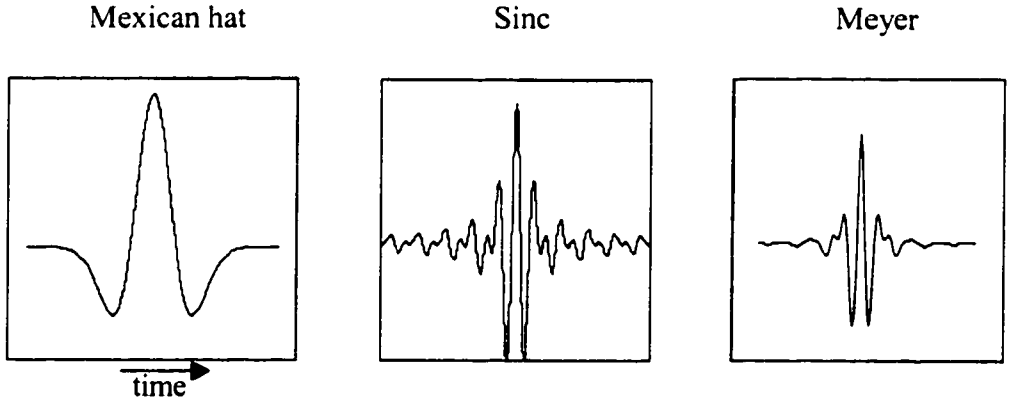

Figure 5.6. Some other continuous wavelets.

5.1.2.2.2 Discrete

To be considered a candidate for the discrete wavelet transform (DWT), the mother wavelet and its dilations must not only meet requirements 1) and 2) listed above, but they must also form an exact orthogonal basis. The wavelets associated with a given mother wavelet are orthogonal if a) the dot product (sum of the product) of any two wavelets from different scales is zero regardless of time-shifting, b) the dot product of a wavelet with itself is finite only at the same time shift, and it is zero otherwise. Although continuous wavelets can approximate this property if they are shifted only by certain discrete intervals, discrete wavelets can meet it exactly, with the exception of the continuous Meyer wavelet which can be made orthogonal (Daubechies 1992). Orthogonality has the advantages that redundancy is eliminated, which significantly reduces computational cost, and that the total energy is exactly conserved (i.e., the sum of the squared modulus of all wavelet coefficients equals the integral of the squared modulus of the original signal). An added criteria, compact support, requires that the wavelet has a finite duration. This serves to further reduce computation costs.

Requirements for both orthogonality and compact support place stringent limitations on the form of the DWT wavelet so that the resulting wavelets often have an unexpected appearance (e.g. Daubechies D4 and Coiflets 2 in Figure 5.7), although some are smoother (e.g., D8 and Coiflets 6 in Figure 5.7). Daubechies pioneered these widely used wavelets which are described in detail in Daubechies (1992). These compactly supported orthogonal wavelets are also restricted to dyadic (power of 2) dilations, further limiting analysis flexibility.

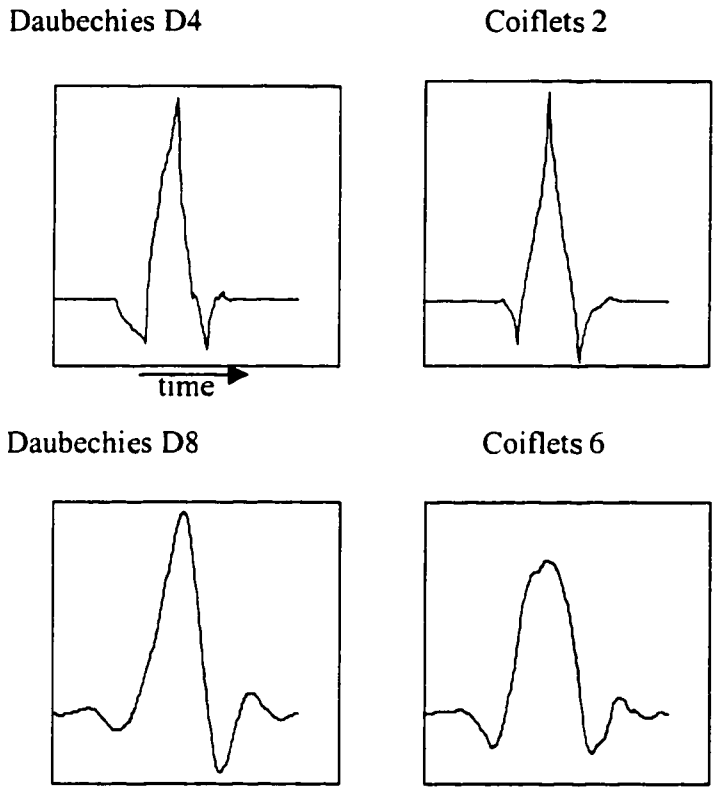

Figure 5.7. Some orthogonal, compactly supported wavelets.

5.1.2.3 Cohen kernel

Energy plots derived from the STFT and the wavelet transform are two extremes of the family of Cohen kernel time frequency transforms (Rioul and Flandrin 1992). Other members of this family of transforms, such as the Wigner-Ville transform, provide a more balanced combination of time/frequency resolution, at the expense of increased computation. Low signal to noise ratios and multiple signals occurring simultaneously, though, can confuse the interpretation of these plots and introduce artifacts in the resulting spectra. For these reasons, vibrometry does not employ these intermediate members of the Cohen kernel transform family.

5.1.2.4 Karhunen-Loève

An alternate non-parametric feature extraction method, the Karhunen-Loève expansion, involves finding data-specific orthogonal bases that form a better match to

the data than the sinusoid of the Fourier methods or the known wavelets. This approach is effective when example data is available in large quantities and the examples exhaustively encompass the signals of interest. This ensures that the selected bases contain all of the features of interest. Since it is prohibitive to obtain large quantities of examples of *in vivo* data (see Appendix A), this method is not suitable for arterial vibrometry.

5.1.2.5 Adaptive filtering

When a signal is slowly time varying, a common means of feature extraction employs what is known as adaptive filtering. Non-parametric adaptive filtering typically relies on a frequency basis such as the Fourier transform. Recently, adaptive filtering was extended to make use of the discrete wavelet transform (Basbug 1993). Implementation of adaptive filtering with the Fourier transform involves the following steps. 1) Some initial assumption is made regarding the Fourier coefficients expected for the data awaiting feature extraction. 2) The data are compared to the inverse Fourier transform of the assumed coefficients. 3) The mean squared error is computed between the true data and the assumed data. 4) The assumed coefficients are altered in an attempt to minimize the mean squared error. Steps 2, 3 and 4 are repeated continuously. The resulting Fourier coefficients converge on the true coefficients of the data, and they will vary as the input data varies. In white noise, or noise with a known spectrum, the Fourier coefficients represent the desired feature extraction. When the noise is colored with an unknown spectrum, however, non-parametric adaptive filtering will model the combination of the signal and noise, but not extract the signal from the noise. Also, for rapidly varying signals, the iterative process will not reach steady state before the signal changes. Thus, adaptive filtering is problematic for arterial vibrometry because the vibrations can be very brief, and the noise spectrum is neither white, nor known.

5.2 Automated detection overview

When the features of interest in a time series have been appropriately extracted, the next step is the detection decision. Modern detection falls into two primary categories.

The first is traditional detection theory and the second includes automated learning systems such as neural networks* or hidden Markov models*. These latter methods use very little *a priori* information about the data sets. Instead, they require a large database of examples with which to learn the nature of the desired signals. Traditional detection theory, alternatively, makes use of all information known about the signals of interest and uses statistics to make the decision, rather than comparison to examples. Since a large database of vascular vibrations is prohibitive, traditional detection theory becomes the method of choice.

It should be mentioned that M. Akay et al. (1992), Y.M. Akay et al. (1993, 1994) and M. Akay (1994) have explored many of the various feature extraction and detector methods for detecting stenoses and coronary artery disease from their acoustic time-frequency signature. Akay's search for a successful detector led to the use of the wavelet transform for feature extraction, but the use of neural networks for the decision process. Because neural networks require large sample databases in order to be trained on the vibrations' possible signatures, the results are very specific to a particular disorder and set of physiological conditions. The goal of vibrometry, however, is to use *a priori* information to develop a detector that applies to a wide range of transient vibrations without the need for prohibitive training.

In detection theory the detection decision is made based on the value of a quantity called the test statistic. This is a result of some manipulation of the output of the feature extractor. The specific manipulation chosen will dictate the distribution of the test statistic values, which is used to determine the detection threshold. That is, the computed value of the test statistic for a given time series is compared to some predetermined threshold, and the detector makes its decision based on the results of this comparison. The difference between various detectors lies primarily in the choice of the test statistic.

---

* See glossary.

Test statistics for detecting signals in noise are found in one of two ways: with knowledge of only the noise probability density function, or with knowledge of the probability density function of both the noise and the desired signal. Detections based on the noise density alone are called non-parametric hypothesis tests. Detectors based on the densities of both the noise and the signal include: the Bayesian detector, the generalized likelihood ratio test, and others. These are described in Kay (1998). Since a concise model of arterial vibrations is not available, the probability density function for the signal is not known. Therefore, non-parametric hypothesis testing is the appropriate detection method.

Non-parametric binary hypothesis testing involves choosing one description that the data is expected to take and determining whether or not that description matches the data to some specified level of confidence. More specifically, a test statistic is formed from the data based on the null hypothesis that only noise is present (H_0). A threshold is then chosen based on the desired level of confidence. If the test statistic exceeds the threshold, then the null hypothesis is rejected and the alternative hypothesis (H_1) is accepted. This results in a signal being detected. Otherwise, the null hypothesis is accepted, and the signal is not detected. For the case in which the parameters are fully known *a priori*, a detection threshold can be determined to assure a fixed, selected false alarm rate. These detectors are given the name CFAR (constant false alarm rate). If some of the parameters of the model for the null hypothesis are not fully known, then they can be estimated. The two key issues in hypothesis testing are: finding an appropriate test statistic, and setting the threshold level.

Since the vibrations are roughly sinusoidal, matched filter theory indicates that a tool employing sinusoidal functions should be a good choice for feature extraction. Due to the fact that in the vibrations of interest the number of sinusoids as well their frequencies and amplitudes are unknown, the non-parametric time-frequency methods should be most useful. For the reasons stated in section 5.1, the non-parametric methods can be reduced to the Fourier and wavelet transforms. Fourier techniques will be discussed here, and wavelet detection will follow.

Other methods are discussed in the literature for detecting sinusoids in noise: the Kalman filter, non-linear least squares fits, and statistical inverse theory. These rely heavily on the availability of a good model for the data under test and would require intense computation for the case in which the number of sinusoids and their parameters are unknown and time-varying.

5.2.1 Fourier detection methods

Many sinusoidal detection methods are based on the discrete Fourier transform (DFT) in the form of the periodogram, the magnitude squared of the DFT of a time series tapered by a rectangular window. Test statistics for these problems typically involve some variant of the magnitude of the ordinates (magnitude values) of the periodogram. Several such detectors are discussed in this section and are summarized in Table 5.1. None of these methods handles all of the challenges of vibrometry, but they illustrate the state of the art in Fourier techniques. Portions of these techniques are applied to the wavelet transform in the design of the detector used for vibrometry (section 6.2).

Table 5.1. Fourier-based sinusoid detection methods.

Decision Method		**Number of stationary sinusoids**	**Able to handle clutter?**	**Noise type** --all require stationarity	**Able to handle unknown variance?**
Maximum likelihood (periodogram) Kay (1998)	**GLRT (generalized likelihood ratio test)**	1	No	Additive white Gaussian	No
Adaptive threshold (CFAR - radar) Ritcey and Holm (1992) El Mashade (1998)	**Mean level detector**	Multiple	No	Colored	Yes
	Order statistic detectors	Multiple	Yes (if low dynamic range)	Colored	Yes
Binary hypothesis (periodogram) Percival and Walden (1993)	**Fisher's test for periodicity**	1	No	Additive white Gaussian	Yes
	Siegel's test	Multiple	No	Additive white Gaussian	Yes
	Thomson's multi-taper test	Multiple	Yes (if low dynamic range)	Colored	Yes

5.2.1.1 Single time series

Consider the specific case of a single stationary sinusoid with unknown frequency, amplitude and phase, in zero mean additive white Gaussian noise with known variance. Kay (1998) describes a detection method for this case which uses the generalized likelihood ratio test (GLRT). The GLRT makes the detection decision based on a known model for the desired signal, a sinusoid. The method estimates the sinusoid's parameters from the periodogram, and then determines whether the data is a better fit to the noise model or the signal model.

The other detectors are based on non-parametric binary hypothesis testing. Percival and Walden (1993) address the same simple case of a single stationary sinusoid with

unknown parameters in additive white Gaussian noise. This method makes use of the fact that the ordinates of the periodogram of the sampled signal are chi-square distributed with two degrees of freedom (also called the exponential distribution). If the variance of the noise (σ_n^2) is known, then detection is made (rejecting the null hypothesis that the data is all noise) when an ordinate of the periodogram exceeds a threshold value that indicates that the data no longer fits the chi-square distribution hypothesis. The threshold value is set by the chosen level of the test as it applies to the chi-square distribution.

Test statistic: peak value of the periodogram

H_0: periodogram ordinates are χ_2^2 distributed

H_1: periodogram ordinates are not χ_2^2 distributed

For the case when the variance of the noise is not known, Fisher derived a test statistic to be used with the periodogram, when the alternative to the null hypothesis is that a single frequency is present. His statistic compares the peak of the periodogram normalized by the sum of periodogram ordinates to a threshold value set by the chosen level of the test.

Test statistic: $$\frac{periodogram \quad peak}{\sum (periodogram \quad ordinates)}$$

H_0: periodogram ordinates are χ_2^2 distributed

H_1: periodogram ordinates are not χ_2^2 distributed

Siegel (1979) derived an expansion to Fisher's test for periodicity that applies when multiple frequencies are present. In this statistic, a sum is formed of all values of the normalized periodogram that exceed a specific fraction of Fisher's threshold.

Test statistic: $$\frac{\sum (periodogram \quad peak) > threshold}{\sum (periodogram \quad ordinates)}$$

H_0: periodogram ordinates are χ^2_2 distributed

H_1: periodogram ordinates are not χ^2_2 distributed

For the case of the stationary sinusoid with zero-mean colored Gaussian noise*, Thomson's multi-taper test for periodicity can be applied (Thomson 1982, Percival and Walden 1993) even if the noise variance is unknown. Thomson's method uses the multi-taper power spectrum estimate mentioned earlier in section 5.1.2.1. He normalizes the spectrum ordinates by their variance to form a ratio of two chi-square variables. The resulting ratio has an F distribution. Detection is then based on a threshold value taken from a chosen level of confidence on the F-distribution. The method requires two additional constraints: the frequency spectrum of the noise is slowly varying in the region of the signal, and the signal frequency is not too close either to zero or the Nyquist frequency. The first constraint would prevent detection when clutter frequencies approach vibration frequencies. Thomson's method applies to multiple sinusoids if they are not closely spaced.

Test statistic:

$$\frac{\dfrac{|\text{multitaper estimate of sinusoid amplitude}|^2}{\text{variance(multitaper estimate of sinusoid amplitude)}}}{(\text{number of multitapers})\dfrac{\text{variance(estimate of noise multitaper spectrum)}}{\text{variance of noise multitaper spectrum}}}$$

H_0: test statistic is F distributed

H_1: test statistic is not F distributed

* See glossary.

Fisher's test for periodicity, which only applies in the case of white noise, and Thomson's test for periodocity in colored noise, remain dependent on direct spectral estimates (such as the periodogram) which are highly biased against the frequencies of interest when strong, low frequency clutter is present (section 5.1.2.1). Thus clutter filtering must be applied before the detector. Care must be taken to ensure that the clutter is sufficiently attenuated, but potential vibrations at the lower frequency range are not affected. These methods also require the arbitrary segmenting of the time series in order to approximate stationarity.

5.2.1.2 Dual time series

If a signal is expected to arise coherently in two time series, then detection will likely improve by testing both time series simultaneously. One such detection method examines the coherence between two stationary time series at a given frequency (Percival 1994). This method computes the multi-taper spectrum of each time series individually followed by their multi-taper cross-spectrum. The coherence is computed as the squared magnitude of the cross-spectrum divided by the product of the individual spectra. Making use of the fact that this coherence estimate follows an F-distribution for Gaussian white noise, Percival shows that the null hypothesis can be rejected at level r when the coherence measure exceeds $1-r^{1/(d-1)}$, where d represents the number of tapers used.

5.2.1.3 Adaptive thresholding in a colored spectrum

Due to the presence of clutter and possibly colored noise, arterial vibrations exist in a colored spectrum, not a white one. The radar community is commonly concerned with target detection amidst clutter and noise. Specifically, their targets appear as peaks in a frequency spectrum which is contaminated by noise and clutter of unknown statistics. Target detection is carried out by employing an adaptive threshold that is set based on estimates of the noise. Many schemes have been developed for estimating this noise as it changes with frequency across the spectrum. Some examples can be found in Ritcey and Holm (1992) and El Mashade (1998). A simple technique, the mean-level detector, estimates the noise level as an average of the

spectrum ordinates over a region of an FFT spectrum centered about the point under test, M, but excluding the test points. For example, the mean level detector noise estimate for point M of Figure 5.8 would be the mean of the ordinates in the frequency region demarcated by the thick black lines. Detection is then determined based on the relative value of M compared to the estimated noise. This method suffers when the data includes either clutter or multiple peaks in the spectrum which represent multiple targets. Therefore, outlier resistant methods were developed. One such method, called the order statistic detector, sorts the ordinates in a region of the spectrum. The noise is then estimated by the k^{th} smallest ordinate. In Figure 5.8 this would mean that the noise estimate for the point M would be the k^{th} smallest FFT magnitude located in the frequency range demarcated by the thick black lines. The value of k must be determined adaptively or from prior information about the expected spectrum. Variations of this, called censoring, include combining multiple small ordinates to form the noise estimate. Further clutter resistance is achieved when two measures of the local noise are computed and the greater of the two is used. For Figure 5.8 this implies that the noise estimate for M would be based on the greater of the mean of the bolded region to the left of point M and the mean of the bolded region to the right of point M. Once a noise estimate has been made, the threshold is set according to knowledge of the distribution of the data.

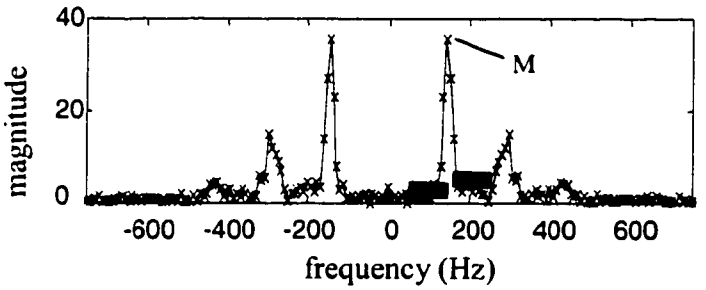

Figure 5.8. Illustration of Fourier adaptive thresholding.
The fast Fourier transform of a time series (-x), a test point (M) and two frequency segments for noise estimation (thick, horizontal, black lines).

5.2.2 Wavelet detection methods

Alongside the rising popularity of the wavelet transform is the wavelet-based detector. As with the wavelet transform, the wavelet-based detector is particularly useful in detecting non-stationary signals that in some way resemble the mother wavelet, such as in duration, bandwidth or shape. The wavelet detector's primary role is in detecting signals that push conventional parametric and Fourier-based detectors beyond their capabilities. Specifically, they are most useful for detection when the expected signal is only roughly known, the noise variance is unknown, the noise is colored, the noise is non-stationary, or some subset of these. This occurs in areas as diverse as biological signals (Englehart et al. 1999, Akin and Sun 1999), biological imaging (Chen et al. 1998), fluid dynamics (Kaspersen and Hudgins 1996), speech (Janer 1998), and underwater acoustics (Campbell 1996, Bailey et al. 1998).

5.2.2.1 Single time series

Wavelet detectors employ all wavelet types and many detection methods. Discussion of wavelet detection with the various methods, for example, can be found in the following: GLRT (generalized likelihood ratio test) (Frisch and Messer 1994, Gallaire and Sayeed 1998), CFAR (constant false alarm rate) (Chen et al. 1998), matched filters (Priebe and Baugh 1994), adaptive filtering in time-varying noise (Basbug 1993) and many others. A few wavelet detectors in the literature are relevant to arterial vibration detection. That is, these detectors apply when both the signal and the stationary noise variance are unknown, as well as when the noise is colored. Specifically, they use the one detection theory method that does not require a signal model: non-parametric binary hypothesis testing for noise only. Table 5.2 summarizes the data types for which the various methods are applicable.

Table 5.2. Comparison of relevant wavelet detectors from the literature.

	Requires Gaussian Noise?	**Accepts non-stationary noise?**	**Accepts colored noise?**	**Accepts unknown noise variance?**
Campbell (1996)	Yes	No	No	No
Torrence and Compo (1998)	Yes	No	Yes, if spectrum known *a priori*	Yes
Ravier and Amblard (1998)	Yes	No	Yes	No
Bailey et al. (1998)	No	Yes, if slowly varying	Yes	Yes
Donoho and Johnstone (1994)	Yes	No	No	No
Johnstone and Silverman (1996)	Yes	No	Yes	Yes

The simplest of the techniques was developed by Campbell (1996). He employed non-parametric binary hypothesis testing to oscillating signals in additive white Gaussian noise with the assumption that the noise variance was known. Torrence and Compo (1998) extended this scenario to include colored stationary Gaussian noise in which the noise variance is estimated with a large sample size and the noise spectrum is known. This result is implicit in their results, not explicit because their aim was not detection, but to develop confidence intervals for time-frequency analysis. In both papers, the detection decision is based on a chi-square null hypothesis of the wavelet power spectra (magnitude squared of the wavelet coefficients) with the thresholds set to a certain level of confidence taken from a chi-square distribution.

Test statistic: $|\text{wavelet coefficient}|^2$

H_0: χ_2^2 distributed

H_1: not χ_2^2 distributed

Ravier and Amblard (1998) followed a similar route in basing their detection decision on the binary hypothesis with stationary Gaussian noise as the null hypothesis, where

the noise variance is known. Their test statistic, however, was the fourth order cumulant (see Papoulis 1991) which should be approximately zero for a Gaussian random variable.

Test statistic: fourth order cumulant of wavelet coefficient

H_0: Gaussian distributed

H_1: not Gaussian distributed

Bailey et al. (1998) attacked the problem of colored non-Gaussian noise, as well as an unknown signal. They use a long noise-only segment of data to compute a density estimate of the noise at individual scales of the discrete wavelet transform. The threshold is chosen at a certain level of confidence of the estimated density.

Test statistic: $mean\left(\sum\left|wavelet\ \ coefficients\right|^2\right)$

H_0: the estimated noise probability density function

H_1: not fitting the estimated noise probability density function

For cases when the noise is stationary and can be estimated, or when there are sufficiently long, known, noise only segments, the methods just listed have been shown to be effective.

The popular "de-noising" use of wavelets can also be considered a form of non-parametric binary hypothesis detection. Donoho and Johnstone pioneered this in what they called "wavelet shrinkage" in 1994. This method applies when the noise is white and stationary with known variance (σ^2). They suggest thresholding the magnitude squared of the wavelet coefficients at $\sigma\sqrt{2\log_e(N)}$, in which N is the length of the time series. Johnstone and Silverman (1997) inform us that this will yield a false alarm rate of approximately 10%. Donoho and Johnstone extend de-noising to the

case in which the noise variance is unknown: the noise variance can be estimated by a scalar multiplied by the median absolute deviation (MAD) from zero across wavelet coefficients at scale 1. The scale factor is 1/0.6745 for a Gaussian distribution. While a false alarm rate of 10% is reasonable for de-noising, it is unacceptably high for arterial vibrometry.

To accommodate colored noise, wavelet detectors can estimate noise on a scale by scale (i.e., frequency range by frequency range) basis. Bailey's method, described above, carries this out implicitly; whereas, scale by scale noise estimates are specifically described in Johnstone and Silverman (1996) in which they use the phrase "level dependent thresholding." The threshold $\sigma\sqrt{2\log_e(N)}$ applies in this situation if the σ value is scale specific. Thus, the noise variance should be estimated at each scale again using the median absolute deviation (MAD) method.

Test statistic: |wavelet coefficient|

H_0: Gaussian distributed

H_1: not Gaussian distributed

5.2.2.2 Dual time series

As mentioned in the Fourier detection section, jointly using two time series with coherent signals can improve detection. Joint wavelet transforms have received some attention in wavelet literature: the wavelet cross spectrum, cross wavelet spectrum, wavelet cross correlation, and wavelet coherence. These are described below and summarized in Table 5.3.

Table 5.3. Wavelet spectra of two time series.

$E\{W_1(t,s)\, W_2^*(t,s)\}$	$W_1(t,s)\, W_2^*(t,s)$	$\frac{E\{W_1(t,s)W_2^*(t,s)\}}{\sqrt{E\{\lvert W_1(t,s)\rvert^2\}E\{\lvert W_2(t,s)\rvert^2\}}}$	$\frac{\lvert W_1(t,s)W_2^*(t,s)\rvert^2}{\lvert W_1(t,s)\rvert^2 \lvert W_2(t,s)\rvert^2}$
Wavelet cross spectrum	Cross wavelet spectrum	Wavelet coherence (A)	Wavelet coherence (B)
Wavelet covariance	Wavelet cross transform	Wavelet cross correlation	
Wavelet cross covariance at lag 0			

5.2.2.2.1 Wavelet cross-spectrum and cross wavelet spectrum

The wavelet cross spectrum is defined as the expected value over time of the product of the wavelet transform of one time series with the conjugate of the wavelet transform of a second time series, in both discrete and continuous time. This can be expressed as:

$$\text{Wavelet cross spectrum (s)} = E\{\, W_1(t,s)\, W_2^*(t,s)\,\}, \qquad (5\text{-}8)$$

where $E\{\}$ stands for the expected value over time, $W_1(t,s)$ and $W_2(t,s)$ are the wavelet transforms of a two different time series, t represents discrete time values, s represents discrete scale values, and ' * ' represents the complex conjugate operator. The result yields one value for each wavelet scale. Because the discussion relates to both the time and frequency domains simultaneously, and the concepts are young, the terminology is not yet unified. The wavelet cross spectrum is referred to by various authors as the wavelet covariance and the wavelet cross-covariance at lag zero (Kumar and Foufoula-Georgiou 1994, Whitcher 1998, Kaspersen and Hudgins 1996, Serroukh and Walden 1998).

The cross wavelet spectrum, on the other hand, is defined in Liu (1994) and referred to in Torrence and Compo (1998) as the product of the wavelet transform of one time series with the conjugate of the wavelet transform of a second time series:

$$\text{Cross wavelet spectrum } (t,s) = W_1(t,s)\, W_2^{*}(t,s). \qquad (5\text{-}9)$$

Hence, these researchers are concerned with the same product as the cross wavelet spectrum without the expectation operator. The cross wavelet spectrum varies in both time and scale since there is no time average. Confusingly, Kaspersen and Hudgins call this same expression the wavelet cross transform.

Torrence and Compo suggest determining the probability density function of the cross wavelet spectrum in order to state the confidence level of signals analyzed with this method. With this suggestion, they have offered a test statistic for binary hypothesis testing. This suggestion forms the basis for dual depth vibrometry discussed in section 6.2.2.

5.2.2.2.2 Wavelet coherence and wavelet cross correlation

Further complicating the terminology, Kaspersen and Hudgin's wavelet coherence function is the continuous time equivalent of Whitcher's wavelet cross correlation: the wavelet cross spectrum divided by the square root of the product of the wavelet power spectra of the two time series:

$$\text{Definition A: wavelet coherence } (s) = \frac{E\{W_1(t,s)W_2^{*}(t,s)\}}{\sqrt{E\{|W_1(t,s)|^2\}E\{|W_2(t,s)|^2\}}}. \qquad (5\text{-}10)$$

Here, the wavelet power spectrum, referred to as the wavelet variance by Whitcher, is defined as the expected value over time of the magnitude squared of the wavelet transform at a given scale.

Liu (1994), followed by Torrence and Compo (1998), define wavelet coherence, however, as the magnitude squared of the cross wavelet spectrum divided by the individual power spectra before application of the expected value operator:

$$\text{Definition B: wavelet coherence } (t,s) = \frac{\left|W_1(t,s)W_2^{*}(t,s)\right|^2}{\left|W_1(t,s)\right|^2\left|W_2(t,s)\right|^2}. \qquad (5\text{-}11)$$

Once again, this result varies in both time and scale, since there is no time average. This definition of coherence, however, is uniquely one at every time and scale, rendering it useless. Some time or scale smoothing must be applied, yet smoothing across the entire time series, such as in definition A, loses the time varying advantage of the wavelet transform. Wavelet coherence has potential to be the foundation for a useful detector when the issue of smoothing is settled.

5.2.3 Specific signal processing challenges

Although detection theory is fairly mature, detection of signals such as the vibrations induced by vascular disorder continue to remain a challenge. Specifically, signals with various combinations of the following characteristics are beyond the limits of the detectors found in the literature and even more so when all are present. The following are all present in arterial vibrations:

- Both broadband and narrowband
- Unknown frequencies, bandwidths, durations, amplitudes, onsets
- Time varying frequencies, bandwidths, amplitudes
- Short duration
- Strong, low frequency clutter
- Non-stationary noise
- Colored noise
- Narrowband noise

5.3 Estimation overview

The goal in vibration estimation is to quantify the vibration characteristics apart from the noise and clutter. As with detection, estimation, in general, can be carried out with either parametric or non-parametric methods. Again, since models of the various vibrations are not available, non-parametric methods must be used. Fourier analysis

has long been the standard means of estimating characteristics such as amplitude, frequency, phase, duration, and arrival time. Methods that rely on Fourier techniques, however, require that the signal is stationary throughout the time window used in the transform. Although this restriction is eliminated with wavelet analysis, parameter estimation with wavelets has received little attention in the literature. The FFT, CWT and DWT approaches will be discussed in the following sections. Serendipitously, the CWT, which has already been shown to be appropriate for detection, is also the appropriate method for estimation. The FFT and DWT approaches are included here for comparison and completeness. Estimation with the CWT will be discussed further in section 6.3.

5.3.1 Fourier analysis

5.3.1.1 Frequency

The frequency is estimated simply as the frequency bin associated with the maximum value of the periodogram. For a single sinusoid, the accuracy of this estimate is on the order of $\frac{1}{N(\Delta t)}$ where N is the number of samples used in the FFT computation (Percival and Walden 1993, Kay 1988). In other words, the frequency estimate falls on the order of the frequency spacing of the spectrum ordinates.

5.3.1.2 Amplitude

The amplitude of a signal at a given frequency bin is computed as the absolute value of the FFT ordinate multiplied by 2/N. For a single sinusoid, this estimation is accurate on the order of 1/N in the same units as the original time series (Percival and Walden, 1993, Kay 1988).

For example, at a sampling rate of 1600 Hz (without zero padding*) a 32 point FFT of a 200 Hz vibration with a constant amplitude of 300 nm would have a magnitude of 4800 at the 5th frequency bin. The amplitude estimate, then, would be (2/32)(4800) which is the correct value of 300 nm. If the vibration frequency were slightly more or

---

* See glossary.

less than 200 Hz so that it did not lie at an exact Fourier frequency bin (i.e. $f_v \neq i\ F_s/N$, for some integer i), then the FFT ordinate would be smaller than 4800. The resulting amplitude estimate would be smaller than 300 nm by an amount on the order of 300/32 nm. Note that the phrase "on the order of" implies that the discrepancy could be some multiple (possibly fairly large) of 300/32 nm. Thus the estimate could be 206.25 nm (300-10*300/32) or even smaller.

5.3.1.3 Phase

The angle of the Fourier cross spectrum provides an estimate of the phase between two stationary time series at the various frequency bins. Its accuracy is not discussed in the literature.

5.3.1.4 Duration

A signal's duration can be estimated with the use of the short time Fourier transform. The duration estimate would simply be the number of time windows for which the signal appeared in the FFT adjusted for the overlap of the windows. The accuracy of this measure would depend on the window length and the amount of overlap between windows.

5.3.1.5 Power (Intensity)

Transient vibrations have a time varying intensity ($I_{inst}(t)$) value which is proportional to the instantaneous power (P(t)):

$$I_{inst}(t) = P(t) / \text{area}. \qquad (5\text{-}12)$$

It is this instantaneous intensity measure that should be estimated for vibrometry. Because the FFT yields one result for an entire time segment, though, it does not offer instantaneous power estimates, P(t). Rather, it provides only time average power estimates associated with each frequency bin. These estimates are computed as the magnitude squared of the FFT ordinates. The sum of the power estimates multiplied by the frequency spacing will exactly equal the energy of the time series over the time

interval included in the FFT calculation, according to Parseval's theorem. The accuracy of the power estimates at a given frequency bin, however, will be affected by the high frequency bias issue mentioned in section 5.1.2.1.

5.3.2 Wavelet analysis

Wavelets are widely used for determining the frequency content and duration of transient, oscillating signals. In addition, they are often used to quantify the variance (power) of stationary signals. However, their utility in quantifying a transient signal's amplitude, phase, and intensity are implied in the literature; yet, rarely explicitly stated. As the following sections indicate, the CWT is more accurate than the DWT for frequency and power estimates, and requires fewer computations for amplitude estimates. Thus, the CWT is the method of choice.

5.3.2.1 Frequency

Frequency estimates of a signal are based on the wavelet scales in which significant magnitudes appear in the wavelet coefficients (e.g., see Figure 5.1 of section 5.1.2.2). The accuracy depends on the frequency spacing of the scales. Due to their orthogonality, DWT wavelets are dyadic, meaning that their center frequencies are spaced by a factor of 2 (e.g. 4000 Hz, 2000 Hz, 1000 Hz, etc.). This results in poor frequency resolution. For the CWT, the frequency spacing can be specifically selected as desired.

5.3.2.2 Amplitude

If its mother wavelet meets the admissibility requirement (section 5.1.2.2.1), then a wavelet transform's coefficients can be combined to reconstruct the original time series (if the DC component is accounted for in some manner). This reconstruction, however, is different for continuous and discrete wavelets. Amplitude estimates are made as a subset of the reconstruction process, and thus also differ for the CWT and the DWT.

5.3.2.2.1 CWT

With the CWT, reconstruction can take place directly in the wavelet domain (Torrence and Compo 1998). Consider the time series x(t) of Figure 5.1, a tapered sinusoid with

a duration longer than the wavelet. If the wavelet is normalized for amplitude estimation, as in the figure, then the amplitude of x(t) can be estimated by the magnitude of the wavelet coefficient at the time location and scale corresponding to the sinusoid provided that the wavelet center frequency is exactly the frequency of the sinusoid (see Appendix D for more details). Figure 5.1 illustrates this in that the peak wavelet coefficient magnitudes of the right column are nearly 1, which are the peak amplitudes of the two sinusoidal components of x(t).

If the vibration is not simply a tapered sinusoid, the amplitude estimate is a sum of the wavelet coefficients across the scales at which the vibration is present, multiplied by a correction factor, C_δ:

$$A_v(t) = \sum_s C_\delta W(t,s), \qquad (5\text{-}13)$$

where $A_v(t)$ is the time varying amplitude of the vibration, s represents the wavelet scales of interest, and t represents the time dimension. The accuracy is affected by the degree to which the true vibration fits within the selected wavelet scales and the frequency spacing of the scales. These concepts, along with a method for determining C_δ for the complex Morlet wavelet, are discussed later in section 6.3.1.1.

5.3.2.2.2 DWT

Unlike the CWT, the DWT is an orthogonal transform which dictates that it can be directly inverted. A vibration amplitude estimate, can therefore be made by inverting only those wavelet scales (frequency ranges) at which the vibration resides. This is accomplished in the following manner. 1) The wavelet transform of the signal is computed. 2) Resulting wavelet coefficients located at the times and scales outside those of interest are set to zero. 3) The inverse transform is computed on the resulting wavelet coefficients. The resulting time domain series is an estimate of the signal. The accuracy is affected by the degree to which the true vibration fits within the frequency range of the selected wavelet scales.

[This method is analogous to determining a sinusoid's amplitude by computing its FFT, zeroing all frequency bins other than the one corresponding to the frequency of the sinusoid, and then computing the inverse Fourier transform. The amplitude of the resulting signal would be the estimate.]

5.3.2.3 Relative phase and time delay

The angle, φ, of the cross wavelet spectrum at the wavelet scale of interest provides an estimate of the time varying phase between two time series (Torrence and Compo 1998). The time delay between vibrations from the individual wavelet power spectra can be computed as the phase values divided by the wavelet center frequencies according to:

$$\text{time\_delay}(t, s) = \frac{\left(\varphi / 2\pi\right)}{f_0}, \qquad (5\text{-}14)$$

where f_0 is the center frequency of the wavelet scale.

5.3.2.4 Duration

A signal's duration is estimated by the number of consecutive wavelet coefficients with significant magnitude in the time dimension. Again, this is evident from the right column in Figure 5.1. The minimum estimate will be affected by the duration of the wavelet itself.

5.3.2.5 Power (intensity)

Like the amplitude estimate, power estimation differs for the DWT and CWT. In both cases the original signal can be reconstructed from the wavelet coefficients, which implies that the signal's energy is conserved in the transformation. The goal in the power estimate (which is related to intensity according to equation (5-12), however, is to obtain instantaneous measures of the power, not just the entire energy of the signal.

5.3.2.5.1 CWT

With the CWT, instantaneous power estimates are made in a manner similar to the amplitude estimates. They can be made from the sum of the magnitudes squared of the CWT coefficients multiplied by a correction factor D_δ:

$$P(t) = \sum_s D_\delta |W(t, s)|^2 . \qquad (5\text{-}15)$$

The method for determining D_δ for the complex Morlet wavelet at its various parameters is discussed later in section 6.3.1.3. As with amplitude estimation, the accuracy is affected by the degree to which the true vibration fits within the selected wavelet scales and the frequency spacing of the scales.

5.3.2.5.2 DWT

The instantaneous power corresponding to the frequency range of one DWT wavelet scale can be estimated by the square of the wavelet coefficients at that scale. Computing instantaneous power estimates for signals spanning multiple, consecutive wavelet scales is problematic, however. Due to the orthogonality of the DWT, in time as well as in frequency, wavelet coefficients at one scale do not align temporally with wavelet coefficients at neighboring scales. This prevents their direct combination which is required for computing an instantaneous power estimate. Localized average power estimates, though, could be obtained by averaging instantaneous power estimates at each scale of interest over some time segment, followed by summing the average power estimates across the scales.

[There is a modification of the DWT which introduces redundancy in the time domain: the MODWT (Percival and Walden 2000). The MODWT coefficients do align in time and can be used for instantaneous power estimation at the expense of increased computational cost.]

6 The detector and estimator

As alluded to earlier, the goal of the detector is to find vibrations amidst clutter and noise by exploiting the known characteristics of the vibrations, clutter and noise that were discussed in section 4.5 above. The vibrations are known to be transient oscillations generally with durations less than ~150 ms and frequencies between approximately 100 and 1000 Hz. Thus, the vibrations are described by their time-frequency characteristics which can be extracted best via a continuous wavelet transform with the complex Morlet wavelet (section 5.1). Because the noise is known to be colored Gaussian, but a concise model of the vibrations is not known, non-parametric binary hypothesis testing is the appropriate detection framework. [The case of narrowband noise will be revisited in section 7.1.4.] Since the noise can be time varying, adaptive thresholding is needed for detection. The vibrometry detector described in this chapter is a combination of 1) an extension to the wavelet detectors that use the chi-square null hypothesis, and 2) a wavelet variation of the Fourier adaptive thresholding techniques. Subsequent to detection, the vibration's parameters can be estimated with the continuous wavelet transform. Each step of the detection and estimation process is described in detail in this chapter. Appendix G offers a discussion of the computational cost of the detector and estimator, along with a comparison to the computational cost of the conventional Fisher's test for periodicity in white noise.

6.1 Feature extraction via the continuous wavelet transform with the complex Morlet Wavelet

6.1.1 Reasons to use the continuous wavelet transform with the Morlet wavelet

The chosen vibration detector is based on feature extraction from the continuous wavelet transform with the Morlet mother wavelet for many reasons.

1. When used in a complete filter bank, Morlet's Gaussian shaped sinusoid acts as a matched filter to condition transient sinusoidal signals in the data.

2. The ability to alter the quality factor of the Morlet wavelet allows this wavelet to approximately match broadband and narrowband vibrations.
3. The wavelet's limited duration poises it to condition the short duration vibration better than its Fourier counterpart.
4. As a non-parametric method, the continuous wavelet transform can be applied with minimal *a priori* information about the vibration.
5. The bandpass nature of the continuous wavelet transform, and the sharp frequency edges of the Morlet wavelet, automatically and superbly perform clutter filtering.
6. As exploited by Johnstone and Silverman in their "denoising" technique, the filter bank approach of the wavelet transform allows for noise estimates at each frequency range (section 5.2.2.1). Thus, colored noise is easily handled. We will see that narrowband noise is also easily neglected. Furthermore, because noise estimates can be made in the joint time-frequency domain, as opposed to just the frequency domain, they can be made on many samples to provide good estimates. The estimates also can be made adaptive to accommodate non-stationary noise.

6.1.2 Methods of transform implementation

To implement the continuous wavelet transform with the complex Morlet wavelet, the wavelets at each scale are first computed from equation (5-5). Each scale will have a different value of ω_0, the center frequency of the wavelet scale. The values of ω_0 are selected for a specific degree of overlap between successive scales. Suppose the spacing between scales is based on a factor, *scalebase*, which is defined by:

$$scalebase = \frac{1 + \frac{1}{2Q_w}}{1 - \frac{1}{2Q_w}}, \qquad (6\text{-}1)$$

where Q_w was defined in equation (5-7). Then, setting

$$\omega_0(s+1) = \omega_0(s) \ (scalebase), \qquad (6\text{-}2)$$

where s is an integer scale number, introduces an overlap in the frequency domain at approximately -8.7 dB (Figure 6.1). Use of a value for *scalebase* that is smaller than that prescribed by equation (6-1) requires more scales to cover the same frequency range, but the scales would overlap with less amplitude dropout between center frequencies. Although most discussion will focus only on positive values of ω_0 for brevity, the wavelet implementation should include negative values of ω_0 whose absolute value equals that of the positive ω_0's. The negative values are required because the quadrature time series are directional. The signal differences between the positive and negative frequencies hold valuable information about the direction of tissue motion.

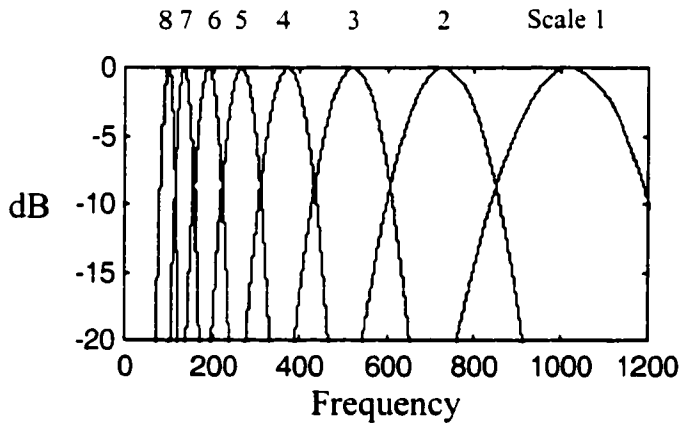

Figure 6.1. A family of Morlet wavelets in the frequency domain. The scale overlap is set at –8.7 dB.

The wavelet coefficients can be computed in several different ways. The first two are common methods: direct implementation with convolution, and indirect implementation with the FFT. The third method, which shall be called the dot product method, exploits the redundancy of the CWT to yield what are equivalently sub-samples of the wavelet coefficients that would have been produced by the convolution

method. The following sections describe these methods. Table 6.1 summarizes their computational costs for one wavelet scale with an arbitrary wavelet length of 100. Notice that, for this example at *tlen* of 8,092, the dot product method requires 20 times fewer computations than the convolution method and 3 times fewer than the FFT method.

Table 6.1. Computation cost for various CWT implementation methods. All values have been computed with: *wlen* = 100.

	tlen	**Convolution**	**FFT**	**Dot product (20% overlap)**
Complex multiplications	8,092	799,200	114,688	40,060
	32,668	3,256,800	524,288	162,940
Complex additions	8,092	791,208	106,496	39,659
	32,668	3,224,232	491,520	161,311
Cost increases with:		*tlen*	*tlen* $\log_2(tlen)$	*tlen*

6.1.2.1 Wavelet implementation with convolution

The most straightforward means of computing the continuous wavelet transform is to simply convolve the wavelet at each scale with the data time series, x(t):

$$W_x(t,s) = \sum_{l=-wlen/2}^{(wlen/2)-1} w(t-l,s)x(l), \quad t = wlen/2 : 1 : tlen - wlen/2, \qquad (6\text{-}3)$$

where w(t) is defined in equation (5-5); *tlen* is the length of x(t); *wlen* is the scale dependent length of w(t); and Δt is the time between data samples. Note that the values of t begin at *wlen/2* and extend to *tlen-wlen/2*. Due to edge effects, wavelet coefficients for the first *wlen/2* and last *wlen*/2 would be inaccurate and are therefore not computed. This summation requires the following number of computations at each scale:

(wlen) (# of wavelet coefficients to compute) complex multiplications and
(wlen-1) (# of wavelet coefficients to compute) complex additions.

Substituting for the number of wavelet coefficients, this becomes:

$$\begin{array}{l} (wlen)\ (tlen - wlen) \text{ complex multiplications and} \\ (wlen-1)\ (tlen - wlen) \text{ complex additions.} \end{array} \qquad (6\text{-}4)$$

This method computes one wavelet coefficient per sample in the data time series at each wavelet scale.

6.1.2.2 Wavelet implementation with the FFT

Since convolution in the time domain corresponds to multiplication in the frequency domain, convolution is commonly accomplished by transforming the two convolution arguments into the frequency domain, taking their product, and then transforming the result back to the time domain. Since the wavelet transform is known *a priori*, its frequency domain equivalent can be computed ahead of time, eliminating computation cost to the detection algorithm. Computation, therefore, includes a) computing the FFT of the data time series, [which need only occur once to be used for all scales], b) multiplying the wavelet Fourier coefficients with the data time series' Fourier coefficients, and c) computing the inverse FFT of this product. Since the length of the convolution of two time series is the sum of the lengths of the individual time series, each time series (wavelet and data) must be zero-padded to attain this longer length. Computing the FFT, or inverse FFT, of a series of length N requires $N*\log_2(N)$ complex multiplications and additions. Hence, at each scale, computation of the continuous wavelet transform will require:

$$\begin{array}{l} (tlen+wlen) + (tlen+wlen)\ \log_2(tlen+wlen) \text{ complex multiplications} \\ (tlen+wlen)\ \log_2(tlen+wlen) \text{ complex additions.} \end{array} \qquad (6\text{-}5)$$

At each wavelet scale, this method computes one wavelet coefficient per sample in the data time series. It is implicitly assumed in equation (6-5) that *tlen+wlen* is a power of two. When this is not true, the computation cost will increase.

It should be clarified that the FFT is used here as a means to accomplish convolution. The FFT is not used for time-frequency feature extraction as described in section 5.1.2.1. Here, one FFT is taken of the entire time series simply to transform it into the frequency domain to be multiplied with the frequency representation of the wavelet.

6.1.2.3 Wavelet implementation with the dot product

Due to the non-zero correlation length of the wavelet (i.e., the wavelet is not orthogonal to its own time shifts), computation of a wavelet coefficient for every sample in the data time series yields redundancy in the resulting wavelet coefficients. For example, approximately 50% of the energy in the wavelet is contained in 20% of its length. So, computing a wavelet coefficient at every fifth data sample point would still yield 50% overlap in the wavelet's energy at each wavelet coefficient. Figure 6.2 illustrates the magnitude of several complex Morlet wavelets spaced with their peaks separated by 20% of the wavelet length.

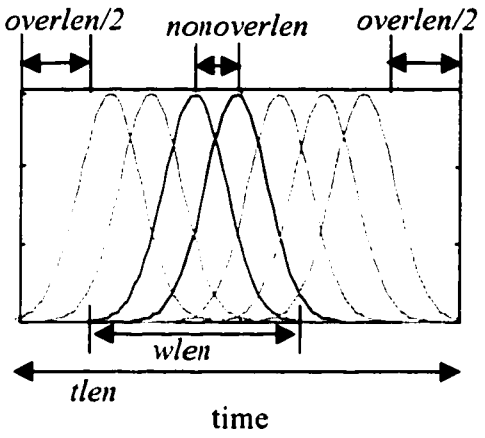

Figure 6.2. Morlet wavelet magnitudes at 20% time separation. *nonoverlen* marks the portion of a wavelet where its magnitude exceeds that of the previous and subsequent wavelets. *nonoverlen* is also the spacing between wavelet coefficients. One wavelet length, *wlen*, is equal to five times *nonoverlen* for 20% time separation. *wlen* is also equal to *nonoverlen* plus *overlen*.

For this method wavelet coefficients are computed at each scale from the dot product of the wavelet with a segment of the time series whose length is the same as the wavelet's, centered about the time location of the wavelet coefficient. The number of wavelet coefficients computed for a given time series is then:

$$\frac{\#\text{ of samples excluding those with edge effects}}{\text{chosen spacing between samples}} = \frac{tlen - overlen}{nonoverlen} \qquad (6\text{-}6)$$

where *tlen*, again, is the length of the time series, *nonoverlen* is the length of the segment of the wavelet between wavelet coefficients, and *overlen* is the difference between *wlen* and *nonoverlen*. For a given scale, the wavelet coefficients computed with a dot product require the following number of computations:

(wavelet length)(# of coefficients to compute)	complex multiplications
(wavelet length - 1)(# of coefficients to compute)	complex additions

Substituting the defined variables, the dot product method requires the following computations for each wavelet scale:

$$\begin{array}{ll} (wlen)\ (\ (tlen - overlen)/nonoverlen\ ) & \text{complex multiplications} \\ (wlen\text{-}1)\ (\ (tlen - overlen)/nonoverlen\ ) & \text{complex additions.} \end{array} \qquad (6\text{-}7)$$

Figure 6.3 illustrates the steps of the dot product method. The left column plots the data time series, x(t), in the first row; the wavelet at a specific scale, w(t), centered about time t = 0 in the second row; and the product of x(t) and w(t) in the third row. Also included in the third row is the scalar value that is the sum of the product of x(t) and w(t). This is the result of the dot product of x(t) and w(t). The bottom row shows the convolution of x(t) with w(t) as the solid curve, and the result of the dot product of x(t) and w(t) at time t = 0 as an '*'. The middle and right columns plot the same information as the first, except that the wavelets have been delayed by 20% and 40% of the wavelet length, respectively. Thus, the dot product can be considered to produce sub-samples of the convolution of x(t) with w(t).

In addition to requiring fewer computations for a given time series length than the other two methods, this method has another advantage over the FFT method. With the dot product method, as well as the convolution method, wavelet coefficients can be computed in a pipeline fashion as new data arrives, rather than requiring the full time series to arrive before computation can be performed.

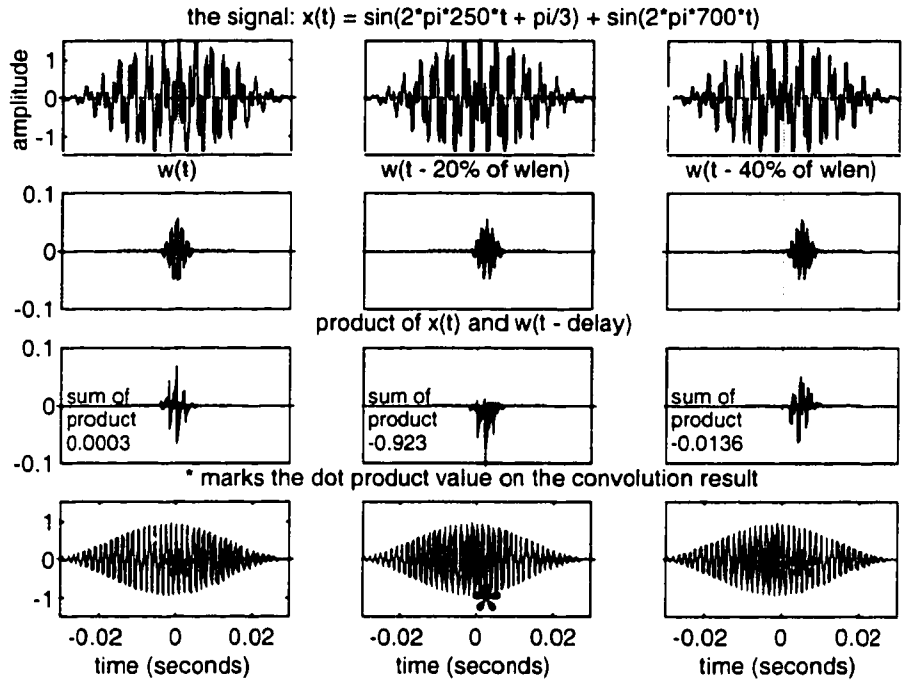

Figure 6.3. Illustration of dot product computation and results.
Upper row (row 1): The signal, x(t), is repeated identically in each column. **Row 2:** The wavelet, w(t) is shown in each column, but successively delayed in time by 20% of its length, *wlen*. **Row 3:** Each column is the product of x(t) and the wavelet of row 2. The sum of this product, which is the dot product of the wavelet with x(t), is displayed with the plots. **Row 4:** The solid line plots the convolution of x(t) with w(t). The large '*' marks the result of the dot product of x(t) with the wavelet of row 2. Thus, the '*' has the value shown in row 3 at the time location corresponding to the center of the wavelet in row 2.

6.2 Detector theory and algorithm

The role of the vibration detector is to estimate the noise level of the data and then to determine when a signal rises above the noise. To accomplish these, the Morlet wavelet transform is first applied to the raw quadrature data one depth at a time in order to condition the data. Then the noise is estimated from the transformed data. The detector further processes the data in order to use a statistical binary hypothesis test to decide the presence or absence of vibration signal.

The test statistic used in the hypothesis test can be based on a time series from one depth in tissue to determine the presence of a vibration. Alternatively, to exploit the *a priori* knowledge that a vibration will have some spatial spread, the test statistic can be formed with time series from two neighboring depths. Since the test statistic for the dual depth method is an extension of the single depth method, the latter will be discussed first.

6.2.1 Morlet wavelet method – single depth

Step 1:

Apply the continuous wavelet transform with the Morlet wavelet to the quadrature data time series from a given depth. Assume that the data consist of Gaussian noise at frequencies above the clutter range in the absence of blood flow, vibration and narrowband noise. For initial simplicity, require that the variance of the real and imaginary parts of the noise is one.

Comments:

a) The requirement for unit variance will be relaxed later.

b) The Gaussian noise assumption was validated in principal in section 4.3.1. It will be validated based on *in vivo* data later in this chapter.

c) Since the wavelet transform is a linear process, the wavelet transform of a Gaussian time series remains Gaussian.

d) If the wavelet is normalized for unit energy at each scale (equation (6-8)), and the original time series has unit variance, the resulting wavelet coefficients will also have unit variance.

$$\sum_t |w(t)|^2 = 1, \qquad w(t) = Be^{-j\omega_0 t - \alpha t^2}, \tag{6-8}$$

where B is chosen to make the left equation true.

e) Since the mean of the wavelet is zero, the series of wavelet coefficients will have zero mean.

f) Any of the three methods described in section 6.1.2 can be used for this wavelet implementation.

Step 2:

Compute the wavelet power spectrum as the magnitude squared of the wavelet coefficients.

Comments:

The resulting time series is chi-square distributed with two degrees of freedom (Kay, 1998) (Figure 6.4).

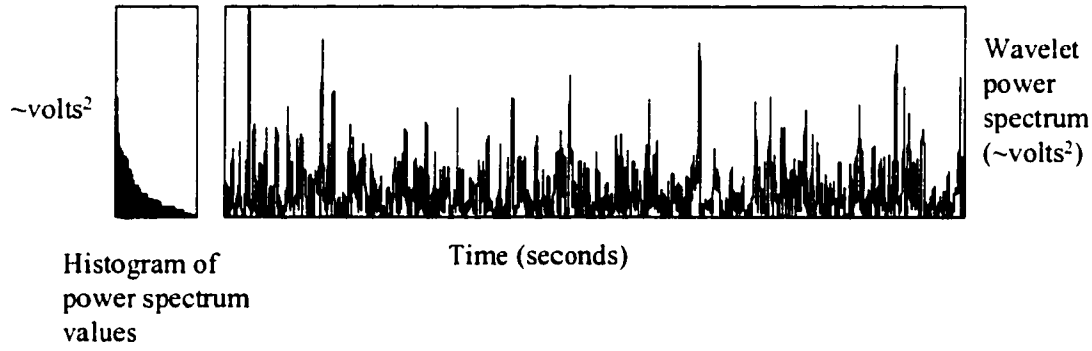

Figure 6.4. Example wavelet power spectrum and histogram.
Right: The wavelet power spectrum coefficients (the magnitude squared of the wavelet transform computed with a complex Morlet wavelet) of a stationary Gaussian time series at an arbitrary wavelet scale. **Left:** The histogram of the wavelet power spectrum coefficients.

Step 3:

At each scale, estimate the noise level at a given data sample, *R*, as follows (Figure 6.5). First, determine the maximum time duration of any expected vibration (e.g., 150 ms). Label this maximum duration *viblen*. Compute the mean of the wavelet power spectrum over the (seg*len* F_s) samples (e.g., 600 samples for F_s = 8000 Hz) beginning at time *(-seglen - viblen/2)* prior to the data sample *R*, and label the result v_L. The value of *seglen* should be greater than or equal to the length of the wavelet, *wlen*.

Now compute the mean of the wavelet power spectrum over the (*seglen* F_s) samples beginning at time *viblen*/2 after the data sample *R*, and label the result v_R. Set the noise estimate, v, equal to the minimum of v_L and v_R. If the time duration of the wavelet, *wlen*/F_s, exceeds *viblen*/2, then set *viblen* to 2*wlen*/F_s.

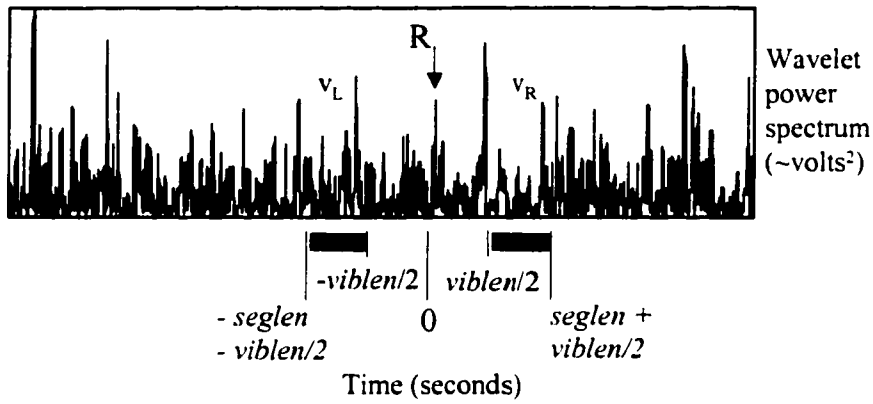

Figure 6.5. Illustration of the process of computing a noise estimate.

Intermediate noise estimates, v_L and v_R, corresponding to the test point, R, are computed as the mean of the wavelet power spectrum coefficients over the time segments marked by the thick, black, horizontal lines. The time segments must be separated by *viblen*, the assumed maximum possible vibration length. The final noise estimate, v, is the minimum of v_L and v_R.

Comments:

a) The quantities v_L and v_R are noise estimates for *R*. Because one of the two noise estimates could have encompassed a vibration signal, the minimum of the two noise estimates is used for vibration detection.

b) Increasing the value of *viblen* for long wavelets is required to ensure that the noise estimates are independent of the sample point R.

c) The value of *seglen* must be greater than or equal to *wlen* in order for the chi-square approximation to be adequate for detection.

d) The question to be answered by the detector is whether the data point R sufficiently exceeds the noise estimate v. If it does, then a vibration is detected. The solution to this question must be decided with knowledge of the statistics of the noise.

If the value of R exceeds the values expected for noise based on the known statistics, then a vibration will be decided.

If the wavelet power spectrum time series samples were independent, then each of the two variance estimates would be chi-square distributed with $2(seglen)Fs$ degrees of freedom; that is, the number of data points used in the averaging, multiplied by two since the values are complex. If this were true, the detector design could continue. However, because the wavelet power spectrum time series represents a convolution of an independent time series with the correlated wavelet, the resulting time series samples are not independent. The distribution of the noise estimates, v_L and v_R, can be approximated by a chi-square distribution, but the number of degrees of freedom must be determined. This determination relies on a concept called "equivalent degrees of freedom" (Priestley 1987, and Percival and Walden 1993). Whitcher (1998), and Percival (1995) applied this concept to estimating the wavelet variance (power spectrum). Assuming that the variance estimate follows a constant multiple of a chi-square distribution, the number of degrees of freedom (η) can be obtained from the following expression, for each wavelet scale:

$$\eta = \frac{2(seglen)F_s}{\sum_{l=-((seglen)F_s-1)}^{(seglen)F_s-1}\left(1-\frac{|l|}{(seglen)F_s}\right)\gamma_l^2}, \qquad (6\text{-}9)$$

where γ_l is the autocovariance of the wavelet at lag l. Selected values of η are listed in Appendix H, Table H.1.

Figure 6.6 plots an empirical density of one of the noise estimates, v_L or v_R, against the chi-square approximation for wavelet scale 1. The samples for the empirical density were generated with the convolution implementation of the wavelet transform. Note the excellent agreement between the two curves.

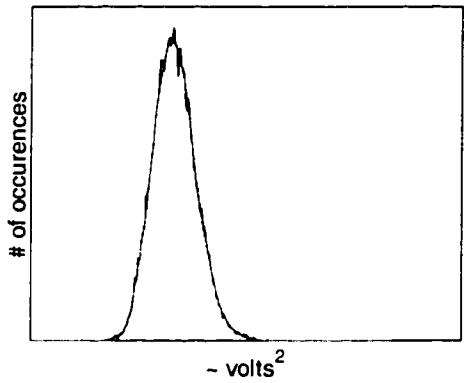

Figure 6.6. Comparison of approximate distribution and empirical density. The approximate density (~ χ_η^2) curve is smooth and the empirical density of v_L (or v_R) is jagged.

The minimum operation further alters the distribution of the final noise estimate. From Papoulis (1991), if $v = \min(v_L, v_R)$ and v_L and v_R are independent and identically distributed, then the probability density function of v is:

$$f_v(v) = 2\, f_{vL}(v)[1-F_{vL}(v)]. \qquad (6\text{-}10)$$

Therefore,

$$f_v(v) = 2\chi_\eta^2(v)\left[1 - \int_0^v \chi_\eta^2(y)dy\right]. \qquad (6\text{-}11)$$

By visual inspection of its graph (Figure 6.7), this probability density function resembles a chi-square density. Thus, this density can also be approximated by some constant, m, multiplied by a chi-square distribution with some number of degrees of freedom, ν:

$$m\chi^2_\nu. \qquad (6\text{-}12)$$

Finding the chi-square density that fits $f_v(v)$ involves matching the first and second moment of the two densities, namely the mean and variance. From Percival and Walden (1993),

$$E\{m\chi^2_\nu\} = m\nu, \text{ and} \qquad (6\text{-}13)$$
$$\text{var}\{m\chi^2_\nu\} = 2m^2\nu,$$

where $E\{\}$ represents the expected value operator and $\text{var}\{\}$ the variance operator. Therefore, after replacing $m\chi^2_\nu$ with $f_v(v)$:

$$\nu = 2\,E^2\{f_v(v)\} \,/\, \text{var}\{f_v(v)\}, \text{ and} \qquad (6\text{-}14)$$
$$m = \text{var}\{f_v(v)\}/(\,2\,E\{f_v(v)\}\,).$$

The density after the min() operation (equation (6-11)) and the density from its chi-square approximation (equation (6-12)) are plotted in Figure 6.7 for the mother wavelet, scale 1. Note the excellent agreement between the two curves which are nearly indistinguishable. Selected values of ν and m are listed in Appendix H, Table H.2 and Table H.3.

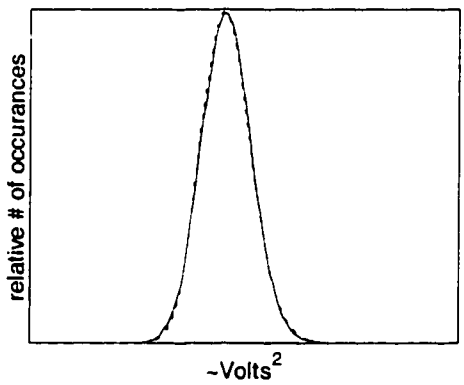

Figure 6.7. Probability density function of the minimum operation and the chi-square approximation.
The minimum operation curve is solid and the chi-square approximation is dotted.

Additional Comments:

aa) Two differences between Whitcher's application and this one (equation (6-9)) should be noted. First, the variance of the noise has been restricted to unity, so Whitcher's wavelet variance terms (which he labels $\nu^2_X(\lambda_j)$) are set to 1. Second, unlike Whitcher's wavelet, the Morlet is complex. Thus each sample is complex valued which introduces two additive terms in the computation of the magnitude of a sample rather than one for real samples. Therefore, the number of degrees of freedom of a power spectrum of complex samples will be twice that of real samples. Hence, a factor of two is introduced in the numerator of equation (6-9).

bb) The agreement of the curves in Figure 6.6 is excellent at the low wavelet scales (high frequency ranges) for all three wavelet implementation methods. As the wavelet length, *wlen*, approaches *seglen*, the approximation worsens for all wavelet implementation methods. It has been discovered heuristically, though, that with a *nonoverlap* percentage of 25% in the dot product method for Q_w of 3, the approximation again becomes quite good out to the 0.9999999 percentile for scales 1 through 9 at -8.7 dB overlap between scales in the frequency domain. Temporal overlaps of 20% and 30% yield considerably worse approximations.

Step 4:

Compute the normalized wavelet power spectrum:

$$\mathrm{NormR}(t,s) = \frac{R(t,s)/2}{v(t,s)/m(s)\nu(s)}. \qquad (6\text{-}15)$$

Compare all values of NormR with thresholds based on a confidence level percentage of the F distribution with 2 and ν degrees of freedom, where ν is scale dependent. Mark the time and scale locations where NormR exceeds these thresholds. These locations are the detected vibrations.

Comments:

a) The ratio of two independent chi-square distributed variables, each divided by its number of degrees of freedom, results in a ratio that has an F distribution. This implies that NormR is F distributed with 2 and ν degrees of freedom, $F(2,\nu)$. The graphs of an example probability density and associated cumulative probability distribution are shown in Figure 6.8. The values of NormR associated with a given time series, x(t), will be referred to as the normalized wavelet power spectrum.

b) The normalized wavelet power spectrum, NormR, is used as the test statistic for the detection decision. Detection is made when the value of NormR falls above the $F(2,\nu)$-distribution for a chosen confidence level. That is, if the chosen level of the test is 0.9, then the threshold level for NormR would be 2.35 for ν of 127 as marked by the cross-hairs in Figure 6.8. If the value of NormR exceeded 2.35, then a vibration would be detected with 90% confidence at the time and scale associated with the point R of Figure 6.5.

The values of m and ν are fixed for a given F_s, *seglen*, Q_w and wavelet scale. Thus, they can be computed in advance. Detector threshold values associated with selected confidence levels are given in Appendix H, Table H.4 through H.7. These thresholds were computed using Waterloo Maple's 'Fratio' function (Waterloo, Ontario, Canada).

c) Note that ν/m in equation (6-15) is chi-square distributed, not ν itself.

d) If the data were not initially unit variance, then the numerator and denominator of the ratio would each be multiplied by the variance. Since the variance would cancel from the ratio, its value is irrelevant to the detector.

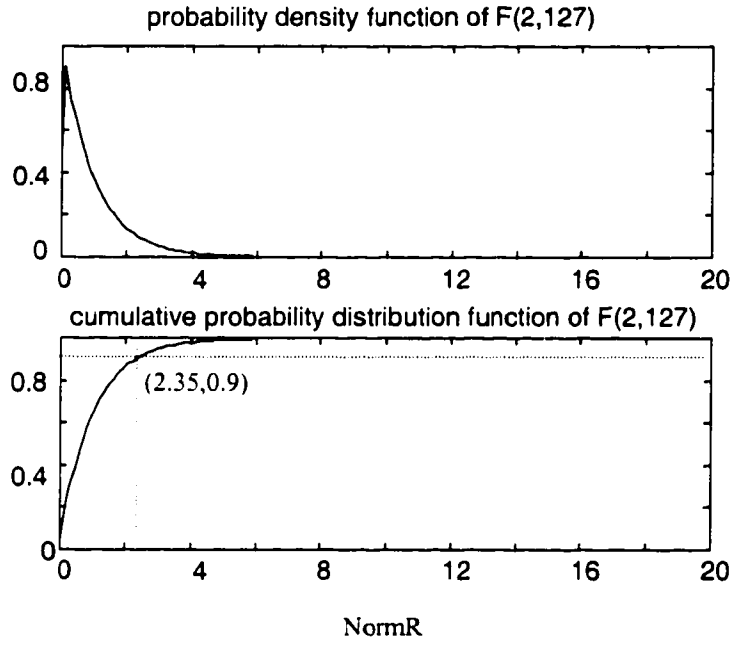

Figure 6.8. Probability plots for the F distribution.
The cross hairs at (2.35, 0.9) in the lower plot illustrate that for the F distribution with 2 and 127 degrees of freedom, 90% of Gaussian noise would generate a value of NormR that is less than or equal to 2.35.

Step 5:

Apply a second detector that only outputs detection hits that occurred at both positive and negative wavelet scales (frequency ranges) simultaneously. These will be referred to as two-sided hits.

Comments:

This step will further limit false alarms by exploiting *a priori* knowledge that vibrations have motion both toward and away from the transducer during a short time period. Thus, they should trigger hits from wavelet spectra with both positive and negative values of ω_0. Detection hits that occur at only positive or negative frequencies, but not both, are either from blood, noise or clutter.

6.2.1.1 Validation of the expected distribution with simulated data

The following steps validated the expected distribution with simulated data. i) A white Gaussian time series was generated in MATLAB® (The Mathworks, Natick, Massachusetts) using the 'randn' function. ii) The continuous wavelet transform of the series was computed using the dot product method with a *nonoverlap* value of 25% of *wlen*. iii) The noise estimates, v_L, v_R and v, were computed. iv) The normalized wavelet power spectrum of equation (6-15) was computed. v) Histograms were made of 100,000 of these normalized wavelet power spectra. vi) Empirical probability density functions were generated by first averaging the histograms at each wavelet scale, and then normalizing the averaged histogram for a cumulative probability of 1. Figure 6.9 shows the right-tail probabilities of the empirical density function against the true probabilities of the expected F-distribution for several wavelet scales. Note the excellent agreement up to the 0.999999 percentile. The slight divergence at scale 10 at the 0.999999 percentile arises because the dot product wavelet implementation computes fewer data points at the higher scales than at the lower scales for a set value of *tlen*. At scale 10 with F_s of 8000 Hz and *tlen* of 8000 samples, 100,000 computations do not provide enough data points (~ 5,000,000) for excellent accuracy out to 0.999999. Also playing a slight roll in the discrepancy is the chi-square assumption of the noise estimate that begins to fail at the large scales (see discussion in section 6.2.1, the additional comments to step 3, part bb).

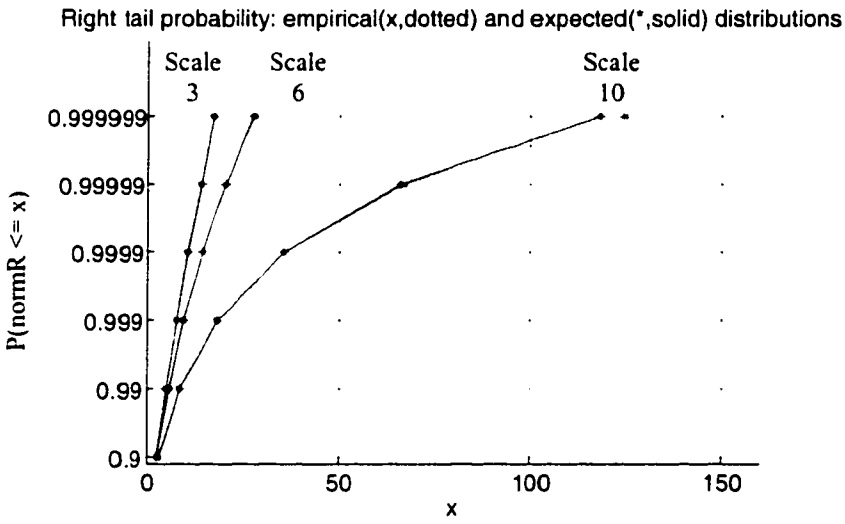

Figure 6.9. Right tail probability comparison of expected and empirical distributions of the normalized wavelet power spectrum.
The empirical distributions were computed with: F_s = 8000 Hz, Q_w = 3, *seglen* = 75 ms, *nonoverlap* percentage of 25%, and *tlen* = 8000. The histograms were computed with dx = 0.1, xmax = 250, and 100,000 generations of simulated noise.

6.2.1.2 Required assumptions and their validity

Use of the vibration detector requires two assumptions about the data from depths, times and wavelet scales in which vibration is absent: 1) the data have a Gaussian distribution in the wavelet domain; and 2) the vibration duration at a given wavelet scale is less than some fixed value.

6.2.1.2.1 Assumption 1: Non-vibration data are Gaussian distributed at each scale of the wavelet transform

An extension of the Gaussian assumption is that these data have a specific F-distribution of the normalized wavelet power spectrum for each wavelet scale. It is known at the outset that some expected events will cause this assumption to fail. First, sharp changes in the "noise level" will cause the normalized wavelet power spectrum distribution to digress from the expected F distribution and will trigger a detection hit. Thus, even if blood flow data are Gaussian distributed, they may trigger a detection hit because they introduce abrupt changes in the wavelet spectrum. It was also mentioned

earlier in section 4.3.2 that occasionally healthy arterial walls move at velocities which generate signals that exceed the minimum frequency expected for vibrations. These will not have a Gaussian distribution.

Finally, it is known *a priori* that narrowband noise is not Gaussian distributed. In fact the variance of the normalized wavelet power spectrum is smaller when narrowband noise is present than when broadband noise is present. An example is shown in Figure 6.10 for simulated narrowband noise. Consider the probability that the value of NormR in Figure 6.10 is 2. This probability is much lower for the narrowband noise (dotted) than the broadband Gaussian noise (solid). This implies that a value of 2 for NormR would be much more likely to appear in broadband noise than in narrowband noise. Therefore, a vibration that yielded a value for NormR of 2 and caused a detection hit under a narrowband assumption, would not cause a hit under the broadband assumption at the same level of confidence. Therefore, the detector, which expects broadband noise, will tend toward both a lower detection rate and a lower false alarm rate for a given signal to noise ratio because the detection threshold will be artificially high.

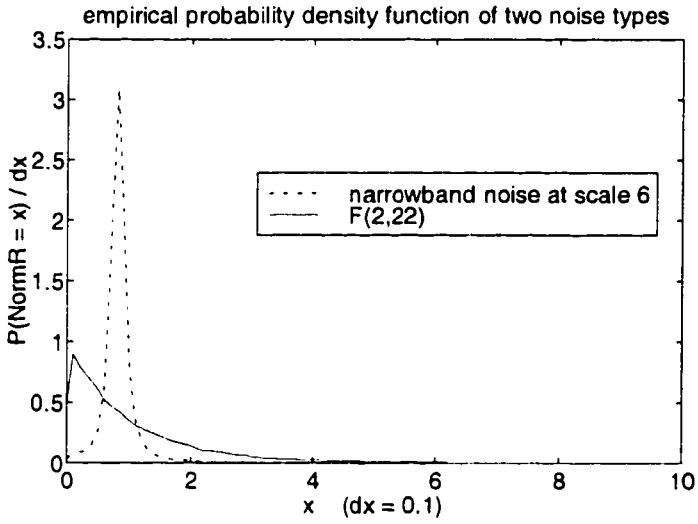

Figure 6.10. Comparison of empirical density functions from normalized wavelet power spectra of simulated narrowband and broadband noise.
These were generated with the FFT wavelet transform implementation from 100 generations of the noise with F_s = 8000 Hz, tlen = 8000 samples, *seglen* = 75 ms, xmax =30, Q_w = 3, scale overlap at -8.7 dB in frequency domain.

Aside from the three caveats just mentioned, blood flow, wall motion and narrowband noise, non-vibration data can be assumed Gaussian. Thus, the wavelet power spectra would be chi-square distributed. To validate this assumption, data has been collected from two human volunteers presumed normal and from four easily accessible anatomical locations where vascular disorders are likely to occur: the carotid artery, the aorta, the superficial femoral artery and the posterior tibial artery. Combined, these data represent time series from more than 150 depths of interest. Empirical probability density functions from these data were computed from the wavelet power spectra, which were not normalized. The resulting distribution should be chi-squared with two degrees of freedom. All data from depths without blood flow, high velocity clutter, vibration or narrowband noise passed the Kolmogoff-Smirnov goodness-of-fit test (described in Appendix E) with a level of confidence greater than 99.9%.

As an example, the following is an empirical probability density function of wavelet power spectra formed from histograms from a normal volunteer. The histograms were

obtained, and averaged, from 23 depths which did not contain blood flow, vibration, high velocity clutter or narrowband noise.

6.2.1.2.2 Assumption 2: Vibration duration at a given wavelet scale is less than some fixed value

According to the literature, as well as *in vivo* vibrations that I have studied to date, vibrations arising from vascular disorders occur primarily during systole, the section of the heart cycle in which peripheral blood velocities are at a maximum, and blood pressures are at a maximum. Since systole encompasses less than 50 percent of each heart beat, most vibrations have less than a 50 percent duty cycle. It is possible that a vascular disorder will cause vibrations that extend into diastole but it is likely that the frequency will be lower in diastole due to the lower blood velocities and blood pressure. If the actual duration at a specific wavelet scale falls within the expected duration, then the detector will operate correctly. The case in which the vibration has a long duration at a particular scale will introduce an artificially large noise estimate, which will reduce the probability of detection.

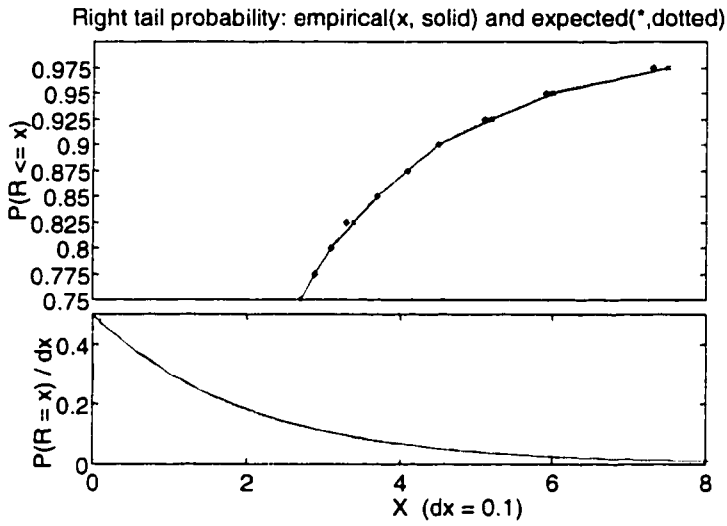

Figure 6.11. Expected (χ_2^2) and empirical probability functions from wavelet power spectra of noise.
Upper plot: Cumulative probability distribution functions. **Lower plot:** Probability density funtions. The data were collected from a normal human volunteer with the HDI 1000 ultrasound instrument. The histograms used in the empirical distribution were obtained, and averaged, from 23 depths which did not contain blood flow, vibration, high velocity clutter or narrowband noise.

6.2.2 Morlet wavelet method – dual depth

The single depth method works well, but fails to make use of the known spatial distribution of the vibration. This spatial distribution in the depth direction can be exploited by basing the test statistic on the cross wavelet spectrum formed with time series from neighboring depths (section 5.2.2.2.1). The magnitude of this spectrum will emphasize the times and scales at which two time series are correlated. The phase of this spectrum reveals the phase shift between the two time series at each time and scale. As with the single depth method, the results of the cross wavelet spectrum for two consecutive depths will be compared with a null hypothesis, the distribution expected from two depths consisting of independent Gaussian noise.

Step 1:

Compute the wavelet transforms ($W_1(t,s)$, $W_2(t,s)$), wavelet power spectra ($|W_1(t,s)|^2 = R_1$, $|W_2(t,s)|^2 = R_2$) and noise estimates ($v_1(t,s)$, $v_2(t,s)$) for each of the two depths of interest.

Comments:

Again, the data are assumed to consist of Gaussian noise at frequencies above the clutter range in the absence of blood flow, vibration and narrowband noise. Initially, for simplicity, the variance of both the real and imaginary components of the noise is restricted to unity, but is unnecessary for the same reasons as in the single depth method (see section 6.2.1, comments to step 4, part d).

Step 2:

Compute the normalized cross wavelet spectrum as:

$$\frac{(W_1(t,s))(\mathrm{conj}(W_2(t,s)))}{\sqrt{v_1(t,s)}\sqrt{v_2(t,s)}}\left(\frac{1/2}{1/m(s)\nu(s)}\right), \qquad (6\text{-}16)$$

where conj() represents the complex conjugate operator.

Subsequently, compute its magnitude squared:

$$NormRcorr(t,s) = \qquad (6\text{-}17)$$

$$\frac{|W_1(t,s)|^2/2}{v_1(t,s)/m(s)\nu(s)}\,\frac{|W_2(t,s)|^2/2}{v_2(t,s)/m(s)\nu(s)} = \frac{R_1(t,s)/2}{v_1(t,s)/m(s)\nu(s)}\,\frac{R_2(t,s)/2}{v_2(t,s)/m(s)\nu(s)}.$$

Comments:

a) As mentioned earlier in the discussion of the single depth detector, for each depth, $\dfrac{R(t,s)/2}{v(t,s)/m(s)\nu(s)}$ is F distributed with 2 and ν degrees of freedom.

b) The squared magnitude of the normalized cross wavelet spectrum is the product of two F-distributed random variables. Its probability density function is derived in Appendix F to be:

$$\int_0^\infty \frac{dy}{y\left[1+\left(\frac{2}{\nu}\right)^2 z+\left(\frac{2}{\nu}\right)\frac{z}{y}+\left(\frac{2}{\nu}\right)y\right]^{1+\nu/2}}. \qquad (6\text{-}18)$$

The graphs of this probability density function and the corresponding cumulative probability distribution function are shown in Figure 6.12.

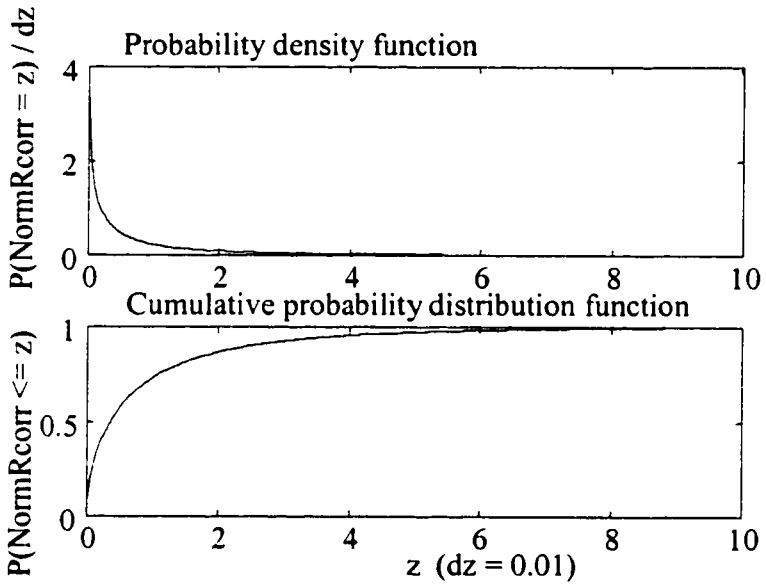

Figure 6.12. The expected probability density of noise from the squared magnitude of the normalized cross wavelet spectrum.
This distribution is the product of two F distributions with 2 and 127 degrees of freedom.

c) In some sense, the normalized cross wavelet spectrum could be considered a variation of the coherence measures of equations (5-10)and (5-11) in section 5.2.2.2.2. Like those definitions, the numerator is the product of wavelet spectra from the two time series, and the denominator is the product of the individual wavelet

power spectra. The wavelet coherence definition A (equation (5-10)), though, computes the expected values of both the cross spectra and the individual spectra; whereas, the wavelet coherence definition B does not involve any averaging (equation (5-11)). The normalized cross wavelet spectrum, alternatively, employs averaging only of the individual spectra. It further differs from the coherence measures in that the individual spectra are computed at time locations that are independent of the cross spectrum computation.

Step 4: Compare the value of *NormRcorr* with a threshold for a given confidence level of the distribution whose density is defined by equation (6-18).

Comments:

Detector threshold values associated with selected confidence levels of this distribution are given in Appendix H, Table H.8 through Table H.11. These were computed using numerical integration of equation (6-18) above.

Step 5:

Apply a second detector that outputs only detection hits that occurred at both positive and negative scales simultaneously.

Comments:

See step 5 of the single depth method.

6.2.2.1 Validation of the expected distribution with simulated data

As done with the single depth detector, the correlated detector was validated with simulated data generated in the same manner as described earlier in section 6.2.1.1, except that two independently generated Gaussian time series were used. Histograms were made of 100,000 normalized cross wavelet spectra. The right-tail probabilities of the empirical density function, computed as the averaged histograms, are plotted against the expected probabilities for various scales. Note the agreement up to the 0.999999 percentile for the low scales. The divergence for scale 10 at the large percentiles is due to the same phenomena as described in section 6.2.1.1 for the single depth detector.

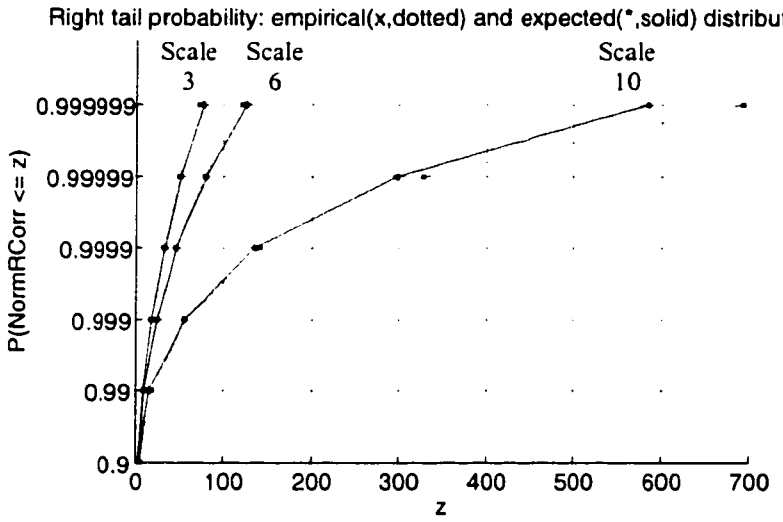

Figure 6.13. Right tail probability comparison of expected and empirical distributions.
The empirical distributions were computed with: F_s = 8000 Hz, Q_w = 3, *tlen* = 8000, and *seglen* = 75 ms. Wavelet spectra were computed with the dot product wavelet implementation with a *nonoverlap* percentage of 25%. The histograms were computed with dz = 0.5, zmax = 1000, and 100,000 generations of simulated noise.

6.2.2.2 Required assumption and its validity

The dual depth detection technique requires the same assumptions as the single depth technique.

6.3 Vibration parameter estimates

The wavelet transform can estimate the amplitude, intensity, frequency and bandwidth of a captured vibration with an accuracy limited only by computational power and SNR. Estimation accuracy of time delay depends greatly on the degree to which the two signals are correlated. Duration estimation accuracy is affected by SNR and the length of the wavelet.

Vibration parameter estimation follows the methods outlined in the literature, as described in section 5.3.2. Where necessary, the methods have been expanded due to insufficient details in the literature.

6.3.1 Algorithms

6.3.1.1 Amplitude

Unlike the detector which operates on the raw quadrature data, the displacement values are computed from the phase of the quadrature data as described in equation (4-13) and section 4.1.3. Therefore, the first step in amplitude estimation is to compute the displacement in the time domain from the ultrasound phase value of quadrature time samples associated with a given depth. The second step is to apply the wavelet transform to the displacement time series in order to condition the data for amplitude estimation. For use with amplitude estimates the wavelet is normalized differently than for detection: by the sum of its Gaussian taper. Mathematically, this is written:

$$w(t) = \frac{e^{-\alpha t^2/2} e^{j\omega_0 t}}{\sum_t e^{-\alpha t^2/2}}, \qquad (6\text{-}19)$$

where α and ω_0 were defined in section 5.1.2.2.

The amplitude estimate of a signal associated with a particular frequency is taken to be the value of the wavelet coefficient at the wavelet scale corresponding to that frequency. The estimate will have a maximum error according to the amplitude level at which the wavelet scales overlap in the frequency domain. For example, if the wavelet scales overlap at -3 dB (*scalebase* = 1.22, see equation (6-2)), then the maximum amplitude error will be 1-1/sqrt(2) or approximately 29.3% (Figure 6.14).

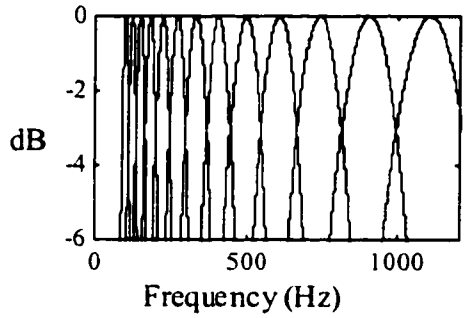

Figure 6.14. A frequency domain plot of wavelet scales.
Q_w=3, -3 dB overlap in frequency (*scalebase = 1.22)*, 13 scales ranging from 100 to 1000 Hz.

The total amplitude of the signal over a given frequency range can be found as the sum of the magnitudes of the scales representing that frequency range. Figure 6.15 shows the sum of the magnitudes of the wavelets themselves. With this method, the amplitude estimates corresponding to frequencies between 100 Hz and 1000 Hz would be underestimated by as much as 1% and overestimated by as much as 6.5% depending on the frequency(ies) of the signal as seen in Figure 6.15. These are the errors expected when C_δ of equation (5-13) is equal to one (section 5.3.2.2.1). Since the errors correspond only to particular frequency values, broadband signals will experience some cumulative effect of these.

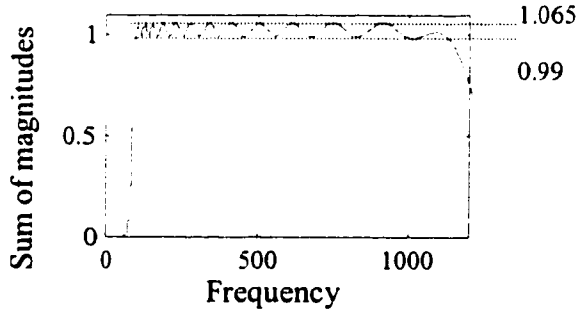

Figure 6.15. A frequency domain plot of the sum of magnitudes of the wavelet scales.
Q_w=3, -3 dB overlap in frequency (*scalebase* = 1.22), 13 scales ranging from 100 to 1000 Hz.

With more overlap between scales, accomplished with closer spacing between the wavelet center frequencies, the error variance will decrease but the overestimation will increase. Therefore, the amplitude error can be reduced by combining increased scale overlap with multiplication by a correction factor to account for the overestimation. The cost comes at the increased number of scales, which could be mitigated by the use of a wavelet with a lower Q_w. For Q_w of 1.5 at –2 dB overlap in frequency (*scalebase* = 1.4), the overestimation would fall between 11% and 14.9% (Figure 6.16). By scaling the result, the error would be an underestimation and overestimation as low as 1.8%. These errors result from a value of C_δ which is the reciprocal of the mean of the magnitudes of Figure 6.16 over the frequency range 100 to 1000 Hz.

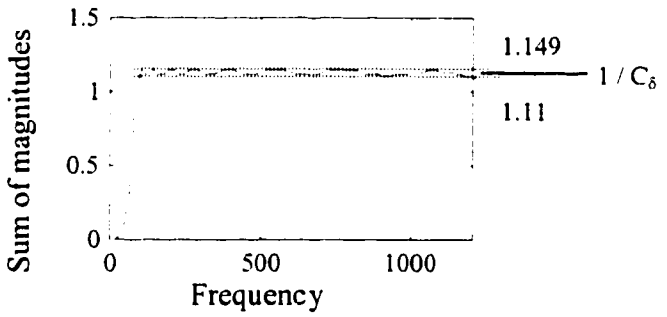

Figure 6.16. A frequency domain plot of the sum of magnitudes of the wavelet scales.
Q_w=1.5, -2 dB overlap in frequency (*scalebase* = 1.4), 9 scales.

If it is unknown *a priori* whether the signal is narrowband or broadband, then broadband must be assumed. Suppose, though, that it were known that the signal consisted of a single, tapered complex sinusoid. If, for some reason, the wavelet had the same center frequency as the signal, but with a much shorter duration, the amplitude estimate would be the maximum value of the wavelet coefficient over the time duration of the signal. For this case, the amplitude estimate error would be negligible. This result can be proven theoretically for the case of a sinusoidally tapered sinusoid (Appendix D), but not for the Gaussian tapered sinusoid of Morlet because the Gaussian taper is not integrable theoretically. The error will increase as the signal

duration shortens, the signal bandwidth broadens within the taper, and the mismatch between the wavelet center frequency and the signal's center frequency increases.

In practice, the amplitude estimate for a signal known to be narrowband can be carried out as follows.

1) Compute a wavelet estimate at each of several wavelet scales spaced for minimal overlap in the frequency domain. Do this using a low Q_w wavelet such as 1 or 1.5 to attempt a wavelet length shorter than the signal duration.
2) Find the wavelet scale which yields the maximum amplitude estimate.
3) Compute wavelet coefficients at wavelet scales more closely spaced over the frequency ranges between the center frequencies of the two scales bordering the maximum amplitude estimate of step 2).
4) The amplitude estimate is the maximum value of these coefficients.
5) Iterate steps 2 through 4 until the algorithm achieves the desired accuracy.

The choice of the scale spacing is an accuracy versus computational cost tradeoff.

6.3.1.2 Frequency and bandwidth

A signal's frequency range can be estimated in different ways. Two straightforward methods involve setting the estimate to the frequency span of the wavelet scales at which either 1) the signal was detected, or 2) the wavelet transform lies above $1/\sqrt{2}$ (the -3 dB points) of its maximum value. These methods are illustrated in Figure 6.17 and Figure 6.18 for an arbitrary broadband signal. Both methods will overestimate the bandwidth of a signal when it does not extend to the extremes of the wavelet scales at which it was detected. In addition, the second method will further overestimate the bandwidth when the peak frequency of the signal falls near the overlap of two wavelet scales. In this instance none of the wavelet transform scales achieves the true peak amplitude value, causing the assumed 3 dB points to actually represent signal as low as the 6 dB points. Alternatively, if the signal had significant frequency content just below its 3 dB level, then its bandwidth may be underestimated depending upon the preferred definition of bandwidth. For both methods, accuracy increases with increasing Q_w, which corresponds to increasingly smaller wavelet bandwidths.

Thus, the frequency error will depend greatly on the minimum spacing between the wavelet center frequencies and on the strength of the signal at the outer edges of the frequency range.

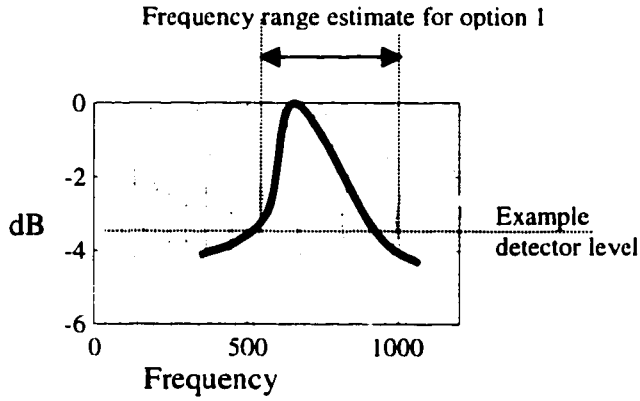

Figure 6.17. Illustration of frequency estimate option 1.
The thick, solid line indicates the frequency domain representation of an example time series. The gray curves are the 13 wavelet scales that span the frequency range of interest. The frequency estimate is made from the frequency span of the wavelet scales that contain wavelet coefficients that surpass the detection threshold (marked by the horizontal dotted line). The two vertical dotted lines indicate the extremes of the wavelet scales with detected coefficients.

For narrowband signals whose duration is longer than the wavelet, the frequency estimate is simply the center frequency of the wavelet scale at which the iterative amplitude estimate of section 6.3.1.1 was maximized.

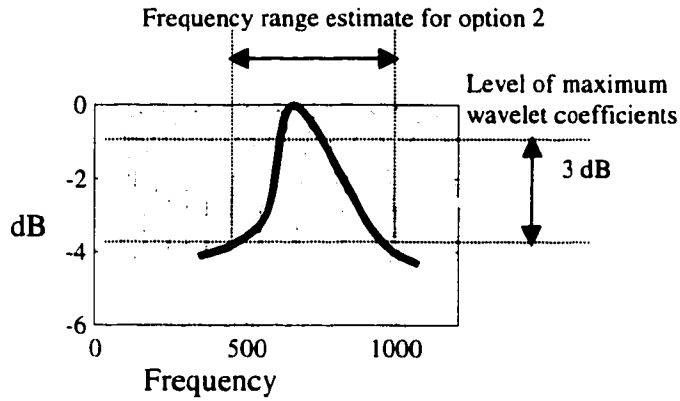

Figure 6.18. Illustration of frequency estimate option 2.
The thick, solid line indicates the frequency domain representation of an example time series. The gray curves are the 13 wavelet scales that span the frequency range of interest. The upper, horizontal, dotted line marks the level of the largest wavelet coefficients, which corresponds to the mean of the product of the wavelet and the signal level within the wavelet's frequency range. The lower, horizontal, dotted line marks the level that is 3 dB down from the upper line. The frequency estimate, marked by the two vertical, dotted lines, is made from the frequency span of the wavelet scales that contain wavelet coefficients whose value surpasses the lower, horizontal line.

6.3.1.3 Intensity

Consider that a signal's intensity measure can be broken down as a sum of intensities over a continuous range of frequencies. An approximate intensity measure, therefore, would be a sum of the intensities at discrete frequency ranges such as the "non-overlapping" wavelet scales (s), where the degree of "non-overlap" is some chosen dB level, possibly -3 dB or -8.7 dB. Consider also that acoustic intensity is found as the particle velocity squared multiplied by the acoustic impedance (Z) in the medium.

$$I_v(t) = vel^2(t)\, Z \qquad (6\text{-}20)$$

The velocity estimates, in turn, can be computed from the first difference of the amplitude estimates divided by the time between wavelet coefficients ($\Delta x/\Delta t$). Therefore, the intensity estimate can be computed as:

$$\sum_s I_{inst}(t,s) = \sum_s \mathrm{vel}(t,s)^2 Z = \sum_s \left(\frac{\Delta x(s)}{\Delta t(s)}\right)^2 Z . \qquad (6\text{-}21)$$

Since the consecutive displacement time samples used in computing Δx will lie at the same frequency location, they will be affected by the same amplitude estimation error (section 6.3.1.1). Hence, the intensity estimation error will be the same as the error associated with the signal's instantaneous power, which is the sum of the wavelet power spectrum across the wavelet scales of interest (i.e., $|W(t,s)|^2$ summed over s)(see section 5.3.2.5.1).

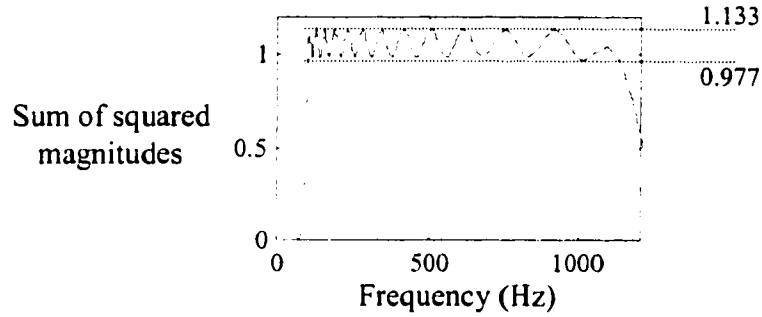

Figure 6.19. A frequency domain plot of the sum of squared magnitudes of the wavelet scales.
Q_w=3, -3 dB overlap, 13 scales.

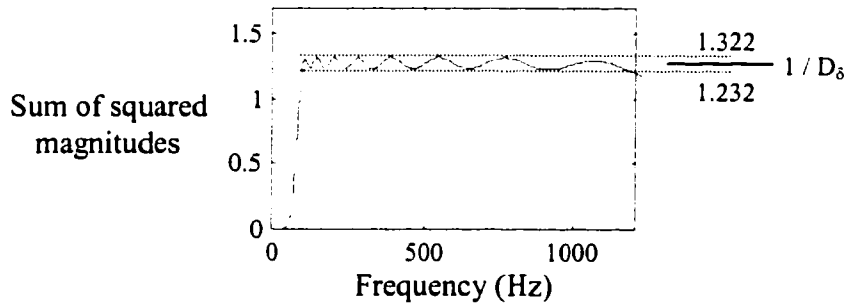

Figure 6.20. A frequency domain plot of the sum of squared magnitudes of the wavelet scales.
Q_w=1.5, -2 dB overlap, 9 scales.

The theoretical accuracy of this estimate is somewhat large when the scales overlap in the frequency domain by approximately -3 dB and when Q_w is 3, as shown in Figure 6.19. This figure reveals that underestimation may be as much as 2.3% and overestimation as much as 13.3%. As with amplitude, with more overlap between scales and a lower Q_w to reduce the number of scales, the error can be improved. For Q_w of 1.5 at -2 dB overlap, the overestimation would fall between 23.2% and 32.2% (Figure 6.20). By scaling the result, the error would be an underestimation and overestimation as low as 3.6%. This corresponds to a correction factor, D_δ (equation (5-15)), equal to the reciprocal of the mean of the squared magnitudes of Figure 6.20 over the frequency range 100 to 1000 Hz.

6.3.1.4 Duration

The duration is estimated as the number of wavelet coefficients exceeding the detection threshold multiplied by the time between wavelet coefficients. The duration estimate accuracy will be dependent on the signal to noise ratio of the data, the chosen level of confidence of the detector and the wavelet length. In a high SNR situation, the minimum duration value at a given scale will be the length of the wavelet.

6.3.1.5 Time delay from phase

The time delay estimate is computed according to equation (5-14). Since phases, and thus time delays, are estimated at each time for a given scale, the final time delay estimate is computed as the mean of the time_delay(t,s) over both time and scale for the time and scale values at which the normalized cross wavelet spectra exceed the detection thresholds.

The phase estimate used in equation (5-14) is the estimate of the phase between signals at consecutive depths at a given wavelet time and scale as obtained from the cross wavelet spectrum. This phase estimate will have negligible error if the signal is narrowband, the wavelet center frequency equals the signal's center frequency, and the SNR is high. If the wavelet center frequency is not the narrowband signal's frequency or if the signal is not narrowband, then phase estimation errors will increase as the signal's strength at the wavelet center frequency decreases. However, since the time

delay estimate uses the frequency at which the phase is estimated, the wavelet center frequency, rather than the true vibration signal center frequency, the time delay estimate would remain accurate, when the SNR is high. Time delay accuracy depends mostly on the degree to which the two signals are correlated.

6.3.2 Confounding factors for vibration parameter estimates

There are three assumptions implicit in estimating vibration parameters from the received ultrasound echoes. 1) The received ultrasound echoes correspond to tissue motion that occurs only along the ultrasound beam. 2) There is no stationary echo superimposed on the tissue motion of interest (see section 4.3.2). 3) The tissue in the sample volume remains fairly consistent.

6.3.2.1 Motion assumed along the ultrasound beam

Due to lack of further information, any motion apparent in the received echoes is automatically assumed to be motion toward or away from the transducer in the direction of the beam. Therefore, if the motion is *not* truly in the direction of the beam, then the amplitude, and, correspondingly, intensity measurements will be underestimated. The frequency, phase and timing measurements, however, would not be affected.

6.3.2.2 Presence of stationary echoes

Estimation errors would also occur if a stationary echo were superimposed on the tissue motion of interest. In the absence of unexpected stationary echoes, the IQ plot of the data at the depth of interest will rotate about the origin. The relative phase changes from sample to sample, then, correctly correspond to the amplitude of motion. When a stationary echo is present, though, the IQ plot rotates about some point other than the origin. In this case, phase changes with respect to this other point would correspond to motion, *not* phase changes with respect to the origin as assumed. Figure 4.18 of section 4.3.2 illustrates this point. Thus, the amplitude and intensity estimates would be erroneous, but other parameter estimates would be unaffected.

6.3.2.3 Blood oscillations rather than solid tissue vibration

Amplitude estimates, as well as intensity estimates, are based on the assumption that the tissue of interest is solid. When blood oscillations are detected, amplitude estimates will be completely false because the motion that was assumed to represent movement of a fixed medium would actually be caused by blood which passes through the sample volume rather than remaining fixed in it. Once again, frequency, phase and timing measurements would remain valid.

7 Results

The detectors and estimators just described have been implemented on data obtained from the various vibration models discussed in section 4.4. In addition, Monte Carlo methods have been used on the simulated data to analyze the accuracy of the detectors and estimators. The simulated data and corresponding results will be discussed first. A discussion of displaying the results follows. The chapter will conclude with results from the other vibration models.

All wavelet-based detections and estimations described here are computed with Q_w of 3, F_s of 8000 Hz, *seglen* of 75 ms, *viblen* of 150 ms, *tlen* of 8000 samples, and *scalebase* of 1.4 for -8.7 dB overlap between scales in the frequency domain, except where noted. The results are all computed from the dot product wavelet transform implementation with a *nonoverlap* value of 25% of the wavelet length (see section 6.1.2.3).

7.1 Simulated detection

7.1.1 Single depth

7.1.1.1 The simulated signal

Signal simulation begins with the selection of all vibration parameters except amplitude, clutter and phase: center frequency, duration (f_v), bandwidth (BW), quality factor (Q_v), and signal-to-noise ratio (SNR). Clutter is not simulated because of the excellent filtering capability of the Morlet wavelet (see Figure 5.5). In addition, any clutter whose frequency content enters the frequency range of interest would be one-sided on a wavelet power spectrum, and therefore would be excluded by the detection scheme that requires a two-sided signal.

The following is the process employed for generating simulated vibration signals. The SNR is first set to a desired value. Vibration quality factor (Q_v), duration and center frequency (f_v) are computed from a uniform distribution across the values that exist *in vivo*, using MATLAB's 'rand' function.

vibration quality factor (Q_v):	0.55 to 10
duration:	10 tol50 ms
center frequency (f_v):	100 to 1000 Hz

7.1.1.1.1 Frequency domain:

1) The number of samples in the signal is determined based on the sample rate and the selected signal duration.
2) Frequency bins are computed as the sampling rate divided by the number of samples.
3) BW is computed as center frequency divided by Q_v.
4) Frequency bins centered around the center frequency extending to a width of BW are set to complex values where each of the real and imaginary components are drawn from a Gaussian distribution of mean zero and variance one using MATLAB's 'randn' function. An example is shown in the upper left corner of Figure 7.1.

7.1.1.1.2 Time domain:

1) The inverse Fourier transform of the frequency domain signal is computed to yield a time domain signal of the requested duration.
2) To mimic *in vivo* vibrations in which there is a gradual amplitude increase at the beginning of the vibration and a gradual decrease at the end, the time domain signal is multiplied by a Hanning taper. An example of the resulting time domain signal is shown in the upper right section of Figure 7.1 above. The Hanning taper of a certain duration will itself have a 3 dB bandwidth of approximately $\frac{\sqrt{2}}{duration}$. This bandwidth will add to the original bandwidth causing the simulated signal to have a larger bandwidth than requested.
3) The total energy in the simulated signal is computed according to $\sum_t abs(x(t))^2$, where x(t) represents the simulated signal in the time domain.

4) A long, zero mean, complex, white noise time series is generated from a Gaussian random number generator, MATLAB's 'randn' function. Its variance is set so that the energy computed for the signal divided by the energy in the noise equals the selected SNR for the time range in which the signal is present.

5) The signal time series is added to the noise in the center of the noise time series (see lower left of Figure 7.1).

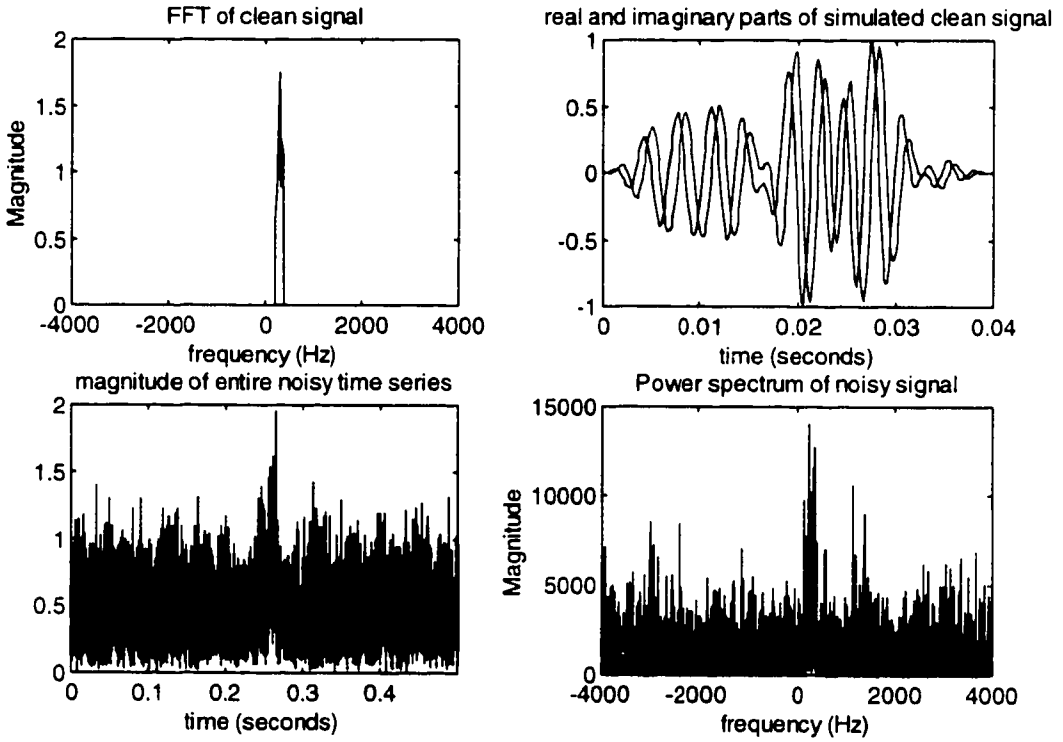

Figure 7.1. Illustration of signal simulation.
This example uses the following parameter values: f_v = 300 Hz, BW = 150 Hz, SNR = 1, Q_v = 2, and signal duration = 40 ms. **Upper left:** A signal with center frequency of 300 Hz (f_v) and a bandwidth of 150 Hz (BW). The values in this frequency range are selected from a white Gaussian distribution. The frequency resolution is selected to correspond to the signal duration of 40 ms (1/duration = 25 Hz). **Upper right:** A Hanning taper multiplied by the inverse FFT of the signal from the upper left. **Lower left:** The time series of the upper right added to a long, stationary, complex, white Gaussian noise time series. The time series from the upper right appears in the center of the noise time series. **Lower right:** The magnitude squared of the FFT of the time series of the lower left.

7.1.1.2 ROC curve

ROC curves were computed as follows. First, a signal is generated as described above for a selected SNR. Then, the detector is applied to both the signal in noise and the noise absent the signal for a selected level of confidence. When applied to the signal in noise, if the detector finds a vibration in the time and frequency range in which the signal was generated, a detection hit is recorded for that signal generation. When the detector finds any hits in the noise-only time series, a false alarm is recorded. One thousand signals are generated for each pair of selected SNR values and detection confidence levels. The ROC curve plots the number of hits divided by the number of signals generated versus the number of false alarms divided by the number of signals generated. Figure 7.2 and Figure 7.3 show the resulting detection curves. Note that the detection rates are excellent even down at the lower SNR values. The simulation indicates that out of 1000 vibration time series with an SNR of 2, 99.3% would be detected at a false alarm rate of 0.2% (Figure 7.3).

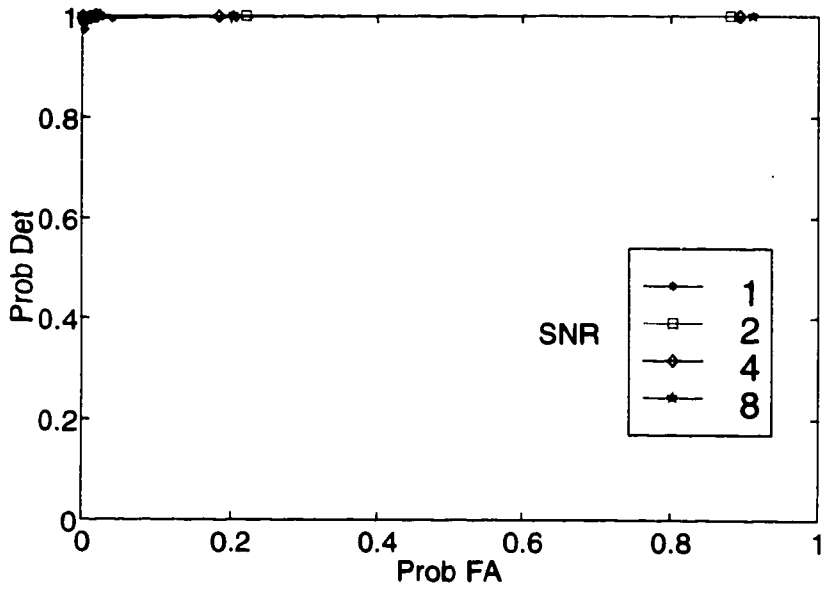

Figure 7.2. ROC curves for single depth detector.
Receiver operating curves for the single depth detector for four signal to noise ratios (SNR), each computed from 1000 different signal and noise realizations. The SNR values are computed as the ratio of the energy of the signal to the energy of the noise when the vibration is present. The simulations were computed with: *seglen* = 75 ms, *viblen* = 150 ms, tlen = 8000, Fs = 8000 Hz, *nonoverlap* percentage of 25%, and wavelet scale overlaps in the frequency domain = -8.7 dB. The confidence levels were 0.999 0.9999 0.99999 0.999995 0.999999 0.9999999.

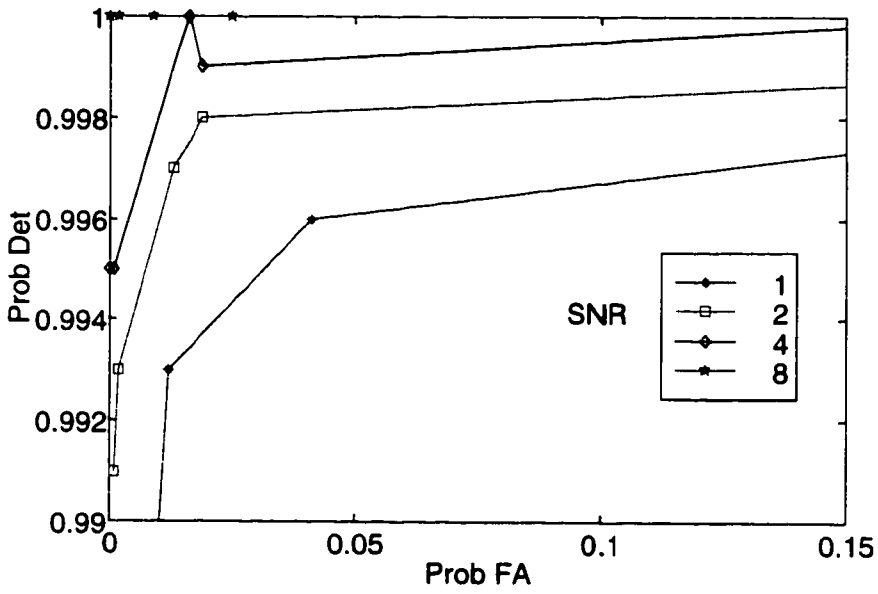

Figure 7.3. ROC curves for single depth detector (zoom).
The same curves as shown in Figure 7.2, but with the axes zoomed to the upper left corner.

To offer some validation for the results, the expected number of false alarms can be computed as:

(1 - confidence level) (# of wavelet coefficients).

The number of wavelet coefficients, then, is found as:

(# of repetitions)(# of wavelet scales)(average number of coefficients per scale).

The average number of coefficients per scale can be found using equation (6-6) and the average wavelet length which is computed in Appendix G, section G.1.1 to be 300 for Q_w of 3. Equation (6-6) stated:

$$(tlen - overlen))/nonoverlen,, \qquad (6.6)$$

where *nonoverlen* is (% nonoverlap)(*wlen)* and *overlen* is (*wlen* - *nonverlen*). Therefore:

average(*nonoverlen*) = (% nonoverlap)(300) and
average(*overlen*) = 300(100% - % nonoverlap).

The expected number of false alarms is now:

$$\frac{(1 - \text{confidence level})(\#\text{of repetitions})(\#\text{of wavelet scales})(tlen - \text{average}(overlen))}{\text{average}(nonoverlen)}.$$

For the largest confidence level used in the simulation, 0.9999999 ($1 - 10^{-7}$), and for 1000 runs of 8000 points each with 25% overlap, and 16 wavelet scales:

$(10^{-7})(1000)(16)(8000 - 0.75(300)) / (0.25(300)) = 0.166.$

Thus, 0.166 false alarms are expected. Correspondingly, a 0.9999 confidence level should involve 166 false alarms. For the ROC curves in Figure 7.3, at the 0.9999999 percentile the average number of false alarms for the four curves is 0.25, which roughly agrees with the expected number of 0.166. At the 0.9999 confidence level, the average number of false alarms is 204. Again, there is rough agreement. Although the ROC curves are expected to be monotonic, the simulated curve at SNR of 4 is not. This lack of monotonicity is not surprising, though, since the change in the number of missed detections between data points is only one simulation out of the total number of simulations (1000).

7.1.2 Dual depth

7.1.2.1 Simulated signal

For evaluation of the dual depth detector, the signal generation scheme outlined in section 7.1.1.1 can be slightly altered to simulate signals that are highly correlated although shifted by a time delay that is less than one time period of their oscillation.

7.1.2.1.1 Frequency domain:

The frequency domain computation of the signal remains the same as in section 7.1.1.1.1 except a second signal is created identical to the first. The second signal is then separated into its magnitude and phase components in the frequency domain. A phase shift that is linear in frequency is applied to the phase components of each non-zero frequency bin, accomplishing a time delay in the time domain version of the signal. The size of the shift is chosen between $-\pi$ and π.

7.1.2.1.2 Time domain:

The first signal is conditioned in the time domain exactly as described in section 7.1.1.1.2. The second signal will have an amplitude adjusted by +/- 30% of the amplitude of the first signal. The noise of the second signal is independent of the noise of the first signal, but with the same variance.

7.1.2.2 ROC curves

The ROC curve for the dual depth detector is computed in the same manner as that of the single depth detector (section 7.1.1.2). Figure 7.4 and Figure 7.5 show the results. The expected number of false alarms for this detector would be the same as the number for the single depth detector: a confidence level of 0.9999999 should correspond to 0.166 false alarms, and a confidence level of 0.9999 should correspond to 166 false alarms. Thus, the results, 0 false alarms and 228, respectively, roughly agree with the expected numbers.

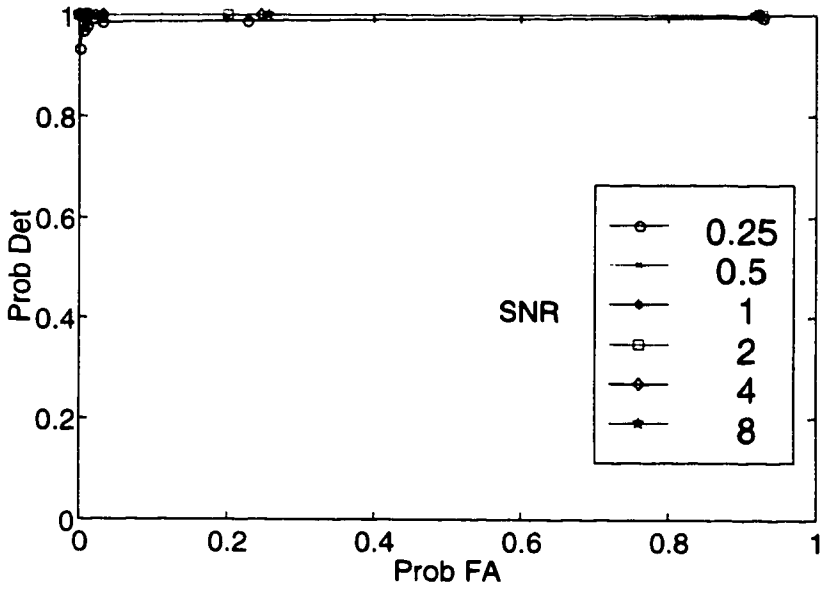

Figure 7.4 ROC curve for dual depth detector.
Receiver operating curves for the dual depth detector for six signal to noise ratios (SNR) computed from 1000 different signal and noise realizations. The SNR values are computed as the ratio of the energy of the signal to the energy of the noise when the vibration is present. The simulations were computed with: *seglen* = 75 ms, *viblen* = 150 ms, tlen = 8000 samples, Fs = 8000 Hz, *nonoverlap* percentage of 25%, and wavelet scale overlaps in the frequency domain = -8.7 dB. The confidence levels were 0.999 0.9999 0.99999 0.999995 0.999999 0.9999999.

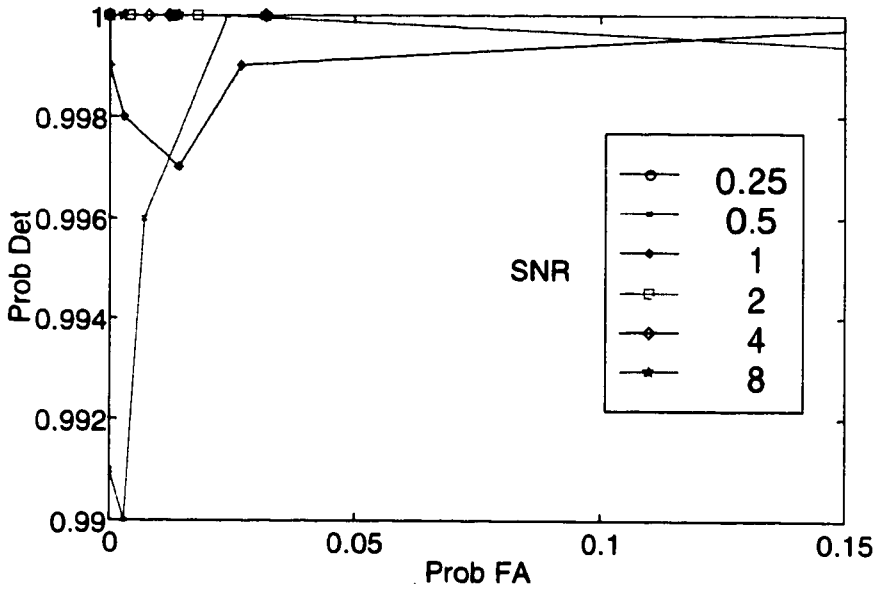

Figure 7.5. ROC curve for dual depth detector (zoom).
The same curves as shown in Figure 7.4, but with the axes zoomed to the upper left corner.

7.1.3 Comparison with Fisher's test for periodicity

Fisher's F-test for periodicity, introduced in section 5.2.1.1, is a common means of detecting simple periodicities (single frequency peaks) in white noise when the noise variance is unknown. Its straightforward method of implementation follows (Percival and Walden, 1993).

1) Compute the periodogram (magnitude squared of the FFT) of the time series under investigation.
2) Set g equal to the maximum value of the periodogram divided by the sum of all periodogram ordinates.
3) Compute the threshold value for the chosen level (φ) of the test:

$$g_F \approx (\varphi / i)^{1/(i-1)},$$

where i = N/2 - 1, and N is the length of the FFT.

4) If g exceeds g_F, then indicate a detection hit.

Fisher's test for periodicity assumes that the data is stationary within the time segment for which the periodogram is computed. A non-stationary time series must be segmented into approximately linear time series. The variable parameters are the number of time samples to use in the FFT calculation and the degree of overlap between segments used for the FFT computation.

Figure 7.6 shows the resulting ROC curves which were generated via Monte Carlo simulation using the same signal generation scheme as outlined in section 7.1.1.1. The FFT's were computed at 50% overlap at N of 1024 because these settings were shown to be superior and/or comparable to other values via a separate Monte Carlo analysis. The results reveal that only at the large SNR values does Fisher's test for periodicity achieve detection rates that approach those of the single depth wavelet based vibrometry detector. [Note the change in the scale of the vertical axis.] Even at the high SNR cases, though, detection rates remain lower than the single depth detection rates (compare with Figure 7.2).

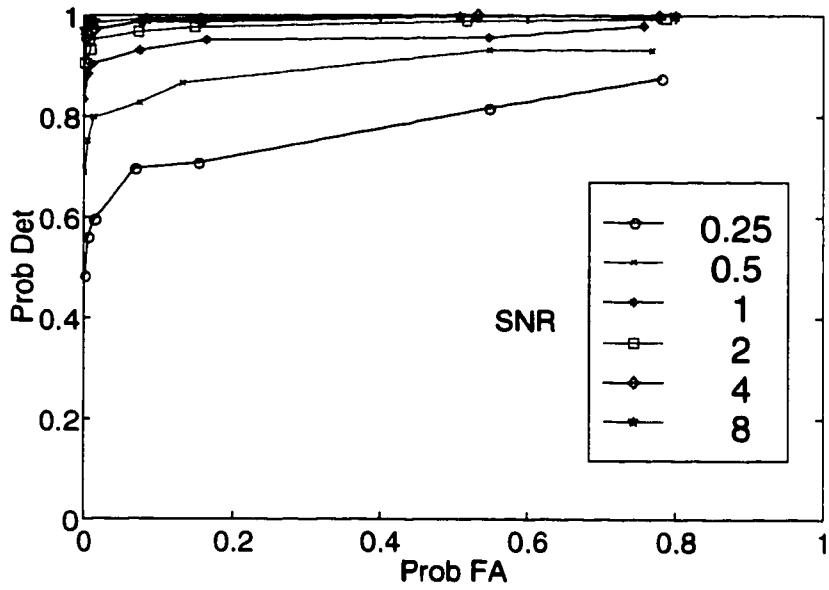

Figure 7.6. ROC curves from Fisher's test for periodicity.
Receiver operating curves for Fisher's test for periodicity for six signal to noise ratios (SNR) computed from 1000 different signal and noise realizations. The SNR values are computed as the ratio of the energy of the signal to the energy of the noise when the vibration is present. Fisher's test for periodicity was computed with 50% overlap between time segments with the following parameters: FFT length = 1024, tlen = 8192. Detection confidence levels were: [0.9 0.95 0.99 0.995 0.999 0.9995 0.9999].

7.1.4 Narrowband noise

It is expected that detection in narrowband noise will be less effective than in Gaussian noise because the detector has been designed for Gaussian noise. To determine the detector's effectiveness, though, ROC curves have been generated with both the single depth and dual depth methods when the noise is narrowband with a frequency that falls into the same wavelet scale as the signal's center frequency. The familiar Gaussian noise is added as well. All of the simulation parameters are the same as in the other ROC curves, except that here, SNR applies to the ratio of the energy in the signal to the energy in the narrowband noise. The energy of the Gaussian noise is a factor of 4 to 10 lower than the narrowband noise.

The ROC curve of Figure 7.7 is from the single depth detector. Notice that the detection percentage is much lower than its Gaussian noise equivalent (Figure 7.3); and the false alarm percentage is slightly lower.

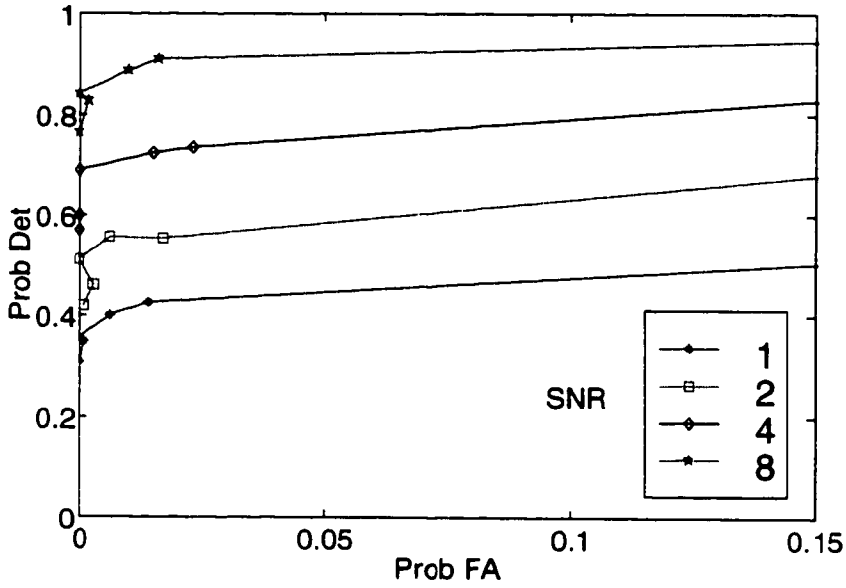

Figure 7.7. Single depth ROC curve with narrowband noise.
Receiver operating curves for the single depth detector with narrowband noise for four signal to noise ratios (SNR) computed from 1000 different signal and noise realizations. The SNR values are computed as the ratio of the energy of the signal to the energy of the noise when the vibration is present. The simulations were computed with: *seglen* = 75 ms, *viblen* = 150 ms, tlen = 8000, Fs = 8000 Hz, *nonoverlap* percentage of 25%, and wavelet scale overlaps in the frequency domain = -8.7 dB. Confidence levels were: 0.9999 0.99999 0.999995 0.999999 0.9999995 0.9999999.

The ROC curve of the dual depth detector in narrowband noise (Figure 7.8) has much better detection rates than its single depth counterpart (Figure 7.7). Notice again, though, that the detection percentage is much lower than its Gaussian noise equivalent (Figure 7.5). Notably, it is slightly higher than the results from Fisher's test for periodicity for single depths in Gaussian noise (Figure 7.6).

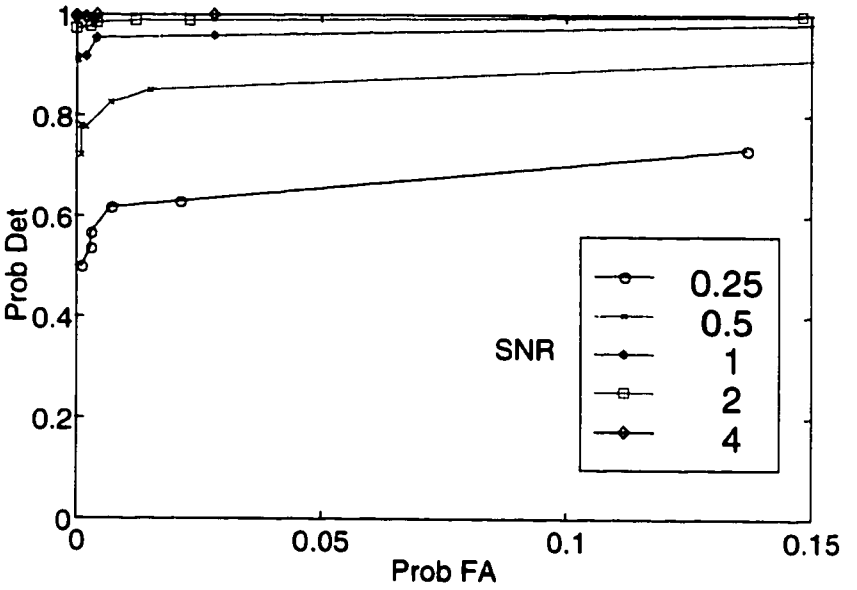

Figure 7.8. Dual depth ROC curve with narrowband noise.
Receiver operating curves for the dual depth detector with narrowband noise for six signal to noise ratios (SNR) computed from 1000 different signal and noise realizations. The SNR values are computed as the ratio of the energy of the signal to the energy of the noise when the vibration is present. The simulations were computed with: *seglen* = 75 ms, *viblen* = 150 ms, tlen = 8000, Fs = 8000 Hz, *nonoverlap* percentage of 25%, and wavelet scale overlaps in the frequency domain = -8.7 dB.Confidence levels: 0.9999 0.99999 0.999995 0.999999 0.9999995 0.9999999.

7.2 Simulated estimation

The expected accuracies of the CWT-based parameter estimation methods discussed in section 6.3 are extremely complicated to predict analytically, so Monte Carlo methods have been employed. Table 7.1 summarizes the results with discussion following. The poor frequency resolution of the wavelet transform at the high frequency ranges is responsible for the large bandwidth errors. The wavelet length is responsible for the large duration errors which will occur when the vibration itself is very short in duration.

All simulations were run with SNR at 10 to reduce the effects of noise. The wavelet scales were chosen to overlap at -3 dB in the frequency domain.

Table 7.1. Accuracy of parameter estimates.
"-" means underestimated and "+" means overestimated. Durations were uniformly distributed over 10 to 150 ms, except for frequency estimation (60 to 150 ms) and bandwidth estimation (20 to 150 ms). SNR = 10. Scales overlap at -3 dB in the frequency domain (*scalebase* = 1.22). 100 signals computed. Q_v values were 4-10 (Frequency), 1.5-10 (Frequency range, amplitude, intensity), 0.55-10 (duration, phase). f_v ranged from 250-750 Hz for frequency estimation, 200-750 Hz for bandwidth estimates, 150-750 Hz for amplitude and intensity measurements, and 100-1000 Hz for duration and time delay estimates.

	Mean % error	**Std % error**	**Max % error**	**Min % error**	**Comments**	**Measure-ment source**
Frequency	0.47	6.4	18	- 12	Narrowband ($Q_v > 4$)	Wavelet coefficients
Bandwidth	- 7.1	40	110	- 74	Both narrow and broad	Wavelet coefficients
Amplitude	- 5.8	6.9	4.6	- 32	Broadband	Wavelet coefficients
Duration	- 2.5	68	410	- 41	0.999999 confidence	Detected cross spectra
Time delay (from phase)	1.2	22	36	-195	0.999999 confidence	Detected cross spectra
Intensity	5.0	2.6	11	- 8.9	Broadband	Wavelet coefficients

7.2.1 Frequency range

Frequency estimate accuracy of narrowband signals was computed using Monte Carlo simulation from the single depth method. Restricting Q_v to fall between 4 and 10 with center frequencies from 250-750 Hz, and the duration to fall between 60 and 150 ms, ensured that the signal bandwidths were less than that of the wavelets. [A Hanning taper of 60 ms duration will have a bandwidth of approximately 24 Hz (see section 7.1.1.1.2). A Q_v of 4 at a vibration frequency of 250 Hz will yield a signal bandwidth of approximately 60 Hz. The largest bandwidth of any Hanning tapered simulated signal, therefore, would be 84 Hz (24 + 60). This is roughly the bandwidth of a Morlet wavelet with a center frequency of 250 Hz and a quality factor of 3.] The true

frequency is the center frequency used in the signal generation. The frequency estimate is the center frequency of the wavelet scale with the maximum wavelet coefficient magnitude. The iterations described in section 6.3.1.1 were not employed. Both the mean and standard deviation of this error are quite acceptable for arterial vibrometry.

Frequency range accuracy estimation of broadband signals was computed using Monte Carlo simulation from the single depth method. Here, Q_v was restricted to the range 1.5 to 10 with center frequencies from 200-750 Hz and durations from 20 to 150 ms to ensure that the entire signal would fall within the wavelet scales spanning 100-1000 Hz. [A Hanning taper of 20 ms duration will have a bandwidth of approximately 60 Hz. The largest signal bandwidth, approximately 140 Hz, would occur at a vibration frequency of 200 Hz and a Q_v value of 1.5. The resulting simulated vibration would have a bandwidth of 200 Hz and would span 100 to 300 Hz. Similarly, a vibration with a center frequency of 750 Hz would have a bandwidth of approximately 500 Hz for a total simulated signal bandwidth of 560 Hz. The simulated signal would span roughly 470 to 1030 Hz.] The estimates were determined with method 2) in section 6.3.1.2. The true bandwidth was taken to be the frequency range over which the signal was generated plus the Hanning bandwidth associated with the signal's duration.

7.2.2 Amplitude

The Monte Carlo simulation was run assuming broadband signals with the single depth method. Amplitude estimates were computed in the manner described in section 6.3.1.1. For these results, C_δ of equation (5-13) was set to 1. The mean and standard deviation of the error of the amplitude estimates is very small with respect to other sources of error (section 6.3.2).

7.2.3 Duration

To perform a Monte Carlo simulation of the accuracy of the duration estimation technique, signals were simulated for the dual depth method again with Q_v from 0.55 to 10 and center frequencies from 100-1000 Hz. Signal duration is estimated as the

length of the signal that exceeds the detection threshold. The true signal duration is taken to be the length of the generated signal.

Underestimation is expected since the signal will be weak at its outer edges and the entire length may not be detected. A small degree of overestimation may occur when the noise estimate is biased slightly low, thus exaggerating the edges of the signal. Overestimation will also occur when the wavelet scale corresponding to the signal frequency has a longer duration than the signal. This accounts for the extreme overestimations occasionally present above.

7.2.4 Time delay (from relative phase)

To perform a Monte Carlo simulation of the accuracy of the time delay estimation technique, the dual depth method was used again with Q_v from 0.55 to 10 and center frequencies from 100-1000 Hz. The nature of the percent error measurement will cause large errors when the true time delay approaches zero.

7.2.5 Intensity

Intensity accuracy is computed in the same manner as amplitude accuracy as described in section 6.3.1.3. Once more, these error ranges are quite low with respect to other sources of error.

7.3 Displaying the results

7.3.1 On M-mode image

New technologies are often best accepted and used by clinicians when they are straightforward extensions of familiar technologies. With that in mind, the chosen means of displaying the results of the vibration detector/estimator is on the familiar M-mode display. [Although, since the data is collected with the Doppler pulse rather than the imaging pulse, the depth resolution of the resulting "M-mode" display will be poorer than the traditional M-mode.] The time and depth location of detected vibrations are marked on a color overlay superimposed on the black and white M-mode display. Additionally, if sufficient processing power is available, color Doppler (CFI) can be implemented as well, and the results displayed to highlight the blood vessel location, blood velocity, and the timing of systole and diastole.

Frequency, amplitude, intensity, phase, duration, and timing in the heart cycle can all be shown on the color overlay; however, frequency, amplitude and intensity cannot be shown simultaneously on the same display. Duration and timing are read directly from the colorized regions. Frequency, amplitude and intensity values at each time-depth location of the M-mode display are revealed using a color scheme other than that used for color Doppler, such as brown through green. The current method is to display the frequency, amplitude or phase at which the maximum amplitude vibration appears (Figure 7.9).

Time delay is best viewed, not with color or with numbers, but with arrows indicating the direction of phase change. Since the phase estimates are made from consecutive depths, the phase arrow appears at each depth pointing in the temporal direction at which the current depth reached its peak with respect to the previous depth. The angle of the arrow indicates what percentage of the frequency cycle corresponds to that time delay. An arrow that is straight up implies that the depths experienced a peak simultaneously (Figure 7.10).

7.3.2 Frequency vs. depth, gated in the cardiac cycle

In order to indicate vibrations at multiple frequencies, as well as to indicate bandwidth, another display can be used either independently or in conjunction with the M-mode display. This display mode would continue to show depth on the vertical axis, but would have frequency (associated with the wavelet scales) on the horizontal axis, with color indicating the vibrations' amplitudes associated with the corresponding frequency range. The display could encompass a fixed time duration, such as during systole when vibrations are most likely. It would update at some adjustable persistence rate. Similarly, it could display some running average across multiple heart beats of vibrations occurring in a specific portion of the heart cycle, such as systole, with an adjustable time duration. An example from the first option is shown in Figure 7.11 for the same time series as shown in Figure 7.9 and Figure 7.10.

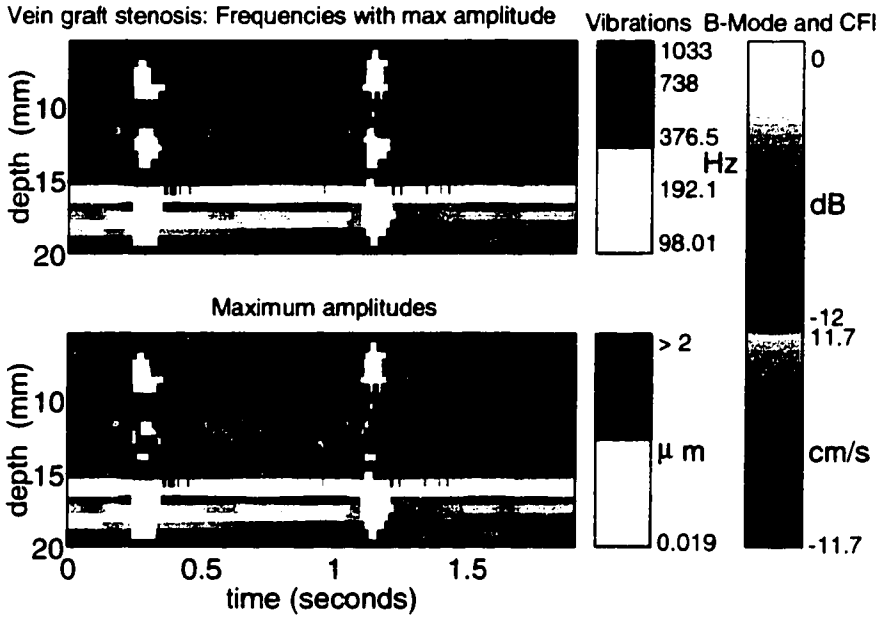

Figure 7.9. Display of vibrations in a human stenosed femoral vein graft. The black and white image is an M-mode display of the vein graft. The red and blue colors represent color Doppler processing and thus display the blood velocities. Detected vibrations are shown in the color scheme that ranges from brown to green. Detection was performed with the dual depth detector at a confidence level of 0.9999999. **Upper plot:** The vibration frequencies (wavelet scale center frequencies) are displayed for which the amplitude of vibration was greater than at other wavelet scales. **Lower plot:** Amplitudes corresponding to the frequencies shown in the upper plot are displayed in the brown-green color scheme.

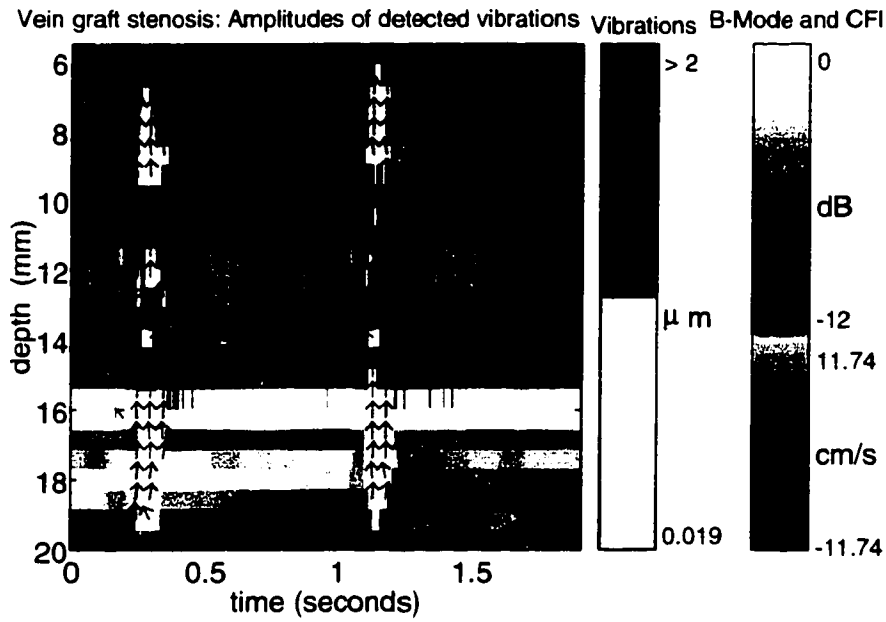

Figure 7.10. Display of phase measurements in a human stenosed femoral vein graft. This is the same display as in Figure 7.9 with the addition of arrows to indicate the phase relationship between the vibrations at neighboring depths. Phase shifts were taken directly from the angle of the cross wavelet spectrum at the wavelet scales with maximum vibration amplitude.

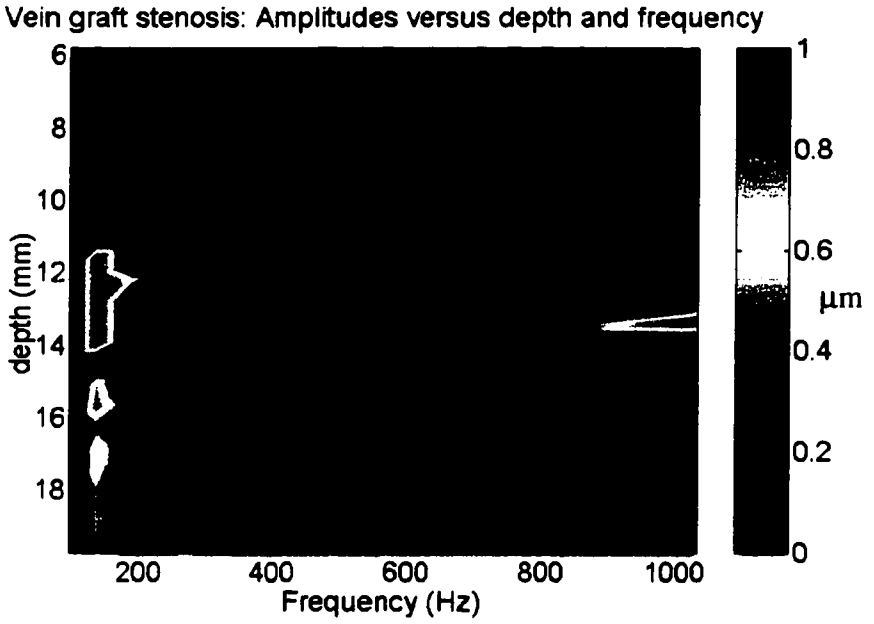

Figure 7.11. Depth versus frequency display of vibrations in a human stenosed femoral artery.
This is an alternative means of displaying the detected vibrations so that all detected frequencies and amplitudes can be shown. The vessel lumen lies between the 9 and 15 mm depths. The amplitude display has been clipped at 1 µm.

7.4 Results from vibration models

Signals have been detected with each of the vibration models described in section 4.4. The following sections present examples of the detected vibrations. All of these images were computed from the dual depth detector with wavelet scale center frequencies ranging from 100 Hz to 1000 Hz at a confidence level of 0.9999999, except where noted. For the cases in which the data were not collected at exactly 8000 Hz, but at PRF values close to 8000 Hz, the detection thresholds used were those computed assuming a PRF of 8000 Hz. The results reveal that the detector is sensitive to vibrations with amplitudes smaller than 30 nm. Furthermore, the detector is effective for both narrowband and broadband signals.

7.4.1 Humming

Humming provides realistic clutter and an ability to look at detection results from data both with and without the humming vibration signal. The following two images were taken at approximately the same location near the carotid artery of a human volunteer within seconds of each other.

The first image, Figure 7.12, is from data collected in which the subject was not humming. Brief vibrations were detected at 140 Hz. This would be the "healthy vibration" discussed in section 4.5.2. This vibration is very short, less than 25 ms, and at a low frequency range. Higher frequencies are detected in the depth ranges 0.8 mm, 2.2 mm, 5 to 7 mm and at 18 mm. These detection hits are suspect because they arise at depth locations which are anechoic (areas of weak ultrasound backscatter), resulting in very low signal to noise ratios. Anechoic regions are also highly sensitive to the reflection effect illustrated in Figure 4.19.

Figure 7.13 illustrates vibration detection while the volunteer rests, hums, and then rests again. Humming is present between time 0.06 and 0.17 seconds at approximately 275 Hz. The low frequency, 100 Hz vibrations at times 0.04 to 0.06 and 0.25 to 0.3 seconds likely arise from healthy arterial motion, or from the volunteer's intention to hum. The brief 385 Hz detection hit at 6 mm depth at time 0.05 to 0.056 seconds is likely a false alarm because it arises at only only depth-time location in the entire data set.

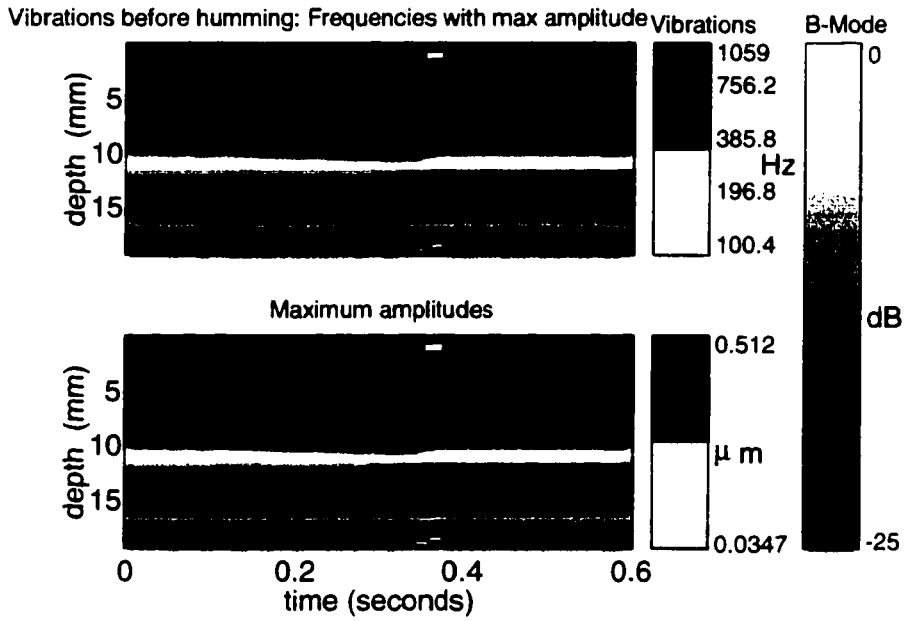

Figure 7.12. Vibrations detected in a human volunteer prior to humming. This data was taken from the carotid artery of a normal human volunteer at rest. The PRF was 8300 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

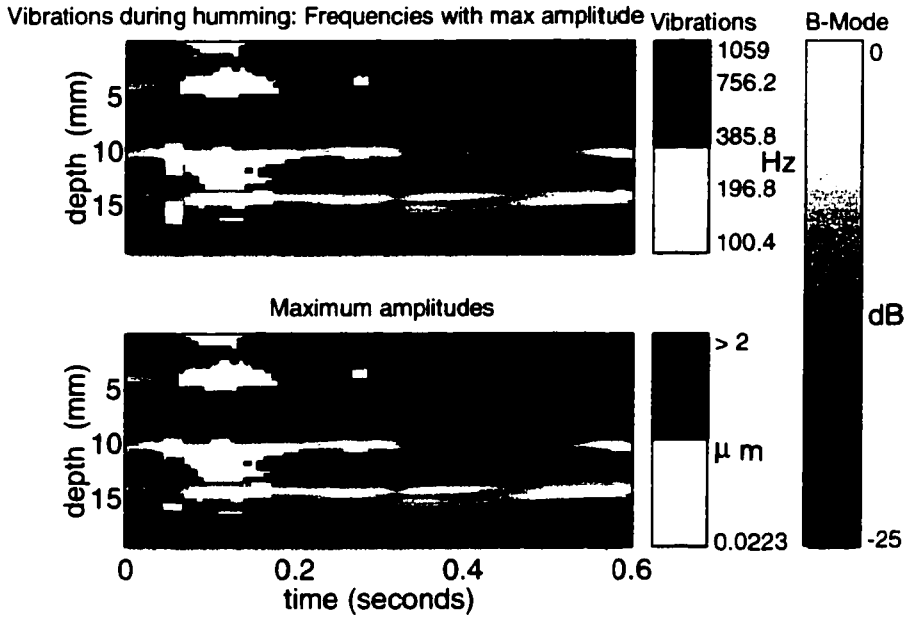

Figure 7.13. Vibration display of a humming human volunteer.
This data was taken from the carotid artery of a normal human volunteer while resting, humming, and the resting again. The PRF was 8300 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

7.4.2 Tissue vibrated with known external source

Images in Figure 7.14 and Figure 7.15 were collected from a region near the superficial femoral artery in a human volunteer. An audio speaker was placed under the volunteer's leg and driven by a signal generator transmitting 72 ms bursts (20 cycles) of a 275 Hz sinusoid at a rate of 1 Hz. Figure 7.14 is comprised of data captured from a very strong vibration from the external source. It shows some presence of lower frequency vibration at both the onset and decay of the vibration. The image in Figure 7.15 is of data collected from a much weaker signal from the external source. Vibrations are detected at some depths, but not all. The depths at which the vibrations were detected are mostly the depths with the strongest returned ultrasound echo and thus have potentially the strongest signal to noise ratio. The detection hits at 0.8 seconds are likely low frequency healthy vibration.

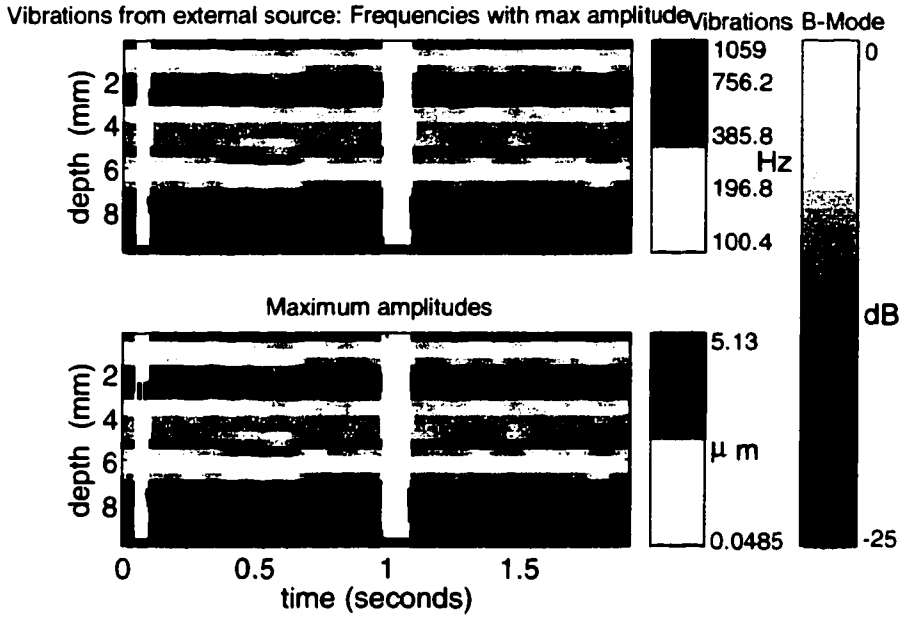

Figure 7.14. Display of strongly externally vibrated tissue.
This data was taken near the superficial femoral artery of a normal human volunteer. An audio speaker was placed beneath the leg and driven by a signal generator transmitting 72 ms bursts (20 cycles) of a 275 Hz sinusoid at a rate of 1 Hz with a large amplitude. The volunteer could feel the vibration and the sound was clearly audible. The PRF was 8300 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

7.4.3 Flow phantom with stenosis

The flow phantom with a stenosis provides vibrations that arise from the same source mechanism as an arterial disorder (section 4.4.2). These vibrations are broadband at unknown frequencies, durations, and amplitudes. Figure 7.16 is an example of vibrations detected in this phantom.

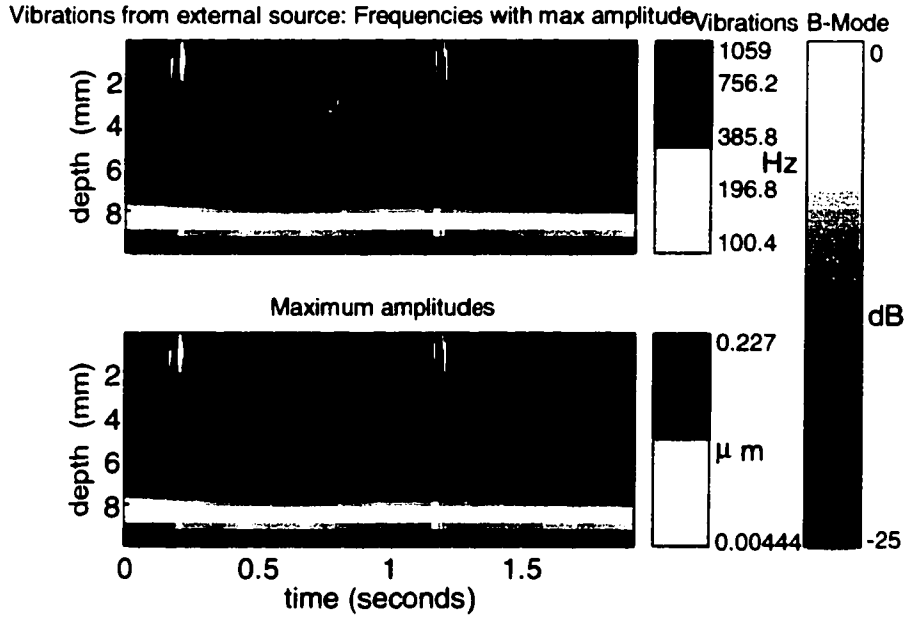

Figure 7.15. Display of vibrations from weakly externally vibrated tissue. This data was taken near the superficial femoral artery of a normal human volunteer. An audio speaker was placed beneath the leg and driven by a signal generator transmitting 72 ms bursts (20 cycles) of a 275 Hz sinusoid at a rate of 1 Hz with a small amplitude. The sound was no longer audible and the volunteer could not feel the vibration. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999. Again, the PRF was 8300 Hz.

7.4.4 *In vivo* Stenosis

The vibrations shown in Figure 7.17 from a human stenosed infrainguinal vein bypass graft have among the highest signal to noise ratios of all *in vivo* vibrations collected. Arterial blood flow is shown in red and blue to highlight the location of the artery. These blood velocities were computed with a method similar to color Doppler processing (section 4.1.4.5).

Vibrations from a different vein bypass graft are shown in Figure 7.18. These vibrations are somewhat broadband which results in low signal to noise ratios at each wavelet scale. Thus, the vibrations are only detected with a low confidence level and many false alarms appear randomly throughout the image.

7.4.5 Pseudoaneurysm

Vibrations detected in a pseudoaneurysm are shown in Figure 7.19. The cross wavelet spectra were generated with data taken from depths spaced 1.8 mm apart to ensure that detected signals are only from vibrations that are spatially distributed, as opposed to noise or blood flow.

7.4.6 Bleed

Figure 7.20 was taken in an undisturbed porcine femoral artery. Subsequently, Figure 7.21 was taken after an arterial puncture, in the same general area, but slightly distal to the location of Figure 7.20. The ultrasound beam was placed near the puncture site. In both figures, the arterial blood flow is displayed in red and blue to highlight the location of the artery. Blood velocities were computed with a method similar to color Doppler processing (section 4.1.4.5).

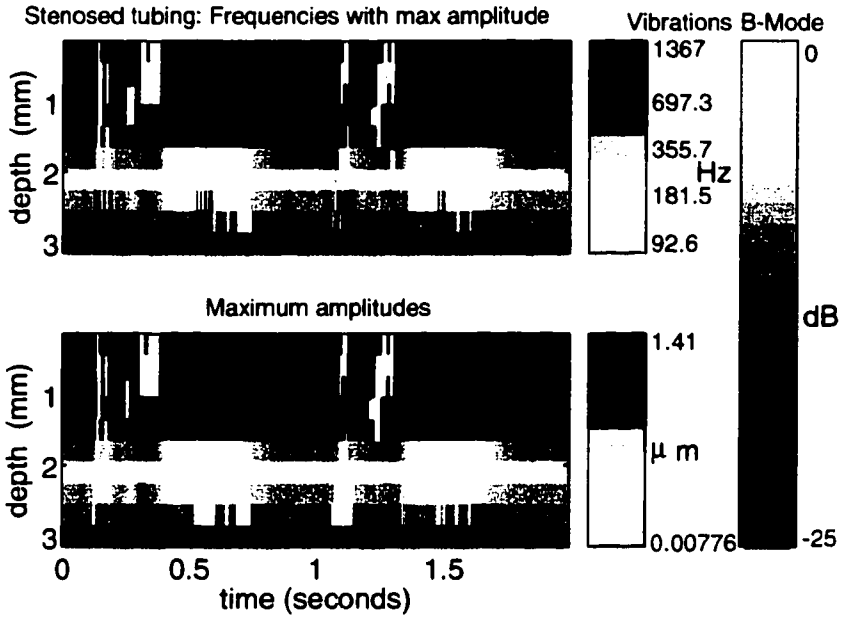

Figure 7.16. Display of vibrations in a stenosis phantom.
This data set was taken from the silicon tubing of the stenosis phantom. The first 2.5 mm are taken from the tubing material, with the fluid beginning at 2.5 mm and deeper. [The tubing itself is anechoic; only the tubing-blood interface produces a strong ultrasound backscatter.] The "vibrations" detected below 2.5 mm are not vibrations of solid tissue, rather they indicate turbulence in the blood flow that generated Doppler shifts in both directions: toward and away from the transducer. The PRF was 15000 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

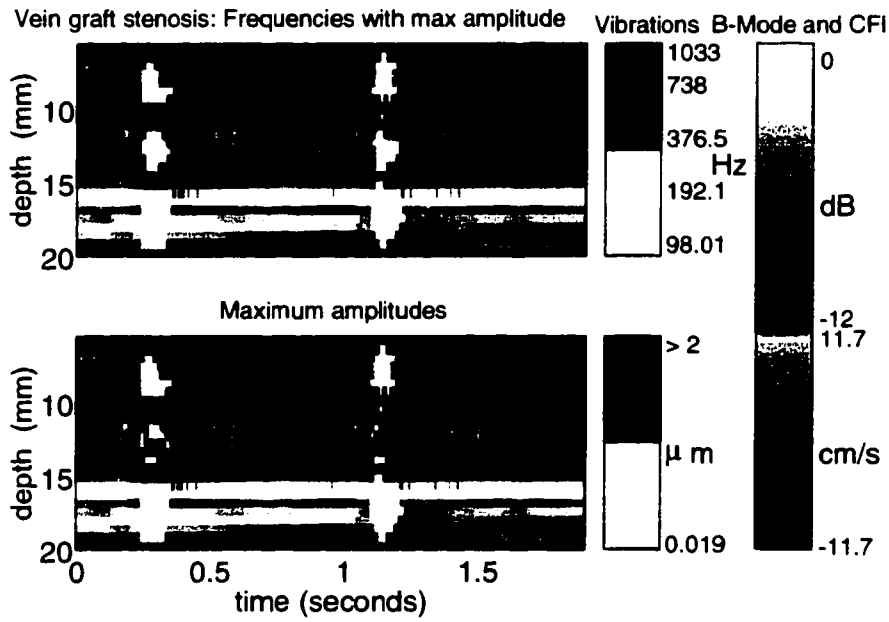

Figure 7.17. Display of vibrations from a human stenosed vein bypass graft. This is the same figure as shown in section 7.3.1. It is repeated here for the reader's convenience. The PRF was 8100 Hz. The detection was computed with a confidence level of 0.9999999.

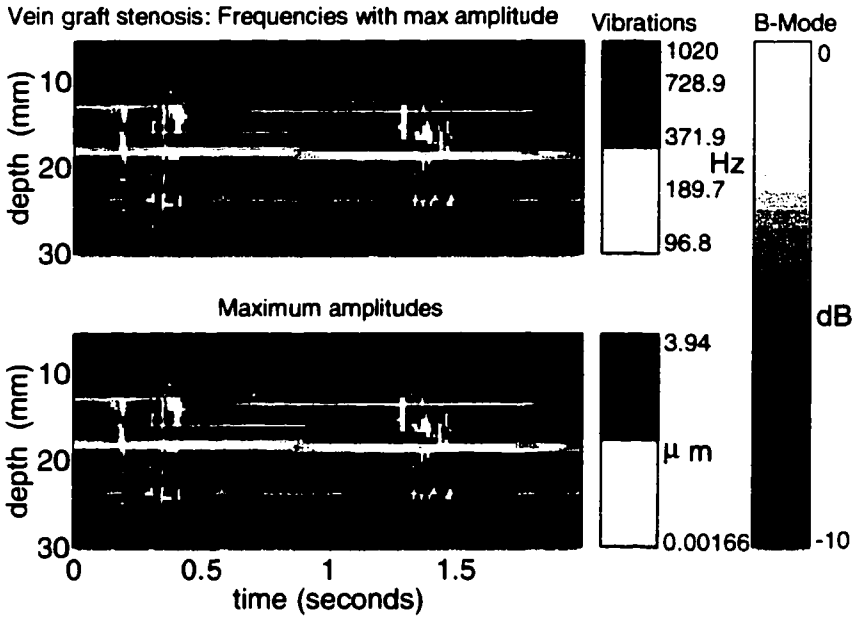

Figure 7.18. Display of vibrations detected from a human vein bypass graft. This data set was taken from another human vein bypass graft. The PRF was 8000 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9995.

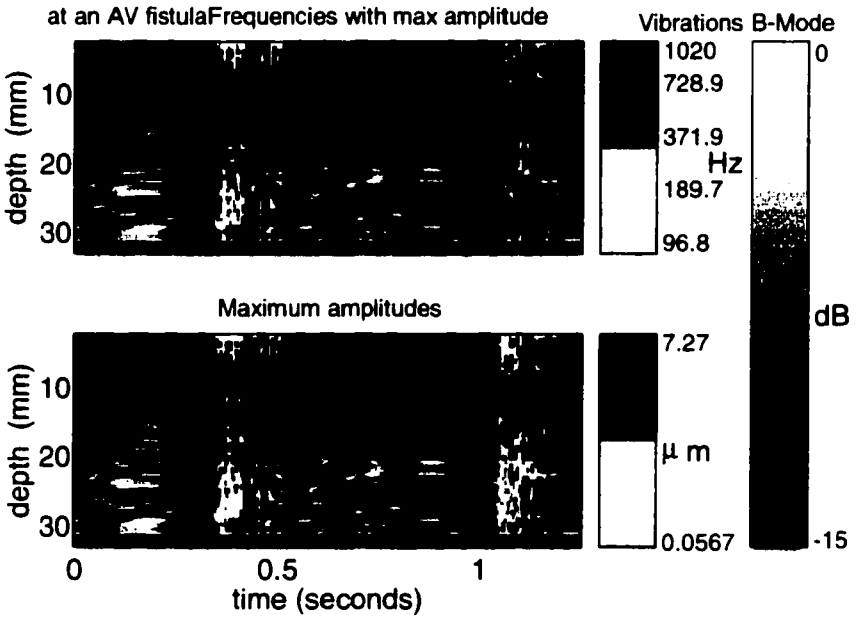

Figure 7.19. Display of vibrations detected in a human pseudoaneurysm.
This data set was taken near a human pseudoaneurysm. The PRF was 8000 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

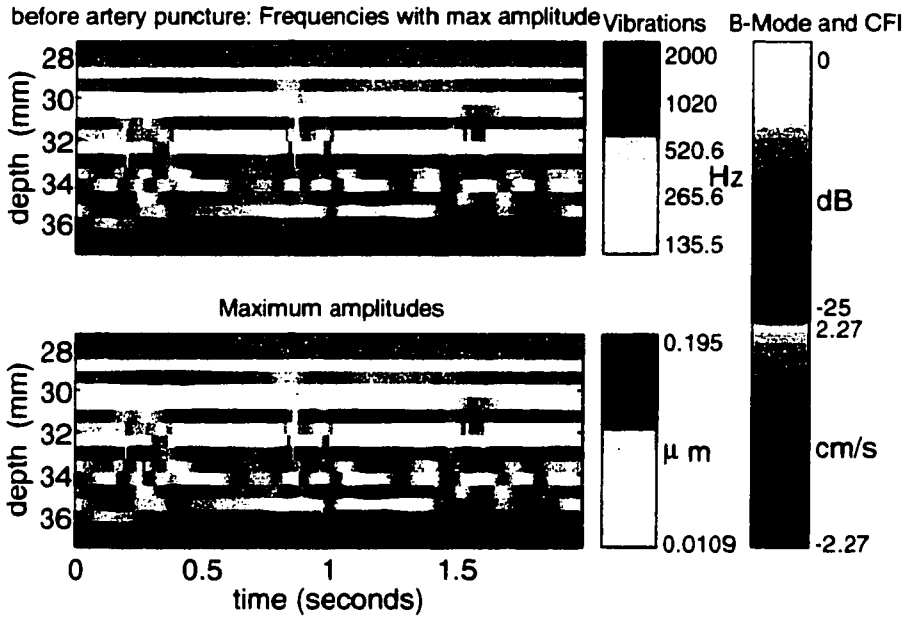

Figure 7.20. Display of vibrations detected in a porcine femoral artery prior to arterial puncture.

This data set was collected from the femoral artery of an anaesthetized pig. These vibrations were detected with wavelet scale center frequencies ranging from 135 to 2000 Hz to eliminate detection of the normal low frequency arterial vibration that occurred up to 100 Hz and to ensure that vibrations above 1000 Hz would be detected if present. The PRF was 8000 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

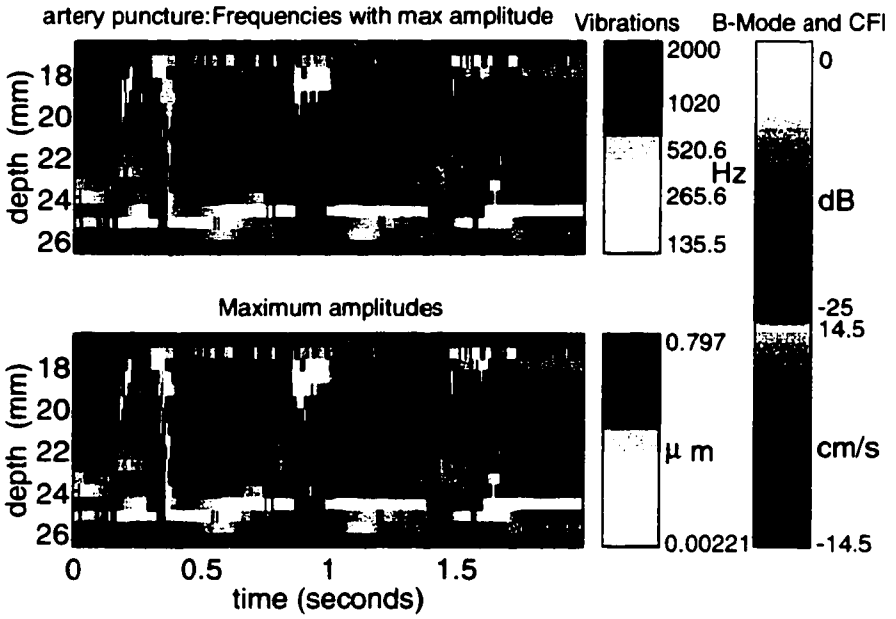

Figure 7.21. Display of vibrations detected in a porcine femoral artery just subsequent to arterial puncture.
This data set was collected from the femoral artery of an anaesthetized pig subsequent to an arterial puncture in the general region of the puncture. These vibrations were detected with wavelet scale center frequencies ranging from 135 to 2000 Hz to eliminate detection of the normal low frequency arterial vibration that occurred up to 100 Hz and to ensure that vibrations above 1000 Hz would be detected if present. The PRF was 8000 Hz. The display was generated in the manner described in section 7.3.1 with a confidence level of 0.9999999.

8 Conclusion

8.1 Evaluation of the method

The detector and estimator can be evaluated from three aspects: data collection, data processing, and clinical applicability. The strengths and weaknesses of each aspect are summarized below.

8.1.1 Strengths of the method

8.1.1.1 Data Collection with Doppler ultrasound

The choice of using ultrasound to collect the vibration data has several advantages besides the ability to capture the data. First, all ultrasound instruments are somewhat portable, and some are extremely portable. Therefore, the instruments can be taken to the patient. This has many clinical benefits such as scheduling flexibility, no time lost in transferring a patient, and applicability to the unstable patient who cannot be transported. Also, compared to other imaging modalities such as magnetic resonance (MR) and computed tomography (CT), ultrasound offers a much lower cost exam; although, neither CT nor MR is feasible for vibrometry because CT only applies to static imaging and functional MR does not have the necessary time resolution. Additionally, since ultrasound is routinely used for complete vascular exams, it is advantageous to use ultrasound for arterial vibrometry as well.

Unlike other ultrasound transmission modes, single line pulsed Doppler ultrasound has the advantage of uninterrupted sampling of the vibration signals at an adjustable sampling rate. Therefore, detection and estimation can be carried out continuously without missing vibrations that may quickly pass by the ultrasound beam. Another advantage of pulsed Doppler ultrasound that was mentioned earlier warrants repeating. Since it interrogates small regions of tissue, signals collected from Doppler ultrasound will be affected minimally by noise that exists in the surrounding tissue regions.

8.1.1.2 Signal processing

Basing the detector and estimator on the wavelet transform affords several advantages. First, the constant quality factor nature of the wavelet transform nicely matches

arterial vibrations reasonably well (section 5.1.2.2). The filter bank nature of the continuous wavelet transform also elegantly filters out noise and clutter from frequency ranges that are not of interest for subsequent detection and noise variance estimation. Thus, complications due to low frequency clutter and colored noise are inherently avoided. In addition, because the wavelet transform does not require stationary signals, it lends itself nicely to adaptive thresholding in non-stationary noise. Furthermore, the mother wavelet can be chosen to approximate the desired signal to accomplish matched filtering. Therefore, the continuous wavelet transform with a tapered sinusoidal shape, specifically the complex Morlet wavelet, has the necessary requirements for application to transient detection of oscillations that resemble sinusoids in low frequency clutter and in noise which is neither white nor stationary. It should also be noted that since the wavelet transform relies on a filter bank, its implementation is easily parallelized.

8.1.1.3 Clinical Applicability

Detection and estimation of arterial vibrations as described in this document have many clinical strengths. Simply stated, these methods provide a non-invasive means of quickly collecting, detecting and displaying vibrations that have been heard with a stethoscope. It allows the clinician to quickly determine the subcutaneous location of the vibration as well as its other parameters, including vibration amplitude and intensity which are not available with any other modality.

8.1.2 Limitations of the method

8.1.2.1 Data collection with pulsed Doppler ultrasound

A key weakness of the use of pulsed Doppler ultrasound for vibrometry is its interrogation of only one spatial line. The consequence is a longer examination time and the potential for missing vibrations that neighbor, but do not intersect, the ultrasound beam.

There is another limitation of the method that should be re-iterated. Amplitude and intensity estimates will only correspond to that component along the ultrasound beam, and may be confounded by unexpected reflection from stationary echoes. True

amplitudes and intensities will almost always exceed the estimate due to this limitation.

8.1.2.2 Signal processing

One potential weakness of the detector with respect to signal processing is its computation cost. With the suggested wavelet transform implementation, the dot product method (see section 8.2), vibrometry detection requires approximately 2,000,000 real additions and 1,000,000 real multiplications per depth location per second (see Appendix G). This value indicates that a modern, high-end instrument capable of processing a large color Doppler ultrasound box (128 ray lines, 400 depths, 10 point ensemble, 10 frames per second) could process approximately 100 spatial locations of vibrometry detection at a PRF of 8000 Hz, and 400 spatial locations at a PRF of 2000 Hz (see Appendix G). Thus, modern, high-end instruments are currently capable of accomplishing vibrometry. Lower cost instruments, however, may not yet have the necessary processing power. For further comparison, the CWT-based vibrometry computational cost requires anywhere from 22% to 117% more real multiplications and 147% to 334% more real additions per depth per second than its short-time Fourier transform counterpart, Fisher's test for periodicity (see Appendix G). The increase depends on the sampling rate and clutter filtering implementation. As computational power continues to rapidly increase with decreasing prices, this potential limitation becomes less of an issue.

The assumption that there is an arbitrary maximum duration of the vibration at each wavelet scale of interest is limiting to detection of continuous and pansystolic arterial vibrations (see section 2). These long vibrations, though, would be more easily detectable with conventional Fourier techniques than the short duration vibration.

8.1.2.3 Clinical Applicability

A fundamental limitation of vibrometry in detecting arterial vibrations is that arterial disorders do not always vibrate, and when they do vibrate, the vibrations do not necessarily vibrate at each heart beat indefinitely. That is, as the patient's physiology

changes over time, vibrations may occur, subside and re-occur. These changes may take place over minutes, hours, or days.

There are several factors that confound the results of the detector. Some relate to the detector itself, others to vibrations. Detection false alarms may arise due to two-sided blood flow which will be detected as vibration when the blood signal is sufficiently strong. Blood flow often appears two-sided when the ultrasound beam is nearly perpendicular to the direction of flow, or when turbulence is present. Another detector issue arises when the vibration signal occurs at the same wavelet scale as narrowband noise. When this occurs, the vibration will be detected only when the signal strength is sufficiently strong with respect to the narrowband noise (Figure 7.8).

The final comment is not exactly a limitation, but reduces the immediate clinical utility of the method: it is not yet clear how the new amplitude and intensity information directly relate to diagnosis or prognosis.

8.1.3 Automation vs. interaction

Although detection and estimation are highly automated in the methods outlined in this document, some interaction of the clinician is still required. Specifically, the confidence level should be lowered until either 1) vibrations are detected at multiple depths at fixed portions of the heart cycle in consecutive heart beats, or 2) false alarms appear randomly in depth and time. This interaction is similar to that required for standard pulsed Doppler and color Doppler ultrasound exams.

It was mentioned earlier that occasionally healthy arteries vibrate at the low frequency range (~100 Hz). If this is clearly occurring, then the clinician may choose to turn off the wavelet scale at this low frequency.

8.2 Suggested implementation

Since the incomplete dot product wavelet transform implementation with a nonoverlap percentage of 25% for Q_w of 3 has an excellent fit to the statistical approximations of section 6.2.1, and relatively fewer computational requirements than other wavelet transform implementations, it is the method of choice. A wavelet quality factor of 3 is suggested for arterial vibrometry because it achieves a balance between two important

tradeoffs. First, the advantages of a low quality factor are needed to minimize the number of scales required to span a certain frequency range. A low quality factor also corresponds to a short wavelet which is needed for good duration estimation. The second tradeoff consideration is that a large quality factor is needed to accomplish the sharp frequency cutoff required to attenuate low frequency clutter while passing the lower frequency vibrations. The Q_w value of 3 has the added advantage of achieving an excellent fit to the expected statistics, despite the approximations used. Finally, a value of three is a compromise for approximately matched filtering both the vibrations with low quality factor (~1) and those with high quality factor (>10).

8.3 Future work

8.3.1 Clinical implementation

Several events need to take place before this detector and estimator can be used in clinical practice.

8.3.1.1 Real time implementation

In order to make use of vibrometry clinically, the clinician must receive the vibrometry results in real time. Although the current algorithm has been implemented in MATLAB for use off-line, its transportation to C or to a DSP would be straightforward for implementation on an ultrasound instrument that works with digital data.

8.3.1.2 Tests to be done with real-time system

Once a real-time system is in effect, two tests should be carried out. First, the method should be applied to a wide range of arteries in normal patients to determine the prevalence of vibration in normals. Second, a comparison should be made of the sensitivity of this detector with the sensitivity of a variety of human observers listening with stethoscopes.

8.3.1.3 Suggestions for clinical trials

After the steps of section 8.3.1.2 have been carried out, a specific disorder should be selected and clinical trials should be done to determine the relevance of vibrometry to

the disorder diagnosis and treatment planning. The goal of the clinical trials would be to 1) determine the prevalence of vibration associated with the disorder, 2) determine the required frequency range for the detector, and 3) explore the relationship between vibration parameters and the significance of the disorder.

8.3.2 Future advances to the detector

8.3.2.1 Multi-line implementation

To rapidly speed up vibration detection, a pulsed Doppler transmission system could be developed to transmit multiple beams in parallel, with the transmission sequences interleaved. For instance, if vibrations from depths less than 8 cm are of interest (corresponding to 106 μs for round trip ultrasound propagation), then the following sequence of events could be used (Figure 8.1):

1) A pulse is transmitted along line 1.
2) This line is sampled, for example at 0.6 mm spacing until 106 μs have elapsed.
3) Advance the ultrasound beam to the next parallel location.
4) A pulse is transmitted along this new line.
5) This line is sampled until 106 μs have elapsed.
6) Repeat steps 3-5 twice more.
7) Wait 76 μs until 500 μs have elapsed since step 1.
8) Repeat steps 1-7 continuously.

With these steps, each line is sampled effectively at a 2 kHz PRF (1/500μs), yet four spatial lines are interrogated with 106 μs delays between them.

Another scheme currently used in commercial ultrasound instruments, called "explososcan" (Shattuck et al. 1984), could be added to increase the number of spatial lines. In this scheme, each transmission would use a somewhat wide ultrasound

transmit beam overlapping multiple, relatively narrower, receive beams (Figure 8.2). These receive beams would be sampled simultaneously.

Both of these options would require increased computational power.

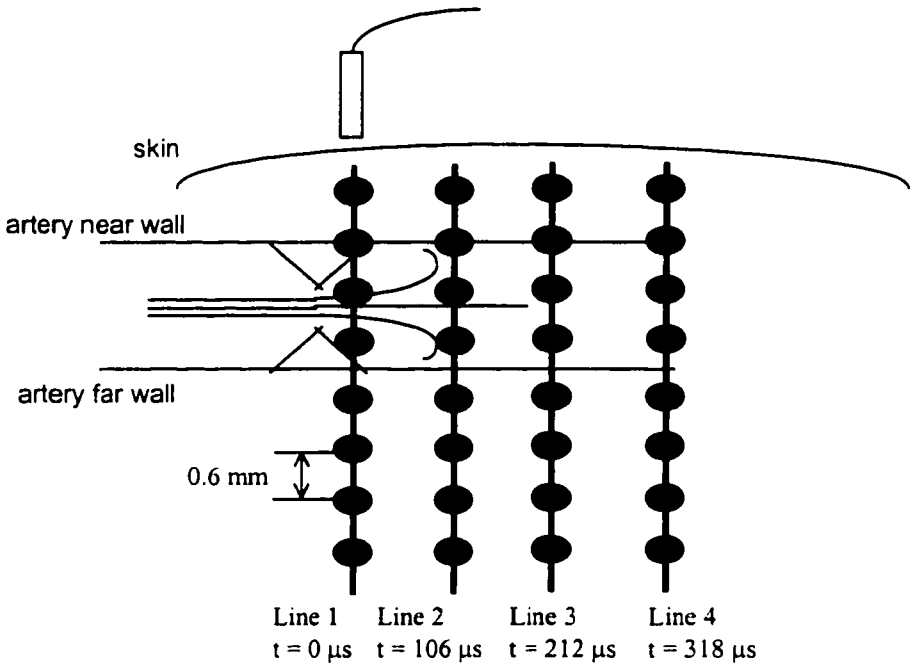

Figure 8.1. Illustration of multi-line vibrometry data collection.
This figure illustrates collection of vibrations from several ultrasound beam lines collected in an interleaved manner. First, a pulse is transmitted along Line 1 and the received echoes digitized. Then, pulses are sent, received and digitized along each of three additional, parallel beam lines. Upon completion of the fourth beam line, the process is repeated continuously.

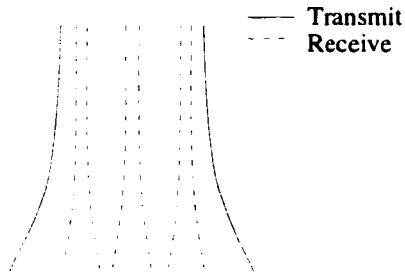

Figure 8.2. Illustration of "explososcan."
This figure illustrates the process of transmitting with a wide beam and receiving simultaneously with several smaller beams.

8.3.2.2 Phase randomness

Phase measurements may offer a means to differentiating typical blood flow from both vibrations and blood flow with eddies. In the absence of eddies, blood flow is expected to yield random phase measurements across both time and depth; whereas, the phase of vibrations and eddies are expected to have some correlated structure. Exploiting this difference may prove to be the required extra detection step to avoid detection hits from typical blood flow.

8.3.2.3 Three dimensional vibration measurement with Vector Doppler

The fact that vibration amplitude and intensity measurement occurs only along the ultrasound beam line is a significant drawback to their clinical utility. The measurement could be improved, however, if it were made from a cross-beam technique such as the Vector Doppler method described in Dunmire (1998 and 1999) and many United States patents (Beach et al. 1998, Vilkomerson 1996, Banjanin and Shahmirian 1995, Beach and Overbeck 1995, Hall and Bernardi 1995). With vector Doppler methods the true three dimensional amplitude and intensity values could be determined. Unfortunately, Vector Doppler methods significantly increase the costs of hardware, transmission, reception and computation.

8.4 Summary

The vibrometry method outlined here uses pulsed ultrasound to capture potential vibrations. Quadrature signals demodulated from the received ultrasound echoes are parsed according to depth in tissue for subsequent processing. The detector operates on either the quadrature time series from a single depth or on quadrature time series from two neighboring depths. The estimator uses the phase of the quadrature time series to make displacement measurements. Both the detector and estimator are based on the use of the continuous wavelet transform as a filter bank with the complex Morlet wavelet used to approximate a matched filter for the transient vibration. A normalized wavelet power spectrum is used for detection in a single time series and a normalized cross wavelet spectrum is used for dual depth detection. Both of these detectors have excellent receiver operating characteristics (ROC curves) as compiled from Monte Carlo simulation: greater than 99.5% detection rates at less than 1 % false alarm rates. These detectors have also been shown to be effective on various vibration models, including examples of vibrations induced by true *in vivo* arterial disorders. Some of the estimates of vibration parameters captured in the data can be made arbitrarily accurate for data with large SNRs with a tradeoff to computational cost: amplitude, frequency, bandwidth and intensity. Duration estimation, however, is limited by the wavelet duration. The estimate of time delay between two neighboring depths depends on the degree to which the signals in the two depths are correlated.

In the clinic, arterial vibrometry lends objectivity to the previously subjective, though routine, practice of auscultating bruits and murmurs with stethoscopes. It is potentially more sensitive than traditional auscultation, but this determination will require additional research. Certainly, though, vibrometry is significantly more resistant to environmental noise than the stethoscope method. Furthermore, unlike the stethoscope, it specifically locates the vibration anatomically. With real-time implementation of the algorithms described here, arterial vibrometry can begin to be

used in the clinic. Implementation of the additional advances described above, however, would greatly enhance its clinical utility.

In the laboratory where arterial disorders are studied, the detector and estimator can make measurements that previously required invasive procedures. Not only will this allow measurements to be taken on animals without disturbing the experiment, but it will allow measurements on human patients in a non-invasive, efficient manner.

Finally, the arterial vibrometry detector and estimator described in this document provide a means of automatically detecting and quantitating oscillations with a thoroughness and sensitivity not previously available. Although these methods have been designed specifically for arterial vibrometry, they apply to any situation in which the goal is to detect and/or estimate transient oscillations occurring in low frequency clutter and Gaussian noise. This Gaussian noise may be colored, non-stationary and have unknown variance.

LIST OF REFERENCES

Aaslid, R. and Nornes, H. (1984) Musical murmurs in human cerebral arteries after subarachnoid hemorrhage. *Journal of Neurosurgery*, 60, 32-36.

Agarwal, R., Agrawal, S.K., Roubin, G.S. , et al. (1993) Clinically guided closure of femoral arterial pseudoaneurysms complicating cardiac catheterization and coronary angioplasty. *Catheterization and cardiovascular diagnosis*, 30, 96-100.

Akay, M., Welkowitz, W., Semmlow, J.L., Akay, Y.M. and Kostis, J. (1992) Noninvasive acoustical detection of coronary-artery disease using the adaptive line enhancer method. *Medical and Biological Engineering and Computing*, 30, 147-154.

Akay, M. (1994) Automated noninvasive detection of coronary artery disease using wavelet-based neural networks. In *Intelligent Engineering Systems Through Artificial Neural Networks, Volume 4*. New York: ASME Press, 517-522.

Akay, Y.M., Akay, M., Welkowitz, W., Semmlow, J.L. and Kostis, J.B. (1993) Noninvasive acoustical detection of coronary-artery disease -- A comparative study of signal-processing methods. *IEEE Transactions on Biomedical Engineering*, 40, 571-578.

Akay, Y.M., Akay, M., Welkowitz, W. and Kostis, J. (1994) Noninvasive detection of coronary artery disease. *IEEE Engineering in Medicine and Biology Magazine*, 13, 761-764.

Akin, A. (1998) *Non-invasive detection of spike activitity of the stomach from cutaneous EGG*. Ph.D. dissertation, Drexel University.

Akin, A. and Sun, H.H. (1999) Time-frequency methods for detecting spike activity of stomach. *Medical and Biological Engineering and Computing*, 37, 381-390.

Bailey, T.C., Sapatinas, T., Powell, K.J. and Krzanowski, W.J. (1998) Signal detection in underwater sound using wavelets. *Journal of the American Statistical Association*, 93, 73-83.

Baker, D.W. and Simmons, V.E. (1968) Phase track techniques for detecting arterial blood vessel wall motion. *Proceedings of the 21st Annual Conference on Engineeringin Medicine and Biology*, Houston, Texas, 8.6.

Banjanin, Z.B. and Shahmirian, V. (1995) *Angle independent doppler in ultrasound imaging*. United States Patent: US5454372.

Basbug, F. (1993) *Wavelet transform-based digital signal processing (detection)*. Ph.D. dissertation, Florida Atlantic University.

Basoglu, C., Managuli, R., York, G. and Kim, Y. (1998) *Computing requirements of modern medical diagnostic ultrasound machines*, 24, 1407-1431.

Bates, B. (1991) *A guide to physical examination and history taking*. Philadelphia: Lipponcott.

Beach, K.W. (1993) Evaluating a pulsed Doppler duplex scanner. In *Vascular diagnosis, fourth edition*, edited by Bernstein, E.F., St. Louis: Mosby-Yeak Book, Inc, 133-148.

Beach, K., McMorrow, G.J. and Barnard, W.L. (1998) *Vector Doppler system for stroke screening*. United States Patent: US5738097.

Beach, K. and Overbeck, J. (1995) *Vector doppler medical devices for blood velocity studies*. United States Patent: US5409010.

Bonnefous, O. and Pesque, P. (1986) Time domain formulation of pulse Doppler ultrasound and blood velocity estimation by cross correlation. *Ultrasonic Imaging*, 8, 73-85.

Bowen, T. (1987) Acoustic radiation temperature for noninvasive thermometry. *Automedica*, 8, 247-267.

Brands, P.J., Hoeks, A.P.G., Ledoux, L.A. and Reneman, R.S. (1997) A radio frequency domain complex cross-correlation model to estimate blood flow

velocity and tissue motion by means of ultrasound. *Ultrasound in Medicine and Biology*, 23, 911-920.

Campbell, T.P. (1996) *Underwater transient detection and localization using the wavelet transform*. Master's thesis, Royal Military College of Canada.

Chapman, B. and Charlesworth, D. (1981) Analysis of the factors influencing the production of sound in tubes. *Vascular diagnosis and therapy*, 2(3), 23-27.

Chen, D., Harris, J.G., Laine, A.F. and Huda, W. (1998) A wavelet CFAR detector for mass detection in mammography. *Proceedings of the SPIE*, 3391, 637-648.

Clark, C. (1977) Turbulent wall pressure measurements in a model of aortic stenosis. *Journal of Biomechanics*, 10, 461-472.

Couch, L.W. (1990) *Digital and analog communication systems, Third edition*. New York: Macmillan Publishing Company.

Daubechies, I. (1992) *Ten lectures on wavelets*. Pennsylvania: Society for Industrial and Applied Mathematics.

Dickinson, R.J. and Hill, C.R. (1982) Measurement of soft tissue motion using correlation between A-scans. *Ultrasound in Medicine and Biology*, 8, 263-271.

Donoho, D.L. and Johnstone, I.M. (1994) Ideal spatial adaptation via wavelet shrinkage. *Biometrika*, 81, 425-455.

Dunmire, B. (1998) *The vector doppler--an ultrasound instrument for 3-D blood velocity measurements: in vitro steady flow studies*. Master's thesis, University of Washington.

Dunmire, B. (2000) Cross beam vector Doppler for angle independent velocity measurements (a review). *Ultrasound in Medicine and Biology*, accepted.

El Mashade, M.B. (1998) Detection analysis of linearly combined order statistic CFAR algorithms in nonhomogeneous background environments. *Signal Processing*, 68, 59-71.

Englehart, K., Hudgins, B., Parker, P.A. and Stevenson, M. (1999) Classification of the myoelectric signal using time-frequency based representations. *Medical Engineering and Physics*, 21, 431-438.

Ezrow, D.H. (1962) Measurement of the thermal-noise spectrum of water. *The Journal of the Acoustic Society of America*, 34, 550-554.

Farge, M. (1992) Wavelet transforms and their applications to turbulence. *Annual Review of Fluid Mechanics*, 24, 395-457.

Ferguson, G.G. (1972) Physical factors in the initiation, growth, and rupture of human intracranial saccular aneurysms. *Journal of Neurosurgery*, 37, 666-677.

Fillinger, M.F., Kerns, D.B., Bruch, E.R., Reinitz, E.R. and Schwartz, R.A. (1990) Does the end-to-end venous anastomosis offer a functional advantage over the end-to-side venous anastomosis in high-output arteriovenous grafts? *Journal of Vascular Surgery*, 12, 676-688.

Foreman, J.E.K. and Hutchison, K.J. (1970) Arterial wall vibration distal to stenoses in isolated arteries of dog and man. *Circulation Research*, 26, 583-590.

Frisch, M. and Messer, H. (1994) Detection of a known transient signal of unknown scaling and arrival time. *IEEE Transactions on Signal Processing*, 42, 1859-1863.

Gabor, D. (1946) Theory of communication. *The Journal of the Institution of Electrical Engineers*, 93-III, 429-457.

Gallaire, J.G. and Sayeed, A.M. (1998) Interference-resistant detection based on time-frequency subspaces. *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, 389-392.

Greenberger, J.J. and Hinthorn, D.R. (1993) *History taking and physical examination*. Chicago: Mosby.

Hall, A.L. and Bernardi, R.B. (1995) *Method for detecting two-dimensional flow for ultrasound color flow imaging*. United States Patent: US5398216.

Hartley, C.J., Hanley, H.G., Lewis, R.M. and Cole, J.S. (1978) Synchronized pulsed Doppler blood flow and ultrasonic dimension measurement in conscious dogs. *Ultrasound in Medicine and Biology*, 4, 99-110.

Hasegawa, J. and Kobayashi, K. (1995) Noninvasive diagnosis of cerebrovascular diseases based on the characteristics of blood flow noise. *IEICE Transactions on Fundamentals*, E78-A, 1634-1639.

Heimdal, A. and Torp, H. (1997) Ultrasound Doppler measurements of low velocity blood flow: Limitations due to clutter signals from vibrating muscles. *IEEE transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 44, 873-881.

Hein, I.A. and O'Brien, Jr, W.D. (1993) Current time-domain methods for assessing tissue motion by analysis from reflected ultrasound echoes -- A review. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 40, 84-102.

Hoeks, A.P.G., Brands, P.J., Smeets, F.A. and Reneman, R.S. (1990) Assessment of the distensibility of superficial arteries. *Ultrasound in Medicine and Biology*, 16, 121-128.

Hokanson, D.E., Strandness, Jr, D.E. and Miller, C.W. (1970) An echo-tracking system for recording arterial-wall motion. *IEEE Transactions on Sonics and Ultrasonics*, SU-17, 130-132.

Holen, J., Waag, R.C. and Gramiak, R. (1985) Representations of rapidly oscillating structures on the Doppler display. *Ultrasound in Medicine and Biology*, 11, 267-272.

Horowitz, P. and Hill, W. (1989) *The art of electronics*. New York: Cambridge University Press.

Huang, S.R., Lerner, R.M. and Parker, K.J. (1992) Time domain Doppler estimators of the amplitude of vibrating targets. *Journal of the Acoustical Society of America*, 91, 965-974.

Ilijevski, N., Radak, D.J., Radevic, B. , et al. (1998) Traumatic arteriovenous fistulas during a period of peace. *Vojnosanit Pregl.*, 55, 19-25.

Islam, M.N., Alimuzzaman, M., Khan, M.N., Bashar, M.A. and Zafar, A. (1996) Ruptured aneurysm of the sinus of Valsalva. *Bangladesh Medical Research Council Bulletin*, 22, 19-26.

Janer, L. (1998) New pitch detection algorithm based on wavelet transform. *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, 165-168.

Jensen, J.A. (1996) *Estimation of blood velocities using ultrasound: A signal processing approach*. New York: Cambridge University Press.

Johnstone, I.M. and Silverman, B.W. (1997) Wavelet threshold estimators for data with correlated noise. *Journal of the Royal Statistical Society, Series B*, 59, 319-351.

Kaklikkaya, I., Ozcan, F. and Kutlu, N. (1999) Double layered autogenous vein graft patch reconstruction of the common carotid-internal jugular fistula caused by gunshot wound. *Journal of Cardiovascular Surgery (Torino)*, 40, 429-433.

Kanai, H., Satoh, H., Hirose, K. and Chubachi, N. (1993) A new method for measuring small local vibrations in the heart using ultrasound. *IEEE Transactions on Biomedical Engineering*, 40, 1233-1241.

Kansky, J.P. (1988) *One-dimensional ultrasonic plethysmography*. Master's Thesis, University of Washington.

Kaspersen, J.H. and Hudgins, L. (1996) Wavelet quadrature methods for detecting coherent structures in fluid turbulence. *Proceedings of the SPIE*, 2825, 540-550.

Kay, S.M. (1988) *Modern spectral estimation*. New Jersey: Prentice Hall.

Kay, S.M. (1998) *Fundamentals of statistical signal processing: detection theory*. New Jersey: Prentice-Hall PTR.

Khoury, G., Sfeir, R., Nabbout, G., Jabbour-Khoury, S. and Fahl, M. (1994) Traumatic arteriovenous fistulae: "The Lebanese war experience". *European Journal of Vascular Surgery*, 8, 171-173.

Kirkeeide, R.L., Young, D.F. and Cholvin, N.R. (1977) Wall vibrations induced by flow through simulated stenoses in models and arteries. *Journal of Biomechanics*, 10, 431-441.

Knox, R., Breslau, P. and Strandness, Jr., D.E. (1981) Quantitative carotid phonoangiography. *Stroke*, 12, 798-803.

Koopmans, L.H. (1974) *The spectral analysis of time series*. New York: Academic Press.

Kumar, P. and Foufoula-Georgiou, E. (1994) Wavelet analysis in geophysics: An introduction. In *Wavelets in Geophysics*, edited by Foufoula-Georgiou, E. and Kumar, P., San Diego: Academic Press.

Laennec, R.T. (1819) *Traite de l'auscultation mediate du diagnostic des maladies des poumons et du coeur, fonde principalement sur ce nouveau moyen d'exploration*. Paris: Brosson, Chaude.

Lees, R.S. and Dewey, Jr., C.F. (1970) Phonoangiography: A new noninvasive diagnostic method for studying arterial disease. *Proceedings of the National Academy of Sciences USA*, 67, 935-942.

Liu, P.C. (1994) Wavelet spectrum analysis and ocean wind waves. In *Wavelets in geophysics*, edited by Foufoula-Georgiou, E. and Kumar, P., San Diego: Academic Press.

Lopez, M.A., Lopez, V.I., Villoria, R., Fernandez, C.G. and Martin, G.J. (1995) Arteriovenous fistulas in the neck: diagnosis with color Doppler sonography. *Journal of Ultrasound in Medicine*, 14, 233-239.

Mast, T.D. and Pierce, A.D. (1995) A theory of aneurysm sounds. *Journal of Biomechanics*, 28, 1045-1053.

McLeod, F.D., Graves, G.M., Dick, D.E. and Miller, C.W. (1977) Doppler measurement of wall motion. *Ultrasound in Medicine, Engineering Aspects*, 3B, 1373-1374.

Meagher Jr., D.P., Defore, W.W., Mattox, K.L. and Harberg, F.J. (1979) Vascular trauma in infants and children. *Journal of Trauma*, 19, 532-536.

Mellen, R.H. (1952) The thermal-noise limit in the detection of underwater acoustic signals. *The Journal of the Acoustic Society of America*, 24, 478-480.

Middleton, W.D., Erickson, S. and Melson, G.L. (1989) Perivascular color artifact: Pathologic significance and appearance on color Doppler US images. *Radiology*, 171, 647-652.

Middleton, W.D., Kellman, G.M., Melson, G.L. and Madrazo, B.L. (1989) Postbiopsy renal transplant arteriovenous fistulas: Color Doppler US charactersitics. *Radiology*, 171, 253-257.

Miller, A., Lees, R.S., Kistler, J.P. and Abbott, W.M. (1980) Effects of surrounding tissue on the sound spectrum of arterial bruits in vivo. *Stroke*, 11, 394-398.

Miyatake, K., Yamagishi, M., Tanaka, N., et al. (1995) New method for evaluating left ventricular wall motion by color-coded tissue Doppler imaging: in vitro and in vivo studies. *Journal of the American College of Cardiology*, 25, 717-724.

Miyatake, K., Yamagishi, M., Tanaka, N. , et al. (1995) New method for evaluating left ventricular wall motion by color-coded tissue Doppler imaging: in vitro and in vivo studies. *JACC*, 3, 717-724.

Myers, R.B. and Morgan, C.D. (1998) Cardiac murmur associated with an abnormal chest x-ray film. *Chest*, 1998, 114, 1459-1462.

Nowicki, A. and Marciniak, A. (1989) Detection of wall vibrations by means of cepstrum analysis. *Ultrasonic Imaging*, 11, 273-282.

Okamoto, M., Hashimoto, M., Sueda, T., Munemori, M. and Yamada, T. (1992) Renal artery aneurysm: The significance of abdominal bruit and use of color Doppler. *Japanese Journal of Medicine*, 31, 1217-1219.

Olinger, C.P. and Wasserman, J.F. (1977) Electronic stethoscope for detection of cerebral aneurysm, vasospasm and arterial disease. *Surgical Neurology*, 8, 298-312.

Oppenheim, A.V. and Schafer, R.W. (1989) *Discrete-time signal processing*. New Jersey: Prentice Hall.

Pansini, R., Longhini, C., Portaluppi, F., Arslan, E., Ravenna, F. and Toselli, T. (1980) Spectral analysis of a "sea gull cry" murmur. *G. Ital. Cardiol.*, 10, 777-780.

Papoulis, A. (1991) *Probability, random variables, and stochastic processes, third edition*. San Francisco: McGraw-Hill.

Pennestri, F., Boccardi, L., Minardi, G., et al. (1989) Doppler study of precordial musical murmurs. *The American Journal of Cardiology*, 63, 1390-1394.

Percival, D.B. (1994) Spectral analysis of univariate and bivariate time series. *Methods of Experimental Physics*, 28, 313-348.

Percival, D.B. (1995) On estimation of wavelet variance. *Biometrika*, 82, 619-631.

Percival, D.B. and Walden, A.T. (1993) *Spectral analysis for physical applications: Multitaper and conventional univariate techniques*. New York: Cambridge University Press.

Percival, D.B. and Walden, A.T. (2000, forthcoming) *Wavelet methods for time series analysis*. Cambridge, U.K.: Cambridge University Press.

Perry, M.O. (1996) Trauma and peripheral vascular disease. In *Peripheral vascular diseases*, edited by Young, J.R., Olin, J.W. and Bartholomew, J.R., St. Louis: Mosby, 599-606.

Plainfosse, M.C., Calonge, V.M., Beyloune-Mainardi, C., Glotz, D. and Duboust, A. (1992) Vascular complications in the adult kidney transplant recipient. *Journal of Clinical Ultrasound*, 20, 517-527.

Priebe, R.D. and Baugh, K.W. (1994) Wavelet based detectors. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 6, IV--105-108.

Ravier, P. and Amblard, P. (1998) Combining an adapted wavelet analysis with fourth-order statistics for transient detection. *Signal Processing*, 70, 115-128.

Reichard, K.W., Hall, J.R., Meller, J.L., Spigos, D. and Reyes, H.M. (1994) Arteriography in the evaluation of penetrating pediatric extremity injuries. *Journal of Pediatric Surgery*, 29, 19-22.

Reneman, R.S. (1981) What measurements are necessary for a comprehensive evaluation of the peripheral arterial circulation? *Cardiovascular Diseases*, 8, 435-454.

Rioul, O. and Flandrin, P. (1992) Time-scale energy distributions: A general class extending wavelet transforms. *IEEE Transactions on signal processing*, 40, 1746-1757.

Ritcey, J.A. and Holm, J.R. (1992) Applications of nonlinear filtering in radar CFAR detection. *IEEE International Symposium on Circuits and Systems*, 6, 2641-2644.

Roach, M.R. (1963) Changes in arterial distensibility as a cause of poststenotic dilatation. *American Journal of Cardiology*, 12, 802-815.

Roach, M.R. and Stockley, D. (1980) The effects of the geometry of a stenosis on poststenotic flow in models and postenotic vibration of canine carotid arteries in vivo. *Journal of Biomechanics*, 13, 623-634.

Sabbah, H.N. and Stein, P.D. (1976) Turbulent blood flow in humans: Its primary role in the production of ejection murmurs. *Circulation Research*, 38, 513-525.

Schmidt-Trucksass, A., Grathwohl, D., Schmid, A. , et al. (1998) Assessment of carotid wall motion and stiffness with tissue Doppler imaging. *Ultrasound in Medicine and Biology*, 24, 639-646.

Sekhar, L.N. and Heros, R.C. (1981) Origin, growth and rupture of saccular aneurysms: A review. *Neurosurgery*, 8, 248-260.

Semmlow, J., Welkowitz, W., Kostis, J. and Mackenzie, J.W. (1983) Coronary artery disease -- correlates between diastolic auditory characterisitics and coronary artery stenoses. *IEEE Transactions on Biomedical Engineering*, 30, 136-139.

Serroukh, A. and Walden, A.T. (1998) The scale analysis of bivariate non-Gaussian time series via wavelet cross-covariance. In *Imperial College of Science Statistics Section, Technical Report TR-98-02.*

Shattuck, D.P., Weinshenker, M.D., Smith, S.W. and vonRamm, O.T. (1984) Explososcan: a parallel processing technique for high speed ultrasound

imaging with linear phased arrays. *Journal of the Acoustical Society of America*, 75, 1273-1282.

Shinozuka, N. and Yamakoshi, Y. (1993) Measurement of fetal movements using multichannel ultrasound pulsed Doppler: autorecognition of fetal movements by maximum entropy method. *Medical and Biological Engineering and Computing*, 31, S59-S66.

Shu, M.C. and Hwang, N.H. (1988) Flow phenomena in compliant and noncompliant arteriovenous grafts. *ASAIO Transactions*, 34, 519-523.

Siegel, A.F. (1979) The noncentral chi-squared distribution with zero degrees of freedom and testing for nonuniformity. *Biometrika*, 66, 381-386.

Simkins, T.E. and Stehbens, W.E. (1973) Vibrational behavior of arterial aneurysms. *Letters in Applied and Engineering Sciences*, 1, 85-100.

Stehbens, W.E., Liepsch, D.W., Poll, A. and Erhardt, W. (1995) Recording unexpectedly high frequency vibrations of blood vessel walls in experimental arteriovenous fistulae of rabbits using a laser vibrometer. *Biorheology*, 32, 631-641.

Strassburg, C.P., Bleck, J.S., Rosenthal, H., Meyer, H.J., Gebel, M. and Manns, M. P. (1996) Diarrhea, massive ascites, and portal hypertension: rare case of a splenic arterio-venous fistula. *Zeitschrift fur Gastroenterol*, 34, 243-249.

Sunagawa, K., Kanai, H., Koiwa, Y., Nitta, K. and Tanaka, M. (1999) Simultaneous measurement of vibrations on the arterial wall downstream and upstream from an atherosclerotic lesions. *1999 IEEE Ultrasonics Symposium. Proceedings.*

Sutherland, G.R., Stewart, M.J., Groundstroem, K.W., et al. (1994) Color Doppler myocardial imaging: A new technique for the assessment of myocardial function. *Journal of the American Society of Echocardiography*, 7, 441-458.

Teich, M.C., Heneghan, C., Khanna, S.M., Flock, A., Ulfendahl, M. and Brundin, L. (1995) Investigating routes to chaos in the guinea-pig cochlea using the continuous wavelet transform and the short-time Fourier transform. *Annals of Biomedical Engineering*, 23, 583-607.

Thomson, D.J. (1982) Spectrum estimation and harmonic analysis. *Proceedings of the IEEE*, 70, 1055-1096.

Torrence, C. and Compo, G.P. (1998) A practical guide to wavelet analysis. *Bulletin of the American Meteorological Society*, 79, 61-78.

Vilkomerson, D. (1996) *Diffracting doppler-transducer*. United States Patent: US5540230.

Vito, R., Tso, W.K. and Schwartz, C.J. (1975) Poststenotic dilatation: arterial wall mechanics in response to vibration. *Canadian Journal of Physiology and Pharmacology*, 53, 998-1004.

Wang, K.Y., Bone, S.N. and Hossack, J.M. (1992) A tool for evaluating Doppler sensitivity. *Journal of Vascular Technology*, 16, 87-94.

Whitcher, B.J. (1998) *Assessing nonstationary time series using wavelets*. Ph.D. thesis, University of Washington.

Wilson, L.S. and Robinson, D.E. (1982) Ultrasonic measurement of small displacements and deformations of tissue. *Ultrasonic Imaging*, 4, 71-82.

Yamakoshi, Y., Sato, J. and Sato, T. (1990) Ultrasonic imaging of internal vibration of soft tissue under forced vibration. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 37(2), 45-53.

Young, J.R., Olin, J.W. and Bartholomew, J.R. (1996) *Peripheral vascular diseases*. St. Louis: Mosby.

Appendix A: Issues related to collecting data from human patients and animal subjects

Data from human patients and animal subjects is difficult to obtain and full of uncertainties. The following two sections enumerate these points.

A.1. Human patients with arterial disorders

Here we consider patients with non-emergent arterial disorders such as stenoses, aneurysms, or AV fistulas. Patients with emergent disorders, such as bleeding, can only be used for research when the method under research has been shown to be clinically useful. Patient safety should never be compromised for research purposes.

A.1.1. Notification

In order to study a patient, the researchers must first be notified that a specific patient has a disorder that meets the researchers' criteria. For this notification to take place, clinicians examining the patient as part of his/her medical care must be aware of the researchers' criteria, mentally relate the patient to the study, and have the time available to contact the researchers. Since clinicians are typically pressed for time and research is secondary to their other responsibilities, notification can be a significant bottleneck in collecting data from human patients.

A.1.2. Access to the patient

After notification, several other factors affect the researchers' access to the patient. First, the patient must be sufficiently conscious, coherent, mobile and willing to sign an informed consent form. Some patients are too sick, sore, tired, medicated, or busy to offer their informed consent. Once the patient is willing, a location must be available for the research to be carried out. If the subject is an in-patient in a hospital, then his/her hospital room is typically the location of choice. For out-patients, however, finding space for carrying out the research at a time when the patient is available can be difficult. To make use of willing patients, the researchers must be continually ready to gather the research data whenever a patient is available. This continual readiness can be very demanding and conflict with other research activities.

A.1.3. Data collection

Once the patient, researchers and data collection instrument have been gathered together in one room, data collection can begin. The patient's time and wishes must be fully accommodated. Thus, equipment setup time and patient discomfort must be completely minimized. The equipment, therefore, must be continually ready for data collection. This priority can conflict with the need to update or modify the experimental equipment.

Although the researchers have been informed that the patient meets their criteria, the possibility remains that either the patient did not in fact meet their criteria, or the patient's physiology has changed so that the patient no longer meets the research criteria. Such a scenario can occur due to blood pressure changes which cause vibrations to cease in the case of stenoses, or medical intervention in the case of pseudo-aneurysms.

Finally, since data analysis occurs only off-line, it is difficult to predict whether the desired data has actually been collected. In this case, the researchers must follow their protocol and wait for future analysis to determine if the data collection was successful.

A.2. Animal subjects

In animal studies used for this research, notification has not been an issue because the disorders (bleeding) were created in the animal. Other challenges occurred in animal studies, however.

Although animal studies can be scheduled in advance and laboratories are set aside for animal experiments, access to the animal is still limited by time, space availability, and finances. Additionally, creation of arterial disorders in an animal requires more personnel than data collection from patients. Specifically, a) technicians must prepare and monitor the animal throughout the experiment, and b) a skilled clinician (e.g., a surgeon) must produce the disorder. This increase in personnel both complicates

scheduling and increases the cost of the experiment. Communication among an increased number of personnel introduces its own challenges and time requirements. Producing the disorder can be a challenge, itself. Communication must occur between the researchers, clinicians and animal care technicians, each of whom has different areas of expertise and does not completely understand what the others are responsible for and can accomplish. Finally, because many steps are involved in animal preparation and disorder creation, it is particularly difficult to predict the time duration of this stage of the process.

If it is desirable that data is collected both before and after creation of the disorder, then further communication is required between the clinician and the researchers in order for data collection to occur just prior to creation of the disorder. As with the human patients, data collection is complicated by changing physiology and off-line analysis.

Appendix B: Human subjects consent form

Human Subjects Division

APR 1 1 1997

UW

UNIVERSITY OF WASHINGTON
CONSENT FORM

NONINVASIVE STUDIES OF BLOOD FLOW THROUGH ORIFICES

Investigators:

Kirk W. Beach, Ph.D., M.D.	Research Professor	Department of Surgery	phone 543 3827
Michael Caps, M.D.	Research Fellow	Department of Surgery	phone 616 8535
Keith A. Comess, M.D.	Associate Professor	Division of Cardiology	phone 731 2289
D Eugene Strandness, Jr.,MD	Professor	Department of Surgery	phone 543 3653

Investigator's Statement

PURPOSE AND BENEFITS

The purpose of this research is to learn more about the ultrasound echoes from blood flow through passages. By analyzing features of the echo signals, we hope to be able to tell the difference between echoes from normal, obstructed, and bleeding arteries.

PROCEDURES

We will attach 3 adhesive ECG electrodes to your arms or chest and place an ultrasound transducer with ultrasonic gel on your skin over one of your arteries to obtain echoes from a portion of that artery. The echo signal will be digitized and stored in a computer along with clinical information about your artery. The entire procedure will take about 5 minutes.

RISK, STRESS OR DISCOMFORT

There is no known risk from diagnostic ultrasound examinations. You should feel no effect or discomfort from this examination. If you feel any effects, please tell us and we will stop the study. Taking off the adhesive electrodes may feel like taking off a Band-Aid. The skin underneath may be red and tender.

OTHER INFORMATION

This examination is done in conjunction with your other medical care and does not substitute for any other examination that you might have. No data will be retained by the investigators about your identity; data will be recorded about the artery from which the echoes came. You will neither be charged nor paid for this study. If you are injured because of this study, the injury will be treated by an investigator or other person providing your care at no charge to you. You may refuse to participate in or stop this study at any time for any reason without any effect on your medical care or your relationship with the investigators, your care providers or the University of Washington.

Signature of Investigator Date

Subject's Statement:

The study described above has been explained to me. I voluntarily consent to take part in this activity. I have had a chance to ask questions. If I have questions in the future about the research or about my rights as a subject, I know I can ask on the researchers.

Signature of Patient Volunteer Date

Copies to: Patient Volunteer
Investigator's file

revised 9 April, 1997

APPROVED FORM

Appendix C: The fiber optic displacement measurement system

The fiber optic system illustrated in Figure 4.24 measures displacement by transmitting light to a reflective tape placed on the tubing, and quantitating the intensity of light that is reflected back. The light is attenuated by the intervening medium (either air or water), so the intensity of the returned light will be a function of the distance between the fiber optic transducer and the reflective tape. The system transduces the received light into a voltage that is proportional to the intensity of the light (Figure C.1). Thus, the voltage level is proportional to tubing displacement provided that a linear region of the curve in Figure C.1 is used.

Calibration is accomplished in the following manner. 1) The micrometer, attached to the fiber optic transducer, is rotated to bring the fiber optic transducer away from the reflective tape so that it is above point A of Figure C.1. 2) The micrometer is rotated monotonically (in only one direction) until point B is reached. The DC voltage value is recorded. 3) The micrometer is again rotated monotonically in this forward direction until points C and D are reached and the DC voltage values recorded. 4) The micrometer is rotated monotonically in the opposite direction until point B has been surpassed. 5) The micrometer is again rotated monotonically in the forward direction until point C is reached. [Steps 4 and 5 are required in order to avoid the backlash of the micrometer.] 6) Voltage fluctuations around point C are amplified and digitized. 7) The digitized voltages are multiplied by the slope of the line between points B and D for calibration, to yield the tubing vibration displacement measurement. Points B, C, and D are thus assumed to be co-linear. Due to the gradations on the available micrometer, points B and D lie much further from point C than the amplitude of the tubing vibrations. Thus, any error in the slope measurement will be amplified when applied to the resulting vibration measurement. Finally, the amplification factor of step 6) is assumed to be its nominal value.

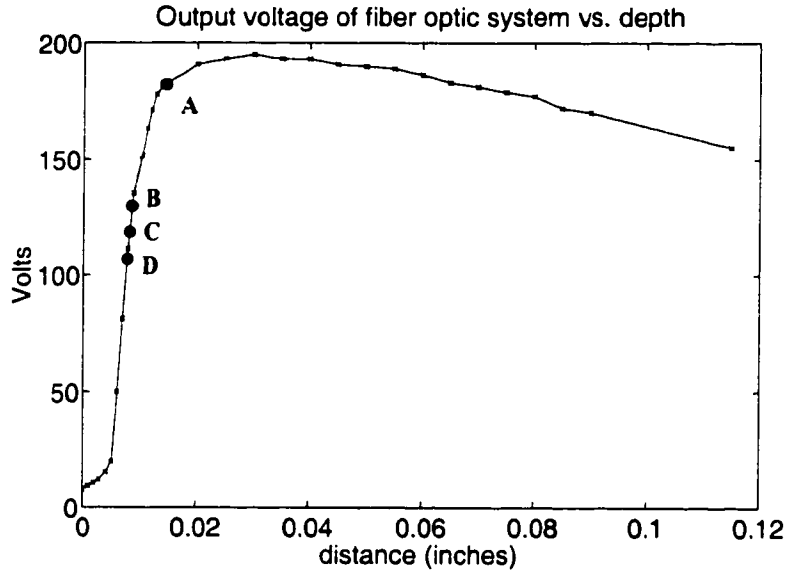

Figure C.1. Calibration curve for the fiber optic system.
This figure illustrates the process of calibrating the fiber optic displacement measurement system. Points B, C and D are assumed to be co-linear. The slope from point B to point D is used to convert voltages measured from the fiber optic system output into distance measures. Measurements should only be taken when the curve is traversed in the direction from A to D. The data points marked with an 'x' were collected with monotonic movement of the micrometer in the direction of the reflective tape.

Appendix D: Amplitude estimation of tapered sinusoid

Consider a signal with complex sinusoid of frequency, ω, windowed by a half wave sinusoid of duration *a*. It is convolved with a wavelet also of complex sinusoid frequency, ω, windowed by a half wave sinusoid of duration *b*.

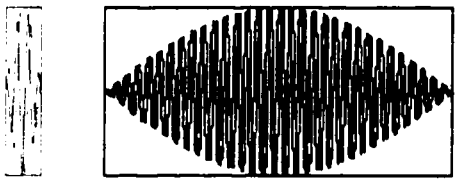

$$\sin\left(\frac{\pi}{b}t\right)e^{j\omega t} \qquad \sin\left(\frac{\pi}{a}t\right)e^{j\omega t}$$

Figure D.1. Illustration of amplitude estimate of a tapered sinusoid.
Left: A sine-tapered sinusoid with center frequency 1/(2b) and duration b. This approximates a Morlet wavelet. **Right:** A sine-tapered sinusoid with center frequency 1/(2a) and duration a. This represents the signal to be estimated.

The wavelet coefficients will be maximum at t = a/2 + b/2 with the following mathematical description:.

$$\int_0^b \sin\left(\frac{\pi}{b}\tau\right)e^{j\omega\tau}\sin\left(\frac{\pi}{a}(t-\tau)\right)e^{j\omega(t-\tau)}d\tau\Bigg|_{t=a/2+b/2} = \frac{2a^2b}{\pi(a^2-b^2)}\cos\left(\frac{b\pi}{2a}\right)e^{j\omega t} \approx \frac{2b}{\pi}e^{j\omega t}, b << a \qquad \text{(D- 1)}$$

Note that $\frac{2b}{\pi} = \int_0^b \sin\left(\frac{\pi}{b}t\right)dt$, which is the integral of the envelope of the wavelet.

Appendix E: Kolmogoroff-Smirnov goodness of fit test

The Kolmogoroff-Smirnov (K-S) test is a means of determining whether a time series fits an expected distribution, $F_0(x)$. To perform the test, it is first necessary to generate an empirical cumulative probability distribution estimate, $\hat{F}(x)$, from the time series. This is easily computed as the cumulative sum of the histogram of the time series divided by the maximum value of the cumulative sum. The test statistic, $q = \max_x \left| \hat{F}(x) - F_0(x) \right|$, indicates how closely the empirical distribution approximates the expected theoretical distribution. It can be shown (Papoulis, 1991) that a long time series is considered to fit the theoretical distribution with confidence level r if:

$$q < \sqrt{\frac{-1}{2N} \ln\left(\frac{r}{2}\right)},$$

where N is the length of the time series.

Appendix F: Dual depth probability density function

It was concluded in section 6.2.2 that the normalized cross wavelet spectrum for the dual depth method is the product of two F distributed variables. Since both variables are computed from two time series that are presumed to be independent and identically distributed, the resulting F-distributed variables will also be independent and identically distributed. According to Papoulis (1991), a random variable z which is the product of random variables x and y will have the following probability density function:

$$f_z(z) = \frac{\int_{-\infty}^{\infty} \frac{1}{|y|} f_{xy}(x, y)\,dy}{\int_{-\infty}^{\infty} \frac{1}{|y|} f_{xy}\left(z/y, y\right)\,dy}, \qquad \text{(F- 1)}$$

where x = z / y. Since x and y are independent, and the joint density function of two independent variables is the product of their individual density functions (Papoulis, 1991), the density function for z can be re-written as:

$$f_z(z) = \int_{-\infty}^{\infty} \frac{1}{|y|} f_x\left(z/y\right) f_y(y)\,dy. \qquad \text{(F- 2)}$$

The density functions $f_x(x)$ and $f_y(y)$ were shown in section 6.2.1 to be F-distributed with 2 and ν degrees of freedom. Therefore, from the definition of the F distribution,

$$f_x(x) = \frac{\left(2/\nu\right)}{\beta\left(1, \nu/2\right)\left(1 + 2/\nu\, x\right)^{1+\nu/2}} \qquad \text{(F- 3)}$$

$$\text{and } f_y(y) = \frac{\left(2/\nu\right)}{\beta\left(1, \nu/2\right)\left(1 + 2/\nu\, y\right)^{1+\nu/2}}, \qquad x, y > 0$$

$$f_x(x) = 0 \text{ and } f_y(y) = 0 \qquad x, y \le 0$$

where $\beta(u, v) = \dfrac{\Gamma(u)\Gamma(v)}{\Gamma(u+v)}$,

and $\Gamma(m+1)$ is the real number interpolation of the integer factorial function. That is, for integer m,

$\Gamma(m+1) = m!$

The density function for z, then becomes:

$$f_z(z) = \int_0^\infty \frac{1}{y}\left(\frac{\left(2/\nu\right)}{\beta\left(1, \nu/2\right)}\right)^2 \frac{dy}{\left(1 + 2/\nu \dfrac{z}{y}\right)^{1+\nu/2}\left(1 + 2/\nu\, y\right)^{1+\nu/2}}. \qquad \text{(F- 4)}$$

Since $\beta\left(1, \nu/2\right)$ evaluates to $\left(2/\nu\right)$, $f_z(z)$ simplifies to:

$$f_z(z) = \int_0^\infty \frac{1}{y} \frac{dy}{\left(1 + \dfrac{2}{\nu}\dfrac{z}{y}\right)^{1+\nu/2}\left(1 + \dfrac{2}{\nu} y\right)^{1+\nu/2}}, \text{ or} \qquad \text{(F- 5)}$$

$$f_z(z) = \int_0^\infty \frac{dy}{y\left(1+\left(\frac{2}{\nu}\right)^2 z+\left(\frac{2}{\nu}\right)\frac{z}{y}+\left(\frac{2}{\nu}\right)y\right)^{1+\nu/2}}. \qquad \text{(F- 6)}$$

An example of the probability density function and cumulative probability distribution function are shown in Figure F.1.

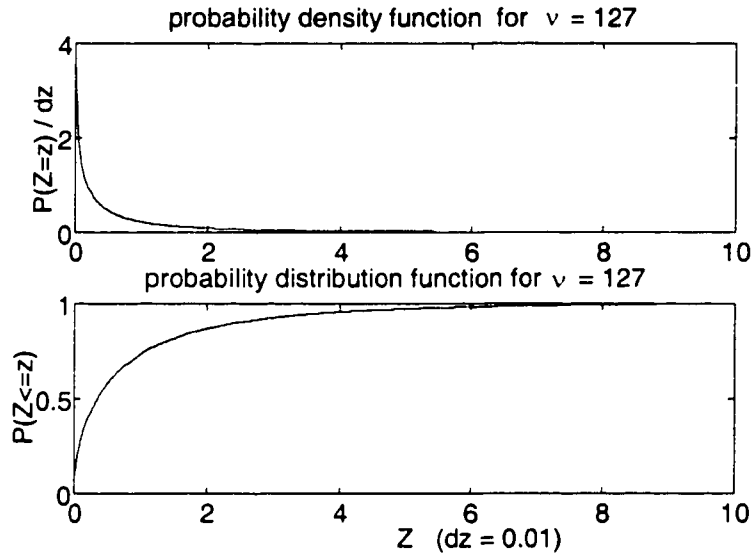

Figure F.1. Example of the probability curves for the dual depth test statistic.

Appendix G: Computation costs

G.1. Detection

The most demanding mode in modern ultrasound instruments is the color Doppler mode. Thus, its computation requirements provide an indicator of the computational capability of modern instruments. Section G.1.2 discusses this topic.

Although Fisher's test for periodicity was shown in chapter 7.1.3 to be incapable of adequate detection of arterial vibrations, it remains a leading detector of simple periodicities (single frequency peaks) in white noise. Therefore, its computation cost serves as a reasonable comparison for the cost of our wavelet-based detector. Table G.1 summarizes the computational cost for the two detector methods for two example sampling rates. The CWT based detector is implemented with the dot product method with a *nonoverlap* percentage of 25%. Sections G.1.1 and G.1.3 provide the details behind these results.

Table G.1. Approximate computation costs for CWT based detector and Fisher's test for periodicity without clutter filtering.
The cost for the wavelet based detector covers the frequency range 100 to 1000 Hz with Q_w of 3 and overlap in the frequency domain at -8.7 dB. All computations are for real numbers.

	Formula		**Fs = 2000 Hz**		**Fs = 8000 Hz**	
	Mults.	**Adds.**	**Mults.**	**Adds.**	**Mults.**	**Adds.**
CWT based detector	128.64 F_s	256 F_s + (1/2350) F_s^2	257,300	513,700	1,029,000	2,076,000
Fisher's test	7 F_s ($\log_2(F_s/8)$ + 0.5)	7 F_s ($\log_2(F_s/8)$ + 0. 5)	118,500	118,500	586,000	586,000
Ratio			2.17	4.34	1.76	3.54

These numbers for Fisher's test for periodicity, however, do not provide for prior signal conditioning which would be required since a vibration's large dynamic range would cause bias in the FFT. Thus, the raw quadrature data first must be high pass filtered to pre-whiten it prior to FFT computation. This can be accomplished in

several ways. For discussion, consider implementing the filter in the time domain with convolution. With a filter length of 16, this would require approximately 32 F_s additional real multiplications and additions per second.

Table G.2. Approximate computation costs for CWT based detector and Fisher's test for periodicity with clutter filtering.
The cost for the wavelet based detector covers the frequency range 100 to 1000 Hz with Q_w of 3 and overlap in the frequency domain at -8.7 dB. All computations are for real numbers.

	Formula		**Fs = 2000 Hz**		**Fs = 8000 Hz**	
	Mults.	**Adds.**	**Mults.**	**Adds.**	**Mults.**	**Adds.**
CWT based detector	128.64 F_s	256 F_s + (1/2350) F_s^2	257,300	513,700	1,029,000	2,076,000
Fisher's test	7 F_s ($\log_2(F_s/8)$+ 0.5)+32F_s	7 F_s ($\log_2(F_s/8)$ +0.5)+32F_s	182,500	182,500	842,000	842,000
Ratio			1.41	2.81	1.22	2.47

Therefore, the CWT based detector provides the ability to detect oscillations in colored noise with far superior detection curves than Fisher test affords for detecting the oscillations in white noise with approximately twice the computation cost.

G.1.1. CWT based detector

Implementation of the detector first requires computation of the continuous wavelet transform of a time series at a given depth at each scale of interest. As stated in section 6.1.2.3, this requires

(wlen) ((tlen - overlen)/nonoverlen) complex multiplications and

(wlen-1)((tlen – overlen) / nonoverlen) complex additions. (G- 1)

Since *tlen* is generally much larger than the wavelet length, the quantity *tlen – overlen* can be conservatively approximated by *tlen*. Furthermore, since *nonoverlen* is a given percentage of *wlen*, the ratio *wlen / nonoverlen* is simply the reciprocal of the given

percentage. Finally, the quantity *wlen-1* can be conservatively approximated by *wlen*. Thus, for each depth time series, the continuous wavelet transform will require approximately

$$(\#\ of\ scales)\ (\ tlen\ )\ /\ (\ nonoverlap\ percentage) \qquad \text{(G- 2)}$$

complex additions and multiplications. For a Morlet wavelet quality factor of 3 with overlaps in the frequency domain of – 8.7 dB, 16 scales are required to cover both positive and negative frequencies between 100 and 1000 Hz. If the sample rate is F_s, and data is collected for one second, and the nonoverlap percentage is 25%, then computation of the continuous wavelet transform requires approximately

$16\ F_s\ /\ (0.25) = 64\ F_s$

complex multiplications and additions per second of data. In general, each complex multiplication requires four real multiplications and 2 real additions. Each complex addition requires two real additions. Since the wavelets at the negative frequencies are the complex conjugates of the wavelets at the positive frequencies, though, the number of real multiplications can be cut in half, but the number of additions remains the same Therefore, computation of the continuous wavelet transform coefficients requires:

$128\ F_s$	real multiplications
$256\ F_s$	real additions.

The detector requires computation of the magnitude squared of each wavelet coefficient to produce the wavelet power spectrum. Conservatively, this involves computing

$(tlen) / (nonoverlen) = F_s / (0.25 wlen(\text{scale}))$

wavelet coefficients per scale, making use of equation (6-6) and the approximation that *tlen-overlen* is approximately equal to *tlen*. Computation of the magnitude squared of each coefficient requires two real multiplications and one real addition. Hence, power spectrum coefficients require:

$2F_s / (0.25 wlen(\text{scale}))$ real multiplications and
$F_s / (0.25 wlen(\text{scale}))$ real additions.

Normalizing the wavelet power spectrum by the noise estimates requires one noise estimate per power spectral coefficient, or:

$(tlen) / (nonoverlen) = F_s / (0.25 wlen(\text{scale}))$

noise estimates and subsequently the same number of real multiplications per scale. Each noise estimate involves computing a mean over two segments of wavelet power spectral coefficients where each segment spans a time duration of *seglen*. This involves approximately:

$2\ seglen\ F_s / (nonoverlen) = 2\ seglen\ F_s / (0.25 wlen(\text{scale}))$

real additions with one comparison and two real multiplications. These multiplications, however, are simply a division by *seglen* which is a fixed value that could be lumped into the detection decision *a priori*.

After including the final comparison which accomplishes the decision stage, detection in one time series for one second at a *nonoverlap* percentage of 25% requires:

$$128\,F_s + \sum_{scale} \frac{3F_s}{0.25 wlen(scale)}$$ real multiplications and

$$256\,F_s + \sum_{scale} \frac{F_s}{0.25 wlen(scale)} \left(\frac{2 seglen\, F_s}{0.25 wlen(scale)} + 1 \right)$$ real additions with

$$\sum_{scale} \frac{2F_s}{0.25 wlen(scale)}$$ comparisons.

The Morlet wavelet can be considered to be time limited with wavelet lengths of:

$$wlen(\text{scale}) \approx 12\, Q_w\, scalebase^{(\text{scale}-1)}, \qquad \text{(G- 3)}$$

where *scalebase* was defined in equation (6-1), and scale 1 corresponds to the center frequency $F_s/4$. For Q_w of 3 and scalebase of 1.4 for -8.7 dB at the frequency overlap of the scales, scales 3 through 10 span the frequency range 100 to 1000 Hz used for vibrometry. Including both positive and negative frequencies requires the use of 16 scales. The average value of *wlen* according to equation (G-3) is approximately 300. Using this approximation, the computation cost of the detector can be approximated as:

$$128\,F_s + (16)\frac{3F_s}{0.25(300)} \approx 128.64F_s$$ real mults

and (G- 4)

$$256\,F_s + (16)\frac{F_s}{0.25(300)} \left(\frac{2(0.075)F_s}{0.25(300)} + 1 \right) \approx 256F_s + \frac{1}{2350} F_s^2$$ real adds

with

$$16\frac{2F_s}{0.25(300)} \approx 0.5F_s$$ comparisons,

where *seglen* was set to 75 ms.

G.1.2. Color Doppler computation costs

Consider that a large color box in the color Doppler mode of a modern ultrasound instrument (128 ray lines, 400 depths, 10 frames per second) would require 512,000 (128*400*10) velocity estimates per second. According to Basoglu et al. (1998), each velocity estimate requires the following computations when based on the autocorrelation method and when applied subsequent to clutter filtering:

4 (ensemble length) real multiplications and
2 (ensemble length – 1) real additions.

The clutter filter is much more costly. Basoglu et al. explain that wall filtering that uses the regression filter matrix method requires:

2 (ensemble length) (ensemble length) real multiplications and
2 (ensemble length) (ensemble length – 1) real additions.

Therefore, for an ensemble length of 10, the color Doppler method with a large color box and an ensemble length of 10 would require:

512,000 (4 (10) + 2 (10) (10)) real multiplications and
512,000 (2 (9) + 2 (10) (9)) real additions.

Thus, 122,880,000 real multiplications and 101,376,000 real additions are needed. Finally, if vibrometry detection requires roughly 1,000,000 real multiplications and 2,000,000 real additions per second per depth for a PRF (F_s) of 8000 Hz, then modern instruments can accomplish vibrometry at roughly 100 tissue depths, since multiplication is more costly than addition. Similarly, if the PRF were reduced to 2000 Hz, then 400 depths could be processed.

G.1.3. Fisher's test for periodicity

Section 5.2.1.1 described Fisher's test for periodicity and section 7.1.3 displayed its detection curves for simulated vibrations. Here we discuss its computational cost. Fisher's test for periodicity first requires computation of the short-time Fourier transform of the time series at a given depth. This results in

$$numsegs = \frac{tlen - FFTLEN}{(1 - overlap\%)(FFTLEN)} \qquad \text{(G- 5)}$$

separate FFT computations per time series.

As with the wavelet-based detector, we will compute the computation cost for one second of data. This sets the value of *tlen* to F_s. It was mentioned in section 7.1.3 that an FFT length of 1024 was desirable for F_s of 8000 Hz, and that an overlap of 50% is sufficient. Making an extrapolation that a good FFTLEN is $F_s/8$, *numsegs* evaluates to 14. When the length of the FFT is a power of 2, each FFT computation requires $N\log_2 N$ complex multiplications and additions, where N is the length of the FFT. The short-time Fourier transform thus requires:

$$numsegs \frac{F_s}{8} \log_2\left(\frac{F_s}{8}\right) = \frac{7}{4} F_s \log_2\left(\frac{F_s}{8}\right)$$

complex multiplications and additions. Since each complex multiplication requires four real multiplications and two real additions, and each complex addition requires two real additions, the short-time Fourier transform requires:

$$7F_s \log_2\left(\frac{F_s}{8}\right) \qquad \text{(G- 6)}$$

real multiplications and additions.

This detector also requires computing the power spectrum, or taking the magnitude squared of each FFT coefficient. This involves two multiplications and one addition per coefficient, or:

$$2\,numsegs(FFTLEN) = 2(14)(F_s / 8) = 3.5F_s \qquad \text{(G- 7)}$$

real multiplications and half as many real additions.

Computing the sum of each power spectrum requires approximately this same number of real additions, 1.75 F_s. Normalizing the peaks of each power spectrum requires a total of *numsegs* (i.e., 14) real multiplications. Finally, a comparison is made at each of these FFT coefficients to make the final detection decisions, requiring:

numsegs (FFTLEN) = 14 (F_s / 8) comparisons.

Hence,

7 $F_s \log_2(F_s/8)$ + 3.5 F_s + 14	complex multiplications
7 $F_s \log_2(F_s/8)$ + 3.5 F_s	complex additions
7/4 F_s	comparisons

are required for detection with Fisher's F-test for periodicity.

G.2. Estimation

Vibration parameter estimates need to be made on the phase of the quadrature data and not on the quadrature data itself. They cannot, therefore, draw only from the detector's computations. Table G.3 summarizes the computation costs for estimation with the wavelet based estimator and the Fourier estimator.

Table G.3. Estimation computation costs for the wavelet based detector and Fisher's test for periodicity.
The wavelet based estimation assumes the signal resides across three wavelet scales. All computations are for real numbers.

	Formula			**Fs = 2000 Hz**			**Fs = 8000 Hz**		
	Mults	**Adds**	**atan**	**Mult**	**Adds**	**atan**	**Mults**	**Adds**	**atan**
CWT based estim-ation	2.6 F_s	2.5 F_s	0.1F_s	5200	5000	200	20,800	20,000	800
STFT based estim-ation	$^3/_8$ F_s + F_s $\log_2(F_s/8)$	$^3/_{16}$ F_s + F_s $\log_2(F_s/8)$	$^3/_{16}F_s$	16,700	16,300	375	82,700	81,200	1500
Ratio				0.31	0.31	0.53	0.25	0.25	0.53

Therefore, the CWT based estimation is both more accurate than the STFT based method (see discussion in section 5.3), and also requires approximately a third the number of computations.

These numbers for the STFT do not provide for prior signal conditioning which would be required since a vibration's large dynamic range would cause bias in the FFT.

G.2.1. CWT based estimation

Vibration parameter estimates need to be made only at the time and scales at which vibrations were detected. As discussed in section 6.3.1, these estimates are best made with the shorter wavelet, $Q_w = 1.5$ at –2 dB overlap in the frequency domain (*scalebase* = 1.4). With the use of the dot product method, wavelet coefficient computation at each time and scale requires *wlen* complex multiplications and *wlen-1* complex additions, or 4*wlen* real multiplications and 2*wlen*+2(*wlen*-1) real additions. All vibration parameters can be drawn from the resulting values.

Consider a vibration that lasts for 100 ms (0.1 s), sampled at 8000 Hz. Displacements must be computed according to equation (4-13) for 0.1 F_s sample points. This requires 0.1 F_s real multiplications and 0.1 F_s arctangent calculations to compute $\phi(t)$ (see equation (4-9)). Subsequently there are (0.1 F_s -1) real additions to form $\Delta\phi$, and (0.1 F_s -1) real multiplications to compute x(t).

Equation (G-2) dictates that the following number of complex additions and multiplications are required to compute the wavelet coefficients for this new time series:

(# of scales) (tlen) / (nonoverlap percentage)

Since the data of interest is the displacement time series, which is real, estimation can use only the positive scales. Using a 25% nonoverlap percentage and assuming a bandwidth of 3 wavelet scales, estimation would require approximately:

$$3\,(0.1\,F_s) / 0.25 = 1.2\,F_s \qquad \text{(G- 8)}$$

complex multiplications and additions. Because the displacement time series are real valued, each complex multiplication requires just two real multiplications and no additions. Each complex addition, though, still requires two real additions. Thus, each estimate requires $2.4F_s$ real multiplications and additions.
Combining the cost of generating the x(t) time series and subsequent wavelet coefficients, the total costs are:

$0.2\,F_s - 1 + 2.4\,F_s$	real multiplications
$0.1\,F_s - 1 + 2.4\,F_s$	real additions
$0.1\,F_s$	arctangents

Note that these numbers are small compared to the computation cost of the detector.

G.2.2. FFT based estimation

As with CWT based estimation, parameter estimates must be made from the displacement measurements, not from the raw quadrature data. Therefore, the short-time Fourier transforms must be re-computed during the time segments at which the vibration is detected. All parameter estimates can then be determined based on the

FFT ordinates. If the 100 ms vibration appears in two FFT segments which overlap by 50%, and the FFT length is $F_s/8$ then 1.5 (F_s / 8) displacement values need to be computed. Using the same argument as outlined in section G.2.1, this requires :

$3/8\ F_s - 1$ real multiplications

$3/16\ F_s - 1$ real additions and

$3/16\ F_s$ arctangents

Subsequently, the FFT computations require:

$2\ F_s/8\ \log_2(F_s/8)$ complex multiplications and additions, or

$F_s\ \log_2(F_s/8)$ real multiplications and additions.

This results in the following total computation cost:

$3/8\ F_s - 1 + F_s\ \log_2(F_s/8)$ real multiplications

$3/16\ F_s - 1 + F_s\ \log_2(F_s/8)$ real additions and

$3/16\ F_s$ arctangents.

Appendix H: Detector values and thresholds

For a sample length of 0.150 $F_s/2$, the complex Morlet wavelet has the following number of equivalent degrees of freedom for the given sampling rate, F_s, and Q_w of 3 according to equation (6-9).

Table H.1. Equivalent degrees of freedom, η, for the complex Morlet wavelet transform.
These apply for a Gaussian time series of length 0.075 F_s.

	$Q_w = 3$	
Scale	**F_s = 8000 Hz**	**F_s = 15000 Hz**
1	89.4	167.08
2	64.0	119.46
3	45.9	85.5
4	33.0	61.3
5	23.8	43.9
6	17.2	31.6
7	12.5	22.7
8	9.1	16.4
9	6.7	11.9
10	5.0	8.7
11	3.9	6.6
12	3.3	5.5

The scale numbers in Table H.1 relate to frequency ranges according to:

$$\frac{F_s}{4 scalebase^{scale-1}},$$

where *scalebase* is set to a value of 1.4 for -8.7 dB at the overlap between scales in the frequency domain. Notice from the table that the numbers of degrees of freedom used for the F distribution at scales 1 through 10 for F_s of 8000 Hz roughly equal the numbers of degrees of freedom for F_s of 15000 Hz at scales 3 through 12. This arises because the frequency ranges of these scales are roughly the same.

The following are the values of m and ν computed from equation (6-14) for F_s of 8000 Hz and *viblen* is the maximum of 150 ms or twice the wavelet time duration. Again, Q_w is 3 and *scalebase* is 1.4.

Table H.2. Chi-square approximation value for the noise estimates, ν, for F_s = 8000 Hz.
These apply for the individual noise estimates ν_L and ν_R computed from Gaussian time series of length 0.075 F_s, with *viblen* = 150 ms.

		F_s = 8000 Hz	
		m	ν
Scale number	**1**	0.0142	127.12
	2	0.0198	89.61
	3	0.0274	63.23
	4	0.038	44.65
	5	0.0524	31.55
	6	0.0716	22.33
	7	0.0970	15.83
	8	0.1302	11.26
	9	0.1722	8.07
	10	0.1722	8.07
	11	0.1722	8.07
	12	0.1722	8.07

Table H.3. Chi-square approximation value for the noise estimates, ν, for F_s = 15000 Hz.
These apply for the individual noise estimates ν_L and ν_R computed from Gaussian time series of length 0.075 F_s, with *viblen* = 150 ms.

		F_s = 15000 Hz	
		m	ν
Scale number	**1**	--	--
	2	0.0160	172.26
	3	0.0148	121.28
	4	0.0206	85.50
	5	0.0288	60.340
	6	0.0398	42.61
	7	0.0546	30.12
	8	0.0746	21.31
	9	0.0505	15.12
	10	0.101	10.77
	11	0.1786	7.72
	12	0.1786	7.72

Note that at the high scales (low frequency ranges), the number of degrees of freedom becomes fixed. This occurs because the wavelet length has reached half of *viblen*. Therefore, *viblen* is set to twice the wavelet time duration as stated in section 6.2.1. Regardless of scale, each wavelet has the same number of degrees of freedom, so when the same number of wavelet lengths is used in the computation, the number of degrees of freedom will be the same.

H.1. Single depth detector threshold values:

Corresponding threshold levels for the single depth detector are given in the following four tables.

Table H.4. Single depth detector threshold levels for F_s = 8000 Hz: the lower confidence levels.
Threshold levels are taken from the F distribution with 2 and ν degrees of freedom, where the values of ν are given in Table H.2. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels						
		0.9	**0.95**	**0.99**	**0.995**	**0.999**	**0.9995**	**0.9999**
Scale	**1**	2.3447	3.0674	4.7760	5.5253	7.2970	8.0740	9.9110
	2	2.3627	3.0981	4.8501	5.6242	7.4686	8.2836	10.2253
	3	2.3884	3.1422	4.9574	5.7681	7.7205	8.5925	10.6924
	4	2.4254	3.2060	5.1145	5.9799	8.0958	9.0552	11.4011
	5	2.4790	3.2990	5.3477	6.2965	8.6672	9.7651	12.5083
	6	2.5572	3.4362	5.7004	6.7806	9.5635	10.8913	14.3117
	7	2.6724	3.6414	6.2472	7.5434	11.0291	12.7629	17.4255
	8	2.8444	3.9545	7.1255	8.7962	13.5685	16.0830	23.2654
	9	3.1044	4.4424	8.5971	10.9644	18.3151	22.5031	35.5043
	10	3.1044	4.4424	8.5971	10.9644	18.3151	22.5031	35.5043
	11	3.1044	4.4424	8.5971	10.9644	18.3151	22.5031	35.5043
	12	3.1044	4.4424	8.5971	10.9644	18.3151	22.5031	35.5043

Table H.5. Single depth detector threshold levels for F_s =8000 Hz: the higher confidence levels.
Threshold levels are taken from the F distribution with 2 and ν degrees of freedom, where the values of ν are given in Table H.2. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels					
		0.99995	**0.99999**	**0.999995**	**0.999999**	**0.9999995**	**0.9999999**
Scale number	**1**	10.7166	12.6214	13.4568	15.4318	16.2980	18.3459
	2	11.0833	13.1273	14.0305	16.1824	17.1332	19.3986
	3	11.6302	13.8888	14.8974	17.3267	18.4115	21.0244
	4	12.4646	15.0654	16.2446	19.1279	20.4352	23.6319
	5	13.7788	16.9531	18.4232	22.0963	23.7974	28.0479
	6	15.9437	20.1477	22.1536	27.3207	29.7861	36.1370
	7	19.7447	25.9816	29.0839	37.4266	41.5762	52.7356
	8	27.0497	37.8591	43.5542	59.8183	68.3834	92.8004
	9	42.9075	65.8365	78.8182	118.1258	139.2040	192.4602
	10	42.9075	65.8365	78.8182	118.1258	139.2040	192.4602
	11	42.9075	65.8365	78.8182	118.1258	139.2040	192.4602
	12	42.9075	65.8365	78.8182	118.1258	139.2040	192.4602

Table H.6. Single depth detector threshold levels for F_s =15000 Hz: the lower confidence levels.
Threshold levels are taken from the F distribution with 2 and ν degrees of freedom, where the values of ν are given in Table H.3. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels						
		0.9	**0.95**	**0.99**	**0.995**	**0.999**	**0.9995**	**0.9999**
Scale number	**1**	--	--	--	--	--	--	--
	2	2.3335	3.0484	4.7304	5.4646	7.1922	7.9463	9.7207
	3	2.3468	3.0709	4.7845	5.5366	7.3165	8.0977	9.9465
	4	2.3656	3.1031	4.8623	5.6405	7.4970	8.3184	10.2776
	5	2.3926	3.1494	4.9752	5.7920	7.7625	8.6441	10.7709
	6	2.4315	3.2165	5.1407	6.0153	8.1591	9.1334	11.5219
	7	2.4879	3.3144	5.3869	6.3500	8.7649	9.8871	12.7010
	8	2.5702	3.4591	5.7603	6.8636	9.7198	11.0892	14.6343
	9	2.6917	3.6761	6.3419	7.6768	11.2919	13.1022	18.0044
	10	2.8735	4.0081	7.2810	9.0215	14.0416	16.7112	24.4113
	11	3.1484	4.5267	8.8637	11.3656	19.2382	23.7790	38.0618
	12	3.1484	4.5267	8.8637	11.3656	19.2382	23.7790	38.0618

Table H.7. Single depth detector threshold levels for F_s =15000 Hz: the higher confidence levels.
Threshold levels are taken from the F distribution with 2 and ν degrees of freedom, where the values of ν are given in Table H.3. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels					
		0.99995	**0.99999**	**0.999995**	**0.999999**	**0.9999995**	**0.9999999**
Scale number	**1**	--	--	--	--	--	--
	2	10.4952	12.3178	13.1133	14.9852	15.8022	17.7249
	3	10.7580	12.6783	13.5212	15.5158	16.3913	18.4632
	4	11.1444	13.2121	14.1268	16.3089	17.2743	19.5772
	5	11.7224	14.0179	15.0449	17.5224	18.6309	21.3050
	6	12.6075	15.2686	16.4781	19.4428	20.7904	24.0936
	7	14.0086	17.2873	18.8111	22.6315	24.4070	28.8587
	8	16.3339	20.7339	22.8435	28.3047	30.9229	37.7013
	9	20.4593	27.1072	30.4361	39.4511	43.9653	56.1901
	10	28.5060	40.3163	46.5960	64.7018	74.3199	101.9483
	11	46.2918	72.0731	86.7808	131.0959	154.1719	207.0102
	12	46.2918	72.0731	86.7808	131.0959	154.1719	207.0102

H.2. Dual depth detector threshold values:

Detector threshold levels are given in the following tables. They were computed from numerical integration of the density function of equation (6-18). Table values less than 50 were computed with a step size of 0.2 with the variables of integration increasing to a maximum of 200. Table values greater than 50 were computed with a step size of 0.5 with variables of integration increasing to a maximum of 1000.

Table H.8. Dual depth detector threshold values for F_s = 8000 Hz: the lower confidence levels.
Threshold levels are taken from the distribution that is the product of two F distributed variables with 2 and ν degrees of freedom, where the values of ν are given in Table H.2. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels						
		0.9	**0.95**	**0.99**	**0.995**	**0.999**	**0.9995**	**0.9999**
Scale number	**1**	2.4	3.8	8.4	10.8	17.6	21.0	30.4
	2	2.4	4.0	8.6	11.0	18.2	21.8	31.6
	3	2.4	4.0	8.8	11.4	18.8	22.8	33.2
	4	2.4	4.2	9.0	12.0	20.0	24.2	35.8
	5	2.6	4.4	9.6	12.8	21.8	26.4	39.8
	6	2.6	4.6	10.0	14.0	24.4	30.0	46.0
	7	2.8	5.0	12.0	16.0	29.2	36.4	59.5
	8	3.2	5.6	13.0	18.8	35.8	45.6	83.5
	9	3.6	6.2	16.0	23.6	55.5	73.5	135.5
	10	3.6	6.2	16.0	23.6	55.5	73.5	135.5
	11	3.6	6.2	16.0	23.6	55.5	73.5	135.5
	12	3.6	6.2	16.0	23.6	55.5	73.5	135.5

Table H.9. Dual depth detector threshold values for F_s = 8000 Hz: the higher confidence levels.
Threshold levels are taken from the distribution that is the product of two F distributed variables with 2 and ν degrees of freedom, where the values of ν are given in Table H.2. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels					
		0.99995	**0.99999**	**0.999995**	**0.999999**	**0.9999995**	**0.9999999**
Scale number	**1**	35.0	47.0	52.5	67.0	74.0	91.0
	2	36.4	49.0	55.0	70.5	78.0	97.0
	3	38.4	52.5	59.0	76.5	84.5	106.0
	4	41.6	57.5	65.5	85.5	95	120.5
	5	46.8	66.0	75.5	100.5	112.5	145.0
	6	56.0	80.5	92.5	126.5	143.5	190.5
	7	71.5	106.5	125.0	178.0	206.0	284.0
	8	102.5	162.0	195.0	295.0	349.5	508.5
	9	173.5	297.0	370.0	586.0	690.5	888.5
	10	173.5	297.0	370.0	586.0	690.5	888.5
	11	173.5	297.0	370.0	586.0	690.5	888.5
	12	173.5	297.0	370.0	586.0	690.5	888.5

Table H.10. Dual depth detector threshold values for F_s = 15000 Hz: the lower confidence levels.

Threshold levels are taken from the distribution that is the product of two F distributed variables with 2 and ν degrees of freedom, where the values of ν are given in Table H.3. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels						
		0.9	**0.95**	**0.99**	**0.995**	**0.999**	**0.9995**	**0.9999**
Scale number	**1**	--	--	--	--	--	--	--
	2	2.4	3.8	8.2	10.6	17.2	20.6	29.6
	3	2.4	3.8	8.4	10.8	17.6	21.2	30.4
	4	2.4	4.0	8.6	11.0	18.2	21.8	31.6
	5	2.4	4.0	8.8	11.4	19.0	23.0	33.6
	6	2.4	4.2	9.2	12.0	20.2	24.4	36.2
	7	2.6	4.4	9.8	13.0	22.2	27.0	40.6
	8	2.8	4.6	10.8	14.4	25.2	31.0	48.0
	9	3.0	5.0	12.2	16.6	30.2	37.8	62.0
	10	3.4	5.8	14.8	20.4	39.4	51.5	88.0
	11	3.8	6.8	19.2	27.4	59.0	79.0	146.5
	12	3.8	6.8	19.2	27.4	59.0	79.0	146.5

Table H.11. Dual depth detector threshold values for F_s = 15000 Hz: the upper confidence levels.

Threshold levels are taken from the distribution that is the product of two F distributed variables with 2 and ν degrees of freedom, where the values of ν are given in Table H.3. Thus, these threshold values hold for: *seglen* = 0.075 F_s, *viblen* = 0.150 F_s, Q_w = 3, *scalebase* = 1.4.

		Confidence Levels					
		0.99995	**0.99999**	**0.999995**	**0.999999**	**0.9999995**	**0.9999999**
Scale number	**1**	--	--	--	--	--	--
	2	34.0	45.4	51.0	64.5	71.5	87.5
	3	35.0	46.8	53.0	67.5	74.5	91.5
	4	36.4	49.2	55.5	71.0	79.0	98.0
	5	38.6	53.0	60.0	77.5	86.0	107.5
	6	42.2	58.5	66.5	87.0	97.0	123.0
	7	47.6	67.5	77.0	103.0	115.5	149.5
	8	57.5	83.0	95.5	131.5	149.5	199.5
	9	74.5	112.0	131.5	188.5	218.5	303.5
	10	109.0	173.5	210.0	320.5	381.0	556.5
	11	189.0	327.5	409.0	642.0	746.5	919.5
	12	189.0	327.5	409.0	642.0	746.5	919.5

VITA

Melani Plett (Shoemaker)
Department of Engineering
Seattle Pacific University
Seattle, Washington 98119

phone (206) 281-2351
email mplett@spu.edu

Education

Undergraduate: B.S. Electrical Engineering 1991
Seattle Pacific University, Seattle, Washington

Graduate: M.S. Electrical Engineering 1993
University of Washington, Seattle, Washington

Coursework: Signal processing, control theory, robotics 1991-1993
Signal processing, electromagnetics, ultrasound 1996-1997
University of Washington, Seattle, Washington

Academic and Professional Employment

Teaching:

Teaching Assistant 1991-1992
University of Washington
Courses: Electronics II and Introduction to Digital Systems and Computers

Instructor 1993 - present (leave of absence 1996-2000)
Electrical Engineering
Seattle Pacific University, Seattle, Washington
Courses taught: Signals and Systems Analysis, Electric Circuits, Computer Organization and Assembly Language

Research:

Research Assistant 1992-1993
Electrical Engineering Biorobotics Lab
University of Washington, Seattle, Washington
-Developed a mathematical model of the mammalian alpha motoneuron and Renshaw Cell to study their combined effect of recurrent

inhibition, which constitutes a basic block in mammalian muscle control.
-Designed, prototyped, and performed PCB layout and testing of an A/D, D/A I/O card to directly interface with a computer system based on Texas Instrument's DSP 320C30 processor.

Faculty Researcher Summers 1994,1995
Phillips Laboratory (PL/LIMI)
United States Air Force
Kirtland Air Force Base, New Mexico
-Participated in the Air Force Office of Scientific Research Summer Faculty Research Program. Applied speckle imaging techniques to images taken of single stars collected by ground-based telescopes with the intent of mitigating the effects of blurring induced by the atmosphere.

Research Assistant 1996-2000
Department of Surgery
University of Washington, Seattle, Washington
-Vibration detection and estimation with ultrasound and the wavelet transform.

Certification

EIT (Engineer in Training), exam passed, Washington State, 1991

Publications

MASTER'S THESIS
M. Shoemaker: A Study and Model of the Role of the Renshaw Cell in Regulating the Transient Firing Rate of the Motoneuron, M.S. Thesis, University of Washington, Seattle, Washington, 1993.

PAPERS AND REPORTS

1. M. Shoemaker, B. Hannaford "A Study and Model of the Role of the Renshaw Cell in Regulating the Transient Firing Rate of the Motoneuron." *Biological Cybernetics* 71:251-262, 1994.

2. M. Shoemaker, "Frequency Domain Analysis of Short Exposure, Photon-Limited Astronomical Data." AFOSR SFRP Final Report, 1994.

3. M. Shoemaker, "Spectral Effects on the Signal-To-Noise-Ratio of Short Exposure Photon-Limited Stellar Data." AFOSR SFRP Final Report, 1995.

4. M. Plett, "Using Doppler Ultrasound to Examine Wall Vibrations and Flow Velocity Fluctuations in Arteries." ASA/ICA conference proceedings, June 1998.

5. M. Plett, "Automated Ultrasonic Arterial Vibrometry: Detection and Measurement." SPIE International Symposium on Medical Imaging 2000 conference proceedings, February 2000.

Honors and Awards

University Scholars award, Seattle Pacific University 1987-1991
Graduation with honors (Summa Cum Laude), Seattle Pacific University 1991
Member of Eta Kappa Nu, electrical engineering honor society, University of Washington 1992
Henry L. Gray Fellowship, School of Engineering, University of Washington 1996-1997